Antitrust and the No-Arbitrage Hypothesis

Robert Fernholz

INTECH
One Palmer Square
Princeton, NJ 08542

December 3, 1997
Revised January 21, 1998

Abstract

Suppose that a continuously traded equity market is compliant with a weak form of antitrust regulation that prevents the concentration of practically all the market capital into a single company. Then the no-arbitrage hypothesis imposes certain conditions on the dividend rates of the stocks in the market. These conditions are observable and can be used to statistically verify the hypothesis. Analysis of U.S. equity market data over the period from 1967 to 1996 indicates that the no-arbitrage hypothesis is likely to be invalid for this market.

Key words: Arbitrage, antitrust, portfolio generating function.
Classification code: G11, C62.

1 Introduction

The no-arbitrage hypothesis is a basic tenet of current mathematical finance (see, e.g., Duffie (1992) and Karatzas (1996)). The hypothesis states that markets do not present opportunities for riskless arbitrage (or "free lunch"), and while there are theoretical examples of markets in which arbitrage exists, these examples appear to be mathematical oddities which do not resemble "real" markets. Faith in market "efficiency" (see Malkiel (1990)) makes no-arbitrage a tempting hypothesis, and it offers a nice setting for mathematical analysis, but whether or not this hypothesis accurately represents actual market conditions has never been established, in part because it depends on technical and unobservable conditions such as the existence of an equivalent martingale measure (see Harrison and Kreps (1979), Harrison and Pliska (1981), and Dybvig and Huang (1988)), and in part because its status as a fundamental axiom has rendered it beyond question.

On what grounds could no-arbitrage be questioned, if it were liable to question? To challenge so fundamental a postulate, three rigorous criteria must be satisfied. First, it must be shown that actual markets are subject to some unconventional constraint that has traditionally been ignored due to its innocuous and apparently irrelevant nature. Second, it must be proved theoretically that, due to this constraint, there exist strategies that violate no-arbitrage. Third, statistical tests must verify that the conditions which permit such strategies are present in some important actual market. Minor, transitory pricing anomalies are not of interest to us; they are quickly "arbitraged" away by clever traders (see Taylor (1986) and Malkiel (1990)). Nor are we interested in theoretical constructs that could not exist in actual markets (e.g., negative-priced options, see Jarrow and Madan (1997)).

In this paper we consider continuously traded equity markets that are compliant with a weak form of antitrust regulation. Antitrust regulation is an unorthodox condition in mathematical finance, and apparently has not been considered in the literature. We show that in a market of stocks that do not pay dividends, this weak antitrust condition makes it possible to construct a well-behaved portfolio that dominates the market portfolio. This portfolio is constructed through the use of a portfolio generating function (see Fernholz (1999b)). The return of the generated portfolio relative to the market portfolio can be expressed as the sum of the change in the value of the generating function plus a monotonically increasing process. There is a positive lower bound on the value of the generating function, and the antitrust condition ensures that the rate of increase of the monotonic process is bounded away from zero. This combination of lower bounds implies that the generated portfolio will dominate the market portfolio.

If the stocks in the market pay dividends, the generated portfolio will dominate the market portfolio unless the difference between the dividend yields of the two portfolios is great enough to offset the increasing process mentioned above. Since these processes are all observable, this permits statistical testing of the no-arbitrage hypothesis. Analysis of data for the U.S. equity market over the 30 year period from 1967 to 1996 indicates that the difference between these dividend yields was not great enough. Hence, for the U.S. equity market, the evidence suggests that the no-arbitrage hypothesis has been invalid.

The example that we consider cannot be dismissed as nothing more than a parlor trick. The weak antitrust condition that we assume depends only on observable parameters and is broadly consistent with the structure of equity markets in industrialized economies. The portfolio construction methodology we use is not merely a mathematical construct, but has been used for actual portfolio management since 1996 (see Fernholz, Garvy, and Hannon (1998)). It follows that the results we obtain are not merely curiosities, but rather show that in certain markets the no-arbitrage hypothesis may be invalid and hence may to lead to incorrect conclusions about market behavior.

We shall use a model of stock price processes represented by continuous semimartingales that is fairly standard in continuous-time financial theory (see, e.g., Karatzas and Shreve (1991)). We shall make certain simplifying assumptions, among them:

1. Companies do not merge or break up, and the total number of shares of a company remains constant. The list of companies in the market is fixed.
2. Dividends are paid continuously rather than discretely.
3. There are no transaction costs, taxes, or problems with the indivisibility of shares.

2 Stochastic portfolio theory

In this section we shall review the basic definitions and results needed in the later sections. Much of the material in this section can also be found in Fernholz (1999a), but since the approach we use may be unfamiliar, it is presented here also. We shall generally follow the definitions and notation used in Karatzas and Shreve (1991) and Karatzas (1996).

Let

$$W = \{W(t) = (W_1(t), \dots, W_n(t)), \mathcal{F}_t, t \in [0, \infty)\}$$

be a standard n-dimensional Brownian motion defined on a probability space $\{\Omega, \mathcal{F}, P\}$ where $\{\mathcal{F}_t\}$ is the augmentation under P of the natural filtration $\{\mathcal{F}_t^W = \sigma(W(s); 0 \leq s \leq t)\}$. We say that a process $\{X(t), \mathcal{F}_t, t \in [0, \infty)\}$ is *adapted* if $X(t)$ is $\mathcal{F}_t$-measurable for $t \in [0, \infty)$. If X and Y are processes defined on $\{\Omega, \mathcal{F}, P\}$, we shall use the notation $X = Y$ if

$$P\{X(t) = Y(t), \quad t \in [0, \infty)\} = 1.$$

For continuous, square-integrable martingales $\{M(t), \mathcal{F}_t, t \in [0, \infty)\}$ and $\{N(t), \mathcal{F}_t, t \in [0, \infty)\}$, we can define the *cross-variation process* $\langle M, N \rangle$. The cross-variation process is adapted, continuous, and of bounded variation, and the operation $\langle \cdot, \cdot \rangle$ is bilinear on the real vector space of continuous, square-integrable martingales. If $M = N$, we shall use the notation $\langle M \rangle = \langle M, M \rangle$; $\langle M \rangle$ is called the *quadratic variation process* of M, and has continuous, nondecreasing sample paths. The Brownian motion process defined above is a continuous, square-integrable martingale, and it is characterized by its cross-variation processes

$$\langle W_i, W_j \rangle_t = \delta_{ij} t, \quad t \in [0, \infty),$$

where $\delta_{ij} = 1$ if $i = j$, and 0 otherwise.

A *continuous semimartingale* $X = \{X(t), \mathcal{F}_t, t \in [0, \infty)\}$ is a measurable, adapted process that has the decomposition,

$$X(t) = X(0) + M_X(t) + V_X(t), \quad t \in [0, \infty), \quad \text{a.s.,} \tag{2.1}$$

where $\{M_X(t), \mathcal{F}_t, t \in [0, \infty)\}$ is a continuous, square-integrable martingale and $\{V_X(t), \mathcal{F}_t, t \in [0, \infty)\}$ is a continuous, adapted process that is locally of bounded variation. It can be shown that this decomposition is a.s. unique (see Karatzas and Shreve (1991)), so we can define the cross-variation process for continuous semimartingales X and Y by

$$\langle X, Y \rangle = \langle M_X, M_Y \rangle,$$

where M_X and M_Y are the martingale parts of X and Y, respectively.

