

Optimal Convergence Trading

Vladislav Kargin*

January 14, 2004

Abstract

This article examines arbitrage investment in a mispriced asset when the mispricing follows the Ornstein-Uhlenbeck process and a credit-constrained investor maximizes a generalization of the Kelly criterion. The optimal differentiable and threshold policies are derived. The optimal differentiable policy is linear with respect to mispricing and risk-free in the long run. The optimal threshold policy calls for investing immediately when the mispricing is greater than zero with the investment amount inversely proportional to the risk aversion parameter. The investment is risky even in the long run. The results are consistent with the belief that credit-constrained arbitrageurs should be risk-neutral if they are to engage in convergence trading.

Myron [Scholes] once told me they are sucking up nickels from all over the world. But because they are so leveraged that amounts to a lot of money.

Merton Miller about the essence of arbitrage.

1 Introduction

Arbitrageurs are people who detect inconsistencies in asset prices and invest in them hoping that the inconsistencies will be eliminated. The waiting time is often uncertain and since the arbitrageur depends on the willingness of other people to lend him money, the irrationality of creditors may lead to great debacles long before prices converge to consistent values. The notorious story of the arbitrage fund LTCM that lost 90 percent of its value on “riskless” deals illustrates the importance of credit constraints. So what policy should the arbitrageur pursue when creditors impose borrowing constraints? In particular, can the arbitrageur allocate the available funds in such a way as to eliminate all the long-run risk?

If this risk elimination is possible, the mispricings should be equally attractive to risk-averse as well as risk-neutral investors. However, a popular view

*Cornerstone Research, 599 Lexington Avenue, New York, NY 10022; slava@bu.edu

asserts that arbitrageurs, unlike other investors, should be risk-neutral. Is there any ground for this belief? The present article provides a justification by solving for arbitrageurs' optimal policies under several types of constraints and showing that under some of them the long-run risk cannot be eliminated. Risk-averse investors that face those constraints are not interested in small mispricings.

The paper investigates two classes of constraints that lead to strikingly different results. Under constraints from the first class, the arbitrageur can only change the leverage slowly. In practice, borrowing additional funds takes time: The arbitrageur must apply for new credit, provide an explanation why he needs it and wait for a decision. Depending on the situation, the process might take from several minutes to several days. In addition a rapid increase in a position adversely affects prices, so in their own interests arbitrageurs must accumulate positions slowly.

For this class of policies, the main result is that the optimal policy is linear in the mispricing, and independent of the coefficient of risk aversion. The variance of the portfolio wealth does not grow with time. Thus, under this constraint the long-run risk can be expunged.

Constraints of the second type are stronger: The arbitrageur cannot change leverage except by closing the position. The motivation is that the arbitrageur is often restricted in his ability to change the leverage – even if the need arises. Higher leverage is mostly needed when mispricing is increasing and the investment account shows negative performance. Unfortunately, this is the worst time to ask for new credit because the creditors hate to invest in accounts with negative performance. As Mark Twain said: “A banker is a fellow who lends you his umbrella when the sun is shining but wants it back the minute it rains.”

For policies in this class, the main result is that the long-run risk cannot be completely removed. Consequently, the arbitrageur will invest an amount that is inversely proportional to his risk aversion.

These two examples suggest that what makes the convergence arbitrage risky in the long run is the inability of the arbitrageur to change the investment amount after the investment is committed. In particular, the results of the second example are consistent with the belief that arbitrageurs should be risk-neutral if they are to engage in convergence trading.

The results of the present paper match closely with results of Grossman and Vila (1992), who study the dynamic investment under a constraint on investment amount. They find that the constraint essentially makes the investor behave as if he were more risk-averse than he actually is. Unlike in the present paper, however, the asset process is not mean-reverting in Grossman and Vila (1992), so the investor could not hope to eliminate the risk completely. Also the constraint is not exogenous as in the present paper but a function of the investor's wealth. Because of these differences it is difficult to conclude whether the similarity of results is incidental or not. Both papers, however, support the view that certain constraints increase long-run riskiness of investment projects.

In a recent paper about convergence trading, Liu and Longstaff (2000) use the Brownian bridge to model the mispricing process, an assumption on the process that requires a fixed horizon at which the mispricing will be effaced. By

the nature of their model, they cannot draw conclusions about long-run risks but they do find that arbitrageurs sometimes cannot eliminate all risk at the end of the arbitrage period. This result is consistent with results of the present paper.

Closely related is the literature on optimal dynamic investments with risky assets. The seminal contributions to this literature were made in Samuelson (1969), Merton (1969), Merton (1971) and Merton (1973). Optimal investment in assets that follow a diffusion process with mean-reverting returns was analyzed in Kim and Omberg (1996), Brennan et al. (1997), Campbell and Viceira (1999), Barberis (2000), Wachter (2002), Campbell et al. (2003). The focus of the present paper is not on general risky investments but on the optimal extraction of profit from near arbitrage opportunities. Consequently, the paper comes to more definite conclusions by using a generalization of the Kelly investment criterion, which emphasizes the long-run behavior of portfolios and especially suitable to modelling objectives of large institutional traders.

In addition, this paper computes the optimal leverage using a new method. While Kim and Omberg ingeniously solve dynamic programming equations by reducing them to a system of non-linear ordinary differential equations, Campbell and Viceira derive an approximate solution by linearization of Euler equations, and Wachter uses the martingale method of Cox and Huang (1989) to separate consumption and financing decisions, the present paper derives the solution by methods of stochastic control, taking the advantage of the asymptotic investment criteria.

