

NBER WORKING PAPER SERIES

NUMERICAL SOLUTION OF DYNAMIC PORTFOLIO OPTIMIZATION WITH
TRANSACTION COSTS

Yongyang Cai
Kenneth L. Judd
Rong Xu

Working Paper 18709
<http://www.nber.org/papers/w18709>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2013

Cai and Judd gratefully acknowledge NSF support (SES-0951576) We thank Walter Murray, Michael Saunders, Sunil Kumar, Benjamin Van Roy and Gerd Infanger for comments on earlier versions of this work. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Yongyang Cai, Kenneth L. Judd, and Rong Xu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Numerical Solution of Dynamic Portfolio Optimization with Transaction Costs
Yongyang Cai, Kenneth L. Judd, and Rong Xu
NBER Working Paper No. 18709
January 2013
JEL No. C61,C63,G11

ABSTRACT

We apply numerical dynamic programming to multi-asset dynamic portfolio optimization problems with proportional transaction costs. Examples include problems with one safe asset plus two to six risky stocks, and seven to 360 trading periods in a finite horizon problem. These examples show that it is now tractable to solve such problems.

Yongyang Cai
Hoover Institution
Stanford University
Stanford, CA 94305
yycai@stanford.edu

Rong Xu
Department of Management Science and Engineering
Stanford University
Stanford, CA 94305, USA
rongxu06@gmail.com

Kenneth L. Judd
Hoover Institution
Stanford University
Stanford, CA 94305-6010
and NBER
kennethjudd@mac.com

1 Introduction

Dynamic management of portfolios is a critical part of any investment strategy by individuals and firms. Standard theory assumes that there is no cost to rebalancing a portfolio but transaction costs are not negligible in real markets. Not only are there transaction fees such as brokerage expenses but the presence of a bid-ask spread creates a transaction cost to a trader. These costs may be small, but the frequency of rebalancing is, theoretically, high. Therefore, any examination of real-world dynamic portfolio management needs to consider these frictions.

Multi-stage portfolio optimization problems with transaction costs assume that there are k risky assets (“stocks”) and/or a riskless asset (“bank account” paying a fixed interest rate r) traded during the period $[0, T]$. In our discrete-time analysis, portfolio adjustments are made at N fixed times in $[0, T]$, $0 = t_0 < t_1 < \dots < t_{N-1} < t_N = T$. Trades are made to maximize the investor’s expected utility over terminal wealth (T is the terminal time) and/or consumption during $[0, T]$. We will examine only the case of proportional transaction costs. If the major transaction cost is the bid-ask spread, then proportional transaction costs are the correct case to study.

Multi-stage portfolio optimization problems with transaction costs have been studied in many papers. The problem with one risky asset has been well studied; see Zabel (1973), Constantinides (1976, 1986), Gennotte and Jung (1994), and Boyle and Lin (1997). The key insight is that transaction costs create a “no-trade region” (NTR); that is, no trading is done if the current portfolio is inside the no-trade region, and otherwise the investor trades to some point on the boundary of the no-trade region. Kamin (1975) considered the case with only two risky assets. Constantinides (1979) and Abrams and Karmarkar (1980) established some properties of the NTR for multiple assets, but present only numerical examples with one safe and one risky asset. Brown and Smith (2011) evaluated some heuristic strategies and their bounds based on simulation, but their method cannot give the optimal portfolios.

In the continuous-time version, there are many papers about the portfolio

optimization problem with transaction costs with one or two risky assets; see Davis and Norman (1990), Duffie and Sun (1990), Akian et al. (1996), Janecek and Shreve (2004), Liu (2004), and Goodman and Ostrov (2010). Muthuraman and Kumar (2006, 2008) gave numerical examples at most three risky assets. Muthuraman and Zha (2008) provided a computational scheme that combines simulation with the boundary update procedure, and presented some computational results with $k \geq 3$. However, the presence of simulation implies that the boundary of the NTR is not precisely computed.

To the best of our knowledge, when the number of correlated risky assets is bigger than three and $T \geq 6$, our DP method is the first one to explicitly give good numerical solutions with transaction costs and general utility functions.

2 Portfolio Models

Assume that there are k risky assets ("stocks") and one riskless asset ("bond") available for investment. The investor's objective is to maximize the expected utility at the terminal time T . In $[0, T]$, there are $N + 1$ stages, $0 = t_0 < t_1 < \dots < t_{N-1} < t_N = T$. At each stage t_i (for $i = 0, \dots, N - 1$), the investor has an opportunity to reallocate the portfolio, which will incur transaction costs. For simplicity, we assume these stages are equally separated with a unit length of time, and then we use $t = i$ as the i -th stage in this paper.

Let $R = (R_1, \dots, R_k)^\top$ be the random one-period return vector of the stocks, and R_f be the return of the bond. The portfolio fraction for asset i at the beginning of period t right before reallocation is denoted x_{ti} , and let $x_t = (x_{t,1}, \dots, x_{t,k})^\top$. Let W_t be the total wealth at the beginning of period t right before reallocation. The difference between the total wealth and the wealth invested in stocks is invested in the bond. Let $\delta_{t,i}W_t$ denote the amount of dollars for buying or selling part of the i -th stock at stage t , expressed as a fraction of wealth, while $\delta_{t,i} > 0$ means buying, and $\delta_{t,i} < 0$ means selling. We assume that $f(\delta_{t,i}W_t) = \tau|\delta_{t,i}W_t|$ with a constant $\tau > 0$ is the transaction cost function for buying or selling part of the i -th stock

using $\delta_{t,i}W_t$ dollars. Let \mathbf{e} denote the column vector of all elements equal to 1. Then the multi-stage portfolio optimization problem can be expressed as

$$\begin{aligned}
V_0(W_0, x_0) &= \max_{\delta_t} \mathbb{E}\{u(W_T)\} \\
\text{s.t. } W_{t+1} &= \mathbf{e}^\top X_{t+1} + R_f(1 - \mathbf{e}^\top x_t - y_t)W_t, \\
X_{t+1,i} &= R_i(x_{t,i} + \delta_{t,i})W_t, \\
y_t &= \mathbf{e}^\top (\delta_t + \tau|\delta_t|), \\
x_{t+1,i} &= X_{t+1,i}/W_{t+1}, \\
t &= 0, \dots, T-1; \quad i = 1, \dots, k,
\end{aligned}$$

where $\mathbb{E}\{\cdot\}$ is the expectation operator, $X_{t+1} = (X_{t+1,1}, \dots, X_{t+1,k})^\top$ is the vector of the amount of dollars invested in the risky assets at stage $t+1$, and $\delta_t = (\delta_{t,1}, \dots, \delta_{t,k})^\top$. Note that all of X_{t+1} , x_{t+1} , W_{t+1} , y_{t+1} and δ_{t+1} are random as R is random. So when R has a continuous distribution, R must be discretized or simulated so that the model is solvable. But the exponential growth of scenarios over the number of periods T allows us to solve this model directly only when T is very limited, typically $T \leq 4$. This limitation makes it necessary to change this model into a DP model discussed in Section 3.

To be more general, we can assume that R_f and the multivariate distribution of R are dependent on a vector discrete time stochastic process θ_t , denoted by $R_f(\theta_t)$ and $R(\theta_t)$ respectively, for $t = 0, \dots, T-1$. Then the above model becomes

$$\begin{aligned}
V_0(W_0, x_0, \theta_0) &= \max_{\delta_t} \mathbb{E}\{u(W_T)\} \\
\text{s.t. } W_{t+1} &= \mathbf{e}^\top X_{t+1} + R_f(\theta_t)(1 - \mathbf{e}^\top x_t - y_t)W_t, \\
X_{t+1,i} &= R_i(\theta_t)(x_{t,i} + \delta_{t,i})W_t, \\
y_t &= \mathbf{e}^\top (\delta_t + \tau|\delta_t|), \\
x_{t+1,i} &= X_{t+1,i}/W_{t+1}, \\
t &= 0, \dots, T-1; \quad i = 1, \dots, k.
\end{aligned}$$

When there is a consumption decision at each stage, a dynamic portfolio optimization problem is to find an optimal portfolio and a consumption decision C_t at each stage t such that we have a maximal expected total utility, i.e.,

$$\max \mathbb{E} \left\{ \beta^T u(W_T) + \sum_{t=0}^{T-1} \beta^t u(C_t) \right\},$$

where u is the given utility function, β is the discount factor, $W_{t+1} = \mathbf{e}^\top X_{t+1} + R_f(1 - \mathbf{e}^\top x_t - y_t)W_t - C_t$, and X_t , x_t and y_t have the same definitions in the equality constraints of the first model.