Definition 2.1. Let X_0 be a positive number. A *stock* $X = \{X(t), \mathcal{F}_t, t \in [0,\infty)\}$ is a process of the form

$$X(t) = X_0 \exp\biggl(\int_0^t \gamma(s)\,ds + \int_0^t \sum_{\nu=1}^n \xi_\nu(s)\,dW_\nu(s)\biggr), \quad t \in [0,\infty), \tag{2.2}$$

where $\gamma = \{\gamma(t), \mathcal{F}_t, t \in [0,\infty)\}$ is measurable, adapted, and satisfies $\int_0^t |\gamma(s)|ds < \infty$, for all $t \in [0,\infty)$, a.s., and for $\nu = 1,\ldots,n$, $\xi_\nu = \{\xi_\nu(t), \mathcal{F}_t, t \in [0,\infty)\}$ is measurable, adapted, and satisfies $\int_0^t \xi_\nu^2(s)ds < \infty$ for all $t \in [0,\infty)$, a.s., and such that there exists a number $\varepsilon > 0$ for which $\xi_1^2(t) + \cdots + \xi_n^2(t) > \varepsilon$, $t \in [0,\infty)$, a.s.

It follows directly from Definition 2.1 that X is adapted, that $X(t) > 0$ for all $t \in [0,\infty)$, a.s., and that X has initial value $X(0) = X_0$. We shall set the initial value X_0 to be the total capitalization of the company represented by X at time $t = 0$, and we shall assume that this total capitalization is positive. This is equivalent to assuming that there is a single share of stock outstanding, and $X(t)$ represents its price at time t. We assume that stock shares are infinitely divisible, so there is no loss of generality in assuming a single share outstanding. The process γ is called the *growth rate (process)* of X and, for each ν, the process ξ_ν represents the sensitivity of X to the ν-th source of uncertainty W_ν.

We shall find it convenient to use a logarithmic representation for stocks (cf. Fernholz and Shay (1982)). Equation (2.2) is equivalent to

$$\log X(t) = \log X_0 + \int_0^t \gamma(s)\,ds + \int_0^t \sum_{\nu=1}^n \xi_\nu(s)\,dW_\nu(s),$$

or, in differential form,

$$d\log X(t) = \gamma(t)\,dt + \sum_{\nu=1}^n \xi_\nu(t)\,dW_\nu(t). \tag{2.3}$$

From this it is clear that $\log X$ is a continuous semimartingale.

Suppose that we have a family of stocks $X_i, i = 1,\ldots,n$,

$$X_i(t) = X_0^i \exp\biggl(\int_0^t \gamma_i(s)\,ds + \int_0^t \sum_{\nu=1}^n \xi_{i\nu}(s)\,dW_\nu(s)\biggr), \quad t \in [0,\infty). \tag{2.4}$$

Consider the matrix valued process ξ defined by $\xi(t) = (\xi_{i\nu}(t))_{1\le i,\nu\le n}$ and define the *covariance process* σ where $\sigma(t) = \xi(t)\xi^T(t)$. The cross-variation processes for $\log X_i$ and $\log X_j$ are related to σ by

$$\langle \log X_i, \log X_j\rangle_t = \int_0^t \sigma_{ij}(s)\,ds, \quad t \in [0,\infty), \quad \text{a.s.} \tag{2.5}$$

Since the processes $\xi_{i\nu}$ are assumed to be square integrable in Definition 2.1, it follows that for all i and j,

$$\int_0^t \sigma_{ij}(s)\,ds < \infty, \quad t \in [0,\infty), \quad \text{a.s.}$$

Definition 2.2. A *market* is a family $\mathcal{M} = \{X_i,\ldots,X_n\}$ of stocks, defined as in (2.4), for which there is a number $\varepsilon > 0$ such that

$$x\sigma(t)x^T \ge \varepsilon\,\|x\|^2, \quad x \in \mathbb{R}^n, t \in [0,\infty), \quad \text{a.s.} \tag{2.6}$$

The *strong nondegeneracy* condition (2.6) is fairly common and can be found, for example, in Karatzas and Shreve (1991) and Karatzas and Kou (1996), and as *uniform ellipticity* in Duffie (1992). Frequently in the literature one finds more restrictions on γ_i and $\xi_{i\nu}$ in order to be able to prove the existence of an equivalent martingale measure (see Harrison and Kreps (1979), Harrison and Pliska (1981), and Dybvig and Huang (1988), as well as Duffie (1992) and Karatzas (1996)). It is known, for example, that if γ_i and $\xi_{i\nu}$ are a.s. bounded on $[0,\infty)$, then there exists an equivalent martingale measure (see, e.g., Karatzas (1996)). But it is not difficult to show that if γ_i and $\xi_{i\nu}$ are bounded, then Definition 4.1 below will not hold.

Definition 2.3. Let $\mathcal{M}$ be a market of n stocks. A *portfolio* in $\mathcal{M}$ is a measurable, adapted process $\pi = \{\pi(t) = (\pi_1(t), \ldots, \pi_n(t)), \mathcal{F}_t, t \in [0,\infty)\}$ such that $\pi(t)$ is bounded on $[0,\infty) \times \Omega$ and

$$\pi_1(t) + \cdots + \pi_n(t) = 1, \quad t \in [0,\infty), \quad \text{a.s.}$$

The processes π_i represent the respective proportions, or weights, of each stock in the portfolio. A negative value for $\pi_i(t)$ indicates a short sale. Suppose $Z_\pi(t)$ represents the value of an investment in π at time t. Then the amount invested in the i-th stock X_i will be

$$\pi_i(t) Z_\pi(t),$$

so if the price of X_i changes by $dX_i(t)$, the induced change in the portfolio value will be

$$\pi_i(t) Z_\pi(t) \frac{dX_i(t)}{X_i(t)}.$$

Hence the total change in the portfolio value at time t will be

$$dZ_\pi(t) = \sum_{i=1}^n \pi_i(t) Z_\pi(t) \frac{dX_i(t)}{X_i(t)},$$

or, equivalently,

$$\frac{dZ_\pi(t)}{Z_\pi(t)} = \sum_{i=1}^n \pi_i(t) \frac{dX_i(t)}{X_i(t)}. \tag{2.7}$$

Since we are interested in the behavior of portfolios, we are interested in solutions to (2.7). The following proposition and corollary are proved in Fernholz (1999a).