The rest of the paper is organized as follows. Section 2 describes the model. Sections 3 and 4 derive the optimal differentiable and threshold policies and describe their properties. Section 5 compares results for differentiable and threshold policies and concludes.

2 Model

An investor can invest in a mispriced asset whose mispricing is measured by $x = \ln(p_1/p_2)$. Here p_1 and p_2 are the asset's actual and "correct" prices, respectively. Mispricing follows the Ornstein-Uhlenbeck process:

$$dx_t = -\alpha x_t dt + \sigma dz_t, \quad (1)$$

where x_t is mispricing at time t , z_t is a standard Wiener process, $\sigma > 0$ and $\alpha > 0$. Parameter α measures the speed of reversion to the correct price: The higher α is, the faster mispricing x drifts towards zero. Parameter σ measures the size of new mispricing shocks introduced into the process. It is also useful to define $\Sigma = \sigma^2/(2\alpha)$, which is the variance of x_t in the long run.

Changes in mispricing induce changes in the arbitrageur's wealth through his choice of the leverage coefficient $f(x)$: The change in the logarithm of wealth is the product of the leverage coefficient and the change in the mispricing,

$$du = f(x)dx. \quad (2)$$

Intuitively, a 1% change in the mispricing results in a $f(x)\%$ change in the investor's wealth. Later we will impose certain restriction on the available leverage.

The arbitrageur's utility is a linear combination of the growth rates in the expectation and variance of the portfolio wealth:

$$U = \liminf_{t \rightarrow \infty} \frac{1}{t} [E(u_t) - \gamma \text{Var}(u_t)], \quad (3)$$

where parameter γ measures risk aversion of the investor. The optimization problem is to choose the leverage function $f(x)$ so that utility U is maximized.

What is the meaning of this maximization criterion? If γ is 0, then the criterion is the same as the criterion of maximizing the portfolio's long-run growth rate – the Kelly criterion. When $\gamma > 0$, it introduces an additional term penalizing deviations from the expected growth rate. This additional term assures the investor that maximizing U protects him against the large deviations in the realized growth of his portfolio from the expectations.

An example may perhaps add some insight into the investment criterion. Suppose that the wealth of the investor follows a geometric Brownian motion with constant parameters μ and σ . Then the utility of the investor is

$$U = \mu - \gamma\sigma^2. \quad (4)$$

This expression shows that the utility depends only on the parameters of the process and on risk aversion but not on the investment horizon.

Another way to get an insight into this criterion is to compare it with the objective under the classical single period Markowitz model. In the Markowitz model the investor maximizes a linear combination of the expectation and the variance of the portfolio return. Therefore, the present model generalizes the Markowitz model to the dynamic setting by substituting the expectation and variance of the single period return with the asymptotic rates of increase in expectation and variance of the investor's wealth.

An important assumption that we adopt in this generalization is that the investor is concerned only with long-run consequences of his policy. This assumption simplifies the analysis considerably and appears to be realistic for small investments by large institutions. In using this assumption we follow Bielecki and Pliska (1999) and Bielecki et al. (2000), who applied it to the analysis of continuous investment policies in a similar situation.

3 Optimal Differentiable Policy

This section is about differentiable policies $f(x)$, for which

$$f \in C^1(-\infty, +\infty), \text{ and} \quad (5)$$

$$|f'(x)| \leq K. \quad (6)$$

The policies from this class will be called **D**-policies. This class excludes policies that prescribe extremely rapid growth of leverage with respect to mispricing. The following theorem is a cardinal ingredient in showing that optimal **D**-policies are linear.

Theorem 3.1 *Linear investment policies are the only **D**-policies such that the variance of the logarithm of wealth u_t is asymptotically constant.*

The proof uses a convenient representation for u : Let

$$g(\xi) =: \int_0^\xi f(\zeta) d\zeta. \quad (7)$$

Then it is easy to check that

$$u_t = u_0 + g(x_t) - g(x_0) - \frac{\sigma^2}{2} \int_0^t f'(x_\tau) d\tau. \quad (8)$$

The intuition behind this representation is simple. The investor can increase his wealth only if he increases his leverage when the mispricing increases. The derivative $f'(x)$ measures the sensitivity of the leverage policy to mispricing, and (8) shows that the change in the logarithm of wealth equals a multiple of the integral of $f'(x)$ plus a stationary process. The more sensitive the leverage policy is to mispricing, the greater the increase in the wealth induced by local variations in mispricing. The addition of $g(x_t) - g(x_0)$ reflects dependence of the wealth on the initial and final conditions.

Proof of Theorem 3.1: Taking the variance of $u_t - u_0$ in (8) gives

$$\text{Var}(u_t - u_0) = \text{Var}(g(x_t)) - \sigma^2 \text{Cov}\left(g(x_t), \int_0^t f'(x_\tau) d\tau\right) + \frac{\sigma^4}{4} \text{Var}\left(\int_0^t f'(x_\tau) d\tau\right). \quad (9)$$

As t increases, all terms except possibly the third one tend to a finite limit. So, asymptotically,

$$\text{Var}(u_t) \sim \text{const} + \frac{\sigma^4}{4} r t, \quad (10)$$

where

$$r =: \lim_{t \rightarrow \infty} \frac{\text{Var} \int_0^t f'(x_\tau) d\tau}{t} \geq 0. \quad (11)$$

The rate $r = 0$ if and only if $\text{Var}(f'(x)) = 0$. Indeed, if $\text{Var}(f'(x)) > 0$ then

$$r(t) = \frac{1}{t} \text{Var}(f'(x)) \int_0^t \int_0^t \text{Corr}(f'(x_\tau), f'(x_s)) d\tau ds. \quad (12)$$

According to Proposition A.1 in Appendix A, $\text{Corr}(f'(x_\tau), f'(x_s)) > 0$, so it follows from (12) that

$$r(t) \geq \frac{1}{t} \text{Var}(f'(x)) \int_0^t 1 d\tau = \text{Var}(f'(x)) > 0. \quad (13)$$

Finally, since $\text{Var}(f'(x)) = 0$ if and only if $f'(x)$ is almost surely constant, the investment policy must be linear if the variance of u_t is not to increase with time. QED.