All these models have their corresponding DP models in the next section to avoid the exponential growth of scenarios over the number of periods T , so that these problems can be solved for a few periods.

3 DP Models for Portfolio Problems with Transaction Costs

For dynamic portfolio problems with transaction costs, we can choose the state variables as the total wealth W_t and allocation fractions $x_t = (x_{t,1}, \dots, x_{t,k})^\top$ invested in the risky assets. Here W_t and x_t are the values right before reallocation at time t . Thus, the DP model becomes

$$V_t(W_t, x_t) = \max_{\delta_t} \mathbb{E} \{V_{t+1}(W_{t+1}, x_{t+1})\},$$

where

$$\begin{aligned} y_t &\equiv \mathbf{e}^\top (\delta_t + \tau |\delta_t|), \\ X_{t+1,i} &\equiv R_i(x_{t,i} + \delta_{t,i})W_t, \\ W_{t+1} &\equiv \mathbf{e}^\top X_{t+1} + R_f(1 - \mathbf{e}^\top x_t - y_t)W_t, \\ x_{t+1,i} &\equiv X_{t+1,i}/W_{t+1}, \end{aligned} \tag{1}$$

for $i = 1, \dots, k$. The terminal value function is $V_T(W, x) = u(W)$ for some given utility function u . Sometimes, the terminal value function is chosen as

$V_T(W, x) = u((1 - \tau \mathbf{e}^\top x)W)$, if we assume that all risky assets have to be converted into the riskless asset before consumption. Later, we just assume that $V_T(W, x) = u(W)$ for simplicity.

In the model, the constraint (1) is nondifferentiable because of the existence of the absolute operator. We can change the nondifferentiable constraint into a linear constraint. Let $\delta_t = \delta_t^+ - \delta_t^-$ with $\delta_t^+, \delta_t^- \geq 0$, we have $|\delta_t| = \delta_t^+ + \delta_t^-$. Here $\delta_t^+ = (\delta_{t,1}^+, \dots, \delta_{t,k}^+)^\top$ is the vector of fractions of wealth W_t for buying stocks, and $\delta_t^- = (\delta_{t,1}^-, \dots, \delta_{t,k}^-)^\top$ is the vector of fractions of wealth W_t for selling stocks. Thus, by using δ_t^+ and δ_t^- instead of δ_t in the models, the constraint (1) is changed into a linear constraint:

$$y_t \equiv \mathbf{e}^\top ((\delta_t^+ - \delta_t^-) + \tau(\delta_t^+ + \delta_t^-)).$$

We will still use the absolute value notation in the rest of the paper to simplify the presentation.

If we do not allow shorting stocks or borrowing cash, then we just need to add the constraints $x_t + \delta_t \geq 0$ and $1 - \mathbf{e}^\top x_t \geq y_t$. And the range of x_t and x_{t+1} is $[0, 1]^k$ while $\mathbf{e}^\top x_{t+1} \leq 1$.

3.1 Portfolio with a CRRA Utility

In economics and finance, we usually assume a CRRA utility function, i.e., $u(W) = W^{1-\gamma}/(1-\gamma)$ for some constant $\gamma > 0$ and $\gamma \neq 1$, or $u(W) = \log(W)$ for $\gamma = 1$. Thus, for $u(W) = W^{1-\gamma}/(1-\gamma)$, if we assume that $V_{t+1}(W_{t+1}, x_{t+1}) = W_{t+1}^{1-\gamma} \cdot g_{t+1}(x_{t+1})$, then

$$V_t(W_t, x_t) = \max_{\delta_t} \mathbb{E} \left\{ W_{t+1}^{1-\gamma} \cdot g_{t+1}(x_{t+1}) \right\},$$

where

$$\begin{aligned}
s_{t+1,i} &\equiv R_i(x_{t,i} + \delta_{t,i}), \\
y_t &\equiv \mathbf{e}^\top (\delta_t + \tau|\delta_t|), \\
\Pi_{t+1} &\equiv \mathbf{e}^\top s_{t+1} + R_f(1 - \mathbf{e}^\top x_t - y_t), \\
W_{t+1} &\equiv \Pi_{t+1}W_t, \\
x_{t+1,i} &\equiv s_{t+1,i}/\Pi_{t+1},
\end{aligned}$$

for $i = 1, \dots, k$. By changing W_{t+1} to $\Pi_{t+1}W_t$ in the objective function, we get $V_t(W_t, x_t) = W_t^{1-\gamma} \cdot g_t(x_t)$, where

$$g_t(x_t) = \max_{\delta_t} \mathbb{E} \left\{ \Pi_{t+1}^{1-\gamma} \cdot g_{t+1}(x_{t+1}) \mid x_t \right\}. \quad (2)$$

Therefore, by induction, from $V_T(W, x) = u(W) = W^{1-\gamma} \cdot 1/(1-\gamma)$, we showed that

$$V_t(W_t, x_t) = W_t^{1-\gamma} \cdot g_t(x_t)$$

for any time $t = 0, 1, \dots, T$, while $g_t(x)$ has the iterative formula given in the above optimization problem and $g_T(x) = 1/(1-\gamma)$ (or $g_T(x) = (1 - \tau\mathbf{e}^\top|x|)^{1-\gamma}/(1-\gamma)$ when we assume that all risky assets have to be converted into the riskless asset before consumption), if we assume a proportional transaction cost and a power utility $u(W) = W^{1-\gamma}/(1-\gamma)$ with a constant relative risk aversion coefficient $\gamma > 0$ and $\gamma \neq 1$.

For $u(W) = \log(W)$, we can also show by induction that

$$V_t(W_t, x_t) = \log(W_t) + g_t(x_t),$$

where

$$g_t(x_t) = \max_{\delta_t} \mathbb{E} \{ \log(\Pi_{t+1}) + g_{t+1}(x_{t+1}) \mid x_t \}, \quad (3)$$

while $g_T(x) = 0$ (or $g_T(x) = \log(1 - \tau\mathbf{e}^\top|x|)$, when we assume that all risky assets have to be converted into the riskless asset before consumption).

Since W_t and x_t are separable for CRRA utilities, we could just do a backward recursion on the functions $g_t(x)$ instead of $V_t(W, x)$.

When there is no riskless asset, we just need to cancel the R_f term and replace m_t with 0 in the above models for g_t or ψ_t , while we should have $\mathbf{e}^\top x_t = 1$. (The state variable vector $x_t = (x_{t,1}, \dots, x_{t,k})$ should be changed to $(x_{t,1}, \dots, x_{t,k-1})$, and there is the same cutoff in x_{t+1}).

If we do not allow shorting stocks or borrowing cash, then the range of x_t is $[0, 1]^k$, and in the models we just need to add the constraints $x_t + \delta_t \geq 0$ and $y_t \leq 1 - \mathbf{e}^\top x_t$, so that $x_{t+1} \in [0, 1]^n$ and $\mathbf{e}^\top x_{t+1} \leq 1$. And we still have the property of separation of W and x in the value functions $V(W, x)$.

We know that there is a NTR, Ω_t , for any $t = 0, 1, \dots, T - 1$. When $x_t \in \Omega_t$, the investor will not trade at all, and when $x_t \notin \Omega_t$, the investor will trade to some point on the boundary of Ω_t . That is, Ω_t is defined as

$$\Omega_t = \{x_t : \delta_t^* = 0\},$$

where δ_t^* is the optimal control for the given x_t . See Kamin (1975), Constantinides (1976, 1979, 1986), Davis and Norman (1990), Muthuraman and Kumar (2006), and so on.

