Discrete time growth optimal investment with costs*

G. Iyengar†

July 16, 2002

Abstract

In this work we ask how should an investor distribute wealth over various assets to maximize the growth rate of the cumulative wealth in a discrete time market with proportional transaction costs. We show that this sequential decision problem has a stationary optimal policy. In addition, we show that for all $\epsilon > 0$ there exists a policy that guarantees a growth rate at most ϵ below optimal on almost every sample path. We also show the existence of an ϵ-optimal control-limit policies – control-limit policies correct the portfolio only it leaves a compact connected no-trade set. For the special case of two-asset markets, we establish that for all $\epsilon > 0$ there exists a control-limit policy that is ϵ-optimal with probability 1.

1 Introduction

In this work, we ask how should an individual distribute funds over various assets to maximize the growth rate of his compounded wealth in a discrete time market with costs. The market dynamics are described by a sequence of price relative vectors $\{\mathbf{X}_n : n \geq 1\}$, where $\mathbf{X}_n$ is vector of the ratio of closing to opening prices over period n. The market levies a transaction cost proportional to the volume of the trades. The details of the market model, the allowed trades and the trading costs are described in the next section.

Growth optimal investment in discrete time markets with proportional transaction costs can be viewed as an extension of the literature that originated with the work of Kelly on horse race markets [15]. Horse race markets are special discrete time markets where in every period one of the assets pays off certain odds and all the others go broke. Kelly showed that it is growth optimum to constantly rebalance to the portfolio that maximizes the expected logarithm of the one step return, i.e. the log-optimum portfolio; and the optimum growth rate is given by $m - H(p)$, where m is the number of assets and $H(p)$ is the entropy of the market distribution. Breiman proved that the log-optimum policy is growth optimum for general independent identically distributed (IID) markets as well [6]. Subsequently, Cover and Algoet established that the conditionally log-optimum investment policy is growth optimum for stationary and ergodic markets and the optimal growth rate is given by $m - H$ where H is the entropy-rate of the market process[3].

All of this earlier work assumes that markets do not have transaction costs. Although policies that ignore costs do not cause the investor to go bankrupt in discrete time markets [14], the growth rate can be substantially improved by designing policies that explicitly account for the cost. In a previous work we investigated the effect of costs in the horse race setting [12]. We establish that the

*Submitted to *Mathematics of Operations Research*. Please do not circulate

†IEOR Department, Columbia University, Email: `garud@ieor.columbia.edu`. Research partially supported by NSF grants CCR-00-09972 and DMS-01-04282.

optimal policy is to rebalance to a portfolio that maximizes a combination of the one-step return and transaction cost. In this paper we extend our work to general IID markets with proportional costs.

Although the rate of return is used as the benchmark in classical investment theory [18], the growth rate of a portfolio is a better indicator of long-term portfolio behavior and is beginning to attract attention [8]. Growth optimal investment in continuous time two-asset Brownian markets with proportional transaction costs was formulated and solved by Taksar, Klass and Assaf [26]. They established that the growth optimal policy keeps the portfolio in a closed interval by trading only when the portfolio hits the boundaries, i.e. the transactions are local times on the boundary. The optimal policy for the closely related problem of optimal consumption and investment in two-asset Brownian markets with proportional costs also has a similar character – the wealth invested in the two assets is kept in a wedge-shaped region by trading only when the portfolio hits the boundary [7] (see also [17]). This local-time character of the optimal trading policy was investigated in detail by Shreve and Soner [24]. The optimal investment-consumption problem for markets with several assets was solved by Akian, Menaldi and Sulem [1]. Iyengar [13] and Akian, Sulem and Taksar [2] investigate growth optimal investment in continuous time markets with several risky assets.

Due to its local-time character the continuous time optimal policy is not implementable in markets that do not allow continuous trades. One remedy is to include a fixed transaction cost; thereby, forcing the investor to take a large step at every transaction instant to cover the fixed costs – thus, trades occur only at discrete instants of time. We investigate this model in an upcoming publication [10]. Another possibility is to explicitly constrain the investor to trade only at discrete instants in time. In this work we take the latter approach and study growth optimal investment in discrete time markets with transaction costs.

Our contributions in this work are as follows:

(a) We formulate growth optimal investment in general discrete time markets with proportional transaction costs and establish the existence of a stationary Markov deterministic optimal policy.

(b) We show that for all $\epsilon > 0$ there exists a stationary ϵ-optimal policy, i.e. a policy with growth rate at most ϵ below optimal, that is a continuous function of the portfolio. Using Weak Feller continuity of the portfolio process induced by an ϵ-optimal continuous policy we show the existence of a policy with a growth rate arbitrarily close to optimal on almost every sample path.

(c) Analogous to the continuous time result [2, 13], we establish that for all $\epsilon > 0$ there exists an ϵ-optimal control-limit policy with the a compact, connected no-trade set $\mathcal{N}$ – the policy does not exercise any control as long as the portfolio lies in $\mathcal{N}$ and corrects to the boundary $\partial\mathcal{N}$ whenever the portfolio drops out of $\mathcal{N}$.

(d) For the special case of two-asset markets, we show that the control-limit policies achieve the corresponding growth rate almost surely. We also present a simple numerical scheme that computes the optimal no-trade interval for a sampled Brownian market. Using this numerical scheme we show that the optimal discrete time growth rate is given by $g = g_0 - g_1\delta$, where δ is the sampling interval, whereas the policy that corrects to the optimal continuous time no-trade interval has a growth rate $g = g_0 - \tilde{g}_1\sqrt{\delta}$.

The rest of this paper are organized as follows. In Section 2 we formulate growth optimal investment in discrete time markets. In Section 3 we discuss results for markets with several assets and

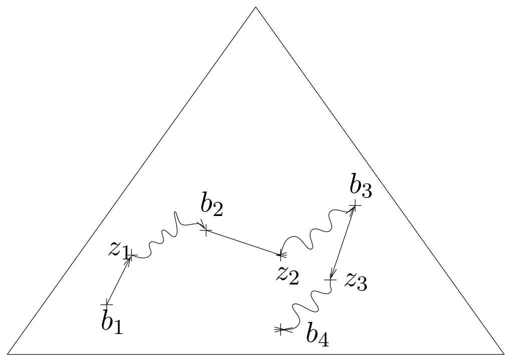

Figure 1: Time evolution of the portfolios (b, z)

general IID market distribution. In Section 4 we present stronger versions of the general results for the special case of two-asset markets. In addition, we present a simple numerical procedure for computing the optimal no-trade set for sampled Brownian markets. Section 5 concludes with some criticisms of the model and avenues of further research. The proofs of the results are in the Appendix.

2 Problem formulation

We consider a market with m assets which opens for trading at discrete instants in time. The price dynamics of the assets is described by the sequence of positive price relative vector $\{\mathbf{X}_n = (X_n(i))_{(1 \leq i \leq m)} : n \geq 1\}$, where $X_n(i)$ is the ratio of the closing to the opening price of asset i over the period n. This sequence of price relative vectors is assumed to be non-negative and IID according to a known distribution.

Cash is the numeraire for the market and mediates all transactions. The market levies proportional transaction costs at a rate $0 \leq \lambda_i < 1$ on asset i, i.e., the sale of a dollar's worth of asset i nets $(1 - \lambda_i)$ dollars in cash and a dollar's worth of asset i costs $1 + \lambda_i$ dollars. Since cash is assumed to mediate all transactions, selling a dollar's worth of asset i and investing the proceeds in the asset j, involves first selling a dollar of asset i to get $(1 - \lambda_i)$ dollars in cash; and then using this capital to buy $\frac{1-\lambda_i}{1+\lambda_j}$ dollars worth of asset j. Although we have assumed symmetric transaction costs, all the results carry over to the case of asymmetric costs.

The investor starts has an initial wealth S_1 that is invested in a portfolio $\mathbf{b}_1 \in \mathbf{R}^m$. (A portfolio vector $\mathbf{b}$ represents the proportion of the total wealth invested in a particular asset, i.e. $\sum_{i=1}^m b(i) = 1$.) The cumulative wealth at the beginning of the n-th market period is denoted by S_n and the portfolio by $\mathbf{b}_n$, i.e. the dollar amount invested in asset i is $S_n b_n(i)$, $i = 1, \ldots, m$. In every investment period, the investor has the option to trade from $\mathbf{b}_n$ to a new portfolio $\mathbf{z}_n$. However, the transaction results in a loss of wealth because of the transaction costs in the market. The net realized wealth when one dollar invested in portfolio $\mathbf{b}_n$ is optimally traded to portfolio $\mathbf{z}_n$ is denoted by $w(\mathbf{b}_n, \mathbf{z}_n)$. In the sequel we shall refer to $w(\mathbf{b}_n, \mathbf{z}_n)$ as the repositioning factor. Since the transaction costs are proportional, wealth S_n invested in portfolio $\mathbf{b}_n$ would net $w(\mathbf{b}_n, \mathbf{z}_n)S_n$ at portfolio $\mathbf{z}_n$.

Subsequently the market reveals a realization $\mathbf{x}_n$ of the random price relative vector $\mathbf{X}_n$. Since a dollar invested in asset i is worth $x_n(i)$ dollars, one dollar invested in portfolio $\mathbf{z}_n$ yields the wealth $\mathbf{z}_n^T \mathbf{x}_n = \sum_{i=1}^m z_n(i) x_n(i)$, and the new portfolio $\mathbf{b}_{n+1} = \mathbf{z}_n \circ \mathbf{x}_n$, where $\mathbf{z} \circ \mathbf{x}$ is defined as follows: for all $i = 1, \ldots, m$,

$$(\mathbf{z} \circ \mathbf{x})(i) = \begin{cases} \frac{z(i)x(i)}{\mathbf{z}^T\mathbf{x}}, & \mathbf{z}^T\mathbf{x} > 0, \\ \frac{1}{m} & \text{otherwise.} \end{cases} \tag{1}$$

(The case $\mathbf{z}^T\mathbf{x} \leq 0$ has to be included for technical reasons.)

This process repeats itself in every market period. Therefore, the cumulative wealth S_n is given by

$$\begin{aligned} S_n &= (\mathbf{z}_{n-1}^T \mathbf{x}_{n-1}) w(\mathbf{b}_{n-1}, \mathbf{z}_{n-1}) S_{n-1}, \\ &= S_1 \prod_{k=1}^{n-1} (\mathbf{z}_k^T \mathbf{x}_k) w(\mathbf{b}_k, \mathbf{z}_k), \end{aligned}$$

and S_n is invested in the portfolio $\mathbf{b}_{n+1} = \mathbf{z}_n \circ \mathbf{x}_n$. The time evolution of the pair of portfolios $\{(\mathbf{b}_n, \mathbf{z}_n) : n \geq 1\}$ is shown in Figure 1. The wiggly lines represent the stochastic portfolio movements corresponding to the market moves and the straight line are the trades that the investor initiates.

In this work, we are interested in maximizing the expected asymptotic growth rate of the wealth of the investor, i.e., very loosely speaking we want to maximize

$$g = \lim_{n\to\infty} \mathbf{E}\Big[\frac{1}{n} \log S_n\Big] = \lim_{n\to\infty} \Big\{\frac{1}{n} \log S_1 + \frac{1}{n} \sum_{k=1}^{n-1} \Big(\mathbf{E} \log w(\mathbf{b}_k, \mathbf{z}_k) + \mathbf{E} \log(\mathbf{z}_k^T \mathbf{X}_k)\Big)\Big\}.$$

Since the costs are proportional, it immediately follows that the growth rate g is independent of the initial wealth S_1 but may depend on the initial portfolio $\mathbf{b}_1$. Therefore, one can assume without loss of any generality that $S_1 = 1$.