Theorem 3.1 shows that all linear strategies eliminate long-run risk. As a natural consequence, the next theorem predicates optimality of linear strategies with respect to the asymptotic investment criterion. The idea is to match any non-linear strategy with an admissible linear strategy that has higher expected return and lower growth in variance. The matching is possible exactly because all linear strategies have zero asymptotic growth in variance.

Theorem 3.2 *For the investor with asymptotic preferences, any \mathbf{D} -policy is dominated by some linear \mathbf{D} -policy.*

Proof: Let the non-linear policy be $f(x)$. According to (8), in the long run

$$E(u_t) = u_0 - \frac{\sigma^2}{2}E(f'(x))t. \quad (14)$$

Take the linear policy $f_L(x) = (E(f'(x)) - \varepsilon)x$ with $\varepsilon > 0$. For a certain ε it is admissible. This is because $|E(f'(x)) - \varepsilon| < K$ follows from $|f'(x)| \leq K$ everywhere and $|f'(x)| < K$ on a set of positive measure, which are both true because f is a non-linear \mathbf{D} -policy. The expectation of the logarithm of wealth under f_L is

$$E(u_t) = u_0 - \frac{\sigma^2}{2}E(f'(x))t + \frac{\sigma^2}{2}\varepsilon t. \quad (15)$$

It is clearly higher than the corresponding expectation for the non-linear policy. From Theorem 3.1 we know that the variance of u_t is asymptotically constant for linear policies and is asymptotically equivalent to rt , where $r > 0$, for non-linear policies. It follows that for sufficiently large T the linear policy f_L will have lower $\text{Var}(u_T)$ than the non-linear policy f . Thus f_L asymptotically dominates f . QED.

It remains to find the optimal policy in the class of linear policies. It turns out that it is the policy that has the maximal sensitivity to mispricing. This is intuitively clear because every linear strategy eliminates the long-run risk, and the policy with the largest sensitivity to mispricing has the highest expected return. Formally, the following theorem holds.

Theorem 3.3 *The optimal linear \mathbf{D} -policy is $f(x) = -Kx$.*

Proof: it is easy to compute

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{E(u_t)}{t} &= \frac{\sigma^2}{2}k, \\ \lim_{t \rightarrow \infty} \frac{\text{Var}(u_t)}{t} &= 0, \\ U &= \frac{\sigma^2}{2}k. \end{aligned}$$

Thus, the utility is maximized by the maximal possible k , from which the theorem follows. QED.

The theorem implies that the arbitrageur should increase the leverage at the maximal possible rate. In particular, the optimal strategy does not depend on the risk aversion parameter or properties of the mispricing process. The intuitive meaning of this conclusion is that the appropriate use of leverage allows the arbitrageur to eliminate all the long run risk. As the next section shows, this conclusion will be reversed if the arbitrageur is more constrained in the use of leverage.

4 Optimal Threshold Policy

When an arbitrageur uses **threshold policies** he keeps his finger on a button that triggers investment while looking at the computer monitor and waiting for a mispricing. If he observes a mispricing that exceeds a threshold, S , he pushes the button and a fixed amount, L , is directed to this opportunity. When the mispricing falls below another threshold, s , he pushes another button and the position closes. Leverage L never changes when the position is opened. **Simple threshold policies** have equal thresholds: $S = s$.

As was said in the Introduction, arbitrageurs use threshold policies because they often cannot secure additional funds for positions they already opened. They also favor threshold policies because these policies allow economizing on transaction costs.

General threshold policies are complicated to analyze. Fortunately, the following theorem shows that it is sufficient to study simple threshold policies.

Theorem 4.1 *Any threshold policy is dominated by a simple threshold policy.*

This theorem is given without proof. Intuitively, for the Markov process of mispricing the optimal investment policy should not depend on the history of investing, and the only threshold policies that pass this selection test are simple threshold policies. Indeed, if $s < S$, and the mispricing is between s and S , then the position is open if the mispricing has fallen from above S but not yet gone below s , and it is closed if the mispricing has risen from below s but not yet gone above S . It follows that investment under the threshold rule with $s \neq S$ depends on history of investment and therefore cannot be optimal.

The relevant properties of the simple threshold policies are described in the next theorem, which uses the following notation:

$$\phi(S) = \frac{1}{\sqrt{2\pi\Sigma}} \exp\left(-\frac{S^2}{2\Sigma}\right), \quad (16)$$

$$\psi(S) = \frac{1}{\sqrt{2\pi\Sigma}} \frac{2}{\alpha} \int_0^1 \frac{1}{\xi} \left[\frac{1}{\sqrt{1-\xi^2}} \exp\left(\frac{S^2}{\Sigma} \frac{\xi}{1+\xi}\right) - 1 \right] d\xi. \quad (17)$$

For $S = 0$, the value of $\psi(S)$ can be computed explicitly: $\psi(0) = 2 \ln 2 / (\alpha \sqrt{2\pi\Sigma})$.