Proposition 2.1. *Let π be a portfolio and let*

$$\gamma_\pi(t) = \sum_{i=1}^n \pi_i(t)\gamma_i(t) + \gamma_\pi^*(t), \tag{2.8}$$

where

$$\gamma_\pi^*(t) = \frac{1}{2}\Big(\sum_{i=1}^n \pi_i(t)\sigma_{ii}(t) - \sum_{i,j=1}^n \pi_i(t)\pi_j(t)\sigma_{ij}(t)\Big). \tag{2.9}$$

Then, for any positive initial value Z_0^π, the process Z_π defined by

$$Z_\pi(t) = Z_0^\pi \exp\Big(\int_0^t \gamma_\pi(s)\,ds + \int_0^t \sum_{i,\nu=1}^n \pi_i(s)\xi_{i\nu}(s)\,dW_\nu(s)\Big), \quad t \in [0,\infty), \tag{2.10}$$

is a strong solution of (2.7).

Corollary 2.1. *Let π be a portfolio and Z_π be its value process. Then for $t \in [0, \infty)$,*

$$d \log Z_\pi(t) = \sum_{i=1}^n \pi_i(t)\, d \log X_i(t) + \gamma_\pi^*(t)\, dt. \tag{2.11}$$

It follows from (2.11) that $\log Z_\pi$ is a continuous semimartingale. The process γ_π in (2.8) is called the *portfolio growth rate (process)* of the portfolio π, and γ_π^* in (2.9) is called the *excess growth rate (process)*. It was proved in Fernholz (1999a) that for portfolios with non-negative weights, the excess growth rate is non-negative, and is positive unless the portfolio consists of a single stock.

For any stock X_i and portfolio π we can consider the quotient process X_i/Z_π defined by

$$\log\bigl(X_i(t)/Z_\pi(t)\bigr) = \log X_i(t) - \log Z_\pi(t). \tag{2.12}$$

This process is a continuous semimartingale with

$$\langle \log(X_i/Z_\pi), \log(X_j/Z_\pi)\rangle_t = \\ \langle \log X_i, \log X_j\rangle_t - \langle \log X_i, \log Z_\pi\rangle_t - \langle \log X_j, \log Z_\pi\rangle_t + \langle \log Z_\pi\rangle_t. \tag{2.13}$$

If we define the process $\sigma_{i\pi}$ by

$$\sigma_{i\pi}(t) = \sum_{j=1}^n \pi_j(t)\sigma_{ij}(t),$$

for $i = 1, \dots, n$, then

$$\langle \log X_i, \log Z_\pi\rangle_t = \int_0^t \sigma_{i\pi}(s)\, ds.$$

Define the *relative covariance (process)* τ^π to be the matrix valued process

$$\tau^\pi(t) = (\tau_{ij}^\pi(t))_{1 \le i,j \le n},$$

where

$$\tau_{ij}^\pi(t) = \sigma_{ij}(t) - \sigma_{i\pi}(t) - \sigma_{j\pi}(t) + \sigma_{\pi\pi}(t), \tag{2.14}$$

for $i, j = 1, \dots, n$, where $\sigma_{\pi\pi}(t) = \pi(t)\sigma(t)\pi^T(t)$, $t \in [0, \infty)$. Then for all i and j,

$$\langle \log(X_i/Z_\pi), \log(X_j/Z_\pi)\rangle_t = \int_0^t \tau_{ij}^\pi(s)\, ds. \tag{2.15}$$

In the case that $i = j$, we know that $\langle \log(X_i/Z_\pi)\rangle_t$ is non-decreasing, so

$$\tau_{ii}^\pi(t) \ge 0, \quad t \in [0, \infty), \quad \text{a.s.}$$

Lemma 2.1. *Let π be a portfolio. Then the rank of $\tau^\pi(t)$ is $n-1$, for $t \in [0, \infty)$, a.s., and the null space of $\tau^\pi(t)$ is spanned by $\pi(t)$.*

Proof. It follows from (2.14) that for any portfolio η, a.s. for $t \in [0, \infty)$,

$$\eta(t)\tau^\pi(t)\eta^T(t) = \bigl(\eta(t) - \pi(t)\bigr)\sigma(t)\bigl(\eta(t) - \pi(t)\bigr)^T.$$

By condition (2.6), a.s. for $t \in [0, \infty)$,

$$\bigl(\eta(t) - \pi(t)\bigr)\sigma(t)\bigl(\eta(t) - \pi(t)\bigr)^T = 0$$

if and only if $\eta(t) = \pi(t)$. ☐

The following lemma expresses the excess growth in terms of the relative covariance process.

Lemma 2.2. *Let π and η be portfolios. Then for $t \in [0, \infty)$,*

$$\gamma_\eta^*(t) = \frac{1}{2}\Big(\sum_{i=1}^n \eta_i(t)\tau_{ii}^\pi(t) - \sum_{i,j=1}^n \eta_i(t)\eta_j(t)\tau_{ij}^\pi(t)\Big).$$

Proof. The proof is a direct calculation using (2.14). □

Lemma 2.3. *Let π be a portfolio. Then there exists an $\varepsilon > 0$ such that for $i = 1, \ldots, n$,*

$$\tau_{ii}^\pi(t) \geq \varepsilon\big(1 - \pi_i(t)\big)^2, \quad t \in [0, \infty), \quad \text{a.s.} \tag{2.16}$$

Proof. For any i and $t \in [0, \infty)$, let $x(t) = (\pi_1(t), \ldots, \pi_i(t) - 1, \ldots, \pi_n(t))$. Then,

$$\begin{aligned}\tau_{ii}^\pi(t) &= \sigma_{ii}(t) - 2\sigma_{i\pi}(t) + \sigma_{\pi\pi}(t) \\ &= x(t)\sigma(t)x^T(t) \\ &\geq \varepsilon\,\|x(t)\|^2, \quad t \in [0, \infty), \quad \text{a.s.,}\end{aligned}$$

where ε is chosen as in (2.6). Since,

$$\|x(t)\|^2 \geq \big(1 - \pi_i(t)\big)^2,$$

the lemma follows. □

Lemma 2.4. *Let π be a portfolio with non-negative weights, and let $\pi_{\max}(t) = \max_{1 \leq i \leq n} \pi_i(t)$. Then there exists an $\varepsilon > 0$ such that*

$$\gamma_\pi^*(t) \geq \varepsilon(1 - \pi_{\max}(t))^2.$$

Proof. If we let $\eta = \pi$ in Lemma 2.2, then Lemma 2.1 implies that

$$\begin{aligned}\gamma_\pi^*(t) &= \frac{1}{2}\sum_{i=1}^n \pi_i(t)\tau_{ii}^\pi(t) \\ &\geq \frac{\varepsilon}{2}(1 - \pi_{\max}(t))^2,\end{aligned}$$

where ε is chosen as in Lemma 2.3, since the $\pi_i(t)$ are non-negative. □

The total capitalization of the market can be represented by a portfolio. Let us assume from now on that the market is $\mathcal{M} = \{X_1, \ldots, X_n\}$, with n stocks.

Definition 2.4. The portfolio

$$\mu = \{\mu(t) = (\mu_1(t), \ldots, \mu_n(t)), \mathcal{F}_t, t \in [0, \infty)\},$$

where

$$\mu_i(t) = \frac{X_i(t)}{X_1(t) + \cdots + X_n(t)}, \tag{2.17}$$

for $i = 1, \ldots, n$, is called the *market portfolio (process)*.