Having motivated the problem, we now rigorously formulate the underlying Markov decision problem. The elements of the model are as follow:

1. The space of admissible portfolios for the investor, i.e. the *state space*, $\mathcal{S}$ is given by

$$\mathcal{S} = \left\{ \mathbf{b} \in \mathbf{R}^m \;\middle|\; \sum_{i=1}^m b(i) = 1, \sum_{i=1}^m (1-\lambda_i) b^+(i) \geq M \sum_{i=1}^m (1+\lambda_i) b^-(i) \right\}, \tag{2}$$

   where $b^+(i) = \max\{b(i), 0\}$, $b^-(i) = \max\{-b(i), 0\}$, $i = 1, \ldots, m$, and $M > 1$. The additional constraint in (2) ensures the long positions are at least a certain *margin* M times the liabilities.

2. From an admissible portfolio $\mathbf{b} \in \mathcal{S}$ a move to portfolio $\mathbf{z} \in \mathcal{S}$ is allowed only if the repositioning factor $w(\mathbf{z}, \mathbf{b}) > 0$. The *action space* $\mathcal{A}(\mathbf{b})$ for the portfolio $\mathbf{b}$ is given by

$$\mathcal{A}(\mathbf{b}) = \{\mathbf{z} \mid w(\mathbf{b}, \mathbf{z}) > 0\}.$$

   Let $\mathcal{A} =$ graph $\mathcal{A}(\cdot) = \{(\mathbf{b}, \mathbf{z}) \mid \mathbf{z} \in \mathcal{A}(\mathbf{b})\}$.

3. The transition probability $P(B \mid \mathbf{z}, \mathbf{b})$, $B \in \mathcal{B}(\mathbf{S})$, $\mathbf{b} \in \mathbf{S}$, $\mathbf{z} \in \mathbf{A}(\mathbf{b})$ is assumed to be a regular conditional probability given by $P(B \mid \mathbf{b}, \mathbf{z}) = \mathbf{P}\{\mathbf{z} \circ \mathbf{X} \in B\}$.

4. The one-step reward $r : \mathcal{S} \times \mathcal{S} \times \mathbf{R}_+^m \to \mathbf{R}$, when the investor moves from portfolio $\mathbf{b}$ to portfolio $\mathbf{z} \in \mathcal{A}(\mathbf{b})$ and the market price relative vector is $\mathbf{x}$, is given by

$$r(\mathbf{b}, \mathbf{z}, \mathbf{x}) = \log w(\mathbf{z}, \mathbf{b}) + \log(\mathbf{z}^T \mathbf{x}).$$

   The function r is continuous in its arguments.

The history spaces $\{\Omega_n : n \geq 0\}$ are defined as follows: $\Omega_0 = \mathcal{S}$, $\Omega_{n+1} = \Omega_n \times \mathcal{A} \times \mathbf{R}_+^m$. A particular element $\omega_n = (\mathbf{b}_1, \mathbf{z}_1, \mathbf{x}_1, \ldots, \mathbf{b}_{n-1}, \mathbf{z}_{n-1}, \mathbf{x}_{n-1}, \mathbf{b}_n) \in \Omega_n$ summarizes all the information available in period n. A randomized policy $\pi = (\pi_1, \pi_2, \ldots)$ is a sequence of transition probabilities from Ω_n to $\mathcal{S}$ such that $\pi_n(\cdot \mid \omega_n)$ is a regular conditional probability and $\pi(\mathcal{A}(\mathbf{b}_n) \mid \omega_n) = 1$, for all $\omega_n \in \Omega_n$. The probability measure induced by a policy π is denoted by $\mathbf{P}^\pi$. A non-random policy $\pi = (\pi_1, \pi_2, \ldots)$ is a sequence of measurable functions from Ω_n to $\mathcal{S}$ such that $\pi_n(\omega_n) \in \mathcal{A}(\mathbf{b}_n)$. We will call a policy Markov if $\pi_n(\omega_n) = \pi_n(\mathbf{b}_n)$, i.e. the decision in period n is a function of $\mathbf{b}_n$ alone, and is independent of the past history.

Since our model does not allow any inflow of capital after the initial investment, we restrict the investor to policies which ensure that $S_n > 0$ and $\mathbf{b}_n \in \mathcal{S}$ for all $n \geq 1$. Define the stopping time τ as follows:

$$\tau = \min\{n \mid \mathbf{b}_n \notin \mathcal{S} \text{ or } S_n \leq 0\}.$$

Then, the set of *admissible* policies $\Pi(\mathbf{b})$ at the initial portfolio $\mathbf{b}$ is given by

$$\Pi(\mathbf{b}) = \{\pi : \tau = \infty, \mathbf{P}^\pi - \text{a.s.}\} \tag{3}$$

It is easy to show that the set $\Pi(\mathbf{b})$ is nonempty. The set of admissible Markov policies is denoted by Π_M.

Given this framework, we are interested in characterizing the *growth rate function* $g(\mathbf{b})$ given by,

$$g(\mathbf{b}) = \sup_{\pi \in \Pi(\mathbf{b})} \liminf_{n \to \infty} \mathbf{E}\Big[\frac{1}{n} \log S_n^\pi(\mathbf{b})\Big], \tag{4}$$

$$= \sup_{\pi \in \Pi(\mathbf{b})} \liminf_{n \to \infty} \Big\{\frac{1}{n} \sum_{k=1}^{n-1} \Big(\mathbf{E} \log w(\mathbf{b}_k, \pi(\mathbf{b}_k)) + W(\pi(\mathbf{b}_k))\Big)\Big\}, \tag{5}$$

where $\Pi(\mathbf{b})$ is set of all policies admissible for $\mathbf{b} \in \mathcal{S}$, $S_n^\pi(\mathbf{b})$ is the wealth at time n generated by policy π, $w(\mathbf{b}, \mathbf{z})$ is the repositioning factor and $W(\mathbf{z}) = \mathbf{E} \log(\mathbf{z}^T \mathbf{X})$. Thus, the *growth rate function* $g(\mathbf{b})$ is the maximum expected interest rate guaranteed for the initial portfolio $\mathbf{b}$.

Since growth rate $g(\mathbf{b})$ is defined in terms of the expected value, a growth optimal policy may not achieve $g(\mathbf{b})$ on every sample path. We show later that a growth rate arbitrarily close to $g(\mathbf{b})$ is achievable with probability 1.

3 Results for general IID markets

In this section we discuss the results for general IID markets. Since the proofs do not contribute much to the understanding of the results, we have relegated them to the Appendix.

We begin with characterizing the repositioning factor $w(\mathbf{b}, \mathbf{z})$ as an optimization problem.

Lemma 1 *The repositioning factor $w(\mathbf{b}, \mathbf{z})$ is the solution of the convex program*

$$\begin{array}{ll} \textit{maximize} & w \\ \textit{subject to} & w + \sum_{i=1}^m \lambda_i \mid w z(i) - b(i) \mid \leq 1. \end{array}$$

Moreover, for all $\mathbf{b}, \mathbf{z} \in \mathcal{S}$

$$w(\mathbf{z}, \mathbf{b}) \geq \left(1 - \frac{1}{M}\right) \left(\frac{1 - \lambda_{\max}}{Z}\right),$$

where $Z = \max_{(\mathbf{z} \in \mathcal{S})}\{1 + \sum_{i=1}^m \lambda_i \, |z(i)|\} < \infty$.

Since $w(\mathbf{b}, \mathbf{z}) > 0$, for all $\mathbf{b}, \mathbf{z} \in \mathcal{S}$, and a single transaction does not affect the growth rate, we have the following corollary.

Corollary 1 *The growth rate $g(\mathbf{b})$ is a constant on the state space $\mathcal{S}$.*

The growth rate function $g(\mathbf{b})$ is defined in (4) as the maximum achievable interest rate when investors are allowed to employ general history dependent policies. However, it is well known from the literature of Markov decision processes that the investor can be restricted to Markov policies without any loss in performance [20, 25].

Lemma 2 *For all initial portfolios $\mathbf{b} \in \mathcal{S}$, the following identity holds:*

$$\sup_{\pi \in \Pi(\mathbf{b})} \liminf_{n \to \infty} \mathbf{E}\Big[\frac{1}{n} \log S_n^{\pi}(\mathbf{b})\Big] = \sup_{\pi \in \Pi_M} \liminf_{n \to \infty} \mathbf{E}\Big[\frac{1}{n} \log S_n^{\pi}(\mathbf{b})\Big].$$

In addition, for all $\mathbf{b} \in \mathcal{S}$ and $\pi \in \Pi(\mathbf{b})$, there exists $\tilde{\pi} \in \Pi_M$ such that

$$\liminf_{n \to \infty} \mathbf{E}\Big[\frac{1}{n} \log S_n^{\pi}(\mathbf{b})\Big] = \liminf_{n \to \infty} \mathbf{E}\Big[\frac{1}{n} \log S_n^{\tilde{\pi}}(\mathbf{b})\Big].$$

The lemma follows immediately from Lemma 4.1 in [25]. The result is established by showing that the conditional distribution induced by a general history dependent policy can be replicated by a Markov policy. For a simple exposition of this fact for finite state and action spaces see Theorem 5.5.1 in [20].

For a Markov policy π the stopping time $\tau = \infty$, $\mathbf{P}^{\pi}$ – a.s., if and only if, for all $\mathbf{b} \in \mathcal{S}$, $\pi(\mathbf{b})^T \mathbf{X} > 0$ and $\pi(\mathbf{b}) \circ \mathbf{X} \in \mathcal{S}$, $\mathbf{P}^{\pi}$-a.s. Therefore, we have the following one-step formulation of admissibility for Markov policies.

Lemma 3 *A Markov policy $\pi = (\pi_1, \pi_2, \ldots)$ is admissible if and only if the transition probability measure $\pi_n(\cdot \mid \mathbf{b}_n)$ puts all its mass on the set*

$$\begin{aligned} \mathcal{K} &= \Big\{\mathbf{z} \in \mathcal{S} \;\Big|\; \mathbf{z}^T \mathbf{X} > 0, \mathbf{z} \circ \mathbf{X} \in \mathcal{S} \ a.s.\Big\}, \\ &= \Big\{\mathbf{z} \in \mathcal{S} \;\Big|\; \sum_{i=1}^{m} ((M-1) + (M+1)\lambda_i) X_i z_i^+ - M \sum_{i=1}^{m} (1+\lambda_i) X_i z_i \le 0, \ a.s.\Big\}, \end{aligned} \tag{6}$$

where (6) follows from the definition of $\mathbf{z} \circ \mathbf{X}$ in (1).

From (6) it follows that $\mathcal{K}$ is a convex compact set. Given the preceding results, we have the following Bellman equation for the investment problem.