Abrams and Karmarkar (1980) showed that NTR is a connected set and that it is a cone when the utility function is assumed to be positively homogeneous (a function $u(x)$ is positively homogeneous if there exists a positive value function $\psi(x)$ such that $u(ax) = \psi(a)u(x)$ for any $a > 0$). Moreover, in the case of proportional transaction costs and concave utility functions, NTR can take on many forms ranging from a simple half-line to a nonconvex set. So we should use numerical methods to compute NTR.

From the separability of W and x , we see that the optimal portfolio rules are independent of wealth W_t . Thus the “no-trade” regions Ω_t are also independent of W_t , for the CRRA utility functions.

3.2 Serially Correlated Asset Returns with Stochastic Parameters

In the previous model, we assume that the interest rate is fixed and the risky asset returns have the same multivariate distribution throughout time, and they are not serially correlated. But in real-life models, the riskless return

R_f and the multivariate distribution of risky asset returns are stochastic and serial-correlated. Assume that they are dependent on some stochastic parameters. Let all these parameters be denoted as a vector θ_t at time t . They could be discrete Markov chains with a given transition probability matrix from the previous stage to the current stage, or continuously distributed, conditional on their previous-stage values. Let us denote $R_f(\theta_t)$ for the bond return and $R(\theta_t)$ for the stocks return vector respectively, for $t = 0, \dots, T - 1$.

With wealth W_t , allocation fractions x_t and parameters θ_t as the state variables, the DP model becomes

$$V_t(W_t, x_t, \theta_t) = \max_{\delta_t} \mathbb{E} \{ V_{t+1}(W_{t+1}, x_{t+1}, \theta_{t+1}) \mid W_t, x_t, \theta_t \},$$

where the terminal value function is $V_T(W, x, \theta) = u(W)$, and

$$\begin{aligned} y_t &\equiv \mathbf{e}^\top (\delta_t + \tau |\delta_t|), \\ X_{t+1,i} &\equiv R_i(\theta_t)(x_{t,i} + \delta_{t,i})W_t, \\ W_{t+1} &\equiv \mathbf{e}^\top X_{t+1} + R_f(\theta_t)(1 - \mathbf{e}^\top x_t - y_t)W_t, \\ x_{t+1,i} &\equiv X_{t+1,i}/W_{t+1}. \end{aligned}$$

Like what we discussed in Section 3.1, when the utility function is $u(W) = W^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$ and $\gamma \neq 1$, we have $V_t(W_t, x_t, \theta_t) = W_t^{1-\gamma} \cdot g_t(x_t, \theta_t)$ where

$$g_t(x_t, \theta_t) = \max_{\delta_t} \mathbb{E} \left\{ \Pi_{t+1}^{1-\gamma} \cdot g_{t+1}(x_{t+1}, \theta_{t+1}) \mid x_t, \theta_t \right\}, \quad (4)$$

where

$$\begin{aligned} s_{t+1,i} &\equiv R_i(\theta_t)(x_{t,i} + \delta_{t,i}), \\ \Pi_{t+1} &\equiv \mathbf{e}^\top s_{t+1} + R_f(\theta_t)(1 - \mathbf{e}^\top x_t - y_t), \end{aligned}$$

with $g_T(x, \theta) = 1/(1-\gamma)$. Similarly, we also have the separability of W and (x, θ) when $u(W) = \log(W)$. Moreover, the separability still holds if we add "no-shorting" and/or "no-borrowing" constraints.

From the separability of W and (x, θ) , we see that the optimal portfolio rules are independent of wealth W_t . Thus, the NTR is also independent of W_t for CRRA utility functions, but will depend on θ_t .

3.3 Portfolio with Transaction Costs and Consumption

If we allow assets to finance consumption during the investment period, then the DP model becomes

$$V_t(W_t, x_t) = \max_{C_t, \delta_t} u(C_t)h + \beta \mathbb{E} \{V_{t+1}(W_{t+1}, x_{t+1})\},$$

where h is the length of one period,

$$\begin{aligned} y_t &\equiv \mathbf{e}^\top (\delta_t + \tau |\delta_t|), \\ X_{t+1,i} &\equiv R_i(x_{t,i} + \delta_{t,i})W_t, \\ W_{t+1} &\equiv \mathbf{e}^\top X_{t+1} + R_f(W_t(1 - \mathbf{e}^\top x_t - y_t) - C_t h), \\ x_{t+1,i} &\equiv X_{t+1,i}/W_{t+1}, \end{aligned}$$

and the terminal value function $V_T(W, x)$ is given.

Let $C_t = c_t W_t$. Like what we discussed in Section 3.1, when the utility function is $u(C) = C^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$ and $\gamma \neq 1$, and the terminal value function is $V_T(W, x) = W^{1-\gamma} \cdot g_T(x)$ for some given $g_T(x)$, we have

$$V_t(W_t, x_t) = W_t^{1-\gamma} \cdot g_t(x_t),$$

where

$$g_t(x_t) = \max_{c_t, \delta_t} u(c_t)h + \beta \mathbb{E} \left\{ \Pi_{t+1}^{1-\gamma} \cdot g_{t+1}(x_{t+1}) \mid x_t, c_t \right\}. \quad (5)$$

Here,

$$\begin{aligned} s_{t+1,i} &\equiv R_i(x_{t,i} + \delta_{t,i}), \\ \Pi_{t+1} &\equiv \mathbf{e}^\top s_{t+1} + R_f(1 - \mathbf{e}^\top x_t - y_t - c_t h). \end{aligned}$$

Similarly, we can also have the separability of W and x when $u(C) = \log(C)$

and $V_T(W, x) = \log(W) + g_T(x_t)$. Moreover, the separability still holds if we add "no-shorting" and/or "no-borrowing" constraints. If R and R_f are dependent on a Markov chain parameter vector θ_t , then we still have the separability of W and (x, θ) .

From the separability of W and x , we see that the optimal portfolio rules are independent of wealth W_t . Thus NTR, Ω_t , is also independent of W_t , for the CRRA utility functions. Here Ω_t is defined as

$$\Omega_t = \{x_t / (1 - c_t^* h) : \delta_t^* = 0\},$$

where c_t^* and δ_t^* are the optimal controls for the given x_t .

4 Numerical DP Algorithms

If state and control variables in a DP problem are continuous, then the value function must be approximated in some computationally tractable manner. It is common to approximate value functions with a finitely parameterized collection of functions; that is, we use some functional form $\hat{V}(x; \mathbf{b})$, where \mathbf{b} is a vector of parameters, and approximate a value function, $V(x)$, with $\hat{V}(x; \mathbf{b})$ for some parameter value \mathbf{b} . For example, \hat{V} could be a linear combination of polynomials where \mathbf{b} would be the weights on polynomials. After the functional form is fixed, we focus on finding the vector of parameters, \mathbf{b} , such that $\hat{V}(x; \mathbf{b})$ approximately satisfies the Bellman equation (Bellman, 1957).

Numerical solutions to a DP problem are based on the Bellman equation:

$$\begin{aligned} V_t(x, \theta) &= \max_{a \in \mathcal{D}(x, \theta, t)} u_t(x, a) + \beta \mathbb{E} \{ V_{t+1}(x^+, \theta^+) \mid x, \theta, a \}, & (6) \\ \text{s.t. } x^+ &= \mathcal{G}_t(x, \theta, a, \omega), \\ \theta^+ &= \mathcal{H}_t(\theta, a, \epsilon), \end{aligned}$$

where x is the continuous state, θ is the discrete state, $V_t(x, \theta)$ is called the value function at time $t \leq T$ (the terminal value function $V_T(x, \theta)$ is given), a is the action variable vector in its feasible set $\mathcal{D}(x, \theta, t)$, x^+ is the next-

stage continuous state with its transition function \mathcal{G}_t at time t , θ^+ is the next-stage discrete state with its transition function \mathcal{H}_t at time t , ω and ϵ are two random variables, and $u_t(x, a)$ is the utility function at time t , β is the discount factor, and $\mathbb{E}\{\cdot\}$ is the expectation operator.