It is clear that the μ_i defined by (2.17) satisfy the requirements of Definition 2.3. If we let

$$Z(t) = X_1(t) + \cdots + X_n(t), \tag{2.18}$$

then $Z(t)$ satisfies (2.7) with weights $\mu_i(t)$ given by (2.17). Hence, the value of the market portfolio represents the combined capitalization of all the stocks in the market. We shall use the notation $Z = Z_\mu$ to represent the market portfolio value process, and $\tau_{ij} = \tau_{ij}^\mu$ to represent its relative covariance processes.

3 Portfolio generating functions

In this section we shall show that certain real-valued functions of the market weights can be used to construct dynamic portfolios. These functionally generated portfolios are of interest because their return relative to the market portfolio is governed by a stochastic differential equation that can be used to establish a dominance relationship between the two portfolios. A general discussion of portfolio generating functions, including examples, can be found in Fernholz (1999b).

We shall consider real-valued functions defined on the open simplex

$$\Delta^n = \{x \in \mathbb{R}^n : x_1 + \cdots + x_n = 1, \quad 0 < x_i < 1, \quad i = 1, \dots, n\}.$$

It will be convenient to use the standard coordinate system in $\mathbb{R}^n$, even though it is not a coordinate system on Δ^n. For this reason we shall consider functions that are defined in an open neighborhood $U \subset \mathbb{R}^n$ of Δ^n. A real-valued function defined on a subset of $\mathbb{R}^n$ is C^2 if it is twice continuously differentiable in all variables. We shall use the notation D_i for the partial derivative with respect to the i-th variable, and D_{ij} for the second partial derivative with respect to the i-th and j-th variables.

Definition 3.1. Let U be an open neighborhood of Δ^n and $\mathbf{S}$ be a positive C^2 function defined in U. Then $\mathbf{S}$ is the *generating function* of the portfolio π if there exists a measurable, adapted process $\Theta = \{\Theta(t), \mathcal{F}_t, t \in [0, \infty)\}$ such that

$$d \log\big(Z_\pi(t)/Z(t)\big) = d \log \mathbf{S}(\mu(t)) + \Theta(t) dt, \quad t \in [0, \infty), \quad \text{a.s.} \tag{3.1}$$

Θ is called the *drift process* corresponding to $\mathbf{S}$.

We shall say that the function $\mathbf{S}$ *generates* π. What follows is the main theorem on portfolio generating functions.

Theorem 3.1. *Let $\mathbf{S}$ be a positive C^2 function defined on a neighborhood U of Δ^n such that for $i = 1, \dots, n$, $x_i D_i \log \mathbf{S}(x)$ is bounded on Δ^n. Then $\mathbf{S}$ generates the portfolio π with weights*

$$\pi_i(t) = \Big(D_i \log \mathbf{S}(\mu(t)) + 1 - \sum_{j=1}^n \mu_j(t)\, D_j \log \mathbf{S}(\mu(t))\Big)\mu_i(t), \quad t \in [0, \infty), \quad \text{a.s.}, \tag{3.2}$$

for $i = 1, \dots, n$, and drift process

$$\Theta(t) = \frac{-1}{2\,\mathbf{S}(\mu(t))} \sum_{i,j=1}^n D_{ij}\mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t), \quad t \in [0, \infty), \quad \text{a.s.} \tag{3.3}$$

Proof. The weight process μ_i is a quotient process with $\mu_i(t) = X_i(t)/Z(t)$ for all t. By (2.15) it follows that

$$d\langle \log \mu_i, \log \mu_j \rangle_t = \tau_{ij}(t)\, dt, \quad t \in [0, \infty), \quad \text{a.s.},$$

so by Itô's Lemma,

$$d\mu_i(t) = \mu_i(t)\, d\log\mu_i(t) + \frac{1}{2}\mu_i(t)\tau_{ii}(t)\, dt, \quad t \in [0,\infty), \quad \text{a.s.,} \tag{3.4}$$

and

$$d\langle\mu_i, \mu_j\rangle_t = \mu_i(t)\mu_j(t)\tau_{ij}(t)\, dt, \quad t \in [0,\infty), \quad \text{a.s.} \tag{3.5}$$

Itô's lemma, along with (3.5), implies that a.s. for all $t \in [0,\infty)$,

$$d\log\mathbf{S}(\mu(t)) = \sum_{i=1}^n D_i \log\mathbf{S}(\mu(t))\, d\mu_i(t) + \frac{1}{2}\sum_{i,j=1}^n D_{ij}\log\mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t)\, dt.$$

Now,

$$D_{ij}\log\mathbf{S}(\mu(t)) = \frac{D_{ij}\mathbf{S}(\mu(t))}{\mathbf{S}(\mu(t))} - D_i\log\mathbf{S}(\mu(t))D_j\log\mathbf{S}(\mu(t)),$$

so, a.s., for all $t \in [0,\infty)$,

$$\begin{aligned} d\log\mathbf{S}(\mu(t)) = &\sum_{i=1}^n D_i\log\mathbf{S}(\mu(t))\, d\mu_i(t) + \frac{1}{2\,\mathbf{S}(\mu(t))}\sum_{i,j=1}^n D_{ij}\mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t)\, dt \\ &- \frac{1}{2}\sum_{i,j=1}^n D_i\log\mathbf{S}(\mu(t))D_j\log\mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t)\, dt. \end{aligned} \tag{3.6}$$

In order for (3.1) to hold, the martingale parts of $\log\mathbf{S}(\mu(t))$ and $\log\big(Z_\pi(t)/Z(t)\big)$ must be equal. Corollary 2.1 implies that for the portfolio π, a.s. for all $t \in [0,\infty)$,

$$\begin{aligned} d\log\big(Z_\pi(t)/Z(t)\big) &= \sum_{i=1}^n \pi_i(t)\, d\log\big(X_i(t)/Z(t)\big) + \gamma_\pi^*(t)\, dt \\ &= \sum_{i=1}^n \pi_i(t)\, d\log\mu_i(t) + \gamma_\pi^*(t)\, dt \\ &= \sum_{i=1}^n \frac{\pi_i(t)}{\mu_i(t)}\, d\mu_i(t) - \frac{1}{2}\sum_{i,j=1}^n \pi_i(t)\pi_j(t)\tau_{ij}(t)\, dt \end{aligned} \tag{3.7}$$

by Lemma 2.2. Suppose that

$$\pi_i(t) = \big(D_i\log\mathbf{S}(\mu(t)) + \varphi(t)\big)\mu_i(t), \tag{3.8}$$

where $\varphi(t)$ is chosen such that $\sum_{i=1}^n \pi_i(t) = 1$. Then, a.s. for all $t \in [0,\infty)$,

$$\begin{aligned} \sum_{i=1}^n \frac{\pi_i(t)}{\mu_i(t)}\, d\mu_i(t) &= \sum_{i=1}^n D_i\log\mathbf{S}(\mu(t))\, d\mu_i(t) + \varphi(t)\sum_{i=1}^n d\mu_i(t) \\ &= \sum_{i=1}^n D_i\log\mathbf{S}(\mu(t))\, d\mu_i(t), \end{aligned}$$

since $\sum_{i=1}^n d\mu_i(t) = 0$. Hence, the martingale parts of $\log\mathbf{S}(\mu(t))$ and $\log\big(Z_\pi(t)/Z(t)\big)$ are equal. If

$$\varphi(t) = 1 - \sum_{j=1}^n \mu_j(t)D_j\log\mathbf{S}(\mu(t)),$$

then $\sum_{i=1}^n \pi_i(t) = 1$, and (3.2) is proved.