Theorem 1 *Let $V : \mathcal{S} \mapsto \mathbf{R}$ be bounded, measurable function and $g \in \mathbf{R}$ such that*

$$V(\mathbf{b}) + g = \sup_{\mathbf{z} \in \mathcal{K}} \{ \log w(\mathbf{b}, \mathbf{z}) + W(\mathbf{z}) + \mathbf{E} V(\mathbf{z} \circ \mathbf{X}) \}, \tag{7}$$

where $W(\mathbf{z}) = \mathbf{E} \log(\mathbf{z}^T \mathbf{X})$ and $w(\mathbf{b}, \mathbf{z})$ is the rebalancing factor defined in Lemma 1. Then g is a upper bound on the achievable growth rate, i.e. for all $\mathbf{b} \in \mathcal{S}$,

$$g \ge \sup_{\pi \in \Pi(\mathbf{b})} \liminf_{n \to \infty} \frac{1}{n} \mathbf{E} \log(S_n^{\pi}(\mathbf{b})). \tag{8}$$

For any (V, g) satisfying (7) with V bounded and measurable, $\pi \in \Pi_M$ and $\mathbf{b}_1 \in \mathcal{S}$, we have that

$$g \geq W(\pi_k(\mathbf{b}_k)) + \mathbf{E}\log w(\mathbf{b}_k, \pi_k(\mathbf{b}_k)) + \Big(\mathbf{E}V(\pi_k(\mathbf{b}_k)\circ \mathbf{X}) - V(\mathbf{b}_k)\Big), \qquad \forall k \geq 1. \tag{9}$$

The bound (8) is established by iterating (9) along a sample path and taking the limit as $n \to \infty$.

Suppose (V_i, g_i), $i = 1, 2$, are two solutions of (7). In addition, suppose there exist measurable functions $\pi_i : \mathcal{S} \mapsto \mathcal{K}$ such that $\pi_i(\mathbf{b}) \in \mathcal{M}(V, \mathbf{b}) \equiv \operatorname{argmax}\{w(\mathbf{b}, \mathbf{z}) + W(\mathbf{z}) + \mathbf{E}V(\mathbf{z}\circ\mathbf{X})\}$. Since functions π_i, $i = 1, 2$, are non-random Markov policies, Theorem 1 together with a simple comparison argument implies that $g_1 = g_2$. Thus, the optimal growth rate can be uniquely characterized by constructing a solution (V, g) of (7) that admits a measurable selection from $\{\mathcal{M}(V, \mathbf{b}) : \mathbf{b} \in \mathcal{S}\}$. The following result formalizes these observations.

Theorem 2 *There exists a unique $g^* \in \mathbf{R}$ such that (V, g^*) is a solution of (7) with $V : \mathcal{S} \mapsto \mathbf{R}$ concave and continuous. Moreover, any such solution (V, g), there exists a measurable function $\pi : \mathcal{S} \mapsto \mathcal{K}$ such that for all $\mathbf{b} \in \mathcal{S}$*

$$V(\mathbf{b}) + g^* = \log w(\mathbf{b}, \pi(\mathbf{b})) + W(\pi(\mathbf{b})) + \mathbf{E}V(\pi(\mathbf{b})\circ\mathbf{X}),$$

i.e. g^* is the optimal growth rate and $\pi^\infty = (\pi, \pi, \ldots)$ is a stationary, non-random, Markov optimal policy.

Remark 1 *Theorem 2 only claims that g^* is unique, i.e. V need not be unique. The uniqueness result is valid even if V is only upper semi-continuous.*

Typically the solution (V, g^*) of the Bellman equation and the associated optimal policy is computed using a combination of value and policy iteration. Without some regularity conditions on the optimal policy π, these numerical methods do not converge. We next show that for all $\epsilon > 0$ there exists a stationary continuous policy π, i.e. $\pi : \mathcal{S} \mapsto \mathcal{K}$ is continuous function, with growth rate no smaller than $g^* - \epsilon$.

Theorem 3 *Let g^* be the optimal growth rate, (V, g^*), V concave and continuous, be a solution of (7), and $\epsilon > 0$. Then there exists a continuous function $\pi_c^\epsilon : \mathcal{S} \to \mathcal{K}$ such that*

$$V(\mathbf{b}) + g \leq \log w(\mathbf{b}, \pi_c^\epsilon(\mathbf{b})) + W(\pi_c^\epsilon(\mathbf{b})) + \mathbf{E}V(\pi_c^\epsilon(\mathbf{b})\circ\mathbf{X}) + \epsilon. \tag{10}$$

Since the potential function V is concave and the space of controls $\mathcal{K}$ convex, this result essentially follows from classical selection theorems. Iterating (10) along a sample path and taking expectation we have

$$\mathbf{E}\log S_n^{\pi_c^\epsilon}(\mathbf{b}) = \mathbf{E}\Big[\sum_{k=1}^{n-1}\Big(\log w(\mathbf{b}_k, \pi_c^\epsilon(\mathbf{b}_k)) + W(\pi_c^\epsilon(\mathbf{b}_k))\Big)\Big] \geq n(g^* - \epsilon) + \mathbf{E}[V(\pi_c^\epsilon(\mathbf{b}_n)\circ\mathbf{X})] - V(\mathbf{b}).$$

Dividing both sides by n and taking the limit as $n \to \infty$ we get the following corollary.

Corollary 2 *The policy π_ϵ^c that satisfies (10) is ϵ-optimal, i.e. the growth rate g_c^ϵ of π_c^ϵ satisfies $g_c^\epsilon \geq g^* - \epsilon$.*

The Markov chain induced by any continuous policy is Weak Feller continuous. Consequently there exists a set $\mathcal{F} \subseteq \mathcal{S}$ such that for all initial portfolios $\mathbf{b} \in \mathcal{F}$ the ϵ-optimal continuous policy achieves the corresponding growth rate almost surely, i.e. if the investor is allowed to perturb the initial portfolio a growth rate arbitrarily close to optimal can be achieved almost surely.

A policy π is defined to be a 1-stationary policy if it is stationary after an initial perturbation. A 1-stationary policy is represented by a pair (π_1, π_2): π_1 is the initial perturbation, i.e. at the beginning of period $n = 1$ the portfolio $\mathbf{b}$ is mapped to $\pi_1(\mathbf{b})$, and, thereafter, π_2 is employed. Thus, the portfolio held during period $n = 1$ is $\pi_2(\pi_1(\mathbf{b}_1))$. Using this notation we have the following result.

Theorem 4 *Let π_c^ϵ be a continuous ϵ-optimal policy with growth rate g_c^ϵ. Then there exists a 1-stationary policy $\pi = (\pi_1, \pi_c^\epsilon)$ such that for all $\mathbf{b} \in \mathcal{S}$*

$$\lim_{n\to\infty} \frac{1}{n} \log S_n^\pi(\mathbf{b}) = g^\epsilon \geq g^* - \epsilon, \quad a.s.$$

If the Markov chain induced by the policy π_c^ϵ has a unique invariant measure, the initial correction is not required.

It is well known that the transactions of the continuous time growth optimal policy is the local time of the portfolio process on the boundary of a closed connected set [2, 13, 26]. We establish that the discrete time growth optimal investment problem has ϵ-optimal policies with similar properties.

Theorem 5 *For all $\epsilon > 0$ there exists an ϵ-optimal policy π_l^ϵ such that*

$$\pi_l^\epsilon(\mathbf{b}) = \begin{cases} \mathbf{b} & \mathbf{b} \in \mathcal{N}, \\ \in \partial\mathcal{N} & \textit{otherwise}, \end{cases} \tag{11}$$

where $\mathcal{N}$ is a closed connected set and $\partial\mathcal{N}$ is the boundary of $\mathcal{N}$.

The no-trade set $\mathcal{N}$ in Theorem 5 is the image of the state space $\mathcal{S}$ under an $\frac{\epsilon}{4}$-optimal continuous policy, i.e. $\mathcal{N} = \pi_c(\mathcal{S})$ for an $\frac{\epsilon}{4}$-optimal continuous policy π_c. The result is established by first showing that there exists functions $f : \mathcal{S} \mapsto \mathcal{K}$ satisfying (11) and (10), and then invoking a classical measurable selection result.

The policy π_l^ϵ does not exercise any control as long as $b \in \mathcal{N}$; however, if $\mathbf{b} \notin \mathcal{N}$ the policy π_l^ϵ maps it to the boundary $\partial\mathcal{N}$. This is analogous to the continuous time solution – take action only if the portfolio threatens to leave a no-trade set and prevent it from leaving. Unfortunately a discrete time policy can only intervene at discrete instants; thus, the best it can do is map the portfolio back if it has left the no-trade set. This result, although plausible, is not obvious. For example, a case may be made that the optimal policy ought to map the portfolio to the interior of $\mathcal{N}$ to maximize the sojourn time in $\mathcal{N}$, and thereby, minimize costs.

The existence of a discrete control-limit ϵ-optimal policy has several implications. The special structure of the policy allows efficient approximation of the solution of the Bellman equation. The structure of control-limit policies opens up the possibility that the discrete time policy that corrects to the continuous time no-trade set performs close to optimal. This is especially attractive since an approximation to the continuous time no-trade set can be efficiently computed [13].

4 Two asset markets

In this section we consider the special case of two-asset markets. This special case is especially important for markets that admit a two-fund separation [23]. Since in such markets all investors hold the same two funds, the results of this section completely describe optimal investment decisions.

We begin this section with stronger versions of the results for general discrete time markets. Then we present a simple numerical scheme for computing the optimal no-trade set for discrete time markets derived by sampling continuous time Brownian markets and compare the performance of the optimal discrete time policy with the policy that corrects to the continuous time optimal no-trade interval.

In two-asset markets the portfolio $\mathbf{b} = [b, 1-b]^T \in \mathbf{R}^2$ is completely parametrized by the fraction b of the total wealth invested in the first asset. Since the portfolio vector is effectively one-dimensional, the following corollary follows from Theorem 5.

Corollary 3 *For every $\epsilon > 0$ there exists a ϵ-optimal control-limit policy π_l^ϵ with no-trade set $[\alpha, \beta] \subset \mathbf{R}$, i.e.*

$$\pi_l^\epsilon(b) = \begin{cases} \alpha, & b \leq \alpha, \\ b, & b \in [\alpha, \beta], \\ \beta, & b \geq \beta. \end{cases}$$

Clearly the control-limit policy π_l^ϵ is continuous. Consequently the Markov chain induced by this policy is Weak Feller continuous, and this in turn implies that the policy achieves its corresponding growth rate with probability one.

Theorem 6 *Fix $\epsilon > 0$ and let π_l^ϵ be an ϵ-optimal control-limit policy with growth rate g_l^ϵ. Then the wealth sequence $\{S_n^{\pi_l^\epsilon} : n \geq 1\}$ satisfies*

$$\lim_{n \to \infty} \frac{1}{n} \log S_n^{\pi_l^\epsilon} = g_l^\epsilon, \qquad a.s.$$

This result essentially follows from Theorem 4. Theorem 6 has important computational implications since it reduces the problem of computing an ϵ-optimal policy to one of estimating the two end-points of a no-trade interval. Another implication of the Theorem 6 is that ϵ-optimal policies can be very effectively approximated even when the distribution of the price relative vectors is unknown (see [11] for details).

Next we present numerical recipe for computing an ϵ-optimal no-trade interval for the discrete time two-asset market resulting from sampling a continuous time two-asset Brownian market. The study of this market is motivated by the fact that a reasonable first approximation for markets that do not allow continuous trading is to assume that the trades are allowed at uniformly distributed instants in time.

Consider a continuous time market with one bond and one stock. Assume that the price processes $(P_1(t), P_2(t))$ of the two assets are given by

$$\begin{aligned} dP_1(t) &= 0, \\ dP_2(t) &= P_2(t)(\mu dt + \sigma dB(t)), \end{aligned} \tag{12}$$

where $B(t)$ is standard Brownian motion. Integrating (12) we get

$$\begin{aligned} P_1(t) &= P_0(0), \\ P_2(t) &= P_1(0) \exp\Big(\tilde{\mu} t + \sigma B(t)\Big), \end{aligned}$$

where $\tilde{\mu} = \mu - \frac{1}{2}\sigma^2$. The price relative vector for the discrete time market obtained by sampling the price processes every δ time units is given by:

$$\begin{aligned} X_1 &= 1, \\ X_2 &= \exp\Big(\tilde{\mu}\delta + \sigma\sqrt{\delta}Z\Big), \end{aligned} \tag{13}$$

where $Z = \pm 1$ with equal probability. This discrete time market model is identical to the binomial model popular in asset pricing and its convergence to its continuous time analog can be put on firm mathematical foundation by invoking the Invariance Principle [5].