[Figure 1 here]

Figure 1: Plots of Optimal Leverage L as Function of Threshold S

[Figure 2 here]

Figure 2: Plots of Optimal Utility U as Function of Threshold S

Theorem 4.2 *For the simple threshold policy with threshold S and leverage L*

$$\lim_{t \rightarrow \infty} \frac{E(u_t - u_0)}{t} = c_1(L, S) \equiv \sigma^2 L \phi(S)$$

$$\lim_{t \rightarrow \infty} \frac{\text{Var}(u_t - u_0)}{t} = c_2(L, S) \equiv (\sigma^2 L \phi(S))^2 \psi(S).$$

The investor's utility is $U(L, S) = c_1(L, S) - \gamma c_2(L, S)$.

The proof is relegated to Appendix B.

The first step in obtaining the optimal policy from this theorem is to calculate reduced utility function that depends only on threshold S :

Corollary 1 *For a fixed threshold S the optimal leverage is*

$$L(S) = \frac{1}{4\gamma\alpha\Sigma\phi(S)\psi(S)}$$

and the corresponding utility is

$$U(S) = \frac{1}{4\gamma\psi(S)}.$$

The functions $L(S)$ and $U(S)$ are illustrated in Figures 1 and 2. We can see that higher long-run variance Σ leads to an increase in both leverage L and utility U . Higher persistence of the process does not change optimal leverage but decreases utility.

The function $\psi(S)$ is increasing in S^2 , and consequently the maximal utility is reached at $S = 0$. The optimal threshold policy is given in the following theorem.

Theorem 4.3 *Utility is maximized for $S = 0$ and $L = \pi/(4\gamma \ln 2)$. The optimal utility is $U = \alpha\sqrt{2\pi\Sigma}/(8\gamma \ln 2)$.*

Predictably, the utility is higher when the convergence is fast (α is high) and the arbitrage opportunity is large (Σ is high). Not so predictable is that the optimal leverage does not depend on the parameters of the process: This leverage optimally balances risk and return for every Ornstein-Uhlenbeck process. What is most important, however, is that the optimal leverage depends on the parameter of the risk aversion γ . The higher γ is, the lower the amount is that the arbitrageur is willing to commit to the arbitrage opportunity: The arbitrageur that uses only threshold strategies is unable to remove the long-run risk and must adjust his behavior.

[Table 1 here]

Table 1: Summary Statistics of Mispricing Factor

[Table 2 here]

Table 2: Results of Estimation of Mispricing Factor Process

5 Empirical Application

This section studies convergent trading in the context of WEBS, which are shares of open-end mutual funds that replicate the price performance of foreign stock market indexes. WEBS trade on a stock exchange like ordinary stock, and their managers try to keep the fund price close to the net asset value (NAV) of their underlying stocks. They buy back shares if the price is less than NAV and issue additional shares if the price is greater than NAV.

As in previous sections, I assume that the investor can hedge the risk of the underlying portfolio. Indeed, trading index futures provides a very good hedge of country exposure. Absence of the perfect hedge limits the implications of my analysis.

I use price and NAV daily data for WEBS that track indices of Australia, Austria, Belgium, Canada, France, Germany, Hong Kong, Italy, Japan, Malaysia, Mexico, Netherlands, Singapore, Spain, Sweden, Switzerland, and the United Kingdom. This is a total of 17 countries. The data start in March of 1996 and end in August of 2000, which gives around 1000 datapoints for each country.

The mispricing factor x_t is computed as the logarithm of the ratio of the price to NAV. Some summary statistics for x_t are given in Table 1. I model the dynamic of the mispricing factor as an AR(1) process, which is the discrete-time counterpart to the Ornstein-Uhlenbeck process:

$$x_t = \beta x_{t-1} + \sigma \epsilon_t, \quad (18)$$

The results of estimation of this process are summarized in Table 2. They show that β is around 0.5, and σ is around 0.01. Durbin-Watson statistic shows that the AR(1) process is a reasonably good approximation to the true mispricing process.

Our first goal is to get an estimate of the order of the coefficient k in the optimal linear strategy. Let us use the following estimates of the order of parameters: $\sigma^2 \sim 10^{-4}$, $\alpha \sim 0.5$. From Theorem 3.3, the long-run variance of the logarithm of wealth is $5 \cdot 10^{-9} k^2$, and the expected change in the logarithm of wealth is $5 \cdot 10^{-5} k$ per day. The average daily change in the logarithm of the S&P500 index has been $8 \cdot 10^{-4} \sim 10^{-3}$ in the last five years. To get this return by convergent trading, the sensitivity k of the linear strategy would have to be set at 20. The corresponding asymptotic variance of the logarithm of wealth would then be $2 \cdot 10^{-8}$. This is considerably smaller than the variance of the deviation of the logarithm of the S&P500 index from its linear trend, which for

[Figure 3 here]

Figure 3: Contour Graph of Mean of Average Daily Returns

[Figure 4 here]

Figure 4: Contour Graph of Standard Deviation of Average Daily Returns

the 5-year period from August 1995 to August 2000 can be estimated at about $3 \cdot 10^{-3}$.

These computations suggest that this market could not be efficient if transactions costs were absent. The goal of the following Monte-Carlo simulations is to analyze properties of threshold strategies and to find the optimal threshold strategy in situations with transaction costs. I assume that the transaction cost c is 0.25%, and that the true process of discounts is AR(1) with $\beta = 0.3$ and $\sigma = 10^{-2}$.

The simulations were organized as follows. One hundred realizations of the mispricing factor process were generated. Each realization had 1250 datapoints. For each realization I simulated the process of investing according to a strategy from a finites set of threshold strategies. Thus, each pair of a strategy and a realization of the mispricing factor process resulted in a realization of wealth process.