If $\pi_i(t)$ satisfies (3.8), then a.s. for all $t \in [0, \infty)$,

$$\begin{aligned}\sum_{i,j=1}^n \pi_i(t)\pi_j(t)\tau_{ij}(t) &= \sum_{i,j=1}^n D_i \log \mathbf{S}(\mu(t)) D_j \log \mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t) \\ &\qquad + 2\varphi(t) \sum_{i,j=1}^n D_i \log \mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t) + \varphi^2(t) \sum_{i,j=1}^n \mu_i(t)\mu_j(t)\tau_{ij}(t) \\ &= \sum_{i,j=1}^n D_i \log \mathbf{S}(\mu(t)) D_j \log \mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t),\end{aligned}$$

since $\mu(t)$ is in the null space of $\tau(t)$ by Lemma 2.1. Hence, a.s. for all $t \in [0, \infty)$,

$$d \log\big(Z_\pi(t)/Z(t)\big) = \sum_{i=1}^n D_i \log \mathbf{S}(\mu(t))\, d\mu_i(t) - \frac{1}{2} \sum_{i,j=1}^n D_i \log \mathbf{S}(\mu(t)) D_j \log \mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t)\, dt.$$

This equation and (3.6) imply that a.s. for all $t \in [0, \infty)$,

$$d \log\big(Z_\pi(t)/Z(t)\big) = d \log \mathbf{S}(\mu(t)) - \frac{1}{2\,\mathbf{S}(\mu(t))} \sum_{i,j=1}^n D_{ij}\mathbf{S}(\mu(t))\mu_i(t)\mu_j(t)\tau_{ij}(t)\, dt,$$

so (3.3) is proved. The process Θ defined by (3.3) is clearly measurable and adapted. □

The next proposition is proved in Fernholz (1999b). It implies that we could have defined generating functions on Δ^n rather than in a neighborhood. We prefer Definition 3.1 because it allows us to use the standard coordinate system on $\mathbb{R}^n$.

Proposition 3.1. *Let $\mathbf{S}_1$ and $\mathbf{S}_2$ be portfolio generating functions. Then $\mathbf{S}_1$ and $\mathbf{S}_2$ generate the same portfolio if and only if $\mathbf{S}_1/\mathbf{S}_2$ is constant on Δ^n.*

4 Antitrust, arbitrage, and dividends

In this section we shall show that if a weak antitrust condition is imposed on the market, then generating functions can be used to construct well-behaved portfolios that dominate the market portfolio. We show that if there is a portfolio that dominates the market portfolio, then arbitrage opportunities exist and the no-arbitrage hypothesis fails. Finally, we show that with the introduction of dividends into our model, antitrust can be made compatible with no-arbitrage.

It has long been accepted that excessive concentration of production or capital in a single company is likely to interfere with competition and be detrimental to the national economy (see, .e.g., Smith (1776) or Blair (1972)). For this reason the U.S. has enacted antitrust legislation to prevent such excessive concentration. Here we are not interested in the economic rationale for antitrust legislation, but rather in the effect such legislation has on the distribution of capital in the equity market. Any credible antitrust regulations will prevent prolonged concentration of practically all the market capital into a single company. From a realistic point of view, in an economy like that of the U.S., it is unlikely that a single company could account for even half of the total market capitalization. The condition we shall impose is a weak consequence of actual antitrust regulations, and any market model bearing even a remote resemblance to the U.S. equity market can safely be assumed to satisfy it.

Definition 4.1. Let $\mu_{\max}(t) = \max_{1\le i\le n} \mu_i(t)$. The market $\mathcal{M}$ is *weakly antitrust compliant* if there exists a number $\delta > 0$ and a time T_0 such that

$$\frac{1}{T}\int_0^T \mu_{\max}(t)\,dt \le 1-\delta, \quad T > T_0, \quad \text{a.s.} \tag{4.1}$$

It is necessary to restrict the class of portfolios we consider in order to prevent the use of "doubling" strategies that will permit unlimited returns at unlimited risk (see Karatzas (1996)). Here we are interested only in portfolios without short sales.

Definition 4.2. A portfolio π is *admissible* if

i) For $i = 1,\dots,n$, $\pi_i(t) \ge 0, \quad t \in [0,\infty)$;

ii) There exists a constant $c > 0$ such that

$$Z_\pi(t)/Z_\pi(0) \ge cZ(t)/Z(0), \quad t \in [0,\infty), \quad \text{a.s.}$$

Admissibility conditions of this nature are not uniform in the literature, and a portfolio that satisfies Definition 4.2 may not be "admissible" in other settings. Condition *ii* implies limited risk with respect to the market as numeraire. The market is a natural numeraire for equity managers whose performance is measured versus the market as benchmark. Note that the market portfolio is admissible.

Definition 4.3. Let η and ξ be portfolios. Then η *strictly dominates* ξ if there is a number $T > 0$ such that for any positive initial values $Z_\eta(0)$ and $Z_\xi(0)$,

$$P\{Z_\eta(T)/Z_\eta(0) > Z_\xi(T)/Z_\xi(0)\} = 1. \tag{4.2}$$

An *arbitrage opportunity* is a combination of investments in portfolios such that the total initial value of the investments is zero and such that the total value will be positive at some given future time T, with probability one (see, e.g., Duffie (1992) or Karatzas (1996)). The no-arbitrage hypothesis states that there exist no arbitrage opportunities composed of investments in admissible portfolios.

Suppose that η and ξ are admissible portfolios that satisfy Definition 4.3. Proposition 2.1 implies that the value of an investment in a portfolio is scalable by setting its initial value. Hence, we can buy one dollar worth of η at time 0, and finance this purchase by selling one dollar worth of ξ short at the same time. Therefore, the total initial value of our portfolio holdings is zero. At time T, the dollar value of our holdings in η will be

$$Z_\eta(T)/Z_\eta(0), \tag{4.3}$$

and the dollar value we owe on the short sale of ξ will be

$$Z_\xi(T)/Z_\xi(0). \tag{4.4}$$

Definition 4.3 implies that, with probability one, (4.3) is greater than (4.4), so the total value of our holdings at time T will be strictly positive. Hence, this combination of investments is an arbitrage opportunity. Therefore, if there exist admissible portfolios that satisfy Definition 4.3, an arbitrage opportunity exists and the no-arbitrage hypothesis fails. The next theorem shows that if an equity market is weakly antitrust compliant, then such portfolios exist.