For ease of computation we parameterize the investor's position by the ratio p of the wealth in the risky asset to that in cash. The portfolio $\mathbf{b}$ is then simply given by $\mathbf{b} = \frac{1}{1+p}[p, 1]^T$. The steps in the proposed numerical recipe are as follows:

1. **Compute optimal ratio p_0^* in the absence of costs**:

$$\begin{aligned} p_0^* &= \text{argmax}_{p\in\mathbf{R}} W(p), \\ &= \text{argmax}_{p\in\mathbf{R}} \left\{ \frac{1}{2}\log\left(\frac{1+pe^{\tilde{\mu}\delta+\sigma\sqrt{\delta}}}{1+p}\right) + \frac{1}{2}\log\left(\frac{1+pe^{\tilde{\mu}\delta-\sigma\sqrt{\delta}}}{1+p}\right)\right\}. \end{aligned}$$

2. **Evaluate growth rate for a given no-trade interval** $[\alpha, \beta]$: The distribution (13) of the price relative vector ensures that the ratio p will leave $[\alpha, \beta]$ in either direction with probability 1. Thus, without loss of generality, one can assume that the initial state $p_1 = \alpha$. Since p is corrected to β if $p > \beta$ and to the ratio α if $p < \alpha$, we have that the state space

$$\mathcal{S} = \left\{\alpha_i \equiv \alpha e^{i(\tilde{\mu}\delta+\sigma\sqrt{\delta})} : -1 \le i \le i_{\max}\right\} \cup \left\{\beta_j \equiv \beta e^{-j(-\tilde{\mu}\delta+\sigma\sqrt{\delta})} : -1 \le j \le j_{\max}\right\},$$

where

$$i_{\max} = \left\lceil \frac{1}{\tilde{\mu}\delta + \sigma\sqrt{\delta}} \log\left(\frac{\beta}{\alpha}\right)\right\rceil \quad \text{and} \quad j_{\max} = \left\lceil \frac{1}{-\tilde{\mu}\delta + \sigma\sqrt{\delta}} \log\left(\frac{\beta}{\alpha}\right)\right\rceil.$$

The transition probability P of the induced Markov chain is given by

$$\begin{array}{l} P(s \mid \alpha_{-1}) = \frac{1}{2}\mathbf{1}(s = \alpha_{-1}) + \frac{1}{2}\mathbf{1}(s = \alpha_2), \\ P(s \mid \alpha_i) = \frac{1}{2}\mathbf{1}(s = \alpha_{i-1}) + \frac{1}{2}\mathbf{1}(s = \alpha_{i+1}), \quad i = 0, \ldots, i_{\max} - 1, \\ P(s \mid \alpha_{i_{\max}}) = \frac{1}{2}\mathbf{1}(s = \beta_{-1}) + \frac{1}{2}\mathbf{1}(s = \beta_2), \\ P(s \mid \beta_{-1}) = \frac{1}{2}\mathbf{1}(s = \beta_{-1}) + \frac{1}{2}\mathbf{1}(s = \beta_2), \\ P(s \mid \beta_j) = \frac{1}{2}\mathbf{1}(s = \beta_{i-1}) + \frac{1}{2}\mathbf{1}(s = \beta_{i+1}), \quad j = 0, \ldots, j_{\max} - 1, \\ P(s \mid \beta_{j_{\max}}) = \frac{1}{2}\mathbf{1}(s = \alpha_{-1}) + \frac{1}{2}\mathbf{1}(s = \alpha_2). \end{array}$$

Since the matrix P is irreducible and finite, there exists a unique invariant measure μ. The one-step reward function $r : \{\alpha_i : -1 \le i \le i_{\max}\} \cap \{\beta_j : -1 \le j \le j_{\max}\} \to \mathbf{R}$ is given by

$$r(s) = \begin{cases} W(s), & s \notin \{\alpha_{-1}, \alpha_{i_{\max}}, \beta_{-1}, \beta_{i_{\max}}\}, \\ W(s) + \log w(s, \alpha_0), & s = \alpha_{-1}, \\ W(s) + \log w(s, \beta_0), & s = \alpha_{i_{\max}}, \\ W(s) + \log w(s, \beta_0), & s = \beta_{-1}, \\ W(s) + \log w(s, \alpha_0), & s = \beta_{i_{\max}}. \end{cases}$$

Thus, the normalized growth rate $g_{\alpha\beta}$ associated with the interval $[\alpha, \beta]$ is given by

$$g_{\alpha\beta} = \sum_{i=1}^{i_{\max}} \mu(\alpha_i) r(\alpha_i) + \sum_{j=1}^{j_{\max}} \mu(\beta_j) r(\beta_j).$$

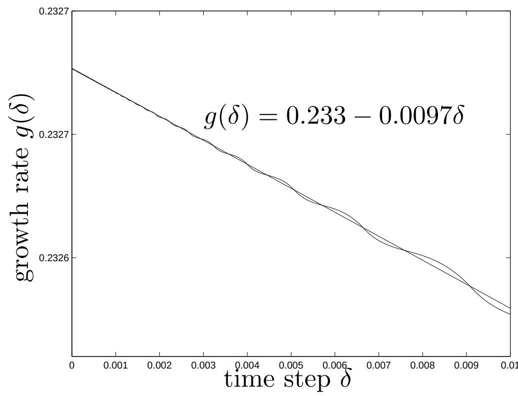

Figure 2: Growth rate vs the time step δ

3. **Optimize over** $[\alpha, \beta]$: We discretized the real line and searched outwards from the no-cost optimal portfolio p_0^*. The search was stopped when the associated growth rate started to decline. The optimal interval $[\alpha^*, \beta^*]$ in terms of the ratio p corresponds to the optimal no-trade interval $\left[\frac{\alpha^*}{1+\alpha^*}, \frac{\beta^*}{1+\beta^*}\right]$ in terms of b.

For our numerical experiments we set the parameters $\mu = \frac{1}{2}\sigma^2 = \frac{1}{2}$, i.e. $P_1(t) = P_1(0)$ and $P_2(t) = P_2(0)\exp(B(t))$. There were no transaction costs on the bond and a 1% transaction cost on the risky asset, i.e. $\lambda_2 = 0.01$. Figure 2 displays the growth rate $g(\delta)$ as a function of the sampling interval δ. Clearly, $g(\delta)$ is approximately a linear function of δ with the best-fit line given by

$$g(\delta) = 0.233 - 0.0097\delta. \tag{14}$$

In our numerical experiments we found that the no-trade interval was always symmetric with respect to the no-cost optimal portfolio $b_0^* = \frac{p_0^*}{1+p_0^*} = 0.5$. Figure 3 displays the length $l(\delta)$ of the no-trade interval as a function of δ. The best-fit curve to $l(\delta)$ is given by

$$l(\delta) = 0.198 - 0.318\delta^{(0.491)}. \tag{15}$$

According to (15) the length of the no-trade interval in the continuous limit, i.e. when $\delta = 0$, ought to be $l(0) = 0.198$. This agrees very closely to the no-trade interval length $l(0) = 2$ calculated by a numerical scheme given by Taksar, Klass and Assaf [26].

We next investigate the dependence of the optimum discrete time growth rate on the transaction cost. For this set of experiments the sampling interval was set to $\delta = (\ln(2))^2 = 0.49$, which resulted in the price relative of the risky asset X_2 taking values $\{2, 1/2\}$ with equal probability. Figure 4 and Figure 5 show the dependence of respectively the optimal growth rate and no-trade interval on the transaction cost λ. The best-fit curve for the growth rate is

$$g(\lambda) = 0.059 - 0.042\lambda^{0.75}$$

and the best-fit curve for the no-trade interval is

$$x(\lambda) = 0.5 \pm 0.273\lambda^{0.5}$$

It is clear that the dependence of the growth rate and the length of the no-trade interval on λ is different from that observed in the continuous time problem (see the Appendix in [24]). We are still

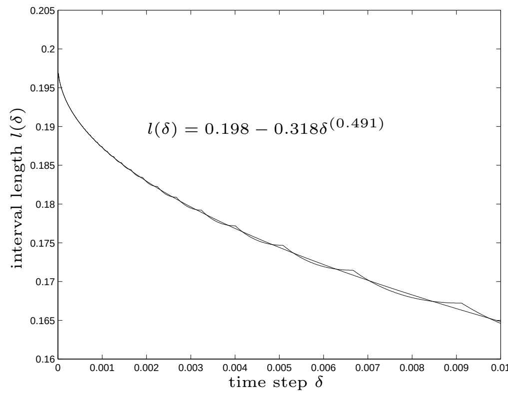

Figure 3: No-trade interval length vs the time step δ

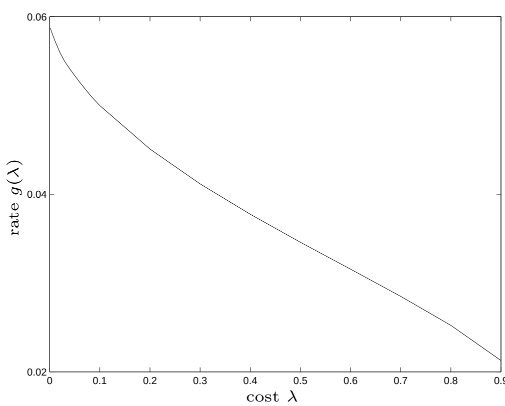

Figure 4: Growth rate vs the transaction cost λ

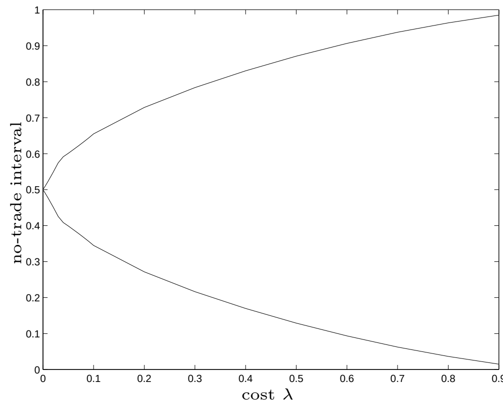

Figure 5: No-trade interval vs the cost λ

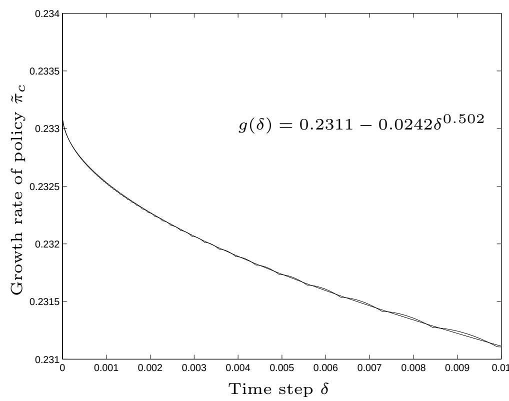

Figure 6: Growth rate of policy $\tilde{\pi}_c$ vs time step δ

investigating analytical expansions that are able to capture the first order behavior of the discrete optimal policies as a function of the sampling interval δ and the transaction cost λ.