For each realization of wealth process, I calculated the average daily return as a difference between logarithm of final wealth and logarithm of initial wealth divided by the length of the realization. After that, I calculated the mean and the variance of this statistic over all realizations of the wealth process that corresponded to a particular strategy. Thus, as a final product I had a function that mapped each strategy into the expectation and variance of the average daily return.

I used the following set of strategies. A rise in the mispricing over a threshold S triggers the opening of the position. When the mispricing returns to the region below $s \in [0, S]$, the arbitrageur closes his position and waits for a new trigger signal. The threshold S was chosen in the range from 0.5% to 2%. The threshold s was chosen in the range from $S + c$ to 0%.

Figure 3 is a contour graph of the mean of the average daily return to a strategy. On the vertical axis of the graph is the low threshold S and on the horizontal axis is the difference between high and low thresholds $S - s$. They are denominated in percentage terms. The lines on the contour graph correspond to the strategies that have the same mean of average daily return.

This graph suggests that the average return is maximized for $s = 0$ and S

[Figure 5 here]

Figure 5: Contour Graph of Sharpe Ratio

set to some $S_0 > c$. Therefore, if the investor wants to maximize average return he should invest only if the mispricing exceeds the transaction cost by some markup.

Figure 4 is a contour graph of the standard deviation of the average daily returns. We can conclude from it that the variance is increasing with an increase in the absolute value of the low threshold and with an increase in the difference between the thresholds.

Figure 5 is a countour graph of the ratio of the mean of average daily return to its standard deviation. It suggests that the ratio is maximized for the strategy that sets $s = 0$ and $S = c$. Thus an investor who uses only threshold strategies and who wants to maximize the Sharpe ratio should invest immediately when the mispricing exceeds his transaction costs.

6 Discussion

Section 3 shows that in the class of differentiable policies with bounded derivative the optimal policy is linear in the mispricing and the coefficient in the linear relationship is the highest possible. The optimal strategy in this case does not depend on the risk-aversion of the arbitrageur, and all the long-run risk can be eliminated.

In contrast, according to the results of Section 4, if only threshold policies are available then the long-run risk is unavoidable and the investment is inversely proportional to risk aversion. This conclusion is consistent with the belief that arbitrageurs are typically risk-neutral. The suggested reason for this belief is that the constraints on flexibility of changes in leverage make the convergence trading risky even in the long run.

A Auxiliary Statistical Result

Suppose that x and y are jointly Gaussian random variables, $\text{Var}(x) = \text{Var}(y) = 1$, and $\text{Cov}(x, y) = \beta$.

Proposition A.1 *If $f \in \mathbf{D}$ and $\text{Var}(f(x)) = 1$, then $\text{Cov}(f(x), f(y)) \in [0, \beta]$.*

Proof: Assume without loss of generality that $Ef(x) = 0$. Since Hermite polynomials are complete in the class of \mathbf{D} -policies, we can use them to approximate f . Then the assertion of Proposition A.1 follows from Proposition A.2.

Proposition A.2 *If f is a polynomial of degree N , $Ef(x) = 0$ and $\text{Var}(f(x)) = 1$, then $\text{Cov}(f(x), f(y)) \in [\beta^N, \beta]$. The maximum and minimum are achieved by $f(x) = x$ and $f(x) = H_N(x)$, respectively, where $H_N(x)$ is the Hermite polynomial of degree N .*

Proof: Represent $f(x)$ as a sum of Hermite polynomials:

$$f(x) = \sum_1^N a_k H_k(x), \quad (19)$$

where by definition

$$H_k(x) = \exp\left(\frac{x^2}{2}\right) \frac{(-1)^k}{k!} \frac{d^k}{(dx)^k} \left[\exp\left(-\frac{x^2}{2}\right) \right]. \quad (20)$$

Hermite polynomials form an orthonormal system with respect to the Gaussian kernel and possess the following useful property:

$$\text{Cov}(H_i(x), H_j(y)) = \beta^i \delta_{ij}. \quad (21)$$

Using this property and orthonormality, we can write

$$\text{Cov}(f(x), f(y)) = \sum_1^N a_k^2 \beta^k \quad (22)$$

$$\text{Var}(f(x)) = \sum_1^N a_k^2. \quad (23)$$

From (22) and (23), the maximum of $\text{Cov}(f(x), f(y))$ is β and it is achieved by $f(x) = H_1(x) = x$. The minimum is β^N and it is achieved by $f(x) = H_N(x)$. QED.

B Proof of Theorem 4.2

Proof: By definition, the threshold policy is

$$f(x) = \begin{cases} -\text{sign}(x)L & \text{if } |x| \geq S, \\ 0 & \text{if } |x| < S. \end{cases} \quad (24)$$

The generalized Ito formula gives

$$u_t = u_0 + g(x_t) - g(x_0) + \frac{\sigma^2}{2} 2L \int_0^t \delta_S(x_\tau) d\tau, \quad (25)$$

where δ_S is the Dirac delta-function and

$$g(\xi) =: \int_0^\xi f(\zeta) d\zeta. \quad (26)$$

The intuition behind this representation is simple: The investor increases his wealth only when he triggers the policy. The number of times the policy is triggered is stochastic and measured by the integral of the delta function. The profit earned at each occasion is proportional to the product of local volatility σ^2 and leverage L . Finally, there is a dependence of wealth on initial and final conditions which is captured by $g(x_t) - g(x_0)$.