The following is the algorithm of parametric DP with value function iteration for finite horizon problems.

Algorithm 1. *Numerical Dynamic Programming with Value Function Iteration for Finite Horizon Problems*

Initialization. *Choose the approximation nodes, $\mathbb{X}_t = \{x_t^i : 1 \leq i \leq m_t\} \subset \mathbb{R}^d$, for every $t < T$, and choose a functional form for $\hat{V}(x, \theta; \mathbf{b})$. Let $\hat{V}(x, \theta; \mathbf{b}^T) \equiv V_T(x, \theta)$. Then for $t = T-1, T-2, \dots, 0$, iterate through steps 1 and 2.*

Step 1. *Maximization Step. Compute*

$$\begin{aligned} v_{i,j} &= \max_{a \in \mathcal{D}(x_i, \theta_j, t)} u_t(x_i, a) + \beta \mathbb{E} \left\{ \hat{V}(x^+, \theta^+; \mathbf{b}^{t+1}) \right\} \\ \text{s.t.} \quad x^+ &= \mathcal{G}_t(x_i, \theta_j, a, \omega), \\ \theta^+ &= \mathcal{H}_t(\theta_j, a, \epsilon), \end{aligned}$$

for each $\theta_j \in \Theta$, $x_i \in \mathbb{X}_t$, $1 \leq i \leq m_t$.

Step 2. *Fitting Step. Using an appropriate approximation method, compute the \mathbf{b}^t such that $\hat{V}(x, \theta_j; \mathbf{b}^t)$ approximates $(x_i, v_{i,j})$ data for each $\theta_j \in \Theta$.*

Algorithm 1 shows that there are three main components in value function iteration for deterministic DP problems: optimization, approximation, and integration. More detailed discussion of numerical DP can be found in Cai (2009), Cai and Judd (2010), Judd (1998) and Rust (2008).

4.1 Approximation

An approximation scheme consists of two parts: basis functions and approximation nodes. Approximation nodes can be chosen as uniformly spaced

nodes, Chebyshev nodes, or some other specified nodes. From the viewpoint of basis functions, approximation methods can be classified as either spectral methods or finite element methods. A spectral method uses globally nonzero basis functions $\phi_j(x)$ such that $\hat{V}(x; \mathbf{b}) = \sum_{j=0}^n b_j \phi_j(x)$ is a degree- n approximation. Examples of spectral methods include ordinary polynomial approximation, Chebyshev polynomial approximation, and shape-preserving Chebyshev polynomial approximation (Cai and Judd, 2012a), and Hermite approximation (Cai and Judd, 2012c). In contrast, a finite element method uses locally basis functions $\phi_j(x)$ that are nonzero over sub-domains of the approximation domain. Examples of finite element methods include piecewise linear interpolation, cubic splines, B-splines, and shape-preserving rational splines (Cai and Judd, 2012b). See Cai (2009), Cai and Judd (2010), and Judd (1998) for more details.

We prefer Chebyshev polynomials when the value function is smooth. Chebyshev polynomials on $[-1, 1]$ are defined as $\mathcal{T}_j(x) = \cos(j \cos^{-1}(x))$, while general Chebyshev polynomials on $[x_{\min}, x_{\max}]$ are defined as $\mathcal{T}_j((2x - x_{\min} - x_{\max})/(x_{\max} - x_{\min}))$ for $j = 0, 1, 2, \dots$. In a d -dimensional approximation problem, let the domain of the value function be

$$\{x = (x_1, \dots, x_d) : x_j^{\min} \leq x_j \leq x_j^{\max}, j = 1, \dots, d\},$$

for some real numbers x_j^{\min} and x_j^{\max} with $x_j^{\max} > x_j^{\min}$ for $j = 1, \dots, d$. Let $x^{\min} = (x_1^{\min}, \dots, x_d^{\min})$ and $x^{\max} = (x_1^{\max}, \dots, x_d^{\max})$. Then we denote $[x^{\min}, x^{\max}]$ as the domain. Let $\alpha = (\alpha_1, \dots, \alpha_d)$ be a vector of nonnegative integers. Let $\mathcal{T}_\alpha(z)$ denote the product $\mathcal{T}_{\alpha_1}(z_1) \cdots \mathcal{T}_{\alpha_d}(z_d)$ for $z = (z_1, \dots, z_d) \in [-1, 1]^d$. Let

$$Z(x) = \left(\frac{2x_1 - x_1^{\min} - x_1^{\max}}{x_1^{\max} - x_1^{\min}}, \dots, \frac{2x_d - x_d^{\min} - x_d^{\max}}{x_d^{\max} - x_d^{\min}} \right)$$

for any $x = (x_1, \dots, x_d) \in [x^{\min}, x^{\max}]$.

Using these notations, the degree- n complete Chebyshev approximation

for $V(x)$ is

$$\hat{V}_n(x; \mathbf{b}) = \sum_{0 \leq |\alpha| \leq n} b_\alpha \mathcal{T}_\alpha(Z(x)), \quad (7)$$

where $|\alpha| = \sum_{j=1}^d \alpha_j$ for the nonnegative integer vector $\alpha = (\alpha_1, \dots, \alpha_d)$. So the number of terms with $0 \leq |\alpha| = \sum_{j=1}^d \alpha_j \leq n$ is $\binom{n+d}{d}$ for the degree- n complete Chebyshev approximation in \mathbb{R}^d .

4.2 Numerical Integration

In the objective function of the Bellman equation, we often need to compute the conditional expectation of $V(x^+ | x, a)$. When the random variable is continuous, we have to use numerical integration to compute the expectation.

One naive way is to apply Monte Carlo or pseudo Monte Carlo methods to compute the expectation. By the central limit theorem in statistics, the numerical error of the integration computed by (pseudo) Monte Carlo methods has a distribution that is close to normal. So there is no bound for the numerical error occurred by (pseudo) Monte Carlo methods. Moreover, the optimization problem often needs hundreds or thousands of evaluations of the objective function. This implies that once one evaluation of the objective function has a big numerical error, the previous iterations to solve the optimization problem may make no sense. Therefore, the iterations may never converge to the optimal solution. Thus it is not practical to apply (pseudo) Monte Carlo methods to the optimization problem generally, unless the stopping criterion of the optimization problem is set very loosely.

Therefore, it will be good to have a numerical integration method with a bounded numerical error. Here we present a common numerical integration method when the random variable is normal.

In the expectation operator of the objective function of the Bellman equation, if the random variable has a normal distribution, then it will be good to apply the Gauss-Hermite quadrature formula to compute the numerical integration. That is, if we want to compute $\mathbb{E}\{f(Y)\}$ where Y has a distribution

$\mathcal{N}(\mu, \sigma^2)$, then

$$\begin{aligned}\mathbb{E}\{f(Y)\} &= (2\pi\sigma^2)^{-1/2} \int_{-\infty}^{\infty} f(y)e^{-(y-\mu)^2/(2\sigma^2)} dy \\ &= (2\pi\sigma^2)^{-1/2} \int_{-\infty}^{\infty} f(\sqrt{2}\sigma x + \mu)e^{-x^2} \sqrt{2}\sigma dx \\ &\doteq \pi^{-\frac{1}{2}} \sum_{i=1}^m \omega_i f(\sqrt{2}\sigma x_i + \mu),\end{aligned}$$

where ω_i and x_i are the Gauss-Hermite quadrature weights and nodes over $(-\infty, \infty)$. See Stroud and Secrest (1966).