The rest of this section studies the performance of the control-limit policy $\tilde{\pi}_c$ that corrects to the optimal continuous time no-trade interval. We are interested in this policy because for small values of δ the complexity of discrete time numerical scheme proposed above is significantly higher than that of the numerical algorithm for computing the optimal continuous time no-trade interval. If it were the case that the performance of the sub-optimal policy $\tilde{\pi}_c$ degrades gracefully with δ, employing this sub-optimal policy would be justified on the basis of the computational cost.

We computed the optimal continuous time no-trade interval using the numerical recipe proposed in Section 7 of [26]. For $\mu = \frac{1}{2}\sigma^2 = \frac{1}{2}$ the optimal no-trade interval is

$$[\alpha, \beta] = [0.4106, 0.5894], \tag{16}$$

and the associated continuous time growth rate

$$g_c = 0.2331.$$

Figure 6 displays the growth rate of the sub-optimal discrete time policy $\tilde{\pi}_c$ that corrects to the optimal continuous time no-trade interval given in (16). The growth rate $\tilde{g}_c(\delta)$ of the sub-optimal policy $\tilde{\pi}_c$ has a square-root decay; thus, $\tilde{g}_c(\delta)$ decays considerably faster than $g(\delta)$. On the more optimistic front, however, $g(0) \approx \tilde{g}(0) > 0$. Thus, both the relative and absolute errors of the sub-optimal policy $\tilde{\pi}_c$ goes to zero as $\delta \to 0$.

5 Conclusion

In this work we introduced growth optimal investment for discrete time markets with proportional transaction costs and IID market structure. Growth optimal investment has been well studied in continuous time markets [2, 13, 26]. The optimal investment policy in continuous time markets has a local-time character, i.e. trades occur only if the portfolio hits the boundary of a no-trade set. Such a policy is clearly infeasible in markets that do not allow continuous trading. In this work we explicitly introduced discrete decision instances and investigated the effect on the growth optimal policies. We established that there exists a stationary growth optimal investment policy and a 1-stationary policy that grows at a rate arbitrarily close to optimal rate almost surely. More significantly, we establish that for all $\epsilon > 0$ there exist ϵ-optimal control-limit policies. A control-limit policy trades only if the portfolio leaves a no-trade set; and in that case the policy corrects the portfolio to the boundary of the no-trade set. Thus, control-limit policies are the discrete analog of the local-time continuous time policies.

Next, we focused on the special case of two-asset markets. For these markets we established that the growth rate of the ϵ-optimal control-limit policies is achieved with probability 1. We developed a simple numerical scheme for calculating the optimal no-trade interval for discrete time markets derived by sampling a Brownian market. Our numerical calculations show that the optimal discrete policy has a growth rate that decreasing linearly with the sampling interval δ, whereas the growth rate of the sub-optimal policy that corrects to the optimal continuous time no-trade interval decays as $\sqrt{\delta}$. Moreover, there is no discontinuity at $\delta = 0$. Thus, for high sampling rates, i.e. small δ, the policy that corrects to the continuous optimal interval performs surprisingly close to optimal.

There are several weaknesses in the market model considered in this work. The transaction costs in a real market are not just proportional – at the very least there is a fixed cost per transaction. In an upcoming publication we study investment in a continuous time market with fixed costs [10]. The optimal investment policy in such a market trades only at discrete instants in time. We believe a discretized version of the policy will be optimal for discrete time markets with fixed costs. Another limitation of our model is that we restrict the investor to policies that ensure that the investor is never leaves the set of admissible states. While this is a reasonable approximation, a more complete model would allow the investor the option of introducing new capital with a penalty. Finally we do not allow for the possibility of investing in assets that have only a finite lifetime, e.g. options.

References

[1] M. Akian, J. L. Menaldi, and A. Sulem. On an investment-consumption model with transaction costs. *SIAM J. Control Optim.*, 34(1):329–364, 1996.

[2] M. Akian, A. Sulem, and M. I. Taksar. Dynamic optimization of long term growth rate for a mixed portfolio with transactions costs. Preprint.

[3] P. H. Algoet and T. M. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *Ann. Prob.*, 16(2):876–898, 1988.

[4] D. P. Bertsekas and S. Shreve. *Stochastic dynamic programming: the discrete time case.* Athena Scientific, Bellmont, MA, 1996.

[5] P. Billingsley. *Probability and measure.* John-Wiley & Sons, Inc., 3rd edition, 1995.

[6] L. Breiman. Optimal gambling systems for favourable games. In *Proc. Fourth Berkeley Symp. Math. Statist. Probab.*, 1. Univ. California Press, 1961.

[7] M. H. A. Davis and A. R. Norman. Portfolio selection with transaction costs. *Math. Oper. Res.*, 15(4):676–713, 1990.

[8] E. R. Fernholz. *Stochastic portfolio theory.* Springer-Verlag, New York, 2000.

[9] A. Hordijk. A sufficient condition for existence of an optimal policy with respect to the average cost criterion in Markov decision processes. In *Trans. Sixth parague Conf. Information Theory, Statistical Decision Functions, Random Processes*, pages 263–174, 1973.

[10] G. Iyengar. Optimal investment in markets with fixed transaction costs. *In preparation.*

[11] G. Iyengar. Robust investment in two-asset markets with transaction costs. *In preparation.*

[12] G. Iyengar and T. M. Cover. Growth optimal investment in horse race markets with costs. *IEEE Trans. Info. Theory*, 46, 2000.

[13] G. N. Iyengar. Analysis of growth rate in continuous time markets with transaction costs. Technical report, Stanford University, 1997.

[14] A. Kalai and A. Blum. Universal portfolios with and without transaction costs. In *Proc. 10th COLT*, pages 309–313, 1997.

[15] J. L. Kelly. A new interpretation of information rate. *Bell Syst. Tech. J.*, 35:917–926, 1956.

[16] A. Kucia and A. Nowak. On ϵ-optimal continuous selectors and their application in discounted dynamic programming. *JOTA*, 54(2):289–302, 1987.

[17] M. J. P. Magill and G. M. Constantinidis. Portfolio selection with transactions costs. *J. Econ. Theory*, 13:245–263, 1976.

[18] H. M. Markowitz. Portfolio selection. *J. Finance*, 7:77–91, 1952.

[19] S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability.* Springer Verlag, 1993.

[20] M. L. Puterman. *Markov Decision Processes: discrete stochastic dynamic programming.* John Wiley & Sons, Inc., New York, NY, 1994.

[21] M. Schäl. Stationary policies in dynamic programming under compactness assumptions. *Math. Oper. Res.*, 8:366–372, 1983.

[22] M. Schäl. Average optimality in dynamic programming. *Math. of Operations Research*, 18:163–172, 1993.

[23] W. Sharpe. *Portfolio theory and capital markets.* McGraw Hill, 1970.

[24] S. E. Shreve and H. M. Soner. Optimal investment and consumption with transaction costs. *Ann. Appl. Probab.*, 4(3):609–692, 1994.

[25] R. Strauch. Negative dynamic programming. *Ann. Math. Stat.*, 37:871–890, 1966.

[26] M.I. Taksar, M. J. Klass, and D. Assaf. Diffusion model for optimal portfolio selection in the presence of brokerage fees. *Math. Oper. Res.*, 13(2):277–294, 1988.

Appendix

Proof of Lemma 1

Let w be the wealth at the portfolio $\mathbf{z}$ realized by any self-financing trade starting from 1 dollar invested in portfolio $\mathbf{b}$. Let c be the associated cost. Then $w + c = 1$.

The cost $c \geq \sum_{i=1}^m \lambda_i |wz(i) - b(i)|$ since the latter sum does not account for any opposing trades. It follows that

$$w + \sum_{i=1}^m \lambda_i |wz(i) - b(i)| \leq 1.$$

This proves that the wealth w realized by any trading strategy is feasible for the convex program

$$\begin{array}{ll} \text{maximize} & w, \\ \text{subject to} & w + \sum_{i=1}^m \lambda_i |wz(i) - b(i)| \leq 1. \end{array} \tag{17}$$

Let w^* be the optimal value of (17). Continuity of the absolute value function implies that

$$w^* + \sum_{i=1}^m \lambda_i |w^* z(i) - b(i)| = 1. \tag{18}$$

Let $S = \{j \mid w^* z(j) < b(j)\}$ and $B = \{j \mid w^* z(j) > b(j)\}$. Then the equality (18) can be rewritten as

$$\sum_{j \in S} (1 - \lambda_j)(b_j - w^* z(j)) = \sum_{j \in B} (1 + \lambda_j)(w^* z(j) - b(j)). \tag{19}$$

Consider the following trade: For $j \in S$ sell $b(j) - w^* z(j)$ dollars worth of the asset j. This results in $\sum_{i \in S} (1 - \lambda_i)(b(i) - w^* z(i))$ dollars in cash. Use this capital to purchase $w^* z(j) - b(j)$ worth of asset j for each $j \in B$. From (19) it follows that this trade is feasible. Moreover, at the end of the trade we have w^* dollars invested in the portfolio $\mathbf{z}$. This completes the proof.

To prove the lower bound on $w(\mathbf{b}, \mathbf{z})$, we employ the following naive trading strategy. First we sell all our holdings to get,

$$\sum_{i=1}^m (1 - \lambda_i) b^+(i) - \sum_{i=1}^m (1 + \lambda_i) b^-(i),$$

in cash. We next invest this wealth in the assets to get the final portfolio $\mathbf{z}$. The final wealth w realized by this trading strategy is given by

$$w \;\; = \;\; \frac{\sum_{i=1}^m (1 - \lambda_i) b^+(i) - \sum_{i=1}^m (1 + \lambda_i) b^-(i)}{1 + \sum_{i=1}^m \lambda_i |z(i)|}.$$

Since $\mathcal{S}$ is compact, it follows that $Z = \max_{(z \in \mathcal{S})} \{1 + \sum_{i=1}^m \lambda_i |z(i)|\} < \infty$. Also, $\mathbf{b} \in \mathcal{S}$ implies that $\sum_{i=1}^m (1 - \lambda_i) b^+(i) \geq M \sum_{i=1}^m (1 + \lambda_i) b^-(i)$. Therefore,

$$\begin{aligned} \sum_{i=1}^m (1 - \lambda_i) b^+(i) - \sum_{i=1}^m (1 + \lambda_i) b^-(i) \;\; &\geq \;\; \left(1 - \frac{1}{M}\right) \sum_{i=1}^m (1 - \lambda_i) b^+(i), \\ &\geq \;\; \left(1 - \frac{1}{M}\right) (1 - \lambda_{\max}) \sum_{i=1}^m b^+(i), \\ &\geq \;\; \left(1 - \frac{1}{M}\right) (1 - \lambda_{\max}) > 0. \end{aligned}$$

Thus, we have that for all $\mathbf{b}, \mathbf{z} \in \mathcal{S}$

$$w(\mathbf{b}, \mathbf{z}) \geq \left(1 - \frac{1}{M}\right)\left(\frac{1-\lambda_{\max}}{Z}\right). \tag{20}$$

■

Proof of Corollary 1

Consider the policy π that initially maps $\mathbf{b}$ to a given portfolio $\mathbf{y}$ and then invests according to a policy π_y admissible for $\mathbf{y}$. Then $\log S_n^\pi(\mathbf{b}) = \log\left(w(\mathbf{b}, \mathbf{y}) S_n^{\pi_y}(\mathbf{y})\right)$. Consequently,

$$\begin{aligned} g(\mathbf{b}) &= \sup_{\pi \in \Pi(\mathbf{b})} \liminf_{n\to\infty} \frac{1}{n} \mathbf{E} \log S_n^\pi(\mathbf{b}) \\ &\geq \liminf_{n\to\infty} \frac{1}{n} \log S_n^\pi(\mathbf{b}) \\ &= \liminf_{n\to\infty} \frac{1}{n} \log w(\mathbf{b}, \mathbf{y}) + \liminf_{n\to\infty} \frac{1}{n} \log S_n^{\pi_y}(\mathbf{y}), \\ &= \liminf_{n\to\infty} \frac{1}{n} \log S_n^{\pi_y}(\mathbf{y}). \end{aligned}$$