Since $g(x_t)$ does not grow with time, the arbitrageur's utility depends only on the moments of the integral of the delta function:

$$\int_0^t \delta_S(x_\tau) d\tau. \quad (27)$$

The first step in the computation of the moments is calculating the expectation and covariance function of the generalized stochastic process $\delta_{t,S} =: \delta_S(x_t)$. The joint density of x_{t_1} and x_{t_2} is

$$p(x_1, x_2) = \frac{1}{2\pi\Sigma\sqrt{1-a(\tau)^2}} \exp \left\{ -\frac{1}{2\Sigma} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}' \begin{pmatrix} 1 & a(\tau) \\ a(\tau) & 1 \end{pmatrix}^{-1} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \right\}, \quad (28)$$

where $\tau = t_2 - t_1$ and $a(\tau) = e^{-\alpha|\tau|}$. The delta function can be approximated by $\frac{1}{\Delta} \chi_{[S, S+\Delta]}$, where χ_A denotes the characteristic function of set A and Δ limits to 0. Then, computing two first moments for $\chi_{[S, S+\Delta]}(x_t)$ and taking the limit $\Delta \rightarrow 0$ give the following formulas:

$$E(\delta_{t_1, S}) = \frac{1}{\sqrt{2\pi\Sigma}} \exp \left(-\frac{S^2}{2\Sigma} \right) \text{ and} \quad (29)$$

$$E(\delta_{t_1, S} \delta_{t_2, S}) = \frac{1}{2\pi\Sigma\sqrt{1-a(\tau)^2}} \exp \left(-\frac{S^2}{\Sigma} \frac{1}{1+a(\tau)} \right). \quad (30)$$

For example, equality (30) can be seen from the following calculation:

$$\begin{aligned}
E(\delta_{t_1,S}\delta_{t_2,S}) &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \chi_{[S,S+\Delta]}(\xi_1)\chi_{[S,S+\Delta]}(\xi_2)p(\xi_1,\xi_2)d\xi_1d\xi_2 \\
&= \lim_{\Delta \rightarrow 0} \frac{1}{2\pi\Sigma\sqrt{1-a(\tau)^2}} \frac{1}{\Delta^2} \int_S^{S+\Delta} \int_S^{S+\Delta} \exp\left[-\frac{\xi_1^2 + \xi_2^2 - 2a(\tau)\xi_1\xi_2}{2\Sigma(1-a(\tau)^2)}\right] d\xi_1d\xi_2 \\
&= \frac{1}{2\pi\Sigma\sqrt{1-a(\tau)^2}} \exp\left(-\frac{S^2}{\Sigma} \frac{1}{1+a(\tau)}\right).
\end{aligned} \tag{31}$$

From (29) and (30) it follows that

$$\text{Cov}(\delta_{t_1,S}, \delta_{t_2,S}) = \frac{1}{2\pi\Sigma\sqrt{1-a(\tau)^2}} \exp\left(-\frac{S^2}{\Sigma} \frac{1}{1+a(\tau)}\right) - \frac{1}{2\pi\Sigma} \exp\left(-\frac{S^2}{\Sigma}\right). \tag{32}$$

Since $\text{Cov}(\delta_{t_1,S}, \delta_{t_2,S})$ depends only on $\tau = t_2 - t_1$, it can be denoted by $\vartheta(\tau, S)$. Next,

$$\begin{aligned}
\text{Var}\left(\int_0^T \delta_{t,S} dt\right) &= 2 \int_0^T dt_1 \int_{t_1}^T \text{Cov}(\delta_{t_1,S}, \delta_{t_2,S}) dt_2 \\
&\text{(substituting } \tau = t_2 - t_1) = 2 \int_0^T dt_1 \int_0^{T-t_1} \vartheta(\tau, S) d\tau \\
&\text{(changing order of integration)} = 2 \int_0^T (T - \tau) \vartheta(\tau, S) d\tau.
\end{aligned} \tag{33}$$

Since $\vartheta(\tau, S) = O(e^{-c\tau})$ for a positive c and $\tau \rightarrow \infty$, and $\vartheta(\tau, S)$ is integrable around $\tau = 0$, it follows

$$\begin{aligned}
\int_0^T \vartheta(\tau, S) d\tau &\rightarrow \text{const and} \\
\int_0^T \tau \vartheta(\tau, S) d\tau &\rightarrow \text{const}
\end{aligned}$$

as $T \rightarrow \infty$.

So, for large T

$$\begin{aligned}
\text{Var}\left(\int_0^T \delta_{t,S} dt\right) &\sim 2T \int_0^\infty \vartheta(\tau, S) d\tau \\
&= 2T \frac{1}{2\pi\Sigma} \exp\left(-\frac{S^2}{\Sigma}\right) \int_0^\infty \left(\frac{1}{\sqrt{1-a(\tau)^2}} \exp\left(\frac{S^2}{\Sigma} \frac{a(\tau)}{1+a(\tau)}\right) - 1\right) d\tau \\
&\text{(substituting } \tau = -\alpha^{-1} \ln \xi) \\
&= 2T \frac{1}{2\pi\Sigma} \exp\left(-\frac{S^2}{\Sigma}\right) \frac{1}{\alpha} \int_0^1 \frac{1}{\xi} \left(\frac{1}{\sqrt{1-\xi^2}} \exp\left(\frac{S^2}{\Sigma} \frac{\xi}{1+\xi}\right) - 1\right) d\xi \\
&= \phi(S)^2 \psi(S) T.
\end{aligned} \tag{34}$$

This implies all the assertions of the theorem. QED.