If Y is log normal, i.e., $\log(Y)$ has a distribution $\mathcal{N}(\mu, \sigma^2)$, then we can assume that $Y = e^X$ where $X \sim \mathcal{N}(\mu, \sigma^2)$, thus

$$\mathbb{E}\{f(Y)\} = \mathbb{E}\{f(e^X)\} \doteq \pi^{-\frac{1}{2}} \sum_{i=1}^m \omega_i f\left(e^{\sqrt{2}\sigma x_i + \mu}\right).$$

If we want to compute a multidimensional integration, we could apply the product rule. For example, suppose that we want to compute $\mathbb{E}\{f(X)\}$, where X is a random vector with multivariate normal distribution $\mathcal{N}(\mu, \Sigma)$ over $(-\infty, +\infty)^d$, where μ is the mean column vector and Σ is the covariance matrix, then we could do the Cholesky factorization first, i.e., find a lower triangular matrix L such that $\Sigma = LL^\top$. This is feasible as Σ must be a positive semi-definite matrix from the covariance property. Thus,

$$\begin{aligned}\mathbb{E}\{f(X)\} &= \left((2\pi)^d \det(\Sigma)\right)^{-1/2} \int_{R^d} f(y)e^{-(y-\mu)^\top \Sigma^{-1} (y-\mu)/2} dy \\ &= \left((2\pi)^d \det(L)^2\right)^{-1/2} \int_{R^d} f\left(\sqrt{2}Lx + \mu\right) e^{-x^\top x} 2^{d/2} \det(L) dx \\ &\doteq \pi^{-\frac{d}{2}} \sum_{i_1=1}^m \cdots \sum_{i_d=1}^m \omega_{i_1} \cdots \omega_{i_d} f\left(\sqrt{2}L_{1,1}x_{i_1} + \mu_1, \right. \\ &\quad \left. \sqrt{2}(L_{2,1}x_{i_1} + L_{2,2}x_{i_2}) + \mu_2, \cdots, \sqrt{2}\left(\sum_{j=1}^d L_{d,j}x_{i_j}\right) + \mu_d\right),\end{aligned}$$

where ω_i and x_i are the Gauss-Hermite quadrature weights and nodes over

$(-\infty, \infty)$, $L_{i,j}$ is the (i, j) -element of L , and $\det(\cdot)$ means the matrix determinant operator.

4.3 Parallelization

Parallelization allows researchers to solve huge problems and is the foundation of modern scientific computation. Our work shows that parallelization can also be used effectively in solving the dynamic portfolio optimization problems using value function iteration. The key fact is that at each maximization step, there are many independent optimization problems, one for each (x_i, θ_j) . In our portfolio problems there are often thousands of such independent problems, and future problems will easily have millions of independent problems. See Cai et al. (2013) for more detailed discussion. In this paper, we will use the parallel DP algorithms to solve a dynamic portfolio problem with seven assets.

5 Numerical DP in Portfolio Problems

All the optimization problems of Section 3 have differential objective functions and constraints after replacing δ_t with $\delta_t^+ - \delta_t^-$, and replacing $|\delta_t|$ by $\delta_t^+ + \delta_t^-$, where $\delta_t^+, \delta_t^- \geq 0$. Therefore, we can apply some fast Newton-type solvers, such as NPSOL optimization package (Gill et al., 1994), to solve the maximization problems in numerical DP algorithms (Algorithm 1). The maximization problems (2) and 4 are formulated in terms of $2k$ control variables (δ_t^+ and δ_t^-), and $2k$ bound constraints ($\delta_t^+, \delta_t^- \geq 0$), where k is the number of risky assets, and other unknowns are expressed in terms of these controls. The problem (5) has one more control at each time, c_t , and one more bound constraint, $c_t \geq 0$. If we do not allow shorting stocks or borrowing cash, then there are $(k + 1)$ more constraints, $x_t + \delta_t^+ - \delta_t^- \geq 0$ and $y_t \leq 1 - \mathbf{e}^\top x_t$ (or $y_t \leq 1 - \mathbf{e}^\top x_t - c_t$ when there is a consumption decision), where $y_t \equiv \mathbf{e}^\top ((\delta_t^+ - \delta_t^-) + \tau(\delta_t^+ + \delta_t^-))$.

If the return for a risky asset i , R_i , is unbounded (e.g., log-normal), then we have $\mathbb{P}(\Pi_t \leq 0) > 0$ when $x_{t,i} + \delta_{t,i} < 0$. It follows that for the

CRRA utility functions, the optimal solution must have $x_{t,i} + \delta_{t,i}^* \geq 0$ unless the asset i can be replicated by other assets, which is a degenerate case. Since the expectation of next-time value function is computed numerically by numerical DP, we must add the constraint $x_{t,i} + \delta_{t,i}^+ - \delta_{t,i}^- \geq 0$ to avoid an unreasonable approximation solution. If returns are unbounded, then for CRRA utility functions, we must have both “no-shorting” and “no-borrowing” constraints: $x_t + \delta_t^+ - \delta_t^- \geq 0$ and $y_t \leq 1 - \mathbf{e}^\top x_t$ (or $y_t \leq 1 - \mathbf{e}^\top x_t - c_t$ when there is a consumption decision). Thus, we can set the domain of x_t as $[0, 1]^k$ for each stage t .

We assume R is log-normal and with correlation, therefore, $\log(R) = (\log(R_1), \dots, \log(R_k))^\top \in \mathbb{R}^k$, has the multivariate normal distribution $\mathcal{N}((\mu - \frac{\sigma^2}{2})h, (\Lambda\Sigma\Lambda)h)$, where h is the length of a period, $\mu = (\mu_1, \dots, \mu_k)^\top$ is the drift, $\sigma = (\sigma_1, \dots, \sigma_k)^\top$ is the volatility, and Σ is the correlation matrix of the log-returns, and $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_k)$. Let the interest rate be r be constant with gross return per period equaling $R_f = \exp(rh)$. We can express the returns in terms of the Cholesky factorization $\Sigma = LL^\top$, where $L = (L_{i,j})_{k \times k}$ is a lower triangular matrix, implying

$$\log(R_i) = (\mu_i - \frac{\sigma_i^2}{2})h + \sigma_i\sqrt{h} \sum_{j=1}^i L_{i,j}z_j,$$

where z_i are independent standard normal random variables, for $i = 1, \dots, k$. Therefore, for the optimization problems (2), (4) and (5), we apply a product Gauss-Hermite quadrature to estimate the conditional expectation of $\Pi_{t+1}^{1-\gamma} \cdot g_{t+1}(x_{t+1})$ while x_t is given.

6 Numerical Examples

In this section, we give several numerical examples for solving the multi-stage portfolio optimization problems with proportional transaction costs and a power utility function $u(W) = W^{1-\gamma}/(1-\gamma)$. In these examples, the length of one period is h , $R_f = \exp(rh)$ for an interest rate r , and R is

always assumed to be log-normal with

$$\log(R) \sim \mathcal{N}\left(\left(\mu - \frac{\sigma^2}{2}\right)h, (\Lambda\Sigma\Lambda)h\right)$$

in \mathbb{R}^k , where $\mu = (\mu_1, \dots, \mu_k)^\top$ is the drift, $\sigma = (\sigma_1, \dots, \sigma_k)^\top$ is the volatility, and Σ is the correlation matrix of the log-returns, and $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_k)$.

In these examples, the number of stages is from 7 to 361, and the number of assets is from 3 to 7. By using the numerical DP method, we computed the “no-trade” regions for each stage in these examples. We applied the NPSOL optimization package (Gill et al., 1994), and the complete Chebyshev approximation method, and the multi-dimensional product Gauss-Hermite quadrature rule in the numerical DP method for these examples.

6.1 Three-Asset Portfolio Problems without Consumption

We give two examples to illustrate our numerical DP method for solving the model (2) and (4) respectively, the multi-stage portfolio optimization problems without consumption. In these examples, the assets available for trading include one bond with a constant interest rate r and $k = 2$ stocks with independent log-normal annual returns. We assume that the utility function at the terminal time is $u(W) = W^{1-\gamma}/(1-\gamma)$ with $\gamma = 3$. In the numerical DP method for these examples, we use the degree-10 complete Chebyshev approximation method with 11^k tensor Chebyshev nodes for k stocks, and the multi-dimensional product Gauss-Hermite quadrature rule with 9 nodes in each dimension.