Theorem 4.1. *If the market $\mathcal{M}$ is weakly antitrust compliant then there exists an admissible portfolio that strictly dominates the market portfolio.*

Proof. Let p be a number $0 < p < 1$, and consider the function $\mathbf{D}_p : \mathbb{R}^n \to \mathbb{R}$ defined by

$$\mathbf{D}_p(x) = \Big(\sum_{i=1}^n x_i^p\Big)^{1/p}. \tag{4.5}$$

By Theorem 3.1, $\mathbf{D}_p$ generates a portfolio π with weights

$$\pi_i(t) = \frac{\mu_i^p(t)}{\big(\mathbf{D}_p(\mu(t))\big)^p}, \tag{4.6}$$

for $i = 1, \dots, n$, and drift process

$$\Theta(t) = (1-p)\gamma_\pi^*(t). \tag{4.7}$$

Let us first show that π is admissible. It is obvious from (4.6) that $\pi_i(t) \geq 0$ for $i = 1, \dots, n$. Lemma 2.4 implies that $\gamma_\pi^*(t) \geq 0$, so,

$$\log\big(Z_\pi(T)/Z_\pi(0)\big) - \log\big(Z(T)/Z(0)\big) \geq \log\big(\mathbf{D}_p(\mu(T))/\mathbf{D}_p(\mu(0))\big), \quad T \in [0,\infty), \quad \text{a.s.}$$

From the definition of $\mathbf{D}_p$ it follows immediately that

$$1 < \mathbf{D}_p(\mu(t)) \leq n^{(1-p)/p}, \quad t \in [0,\infty), \quad \text{a.s.}, \tag{4.8}$$

so,

$$Z_\pi(T)/Z_\pi(0) \geq n^{(p-1)/p} Z(T)/Z(0), \quad T \in [0,\infty), \quad \text{a.s.},$$

and hence π is admissible.

Now we must show that π strictly dominates μ. By (4.6), $\pi_{\max}(t) \leq \mu_{\max}(t)$, for all $t \in [0,\infty)$, a.s., and since $\mathcal{M}$ is weakly antitrust compliant there is a $\delta > 0$ such that

$$\frac{1}{T}\int_0^T \big(1 - \pi_{\max}(t)\big)\, dt \geq \delta, \quad T > T_0, \quad \text{a.s.},$$

where T_0 is chosen as in Definition 4.1. By Schwarz's inequality,

$$\frac{1}{T}\int_0^T \big(1 - \pi_{\max}(t)\big)^2\, dt \geq \delta^2, \quad T > T_0, \quad \text{a.s.}$$

By Lemma 2.4 there exists an $\varepsilon > 0$ such that

$$\gamma_\pi^*(t) \geq \varepsilon\big(1 - \pi_{\max}(t)\big)^2,$$

so,

$$\frac{1}{T}\int_0^T \gamma_\pi^*(t)\, dt \geq \varepsilon\delta^2, \quad T > T_0, \quad \text{a.s.} \tag{4.9}$$

It follows that if $T > \max(T_0, \log n/p\varepsilon\delta^2)$, then $P\{Z_\pi(T)/Z_\pi(0) > Z(T)/Z(0)\} = 1$. □

Remark. Generating functions that are symmetric and concave, as is $\mathbf{D}_p$, are called *measures of diversity*. Portfolios generated by measures of diversity have non-decreasing drift processes and positive weights π_i for which the weight ratios π_i/μ_i decrease with increasing μ_i. See Fernholz (1999b) for details.

In the mathematical finance literature, for reasons of convenience it is often assumed that companies pay no dividends. Of course, this assumption is not realistic. Let us now formally introduce dividends into our model, and see how this affects the conclusions of Theorem 4.1.

A *dividend rate (process)* is a process $\delta = \{\delta(t), \mathcal{F}_t, t \in [0,\infty)\}$ that is measurable, adapted, and satisfies $\int_0^t |\delta(s)| ds < \infty$, for all $t \in [0,\infty)$, a.s. Usually dividend rates are assumed to be non-negative, but this assumption is not necessary. We shall henceforth assume that all stocks have dividend rates, even if in some cases the dividend rate is identically zero.

We define the *total return process* $\widehat{X}$ for a stock X by

$$\widehat{X}(t) = X(t) \exp\Big(\int_0^t \delta(s)\, ds\Big). \tag{4.10}$$

If $\delta = 0$, then $\widehat{X} = X$. It follows from (4.10) that $\widehat{X}(0) = X(0)$ and that

$$d \log \widehat{X}(t) = d \log X(t) + \delta(t)\, dt.$$

Let $\delta_1, \dots, \delta_n$ be the respective dividend rates of the stocks $X_1, \dots, X_n$ in the market $\mathcal{M}$. For any portfolio π, we define the *portfolio dividend rate (process)* δ_π by

$$\delta_\pi(t) = \sum_{i=1}^n \pi_i(t)\delta_i(t), \quad t \in [0,\infty),$$

and the *total return (process)* $\widehat{Z}_\pi$ of π by

$$\widehat{Z}_\pi(t) = Z_\pi(t) \exp\Big(\int_0^t \delta_\pi(s)\, ds\Big). \tag{4.11}$$

As for individual stocks,

$$d \log \widehat{Z}_\pi(t) = d \log Z_\pi(t) + \delta_\pi(t)\, dt.$$

The process $\widehat{Z}_\pi$ represents the value of a portfolio with the same weights $\pi_1(t), \dots, \pi_n(t)$ as π, but in which all dividends are reinvested proportionally across the entire portfolio according to the weight of each stock. Hence the reinvestment of the dividends modifies the value of $\widehat{Z}_\pi$ while preserving the weights of the portfolio π.

We shall use the notation $\widehat{Z}$ to represent the total return process of the market portfolio μ, and δ_μ to represent its dividend rate. We must extend Definition 4.3 so that it applies to stocks with dividends.

Definition 4.4. Let η and ξ be portfolios. Then η *strictly dominates* ξ if there is a number $T > 0$ such that for any positive initial values $\widehat{Z}_\eta(0)$ and $\widehat{Z}_\xi(0)$,

$$P\{\widehat{Z}_\eta(T)/\widehat{Z}_\eta(0) > \widehat{Z}_\xi(T)/\widehat{Z}_\xi(0)\} = 1. \tag{4.12}$$

This definition coincides with Definition 4.3 for portfolios of stocks which pay no dividends. The following proposition generalizes (3.1) to the total return processes.

Proposition 4.1. *Suppose that* $\mathbf{S}$ *generates the portfolio* π *with drift process* Θ*. Then*

$$d \log\big(\widehat{Z}_\pi(t)/\widehat{Z}(t)\big) = d \log \mathbf{S}(\mu(t)) + \big(\delta_\pi(t) - \delta_\mu(t) + \Theta(t)\big) dt, \quad t \in [0,\infty), \quad \text{a.s.} \tag{4.13}$$

Proof. This follows immediately from Theorem 3.1 and (4.11). □

Theorem 4.1 is not valid for π and μ because the differential dividend rate $\delta_\pi(t) - \delta_\mu(t)$ in (4.13) can offset the drift process $\Theta(t)$ and prevent dominance from occurring. Instead we have

Proposition 4.2. *Suppose that* $\mathbf{S}$ *generates the portfolio* π *with drift process* Θ. *If there exist positive constants* c_1 *and* c_2 *such that a.s. for all* $t \geq 0$,

$$\mathbf{S}(\mu(t)) > c_1 \tag{4.14}$$

and

$$\big(\delta_\pi(t) - \delta_\mu(t) + \Theta(t)\big) > c_2, \tag{4.15}$$

then π *strictly dominates* μ.