References

- Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance* 55(1), 225–264.
- Bielecki, T. R. and S. R. Pliska (1999). Risk sensitive dynamic asset management. *Journal of Applied Mathematics and Optimization* 39, 337–360.
- Bielecki, T. R., S. R. Pliska, and M. Sherris (2000). Risk sensitive asset allocation. *Journal of Economic Dynamics and Control* 24, 1145–1177.
- Brennan, M. J., E. S. Schwartz, and R. Lagnado (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Campbell, J. Y., Y. L. Chan, and L. M. Viceira (2003). A multivariate model of strategic asset allocation. *Journal of Financial Economics* 67(1), 41–80.
- Campbell, J. Y. and L. Viceira (1999). Consumption and portfolio decisions when expected returns are time-varying. *Quarterly Journal of Economics* 114, 433–495.
- Cox, J. and C. Huang (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory* 39, 33–83.
- Grossman, S. J. and J.-L. Vila (1992, June). Optimal dynamic trading with leverage constraints. *Journal of Financial and Quantitative Analysis* 27(2), 151–168.
- Kim, T. S. and E. Omberg (1996, Spring). Dynamic nonmyopic portfolio behavior. *The Review of Financial Studies* 9(1), 141–161.
- Liu, J. and F. A. Longstaff (2000, June). Losing money on arbitrages: Optimal dynamic portfolio choices in markets with arbitrage opportunities. Anderson School at UCLA, Working Paper.
- Merton, R. C. (1969, August). Lifetime portfolio selection under uncertainty: the continuous time case. *The Review of Economics and Statistics* 51(3), 247–257.
- Merton, R. C. (1971). Optimal consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3, 213–214.
- Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica* 41, 867–887.
- Samuelson, P. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics* 51, 239–246.
- Wachter, J. A. (2002). Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. *Journal of Financial and Quantitative Analysis* 37(1), 63–91.

Table 1. Descriptive Statistics for Mispricing Factor

Country	Number of datapoints	MMF	VMF
Australia	1078	0.64	1.11
Austria	901	0.61	1.28
Belgium	652	0.19	1.65
Canada	992	0.28	0.98
France	1065	0.26	0.80
Germany	1083	0.46	0.96
Hong Kong	1089	0.29	2.09
Italy	1081	0.22	0.87
Japan	1096	0.46	1.17
Malaysia	1100	2.33	7.49
Mexico	709	0.46	0.67
Netherlands	1082	-0.12	1.43
Singapore	1093	0.72	2.28
Spain	1012	0.37	0.79
Sweden	992	0.44	0.84
Switzerland	1037	0.61	0.95
United Kingdom	1034	0.60	0.92

$$MMF = 100 \times \frac{1}{T} \sum_{t=1}^T \ln(price_t / NAV_t)$$

$$VMF = 100 \times \left(\frac{1}{T} \sum_{t=1}^T \ln^2(price_t / NAV_t) - MMF^2 \right)$$

Table 2. Estimates of AR(1) Model For Mispricing Factor Process

Country	BETA	SIGMA	Durbin-Watson
Australia	0.28	0.011	2.1
Austria	0.38	0.012	2.14
Belgium	0.08	0.016	2.00
Canada	0.22	0.010	2.07
France	0.20	0.008	2.04
Germany	0.13	0.010	2.01
Hong Kong	0.39	0.019	2.15
Italy	0.24	0.008	2.04
Japan	0.33	0.011	2.08
Malaysia	0.93	0.028	2.43
Mexico	0.07	0.007	2.01
Netherlands	0.54	0.012	2.25
Singapore	0.61	0.018	2.22
Spain	0.13	0.008	2.02
Sweden	0.16	0.008	2.04
Switzerland	0.14	0.009	2.02
United Kingdom	0.23	0.009	2.06