6.1.1 Example 1

The first example is a simple three asset portfolio problem of the model (2) with quarterly time periods and 10-year horizon, i.e., $h = 0.25$ year and $T = 40$ periods. We let $r = 0.03$, $\mu = (0.07, 0.07)^\top$ and $\sigma = (0.2, 0.2)^\top$. Figure 1 shows NTR’s with $\tau = 0.5\%$ at stages $t = 0, 37, 38, 39$, and Figure 2 shows NTR’s at stage $t = 0$ with various $\tau = 0.1\%, 0.2\%, 0.5\%, 1\%, 2\%, 4\%$.

Figure 1: No trade regions of Example 1 at various stages





Figure 2: No trade regions of Example 1 at various transaction costs



The circle point located inside the NTR's is the Merton point, that is, the portfolio that would be chosen if there were no transaction costs and the horizon were infinite (Merton 1969, 1971). These NTR's are symmetric w.r.t. the 45 degree line, which is expected because the two stocks are i.i.d. Also, each NTR is close to being square.

Figure 1 shows that the NTR's depend on the trading horizon, but the sensitivity is small. Only the $T = 39$ case is significantly different than the others, due to the fact that the $T = 39$ is the last period before liquidation. The NTR's are otherwise nearly the same for all stages less than 37. This indicates that the trading strategies for most periods are the same as an infinite-horizon problem.

Figure 2 displays the sensitivity of the NTR to transaction costs τ . Note

that a very small transaction cost, $\tau = 0.1\%$, creates a nontrivial NTR, but that the NTR's grow more slowly in size than τ grows as we increase τ . The width of the NTR's is expanding in a rate of about $\tau^{1/3}$. For example, the width of the NTR of $\tau = 4\%$ is a nearly $(0.04/0.001)^{1/3} \approx 3.4$ times of the width of $\tau = 0.1\%$. This reflects the theoretical asymptotic results of Goodman and Ostrov (2010) for an infinite-horizon problem.

6.1.2 Example 2

In this example, we solve the three asset portfolio model (4) with one-year time periods and 6-year horizon, i.e., $h = 1$ year and $T = 6$ periods. The drift terms of stocks, $\mu_t = (\mu_{t,1}, \mu_{t,2})^\top$, are discrete Markov chains and independent of each other. The transition probability matrix

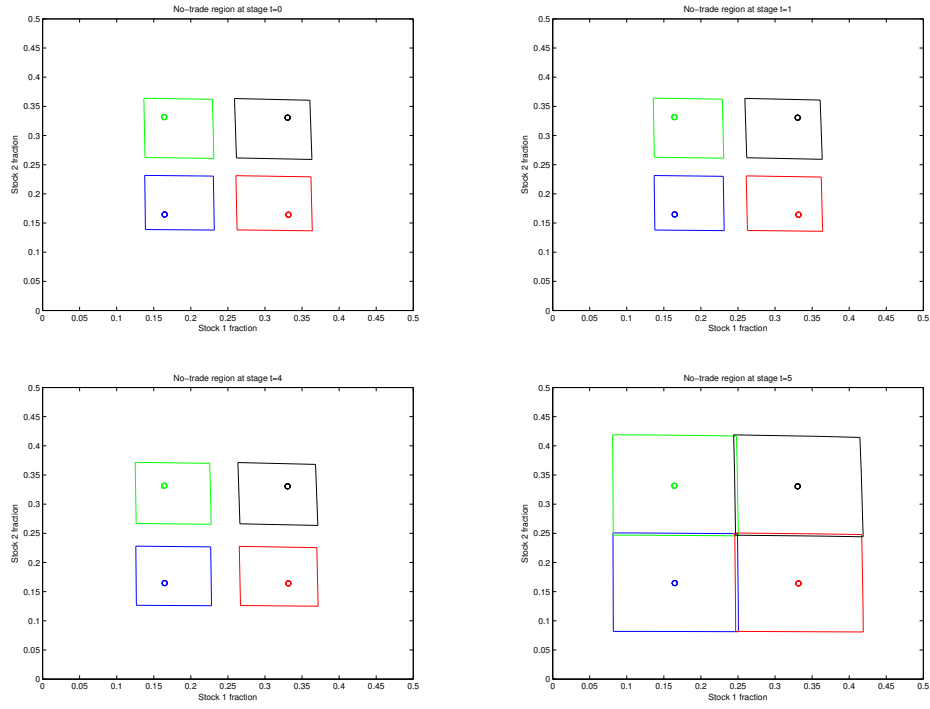
$$\mathbb{P}(\mu_{t+1,i} | \mu_{t,i}) = \begin{bmatrix} 0.75 & 0.25 \\ 0.25 & 0.75 \end{bmatrix}$$

for $i = 1, 2$, while $\mu_{t,1}$ is independent of $\mu_{t,2}$. We let $r = 0.04$ and $\sigma = (0.2, 0.2)^\top$. The proportional transaction cost ratio is $\tau = 1\%$ for buying or selling stocks.

Figure 3 displays NTR's for four possible discrete states of $(\mu_{t,1}, \mu_{t,2})$ at stages $t = 0, 1, 4, 5$. These NTR's are close to be square as the two stocks are i.i.d.. The top-right squares are NTR's for the state $(\mu_{t,1}, \mu_{t,2}) = (0.08, 0.08)$, the bottom-left squares are the NTR for the state $(\mu_{t,1}, \mu_{t,2}) = (0.06, 0.06)$, and the top-left and the bottom-right squares are respectively the NTR for the states $(\mu_{t,1}, \mu_{t,2}) = (0.06, 0.08)$ and $(\mu_{t,1}, \mu_{t,2}) = (0.08, 0.06)$. The circle points inside the squares are the optimal allocation fractions given the discrete states $(\mu_{t,1}, \mu_{t,2})$ if there were no transaction costs for trading stocks. We see that in the case with serially correlated parameters, the sensitivity of NTR's on the trading horizon is still small, and only the last period is significantly different.

In Figure 3, the circle points representing $\tau = 0$ are not in centers of their corresponding squares when $t < 5$. Each circle point stays away from the center of its square at each stage, and the "outside edges" of its NTR are closer

Figure 3: NTR for 2 stocks with stochastic μ and 1 bond



to the circle point than the “inside edges”. Suppose that the discrete state at time t is $(\mu_{t,1}, \mu_{t,2}) = (0.06, 0.06)$, i.e., the NTR is the bottom-left square, a portfolio before rebalance locates near the center of all the four squares (i.e., a point in the northeast direction of the circle in the bottom-left square) tends to have less trade than another portfolio having the same distance from the circle but away from the center (i.e., a point in the southwest direction of the circle in the NTR). This confirms the intuition that if the portfolio is in a position closer to the center of all NTR’s, there is little or no incentive to trade, because the expected direction of next period’s trade is close to zero.

We ran Example 1 and 2 on a single core of a Mac laptop with a 2.5 GHz processor. Every case in both examples took only several seconds.

6.2 Portfolio Problems with Consumption

We gives two examples to illustrate our numerical DP method for solving the model (5), the multi-stage portfolio optimization problems with consumption. In these examples, the assets available for trading include one bond with a constant interest rate r and two or three stocks with correlated log-normal annual returns. Table 1 lists the other parameters and running times on a single core of a Mac laptop with a 2.5 GHz processor.