Taking the supremum over all feasible policies π_y we get

$$g(\mathbf{b}) \geq \sup_{\pi_y \in \Pi(\mathbf{y})} \liminf_{n\to\infty} \frac{1}{n} \log S_n^{\pi_y}(\mathbf{y}) = g(\mathbf{y})$$

Since $\mathbf{b}$ and $\mathbf{y}$ were arbitrary, it follows that $g(\mathbf{b})$ is constant over $\mathcal{S}$. ■

Proof of Theorem 1

Let $\pi \in \Pi(\mathbf{b})$ and let $\{\mathbf{b}_n : n \geq 1\}$ be a sample path generated by π. Then (7) implies that

$$\mathbf{E}[V(\mathbf{b}_k)] + g \geq \mathbf{E}\Big[\log w((\mathbf{b}_k, \pi(\mathbf{b}_k)) + W(\pi(\mathbf{b}_k))\Big] + \mathbf{E}V(\mathbf{b}_{k+1}), \qquad \forall k \geq 1. \tag{21}$$

Iterating (21) n times we get

$$\mathbf{E}\left[\frac{1}{n}\sum_{k=1}^{n-1}\Big(\log w(\mathbf{b}_k, \pi(\mathbf{b}_k)) + W(\pi(\mathbf{b}_k))\Big)\right] \leq g + \frac{1}{n}\left(V(\mathbf{b}) - \mathbf{E}V(\mathbf{b}_n)\right).$$

Since V is bounded, it immediately follows that g is an upper bound on the achievable growth rate.

■

Proof of Theorem 2

We first prove the uniqueness result and then establish that there exists a solution (V, g^*) of (7) with V concave and continuous.

Suppose (V_i, g_i), $i = 1, 2$, are any two solutions of (7) with V_i, $i = 1, 2$, continuous and concave. It is easy to establish that V continuous implies that $f(\mathbf{b}, \mathbf{z}) = \log w(\mathbf{b}, \mathbf{z}) + W(\mathbf{z}) + \mathbf{E}V(\mathbf{z} \circ \mathbf{X})$ is continuous on $\mathcal{S} \times \mathcal{K}$. The continuity of $\log(w(\mathbf{b}, \mathbf{z}))$ follows from a triangle argument, whereas the continuity of $W(\mathbf{z})$ and $\mathbf{E}V(\mathbf{z} \circ \mathbf{X})$ follows from the definition of $\mathcal{K}$.

Thus, Proposition 7.33 in [4] implies the existence of measurable functions $\pi_i : \mathcal{S} \mapsto \mathcal{K}$, $i = 1, 2$, such that, for all $\mathbf{b} \in \mathcal{S}$,

$$\pi_i(\mathbf{b}) \in \operatorname*{argmax}_{\mathbf{z}\in\mathcal{K}} \Big\{ w(\mathbf{b}, \mathbf{z}) + W(\mathbf{b}) + \mathbf{E}[V_i(\mathbf{z} \circ \mathbf{X})\Big\}.$$

Then an argument similar to one employed in the proof of Theorem 1 implies that for all $\mathbf{b} \in \mathcal{S}$,

$$g_i = \liminf_{n\to\infty} \mathbf{E}\Big[\frac{1}{n} \log(S_n^{\pi_i}(\mathbf{b}))\Big], \qquad i = 1, 2. \tag{22}$$

Since Theorem 1 implies that each g_i, $i = 1, 2$, is an upper bound on the achievable growth rate, we have from (22) that $g_1 = g_2$. This completes the proof of uniqueness. Note that we only needed the existence of a measurable selection to establish uniqueness. Since such a selection exists if V is upper semi-continuous, the uniqueness result extends to this case.

Next, we establish that the Bellman equation (7) has a solution (V, g) with V concave and continuous. Define the discounted reward function J_β as follows:

$$J_\beta(\pi, \mathbf{b}) = \mathbf{E}\left[\sum_{k\geq 1} \beta^{k-1} \Big(\log w(\mathbf{b}, \pi(\mathbf{b})) + W(\pi(\mathbf{b}))\Big)\right], \qquad 0 < \beta < 1. \tag{23}$$

Define the discounted potential $V_\beta(\mathbf{b})$ as follows:

$$V_\beta(\mathbf{b}) = \sup_{\pi\in\Pi(\mathbf{b})} J_\beta(\pi, \mathbf{b}).$$

Let $m_\beta = \sup_{\mathbf{x}\in\mathcal{S}} V_\beta(\mathbf{x})$ and $\overline{g} = \limsup_{n\to\infty}(1-\beta)m_\beta$. Since $\overline{g} \leq \max_{\mathbf{b}\in\mathcal{S}} W(\mathbf{b}) < \infty$, it is well defined. From the well-known Tauberian relationship (c.f. [9]) we have

$$g = \sup_{\pi\in\Pi(\mathbf{b})} \liminf_{n\to\infty} \frac{1}{n}\mathbf{E}\log S_n^\pi \leq \liminf_{\beta\to 1}(1-\beta)m_\beta \leq \limsup_{\beta\to 1}(1-\beta)m_\beta = \overline{g}$$

We next show the existence of a convex continuous function V and a stationary, non-random policy π^* such that

$$V(\mathbf{b}) + \overline{g} \leq w(\mathbf{b}, \pi^*(\mathbf{b})) + W(\pi^*(\mathbf{b})) + \mathbf{E}V(\pi^*(\mathbf{b}) \circ \mathbf{X})$$

From the conclusion of Theorem 1, it follows that $g = \overline{g}$ and the policy π^* is optimal. This proof technique closely parallels that of Schäl [22].

We first prove that the discounted problem has a stationary optimal policy.

Lemma 4 *For all $0 < \beta < 1$ there exists a discount optimal policy $\pi_\beta^* : \mathcal{S} \mapsto \mathcal{K}$ which satisfies*

$$V_\beta(\mathbf{x}) = \log w(\mathbf{x}, \pi_\beta^*(\mathbf{x})) + W(\pi_\beta^*(\mathbf{x})) + \beta\mathbf{E}V_\beta(\pi_\beta^*(\mathbf{x}) \circ X), \qquad \forall \mathbf{x} \in \mathcal{S}. \tag{24}$$

Proof: Since the action space $\mathcal{A}(\mathbf{x}) = \mathcal{K}$ for all $\mathbf{x} \in \mathcal{S}$ and $w(\mathbf{x}, \mathbf{z}) + W(\mathbf{z})$ is continuous on $\mathcal{S} \times \mathcal{K}$, the existence of π_β would immediately follow from the results in [21] if the transition distribution is continuous with respect to weak convergence.

For $\mathbf{b} \in \mathcal{S}$ and $\mathbf{z} \in \mathcal{K}$ let $\mathbf{P}(\cdot \mid \mathbf{b}, \mathbf{z})$ be the distribution of $\mathbf{z} \circ \mathbf{X}$, i.e. the one-step transition distribution when the investor readjusts the portfolio to $\mathbf{z}$. Let $(\mathbf{b}_k, \mathbf{z}_k) \to (\mathbf{b}, \mathbf{z})$ and f be a continuous bounded function on $\mathcal{S}$. Then $\int_{\mathcal{S}} \mathbf{x}\mathbf{P}(d\mathbf{x} \mid \mathbf{b}_k, \mathbf{z}_k) = \mathbf{E}f(\mathbf{z}_n \circ \mathbf{X})$.

Fix $\mathcal{X} \subset \mathbf{R}^+$ such that $\mathbf{P}(\mathcal{X}) = 1$ and

$$\sum_{i=1}^{m}(1+\lambda_i)X_i z_k^-(i) \leq M \sum_{i=1}^{m}(1-\lambda_i)X_i z_k^+(i), \qquad \mathbf{X} \in \mathcal{X},\ k \geq 1. \tag{25}$$

The bound (25) implies that $\mathbf{z}_k^T\mathbf{X} > 0$ for all $\mathbf{X} \in \mathcal{X}$ and $k \geq 1$. Also $\{g_k(\mathbf{X}) = f(\mathbf{z}_k \circ \mathbf{X}) : k \geq 1\}$ is uniformly bounded and $\lim_{k\to\infty} g_k(\mathbf{X}) = f(\mathbf{z} \circ \mathbf{X})$ for all $X \in \mathcal{X}$.

From the bounded convergence theorem it follows that

$$\begin{aligned}
\lim_{k\to\infty} \mathbf{E} f(\mathbf{z}_k \circ \mathbf{X}) &= \lim_{k\to\infty} \mathbf{E}[g_k(\mathbf{X}); \mathcal{X}] \\
&= \mathbf{E}[f(\mathbf{z} \circ \mathbf{X}); \mathcal{X}] \\
&= \int_{\mathcal{S}} x\mathbf{P}(d\mathbf{x} \mid \mathbf{b}, \mathbf{z}),
\end{aligned}$$

i.e. the transition density is continuous with respect to weak convergence. ■

The regularity properties of V_β needed to establish the existence of the potential function V are collected together in the following lemma.

Lemma 5 *The optimal discounted reward function V_β has the following properties:*

1. *$(1-\beta)V_\beta$ is uniformly bounded for all $\beta \in (0,1)$,*
2. *V_β is a concave function for all $\beta \in (0,1)$,*
3. *$\Big\{(1-\beta)V_\beta : 0 < \beta < 1\Big\}$ is an equi-continuous family.*

Proof: Let $R = \max_{\mathbf{b}\in\mathcal{S}} W(\mathbf{b}) < \infty$. Since $w(\mathbf{b},\mathbf{z}) \leq 1$ for all $\mathbf{b},\mathbf{z} \in \mathcal{S}$, (23) implies that

$$J_\beta(\pi, \mathbf{b}) \leq \frac{R}{1-\beta}.$$

Thus, $(1-\beta)V_\beta(\mathbf{b}) \leq R$ for all $\mathbf{b} \in \mathcal{S}$.

To prove a uniform lower bound on $(1-\beta)V_\beta$ consider the policy π_x that rebalances the portfolio to a fixed portfolio $\mathbf{x}$ at the beginning of every market period. Then, using the uniform lower bound (20) on the rebalance factor $w(\mathbf{b},\mathbf{z})$ we get

$$\begin{aligned}
V_\beta(\mathbf{b}) \geq J_\beta(\pi_x, \mathbf{b}) &= \mathbf{E}\Big[\sum_{k=1}^{\infty} \beta^{k-1}\Big(\log w(\mathbf{b}_k, \mathbf{x}) + W(\mathbf{x})\Big)\Big], \\
&\geq \sum_{k=1}^{\infty} \beta^{k-1}\bigg(\log\Big(\Big(1 - \frac{1}{M}\Big)\Big(\frac{1-\lambda_{\max}}{Z}\Big)\Big) + W(\mathbf{x})\bigg), \\
&= \frac{1}{1-\beta}\bigg(\log\Big(\Big(1 - \frac{1}{M}\Big)\Big(\frac{1-\lambda_{\max}}{Z}\Big)\Big) + W(\mathbf{x})\bigg).
\end{aligned}$$

The convexity of V_β follows from the concavity of the logarithm and the convexity of the space of admissible policies.