Proof. It is easily seen that if $T > \log\big(c_2/\mathbf{S}(\mu(0))\big)/c_2$, then $P\{\widehat{Z}_\pi(T)/\widehat{Z}_\pi(0) > \widehat{Z}(T)/\widehat{Z}(0)\} = 1$. □

This proposition shows that the validity of the no-arbitrage hypothesis depends on the behavior of the processes $\mathbf{S}(\mu)$, δ_π, δ_μ, and Θ. The processes $\mathbf{S}(\mu)$, δ_π, and δ_μ are all observable, and Θ can be calculated from observable processes using (4.13). Hence Proposition 4.2 provides the basis for statistical testing of the no-arbitrage hypothesis.

5 Statistical testing of the no-arbitrage hypothesis

We can think of at least two reasons why the no-arbitrage hypothesis has become dogma in mathematical finance. First, on the surface, it appears to be a satisfactory representation of actual markets—Karatzas (1996) characterizes it as "a basic tenet of the reality available to most of us." Arbitrage opportunities are probability-one events, and outside mathematics there are no probability-one events (except for death and taxes, of course), so one could argue that no-arbitrage holds by default. But we must remember that no-arbitrage is a strong hypothesis and it has been levered into strong conclusions, particularly regarding market equilibrium (see, e.g., Sharpe (1964), Merton (1971) and Karatzas, Lehoczky, and Shreve (1990)). Mathematical models should provide a faithful representation of reality—are these conclusions faithful to reality?

Second, until now it has been difficult, if not impossible, to test the hypothesis empirically. In the literature, no-arbitrage frequently follows from the assumed existence of an equivalent martingale measure, and the existence of such a measure is not amenable to statistical verification (see Harrison and Kreps (1979), Harrison and Pliska (1981), and Dybvig and Huang (1988)). While statistical tests of various versions of the *efficient market hypothesis* have appeared, none of these constitute a test of no-arbitrage (see, e.g., Taylor (1986) and Malkiel (1990)).

Proposition 4.2 makes it possible to test the no-arbitrage hypothesis in actual markets. Here we test it in the U.S. equity market over the 30 year period from 1967 to 1996. The analysis we present is exploratory in the sense of Tukey (1977), and our conclusions should not be considered definitive. Nevertheless, our results provide some insight into the validity of this hypothesis which until now has never been subjected to scientific verification.

Since the composition of actual equity markets is continually changing, certain modifications must be made to Proposition 4.2 in order that it be applicable. Changes in the value of the generating function $\mathbf{S}(\mu(t))$ in (4.13) that are caused by changes in the composition of the market or by corporate actions such as takeovers and breakups must be excluded since they have no effect on the returns of the portfolios π and μ. Hence we include only return induced changes in the generating function. This exclusion has the same purpose as adjustments made to the "divisor" of a stock index. However, this also means that any natural lower bound on $\mathbf{S}(\mu(t))$ such as (4.8) will no longer be valid, and some other argument will be needed to establish (4.14) in Proposition 4.2.

The data used for our analysis were the monthly stock data from the Center for Research in Security Prices at the University of Chicago (CRSP). The market portfolio consisted of all the U.S. exchange traded and NASDAQ stocks with market weight greater than one half of a basis point, after the removal of REITs, closed-end funds, and ADRs not included in the S&P 500 Index. There were about 6000 stocks in this market. The generating function of the portfolio π was $\mathbf{D}_p$ defined in Theorem 4.1 with $p = .5$. The market weights were calculated each month using the monthly closing prices and shares outstanding, and π was generated from the market weights.

Table 1: Relative return decomposition, 1967–1996
Annual means (logarithmic)

Relative return	Change in $\mathbf{D}_p$	Drift process	Differential dividend rate
1.57%	−.45%	2.45%	−.43%

The results of the data analysis are presented in Table 1. All the values are logarithmic and are in the form of annual averages. The relative return is equal to the difference of the average annual logarithmic return of the portfolio π minus the average annual logarithmic return of the market portfolio. The three components add up to this difference (up to roundoff error). As can be seen, the contribution of the drift process is substantially greater than either of the other two components of the relative return.

Figures 1 through 4 show the cumulative values of the components of the relative return. Figure 1 shows that π generated about 47% higher logarithmic return than the market over the 30 year period, and this resulted in the 1.57% average relative return in Table 1.

The change in $\mathbf{D}_p$ in Figure 2 was responsible for almost all of the volatility of the relative return, but over the 30 years total change was only about −14%. The negative change in $\mathbf{D}_p$ is small compared to total variation over the period ($t = -.58$), and the time series appears to be mean-reverting and stable. $\mathbf{D}_p$ is a measure of diversity, and there is reason to believe that the diversity of a large equity market like that of the U.S. is likely to remain stable over time (see Fernholz (1999a) and Fernholz (1999b)).

The drift process in Figure 3 is close to a pure trend process ($t = 54.$) that added more than 73% to the relative return of π over the period studied. Since $\Theta(t) = (1-p)\gamma^*_\pi(t)$, the slope of this process depends on the relative variances of the stocks in the market. From the looks of Figure 3, these relative variances appear to have been fairly stable over the period.

The contribution of the differential dividend rate in Figure 4 was significantly less than that of the drift process (two-sample $t = 44.$). Although the differential dividend rate was flat for the first few years, after 1975 it averaged about .57% a year in favor of the market portfolio. But to neutralize the drift process, it would have to have been more than four times greater even after 1975.

The evidence from this exploratory analysis indicates that the change in $\mathbf{D}_p$ is mean-reverting and stable with no significant trend, and that the rate of increase of the drift process is significantly greater than the rate of decrease of the differential dividend process. If we accept the results of this analysis, then over the period studied, the conditions of Proposition 4.2 are satisfied, and we must conclude that the no-arbitrage hypothesis did not accurately represent the U.S. equity market.

6 Conclusion

The no-arbitrage hypothesis must be subjected to the same verification process as any other scientific principle. We have presented prima facie evidence that the no-arbitrage hypothesis has been invalid for the U.S. equity market. We have shown theoretically that no-arbitrage fails under certain conditions, and the evidence indicates that these conditions have prevailed in the U.S. equity market.

References

Blair, J. M. (1972). *Economic Concentration*. New York: Harcourt Brace.

Duffie, D. (1992). *Dynamic Asset Pricing Theory*. Princeton, NJ: Princeton University Press.

Dybvig, P. and C. F. Huang (1988). Nonnegative wealth, absence of arbitrage, and feasible consumption plans. *Review of Financial Studies 1*, 377–401.

Fernholz, R. (1999a). On the diversity of equity markets. *Journal of Mathematical Economics 31*(3), 393–417.