6.2.1 Example 3

This example assumes that the total investment horizon is 30 years and the portfolio is rebalanced each month. This is a finite-horizon discrete time approximation for the infinite-horizon continuous-time dynamic portfolio problem given by Discussion 1 in Muthuraman and Kumar (2006). To approximate the infinite-horizon problem, we assume that after the terminal stage $T = 360$, the transaction cost becomes 0 and the portfolio can be traded at any time, so the investor will always choose the Merton ratio,

$$x^* = \frac{(\Lambda \Sigma \Lambda)^{-1}(\mu - r)}{\gamma},$$

Table 1: Parameters and Running Times for Examples of Portfolio Problems with Consumption

	Example 3	Example 4
k	2	3
h	1/12 (1 month)	1
T	360	6
γ	2	3
β	$\exp(-0.1h)$	0.95
r	0.07	0.04
μ^\top	(0.15, 0.15)	(0.07, 0.07, 0.07)
σ^\top	$(\sqrt{0.17}, \sqrt{0.17})$	(0.2, 0.2, 0.2)
Σ	$\begin{bmatrix} 1 & 0.4706 \\ 0.4706 & 1 \end{bmatrix}$	$\begin{bmatrix} 1 & 0.4 & 0.4 \\ 0.4 & 1 & 0.16 \\ 0.4 & 0.16 & 1 \end{bmatrix}$
Running time	3.5 minutes	8 minutes

and the optimal portfolio and the optimal consumption rate is

$$c^* \equiv \frac{1}{\gamma} \left[\rho - (1 - \gamma) \left(\frac{(\mu - r)^\top x^*}{2} + r \right) \right] = 9.14\%,$$

where $\rho = -\log(\beta)/h = 0.1$ is the continuously compounded discount rate corresponding to the discretely compounded discount factor β . See Merton (1969, 1971) for more details. Therefore, the terminal value function is $V_T(W, x) = W^{1-\gamma} \cdot g_T(x)$, with

$$g_T(x) = \int_0^\infty e^{-\rho t} u(c^*) dt = \frac{(c^*)^{1-\gamma}}{\rho(1-\gamma)}.$$

In the numerical DP method for this example, we use the degree-14 complete Chebyshev approximation method with 15^2 tensor Chebyshev nodes for two stocks, and the multi-dimensional product Gauss-Hermite quadrature rule with 9 nodes in each dimension.

Figure 4 displays NTR's at the stage $t = 0, 320, 357, 358, 359$, when $h = 1/12$ year and the transaction cost proportion is $\tau = 1\%$. The circle point located inside NTR is the Merton point (0.16, 0.16). We found that

Figure 4: NTR for 2 correlated stocks and 1 bond with consumption when $\tau = 1\%$ (Example 2)





Figure 5: Initial-time NTR for 2 correlated stocks and 1 bond with consumption at various transaction costs (Example 2)



NTR at stage $t = 0$ is close to the solution given in Muthuraman and Kumar (2006).

We also ran the same example with $h = 0.25$ year and various τ using numerical DP algorithm. Figure 5 shows NTR's at $t = 0$ with $h = 0.25$ year and various transaction cost proportions $\tau = 0.1\%, 0.2\%, 0.5\%, 1\%, 2\%$. We can also see that the NTR with larger τ contains the NTR with smaller τ . It also shows that the NTR's at time $t = 0$ converges to a point $(0.159, 0.159)$ which is very close to the Merton's point, as $\tau \rightarrow 0$.

6.2.2 Example 4

In this example, we have one bond and three stocks available for trading, with a transaction cost proportion $\tau = 1\%$. The terminal value function is $V_T(W, x) = W^{1-\gamma}/(1-\gamma)$. In the numerical DP method for this example, we use the degree-10 complete Chebyshev approximation method with 11^3 tensor Chebyshev nodes for three stocks, and the multi-dimensional product Gauss-Hermite quadrature rule with 9 nodes in each dimension.

Figure 6 displays NTR's for $t = 0, 1, 4, 5$. Notice that the faces of the NTR's seem to be flat, but in fact there are small perturbation on the faces, which might be due to numerical errors or the possibility that the exact NTR might have curvy faces. The NTR's are tilted as the stocks are correlated. We see that in the case with 3 correlated stocks and 1 bond and consumption, the sensitivity of NTR's on the trading horizon is still small, and only the last period is significantly different.

6.3 Seven-Asset Portfolio Problem

Our last example solves the model (2), where the assets available for trading include one bond with a constant interest rate $r = 0.03$ and six stocks with independent log-normal annual returns. We assume that

$$\begin{aligned}\mu &= (0.06, 0.064, 0.068, 0.072, 0.076, 0.08)^\top, \\ \sigma &= (0.2, 0.22, 0.24, 0.26, 0.28, 0.3)^\top,\end{aligned}$$

$\gamma = 5$, $\tau = 1\%$, and $h = 1$ year.

We applied the degree-4 complete Chebyshev approximation method with 5 Chebyshev nodes in each dimension, and the multi-dimensional product Gauss-Hermite quadrature rule with 5 nodes in each dimension. Thus, for each value function iteration, it has $5^6 = 15,625$ optimization problems in the maximization step. Moreover, the evaluation of the objective function of each optimization problem is also time-consuming: it needs $5^6 = 15,625$ evaluations of degree-4 complete Chebyshev polynomials. This is a computationally intensive problem. But we can solve it by parallelization.

Figure 6: NTR for 3 stocks with higher correlation and 1 bond with consumption (Example 4)

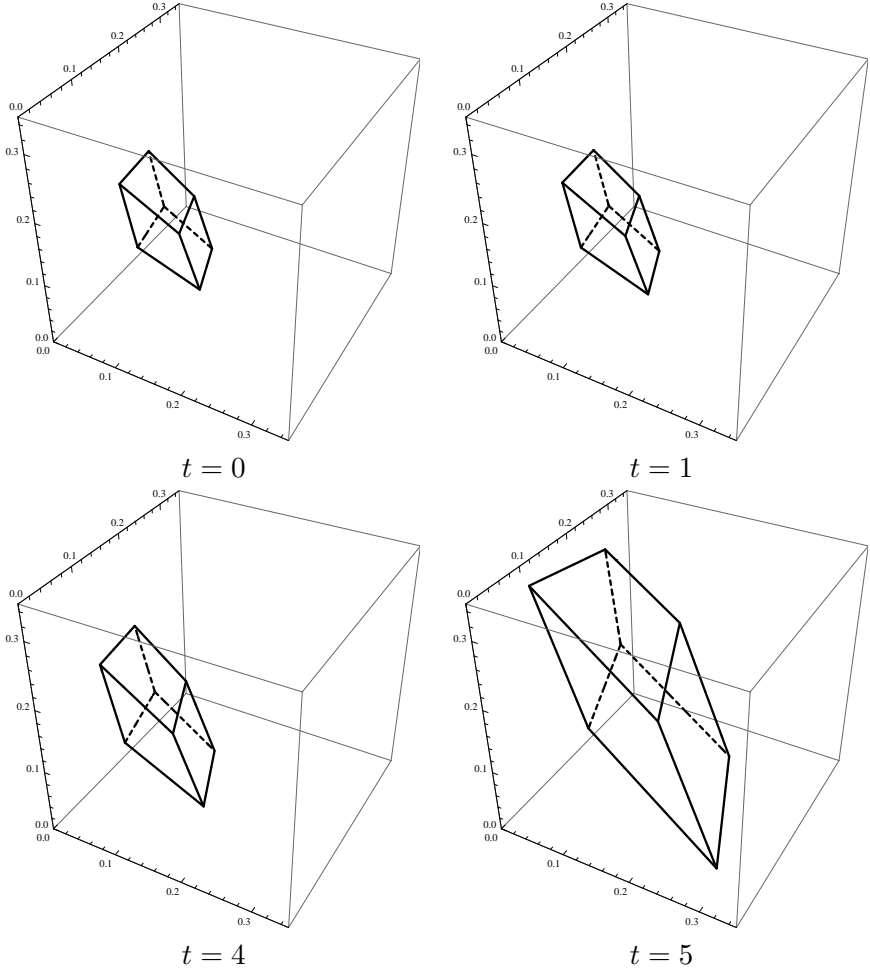
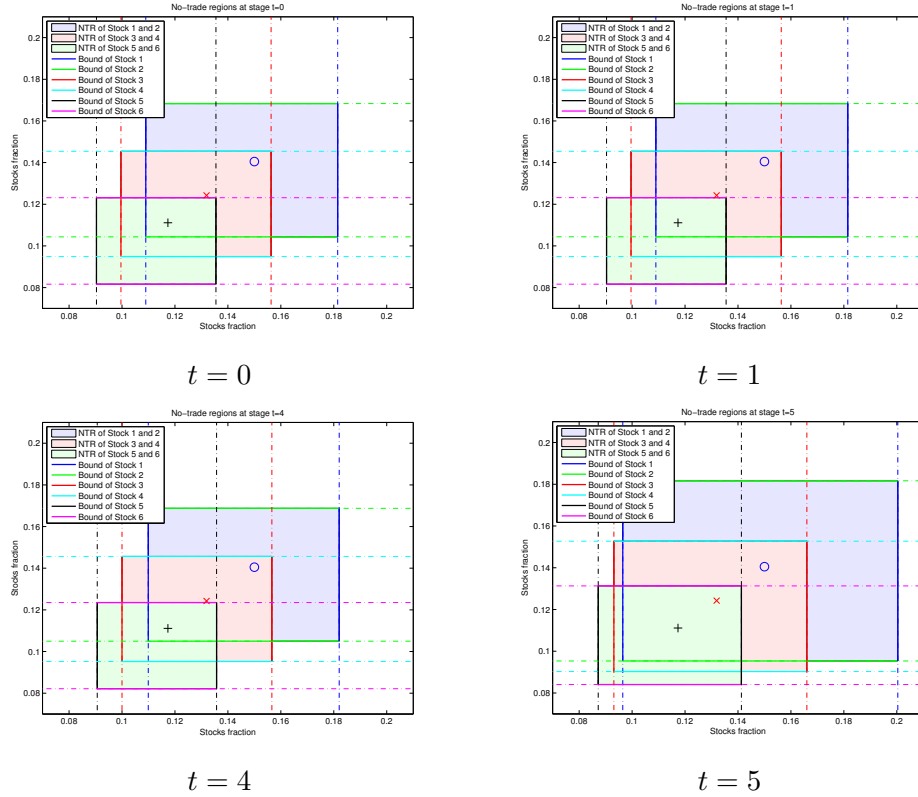