Figure 1. Optimal Leverage K as a Function of Threshold




Figure 2. Optimal Utility U as a Function of Threshold




Figure 3. Contour Graph of Mean of Average Daily Returns

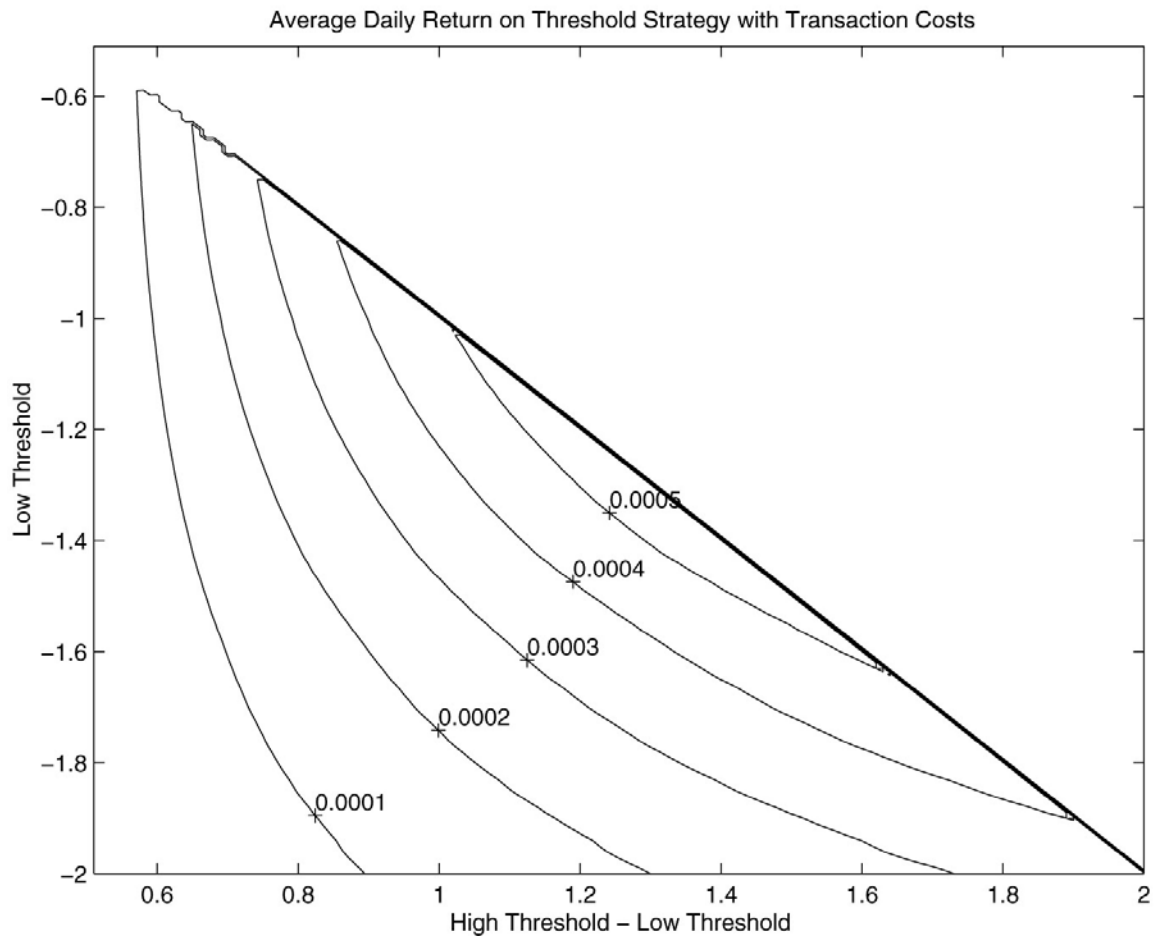


Figure 4. Contour Graph of Mean of Average Daily Returns




Figure 5. Contour Graph of Standard Deviation of Average Daily Return

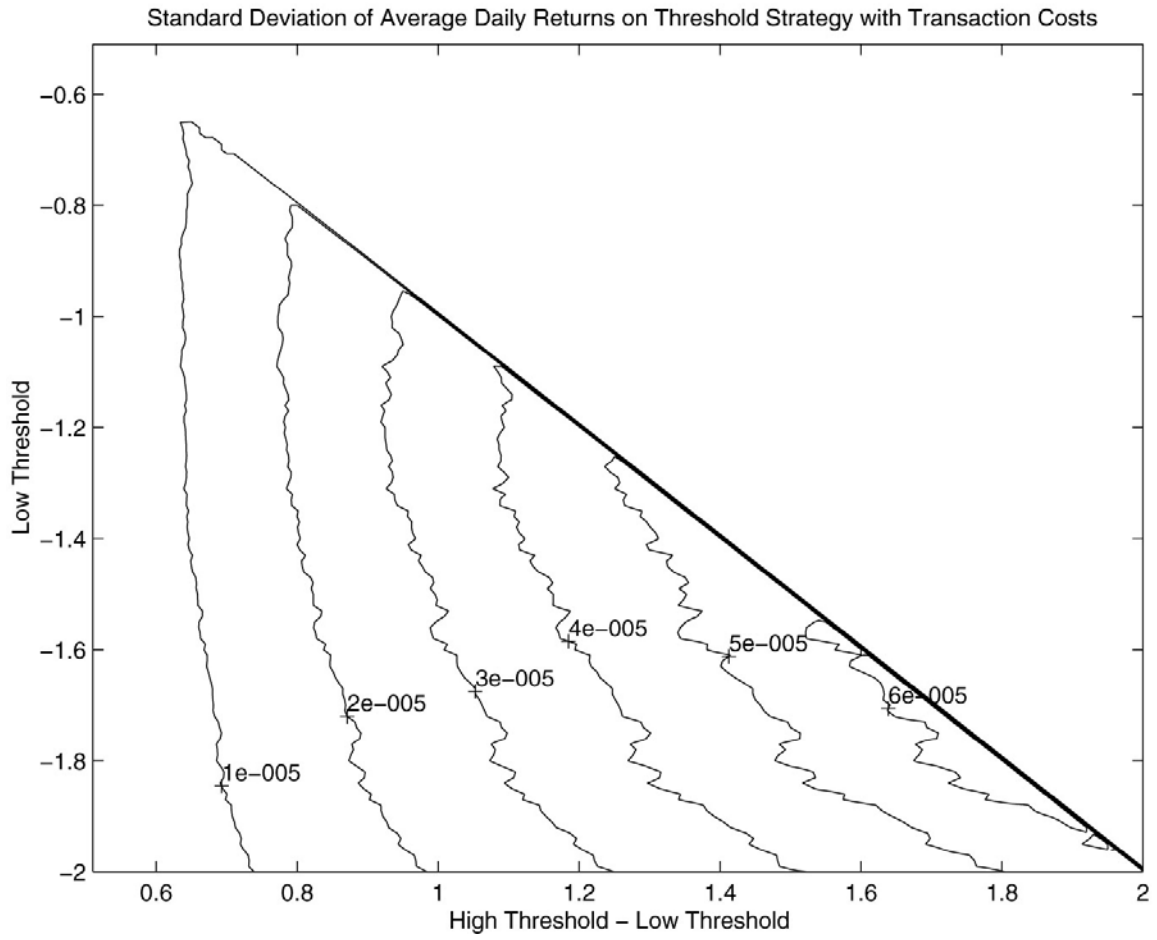


Figure 6. Contour Graph of Sharpe Ratio