An argument similar to one in the proof of Corollary 1 establishes that

$$|V_\beta(\mathbf{b}) - V_\beta(\mathbf{z})| \leq \max\{-\log w(\mathbf{b},\mathbf{z}), -\log w(\mathbf{z},\mathbf{b})\}.$$

The equi-continuity of the family $\{V_\beta\}$ immediately follows from the fact that $w(\cdot,\cdot)$ is uniformly continuous on $\mathcal{S} \times \mathcal{S}$. ■

Define $\tilde{V}_\beta(\mathbf{b}) = V_\beta(\mathbf{b}) - m_\beta$, $\mathbf{b} \in \mathcal{S}$. From Lemma 4, it follows that there exists a measurable function π_β such that

$$\tilde{V}_\beta(\mathbf{b}) + (1-\beta)m_\beta = w(\mathbf{b}, \pi_\beta(\mathbf{b})) + W(\pi_\beta(\mathbf{b})) + \beta \mathbf{E}[\tilde{V}_\beta(\pi_\beta(\mathbf{b}) \circ X)].$$

Fix a sequence $\beta_k \to 1$ such that

$$(1-\beta_k)m_{\beta_k} \to \overline{g}, \tag{26}$$

and

$$\tilde{V}_{\beta_k} \to V, \tag{27}$$

where V is a bounded continuous function and the convergence is uniform. Since each $\tilde{V}_{\beta_k}$ is concave, it follows that V is concave.

Since $\mathcal{K}$ is compact, it follows from [21] Lemma 4 that there exists a measurable function $\pi : \mathcal{S} \to \mathcal{K}$ such that $\pi(\mathbf{b})$ is an accumulation point of $\{\pi_{\beta_k}(\mathbf{b})\}$ for all $\mathbf{b} \in \mathcal{S}$. Fix $\mathbf{b} \in \mathcal{S}$ and choose a subsequence $\{k'\} \subset \{k\}$ such that $\pi_{\beta_{k'}}(\mathbf{b}) \to \pi(\mathbf{b})$. Since $\log w(\cdot,\cdot) + W(\cdot)$ is continuous, it follows that

$$\log w(\mathbf{b}, \pi_{\beta_{k'}}(\mathbf{b})) + W(\pi_{\beta_{k'}}(\mathbf{b})) \to \log w(\mathbf{b}, \pi(\mathbf{b})) + W(\pi(\mathbf{b}). \tag{28}$$

In Lemma 4 it was established that the transition density is continuous with respect to weak convergence; therefore

$$\limsup_{n\to\infty} \mathbf{E}[\tilde{V}_{\beta_{k'}}(\pi_{\beta_{k'}}(\mathbf{b}) \circ X)] = \mathbf{E}[V(\pi(\mathbf{b}) \circ \mathbf{X})]. \tag{29}$$

From (26)-(29) it follows that

$$V(\mathbf{b}) + \overline{g} \le \log w(\mathbf{b}, \pi(\mathbf{b})) + W(\pi(\mathbf{b})) + \mathbf{E}[V(\pi(\mathbf{b}) \circ \mathbf{X})].$$

This completes the proof. ■

Proof of Theorem 3

Fix $\epsilon > 0$. Define conic embeddings of the sets $\mathcal{K}$ and $\mathcal{S}$ in $\mathbf{R}^{m+1}$ as follows:

$$\widehat{S} \equiv \{w\mathbf{b} \mid w \ge 0, \mathbf{b} \in \mathcal{S}\} = \Big\{\widehat{\mathbf{b}} \;\Big|\; \sum_{i=1}^m (1-\lambda_i)\widehat{b}^+(i) \ge M \sum_{i=1}^m (1+\lambda_i)\widehat{b}^-(i)\Big\},$$

$$\widehat{\mathcal{K}} \equiv \{w\mathbf{z} \mid w \ge 0, \mathbf{z} \in \mathcal{K}\} = \Big\{\widehat{\mathbf{z}} \;\Big|\; \sum_{i=1}^m (1-\lambda_i)X_i\widehat{z}^+(i) \ge M \sum_{i=1}^m (1+\lambda_i)X_i\widehat{z}^-(i) \text{a.s.}\Big\}.$$

For all $\mathbf{b} \in \widehat{S}$ and $0 < \beta < 1$, let

$$\begin{aligned} \widehat{V}_\beta(\mathbf{b}) &= \sup_{\pi \in \Pi_M(\mathbf{b})} \mathbf{E}\left\{\sum_{k=1}^\infty \beta^{k-1} \log\left(\frac{S^\pi_{k+1}}{S^\pi_k}\right) + \log(S^\pi_1)\right\}, \\ &= \sup_{\pi \in \Pi_M(\mathbf{b})} \mathbf{E}\left\{(1-\beta)\sum_{k=1}^\infty \beta^{k-1} \log S^\pi_{k+1}\right\}. \end{aligned}$$

It is easy to establish that the value function $\widehat{V}_\beta$ is concave and

$$\widehat{V}_\beta(\mathbf{b}) = V_\beta\left(\frac{\mathbf{b}}{\mathbf{1}^T\mathbf{b}}\right) + \log(\mathbf{1}^T\mathbf{b}), \qquad \forall \mathbf{b} \in \widehat{S}.$$

Thus, for all $\beta \in (0,1)$, $\widehat{V}_\beta(\mathbf{b}) - m_\beta = \tilde{V}_\beta\left(\frac{\mathbf{b}}{\mathbf{1}^T\mathbf{b}}\right) + \log(\mathbf{1}^T\mathbf{b})$ is concave on $\widehat{S}$. By taking the limit over a suitable subsequence $\beta_k \to 1$ (independent of $\mathbf{b}$), we have that

$$\widehat{V}(\mathbf{b}) = V\left(\frac{\mathbf{b}}{\mathbf{1}^T\mathbf{b}}\right) + \log(\mathbf{1}^T\mathbf{b}) \tag{30}$$

is also concave on $\widehat{S}$.

From (30) it follows that $\log w(\mathbf{b},\mathbf{z}) + W(\mathbf{z}) + \mathbf{E}V(\mathbf{z}\circ\mathbf{X}) = \mathbf{E}\widehat{V}(\mathbf{z}\cdot\mathbf{X})$, where $(\mathbf{z}\cdot\mathbf{X})(i) = z(i)X_i$, $i = 1,\ldots,m$. Thus, the Bellman equation (7) can be rewritten as

$$V(\mathbf{b}) + g = \max_{\{\mathbf{z}\in Z(\mathbf{b})\}} \mathbf{E}\widehat{V}(\mathbf{z}\cdot\mathbf{X}),$$

where $Z(\mathbf{b}) = \{\mathbf{z} \in \widehat{\mathcal{K}} \mid \sum_{i=1}^m (z_i - b_i) + \sum_{i=1}^m \lambda_i |z_i - b_i| \le 0\}$.

Since $\widehat{V}$ is concave and $\mathbf{X} \ge 0$, it follows that $\mathbf{E}\widehat{V}(\mathbf{z}\cdot\mathbf{X})$ is also concave on $\widehat{S}$. Therefore, the existence of an ϵ-optimal continuous selector will immediately follow from Theorem 3.2 in [16] provided $Z(\mathbf{b})$ is a continuous set-valued function. We prove this in the lemma below. ■

Lemma 6 *Let $Z : \widehat{S} \to \widehat{\mathcal{K}}$ be a set-function defined as follows*

$$Z(\mathbf{b}) = \left\{\mathbf{z} \in \widehat{\mathcal{K}} \mid \mathbf{1}^T(\mathbf{z}-\mathbf{b}) + \boldsymbol{\lambda}^T|\mathbf{z}-\mathbf{b}| \le 0\right\},$$

where $\boldsymbol{\lambda} = (\lambda_i)_{(1\le i\le m)}$. Then Z is a continuous set function.

Proof: Let $A \subset \widehat{\mathcal{K}}$ be a closed set. Then

$$\begin{aligned} Z^{-1}(A) &= \{\mathbf{b} \in \widehat{S} \mid Z(\mathbf{b}) \cap A \neq \emptyset\}, \\ &= \{\mathbf{b} \in \widehat{S} \mid \exists \mathbf{z} \in A \text{ such that } \mathbf{1}^T(\mathbf{z}-\mathbf{b}) + \boldsymbol{\lambda}^T|\mathbf{z}-\mathbf{b}| \le 0\}. \end{aligned}$$

Let $\mathbf{b}_k \to \mathbf{b}_0$ such that $\mathbf{b}_k \in Z^{-1}(A)$ for all $k \ge 1$. Choose a sequence $\mathbf{z}_k \in A \cap Z(\mathbf{b}_k)$ such that

$$\mathbf{1}^T(\mathbf{z}_k - \mathbf{b}_k) + \boldsymbol{\lambda}^T|\mathbf{z}_k - \mathbf{b}_k| \le 0. \tag{31}$$

Since $\cap_{k\ge1} Z(\mathbf{b}_k)$ is compact and A is closed, it follows that there exists a subsequence $\mathbf{z}_{k'}$ such that $\mathbf{z}_{k'} \to \mathbf{z}_0 \in A$. By taking the limit of (31) along the subsequence $\{k'\}$, we have

$$\mathbf{1}^T(\mathbf{z}_0 - \mathbf{b}_0) + \boldsymbol{\lambda}^T|\mathbf{z}_0 - \mathbf{b}_0| \le 0,$$

i.e. $\mathbf{z}_0 \in Z(\mathbf{b}_0)$. Therefore, $\mathbf{b}_0 \in Z^{-1}(A)$. This proves the upper semi-continuity of the function Z.

To prove lower semi-continuity, one needs to show that $Z^{-1}(A)$ is open if $A \subset \widehat{\mathcal{K}}$ is open. Let $\mathbf{b}_0 \in Z^{-1}(A)$ and $\mathbf{z}_0 \in A$ such that

$$\mathbf{1}^T(\mathbf{z}_0 - \mathbf{b}_0) + \boldsymbol{\lambda}^T|\mathbf{z}_0 - \mathbf{b}_0| \le 0.$$

There are the following two possible cases:
Case(i): $\mathbf{1}^T(\mathbf{z}_0 - \mathbf{b}_0) + \boldsymbol{\lambda}^T|\mathbf{z}_0 - \mathbf{b}_0| < 0$. In this case, there exists $\epsilon > 0$ such that for all $\mathbf{b} \in \widehat{S}$ with $\|\mathbf{b} - \mathbf{b}_0\| \le \epsilon$

$$\mathbf{1}^t(\mathbf{z}_0 - \mathbf{b}) + \boldsymbol{\lambda}^T|\mathbf{z}_0 - \mathbf{b}| \le 0,$$

i.e. $\mathbf{b} \in \widehat{S}$ such that $\|\mathbf{b} - \mathbf{b}_0\| \le \epsilon$ implies that $\mathbf{b} \in Z^{-1}(A)$.