Figure 7: NTR for 6 independent stocks and 1 bond



At first, we applied the parallel DP algorithm in the HTCondor-MW system, a high-throughput distributed computing framework with high-latency on a cluster of computers. It took 1.3 hours to get the solutions by requesting 100 worker processes in the system. More details about the DP parallelization in the HTCondor-MW system are discussed in Cai et al. (2013). The parallel DP algorithm can also be extended to be applied in high-performance supercomputing with low-latency. We used a supercomputer in Computation Institute of University of Chicago, called Beagle (see <http://beagle.ci.uchicago.edu> for more details), to solve the problem. The parallel DP algorithm used 480 cores and took only 16 minutes to get the solutions.

Figure 7 displays NTR's for $t = 0, 1, 4, 5$, and we see that the regions

are close to hypercubes as the six stocks are uncorrelated. In the figure, the circle, the mark and the plus are, respectively, the Merton points for stock 1 and 2, stock 3 and 4, and stock 5 and 6. The top-right squares, the middle squares, and the bottom-left squares are the cross-sections of the hypercube NTR's for stock 1 and 2, stock 3 and 4, and stock 5 and 6, respectively. We see that in the case with 6 stocks and 1 bond, the sensitivity of NTR's on the trading horizon is still small, and only the last period is significantly different.

There is no problem to apply parallel DP algorithms to solve a problem with 6 correlated stocks and 1 bond, and clearly the parallel DP algorithms can be also applied to solve a problem with more stocks as the supercomputers typically now have more than 100,000 cores.

7 Conclusion

This paper has shown that numerical value function iteration can solve multi-stage portfolio optimization problems with multiple assets and transaction costs in an efficient and accurate manner. We illustrate the trading strategies by describing the no-trade regions for various choices of asset returns and transaction costs. The numerical DP algorithms may be computational intensive for large portfolio optimization problems, but modern hardware and parallel DP algorithms make it possible to now solve such big problems.

References

- [1] Abrams, R.A., and U.S. Karmarkar (1980). Optimal multiperiod investment-consumption policies. *Econometrica*, 48(2):333–353.
- [2] Akian, M., J.L. Menaldi, and A. Sulem (1996). On an investment-consumption model with transaction costs. *SIAM Journal of Control and Optimization*, 34(1):329–364.
- [3] Boyle, P.P., and X. Lin (1997). Optimal portfolio selection with transaction costs. *North American Actuarial Journal*, 1(2):27–39.
- [4] Brown, D.B., and J.E. Smith (2011). Dynamic portfolio optimization with transaction costs: heuristics and dual bounds. *Management Science*, 57(10):1752–1770.
- [5] Cai, Y. (2009). *Dynamic Programming and Its Application in Economics and Finance*. PhD thesis, Stanford University.
- [6] Cai, Y., and K.L. Judd (2010). Stable and efficient computational methods for dynamic programming. *Journal of the European Economic Association*, Vol. 8, No. 2-3, 626–634.
- [7] Cai, Y., and K.L. Judd (2012a). Dynamic programming with shape-preserving rational spline Hermite interpolation. *Economics Letters*, Vol. 117, No. 1, 161–164.
- [8] Cai, Y., and K.L. Judd (2012b). Shape-preserving dynamic programming. *Mathematical Methods of Operations Research*, DOI: 10.1007/s00186-012-0406-5.
- [9] Cai, Y., and K.L. Judd (2012c). Dynamic programming with Hermite approximation. NBER working paper No. 18540.

- [10] Cai, Y., K.L. Judd, G. Thain, and S. Wright (2013). Solving dynamic programming problems on a computational grid. Hoover working paper.
- [11] Constantinides, G.M. (1976). Optimal portfolio revision with proportional transaction costs: extension to HARA utility functions and exogenous deterministic income. *Management Science*, 22(8):921–923.
- [12] Constantinides, G.M. (1979). Multiperiod consumption and investment behavior with convex transaction costs. *Management Science*, 25:1127–1137.
- [13] Constantinides, G.M. (1986). Capital market equilibrium with transaction costs. *Journal of Political Economy*, 94(4):842–862.
- [14] Davis, M.H.A., and A.R. Norman (1990). Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713.
- [15] Duffie, D., and T.S. Sun (1990). Transaction costs and portfolio choice in a discrete-continuous-time setting. *Journal of Economic Dynamics and Control*, 14:35–51.
- [16] Gennotte, G., and A. Jung (1994). Investment strategies under transaction costs: the finite horizon case. *Management Science*, 40(3):385–404.
- [17] Gill, P., W. Murray, M.A. Saunders and M.H. Wright (1994). User’s Guide for NPSOL 5.0: a Fortran Package for Nonlinear Programming. Technical report, SOL, Stanford University.
- [18] Goodman, J., and D.N. Ostrov (2010). Balancing small transaction costs with loss of optimal allocation in dynamic stock trading strategies. *SIAM Journal of Applied Mathematics*, 70(6):1977–1998.

- [19] Janecek, K., and S.E. Shreve (2004). Asymptotic analysis for optimal investment and consumption with transaction costs. *Finance Stochast.*, 8(2):181–206.
- [20] Kamin, J.H. (1975). Optimal portfolio revision with a proportional transaction cost. *Management Science*, 21(11):1263–1271.
- [21] Liu, H. (2004). Optimal consumption and investment with transaction costs and multiple risky assets. *Journal of Finance*, 59:289–338.
- [22] Merton, R. (1969). Lifetime portfolio selection under uncertainty: the continuous time case. *The Review of Economics and Statistics*, 51(3):247–257.
- [23] Merton, R. (1971). Optimum consumption and portfolio rules in a continuous time model. *Journal of Economic Theory*, 3:373–413.
- [24] Muthuraman, K., and S. Kumar (2006). Multidimensional portfolio optimization with proportional transaction costs. *Mathematical Finance*, 16(2):301–335.
- [25] Muthuraman, K., and S. Kumar (2008). Solving free-boundary problems with applications in finance. *Foundations and Trends in Stochastic Systems*, 1(4):259–341.
- [26] Muthuraman, K., and H. Zha (2008). Simulation-based portfolio optimization for large portfolios with transaction costs. *Mathematical Finance*, 18(1):115–134.
- [27] Rust, J. (2008). Dynamic Programming. In: *New Palgrave Dictionary of Economics*, ed. by Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan, second edition.
- [28] Zabel, E. (1973). Consumer choice, portfolio decisions, and transaction costs. *Econometrica*, 41(2):321–335.