Fernholz, R. (1999b). Portfolio generating functions. In M. Avellaneda (Ed.), *Quantitative Analysis in Financial Markets*, River Edge, NJ. World Scientific.

Fernholz, R., R. Garvy, and J. Hannon (1998, Winter). Diversity-weighted indexing. *Journal of Portfolio Management 24*(2), 74–82.

Fernholz, R. and B. Shay (1982). Stochastic portfolio theory and stock market equilibrium. *Journal of Finance 37*, 615–624.

Harrison, J. M. and D. Kreps (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory 20*, 381–408.

Harrison, J. M. and S. Pliska (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and Their Applications 11*, 215–260.

Jarrow, R. and D. B. Madan (1997). Is mean-variance analysis vacuous: or was beta still born? *European Finance Review 1*, 14–24.

Karatzas, I. (1996). *Lectures on the Mathematics of Finance*. Providence, RI: American Mathematical Society.

Karatzas, I. and S. G. Kou (1996). On the pricing of contingent claims under constraints. *The Annals of Applied Probability 6*, 321–369.

Karatzas, I., J. P. Lehoczky, and S. E. Shreve (1990). Existence and uniqueness of multi-agent equilibrium in a stochastic, dynamic consumption/investment model. *Mathematics of Operations Research 15*, 80–128.

Karatzas, I. and S. Shreve (1991). *Brownian Motion and Stochastic Calculus*. New York: Springer-Verlag.

Malkiel, B. (1990). *A Random Walk Down Wall Street*. New York: Norton.

Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*, 373–413.

Sharpe, W. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance 19*, 425–442.

Smith, A. (1776). *The Wealth of Nations.* Modern Library Edition (1994). New York: Random House.

Taylor, S. (1986). *Modelling Financial Time Series*. Chichester: John Wiley & Sons.

Tukey, J. W. (1977). *Exploratory Data Analysis*. Reading, MA: Addison-Wesley.

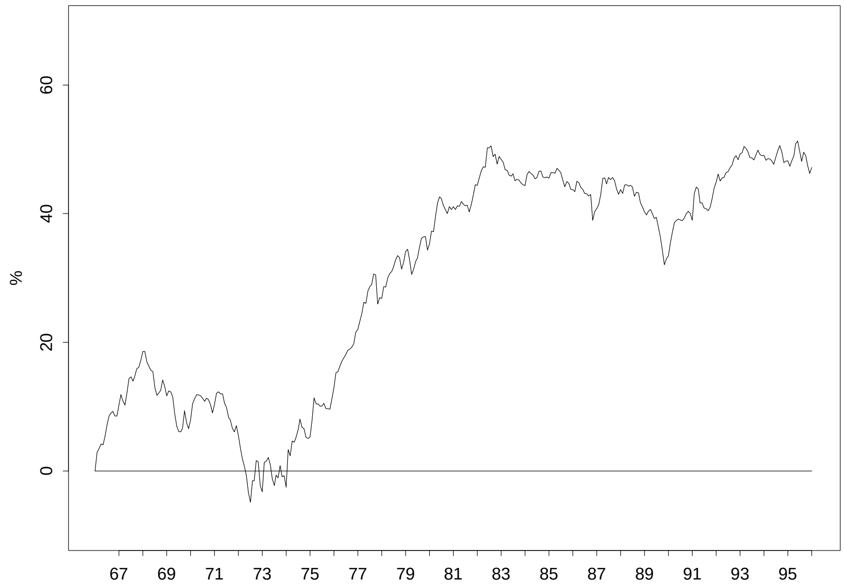

Figure 1: Relative return (logarithmic), 1967–1996

Figure 2: Change in $\mathbf{D}_p$, 1967–1996

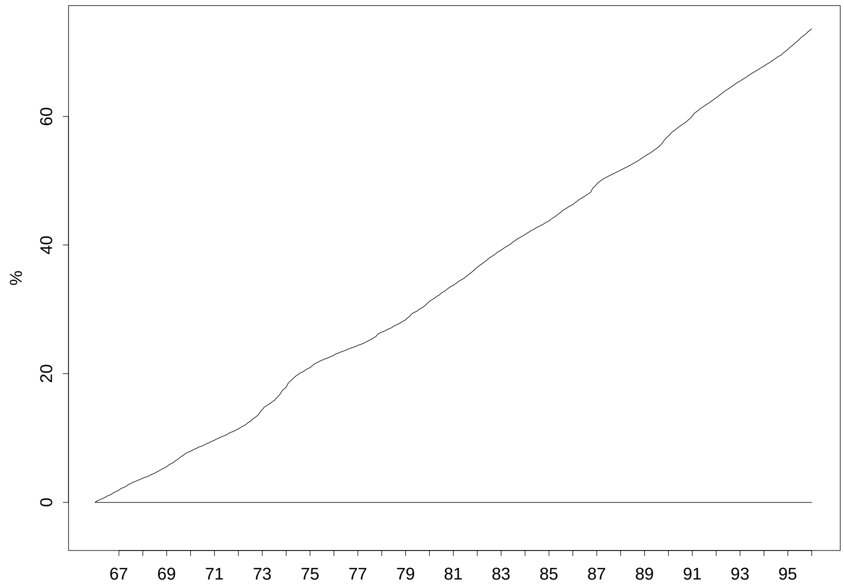

Figure 3: Drift process, 1967–1996

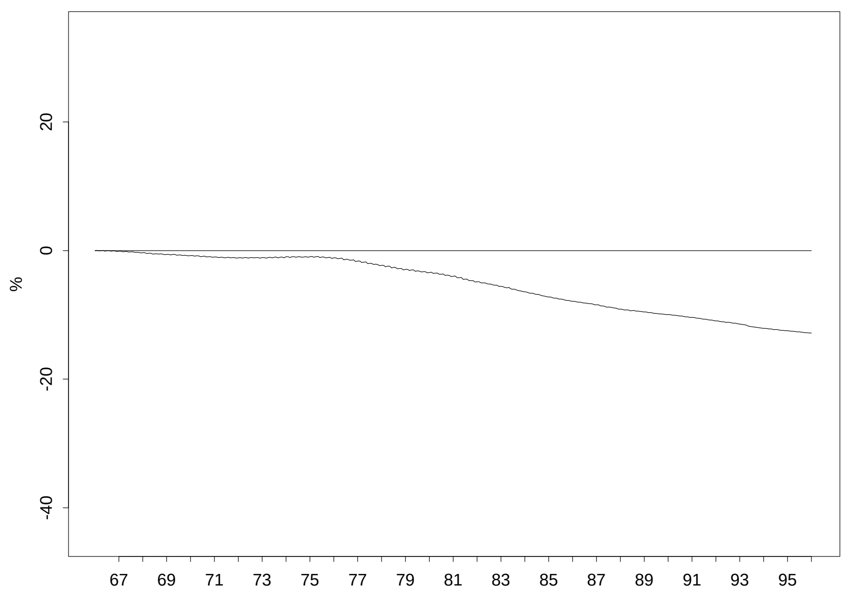

Figure 4: Differential dividend rate, 1967–1996