Case(ii): $\mathbf{1}^T(\mathbf{z}_0 - \mathbf{b}_0) + \boldsymbol{\lambda}^T |\mathbf{z}_0 - \mathbf{b}_0| = 0$. Suppose first that $\mathbf{z}_0 \neq 0$. Then, A open implies that exists $w < 1$ such that $w\mathbf{z}_0 \in A$. Therefore,

$$\begin{aligned}\mathbf{1}^T(w\mathbf{z}_0 - \mathbf{b}_0) + \boldsymbol{\lambda}^T |w\mathbf{z}_0 - \mathbf{b}_0| &\leq [\mathbf{1}^T(z_0 - b_0) + \boldsymbol{\lambda}^T |\mathbf{z}_0 - \mathbf{b}_0|] + (w-1)(\mathbf{1}^T\mathbf{z}_0 - \boldsymbol{\lambda}^T |\mathbf{z}_0|) \\ &\leq (w-1)(\mathbf{1}^T\mathbf{z}_0 - \boldsymbol{\lambda}^T |\mathbf{z}_0|).\end{aligned}$$

Since $\mathbf{z}_0 \in \widehat{S}$ and $\mathbf{z}_0 \neq 0$, it is easy to show that $\mathbf{1}^T\mathbf{z}_0 - \boldsymbol{\lambda}^t |\mathbf{z}_0| > 0$, i.e.

$$\mathbf{1}^T(w\mathbf{z}_0 - \mathbf{b}_0) + \boldsymbol{\lambda}^T |w\mathbf{z}_0 - \mathbf{b}_0| < 0.$$

We are back in case(i), so one can conclude that there exists $\epsilon > 0$ such that for all $\mathbf{b}$ such that $\|\mathbf{b} - \mathbf{b}_0\| \leq \epsilon$ we have $\mathbf{b} \in Z^{-1}(A)$.

Finally suppose $\mathbf{z}_0 = 0$. Since $0 \in Z(\mathbf{b})$ for all $\mathbf{b} \in \widehat{S}$ we have that $\widehat{S} \subset Z^{-1}(0) \subset Z^{-1}(A)$. This completes the proof. ■

Proof of Theorem 4

The proof proceeds as follows. First we establish that the Markov chain induced by any continuous policy is Weak Feller continuous. Next we establish that this Markov chain has a invariant probability measure. Finally invoking Theorem 17.1.2 in [19] we complete the proof.

Fix an ϵ-optimal continuous policy π_c and let $h : \mathcal{S} \times \mathcal{S}$ be a bounded continuous function. Define

$$\mathcal{P}h(\mathbf{b}) \triangleq \int_{\mathcal{S}} P^{\pi_c}(\mathbf{b}, d\mathbf{y})h(\mathbf{y}) = \mathbf{E}h(\pi_c(\mathbf{b}) \circ X),$$

where P^{π_c} is the measure induced by the policy π_c. Since h is bounded it immediately follows that $\mathcal{P}h$ is bounded.

Let $\mathbf{b}_n \to \mathbf{b}$ be any sequence converging to $\mathbf{b}$. Since h is bounded and continuous, it immediately follows that

$$\mathcal{P}h = \mathbf{E}h(\pi_c(\mathbf{b}) \circ X) = \lim_{n\to\infty} \mathbf{E}h(\pi_c(\mathbf{b}_n) \circ X),$$

i.e. $\mathcal{P}h$ is continuous function. Thus, the Markov chain induced by π is Weak Feller continuous.

Since the portfolio $\mathbf{b}_n \in \mathcal{S}$, $n \geq 1$ and $\mathcal{S}$ is compact, it follows that the set of transition kernels $\{(P^{\pi_c})^k(\mathbf{b}, \cdot) : k \geq 1\}$ is tight. From Theorem 12.1.2 of [19] it follows that the Markov chain has an invariant measure μ.

From Theorem 17.1.2 in [19] it follows that there is a set $\mathcal{F} \subseteq \mathcal{S}$ such that $\mu(\mathcal{F}) = 1$ and for all $\mathbf{b} \in \mathcal{F}$,

$$\lim_{n\to\infty} \frac{1}{n} \log S_n^{\pi_c}(\mathbf{b}) = g_c, \quad P^{\pi_c} - \text{a.s.},$$

where g_c is the growth rate of the policy π_c.

Fix $\mathbf{b}_0 \in \mathcal{F}$ and define the 1-stationary policy $\pi = (\pi_1, \pi_2)$ as follows:

$$\pi_1(\mathbf{b}) = \mathbf{b}_0, \qquad \pi_2(\mathbf{b}) = \pi_c(\mathbf{b}), \qquad \forall \mathbf{b} \in \mathcal{S}.$$

It is easy to show that for all $\mathbf{b} \in \mathcal{S}$

$$\lim_{n\to\infty} \frac{1}{n} \log S_n^{\pi} = g_c, \quad P^{\pi} - \text{a.s.}$$

Since π_c was ϵ-optimal, this establishes the result. ■

Proof of Theorem 5

Fix $\epsilon > 0$ and an $\frac{\epsilon}{4}$-optimal continuous policy π. Since π is continuous, set $\mathcal{N} = \pi(\mathcal{S})$ is a compact and connected set. Define a policy $\tilde{\pi}$ as follows

$$\tilde{\pi}(\mathbf{b}) = \begin{cases} \pi(\mathbf{b}), & \mathbf{b} \in \mathcal{S}\backslash\mathcal{N}, \\ \mathbf{b}, & \mathbf{b} \in \mathcal{N}, \end{cases}$$

i.e. the policy $\tilde{\pi}$ follows the policy π outside the set $\mathcal{N}$ and does nothing within the set.

Since π is an $\frac{\epsilon}{4}$-optimal policy we have that

$$\begin{aligned} V(\mathbf{b}) + g &\geq \log w(\mathbf{b}, (\pi)^2(\mathbf{b})) + W((\pi)^2(\mathbf{b})) + \mathbf{E}V((\pi)^2(\mathbf{b}) \circ \mathbf{X}), \\ &\geq \log w(\mathbf{b}, (\pi)^2(\mathbf{b})) + \left[V(\pi(\mathbf{b})) + g - \log w(\pi(\mathbf{b}), (\pi)^2(\mathbf{b})) - \frac{\epsilon}{4}\right], \end{aligned} \tag{32}$$

where $(\pi)^2(\mathbf{b}) = \pi(\pi(\mathbf{b}))$. Also, one has that

$$\log w(\mathbf{b}, (\pi)^2(\mathbf{b})) \geq \log w(\mathbf{b}, \pi(\mathbf{b})) + \log w(\pi(\mathbf{b}), (\pi)^2(\mathbf{b})). \tag{33}$$

Substituting (33) into (32) we get

$$V(\mathbf{b}) \geq \log w(\mathbf{b}, \pi(\mathbf{b})) + V(\pi(\mathbf{b})) - \frac{\epsilon}{4}. \tag{34}$$

Since $\tilde{\pi}(\mathbf{b}) = \pi(\mathbf{b})$ for all $\mathbf{b} \in \mathcal{S}\backslash\mathcal{N}$ and π is $\frac{\epsilon}{4}$-optimal, we have that for all $\mathbf{b} \in \mathcal{S}\backslash\mathcal{N}$

$$V(\mathbf{b}) + g \leq \log w(\mathbf{b}, \tilde{\pi}(\mathbf{b})) + W(\tilde{\pi}(\mathbf{b})) + \mathbf{E}V(\tilde{\pi}(\mathbf{b}) \circ \mathbf{X}) + \frac{\epsilon}{4}. \tag{35}$$

For all $\mathbf{b} \in \mathcal{N}$, then there exists $\mathbf{z} \in \mathcal{S}$ such that $\pi(\mathbf{z}) = \mathbf{b}$. Therefore,

$$V(\mathbf{z}) + g \leq \log w(\mathbf{z}, \pi(\mathbf{z})) + W(\pi(\mathbf{z})) + \mathbf{E}V(\pi(\mathbf{z}) \circ X) + \frac{\epsilon}{4}.$$

Substituting the lower bound for $V(\mathbf{z})$ from (34) and collecting terms we get

$$V(\mathbf{b}) + g \leq W(\mathbf{b}) + \mathbf{E}V(\mathbf{b} \circ X) + \frac{\epsilon}{2}. \tag{36}$$

Thus (35) and (36) imply that the policy $\tilde{\pi}$ is $\frac{\epsilon}{2}$-optimal.

The policy $\tilde{\pi}$ is not yet a control-limit policy since π could map portfolios outside of $\mathcal{N}$ to portfolios in the interior of $\mathcal{N}$. The rest of the proof corrects this.

For each portfolio $\mathbf{b} \in \mathcal{S}\backslash\mathcal{N}$ fix a trading strategy that takes the portfolio $\mathbf{b}$ to the portfolio $\pi(\mathbf{b})$ in a continuous manner. For a portfolio $\mathbf{b} \in \mathcal{S}\backslash\mathcal{N}$ and let $\mathbf{z}_b \in \partial\mathcal{N}$ be the portfolio where the continuous trading strategy intersects the boundary. Then

$$\begin{aligned} V(\mathbf{b}) + g &\leq \log w(\mathbf{b}, \pi(\mathbf{b})) + W(\pi(\mathbf{b})) + \mathbf{E}V(\pi(\mathbf{b}) \circ X) + \frac{\epsilon}{2}, \\ &\leq \log w(\mathbf{b}, \pi(\mathbf{b})) - \log w(\mathbf{z}_b, \pi(\mathbf{b})) + V(\mathbf{z}_b) + g + \frac{\epsilon}{2}, \qquad (37) \\ &= \log w(\mathbf{b}, \mathbf{z}_b) + V(\mathbf{z}_b) + g + \frac{\epsilon}{2}, \qquad (38) \\ &\leq \log w(\mathbf{b}, \mathbf{z}_b) + W(\mathbf{z}_b) + \mathbf{E}V(\mathbf{z}_b \circ \mathbf{X}) + \epsilon, \qquad (39) \end{aligned}$$

where (37) follows from the simple bound

$$V(\mathbf{z}_b) + g \geq \log w(\mathbf{z}_b, \pi(\mathbf{b})) + W(\pi(\mathbf{b})) + V(\pi(\mathbf{b}) \circ \mathbf{X}),$$

(38) follows from the fact that $\mathbf{z}$ lies in the path from $\mathbf{b}$ to $\pi(\mathbf{b})$ and (39) follows by recognizing that for all $\mathbf{z} \in \mathcal{N}$ it is $\frac{\epsilon}{2}$-optimal not to move.

Unfortunately one cannot prove that function $f : \mathcal{S} \mapsto \mathcal{K}$, that maps $\mathbf{b} \in \mathcal{N}$ to itself and $\mathbf{b} \in \mathcal{S} \backslash \mathcal{N}$ to the portfolio $\mathbf{z}_b \in \partial\mathcal{N}$, is measurable and, therefore, admissible. We circumvent this problem by invoking a simple measurable selection result given in Proposition 7.33 in [4].

For every portfolio $\mathbf{b} \in \mathcal{S} \backslash \mathcal{N}$ define

$$Z(\mathbf{b}) = \{z \in \partial\mathcal{N} \mid V(\mathbf{b}) + g \leq \log w(\mathbf{b}, \mathbf{z}) + W(\mathbf{z}) + V(\mathbf{z} \circ \mathbf{X}) + \epsilon\},$$

i.e. $Z(\mathbf{b})$ are portfolios on $\partial\mathcal{N}$ that are ϵ-optimal. From (39) it follows that the set $Z(\mathbf{b})$ is nonempty and closed. Then Proposition 7.33 in [4] states that there exists a measurable selector π_l^ϵ such that

$$\pi_l^\epsilon(\mathbf{b}) = \begin{cases} \mathbf{b} & \mathbf{b} \in \mathcal{N}, \\ \mathbf{b} \in \operatorname{argmax}_{\mathbf{z} \in Z(\mathbf{b})} \{ \log w(\mathbf{b}, \mathbf{z}) + W(\mathbf{z}) + \mathbf{E}V(\mathbf{z} \circ \mathbf{X})\}, & \mathbf{b} \in \mathcal{S} \backslash \mathcal{N}, \end{cases} \tag{40}$$

i.e. π_l^ϵ is ϵ-optimal control-limit policy. ■