Optimal arbitrage trading

Elena Boguslavskaya[*], Michael Boguslavsky[†]

September 2003

Abstract

We consider the position management problem for an agent trading a mean-reverting asset. This problem arises in many statistical and fundamental arbitrage trading situations when the short-term returns on an asset are predictable but limited risk-bearing capacity does not allow to fully exploit this predictability. The model is rather simple; it does not require any inputs apart from the parameters of the price process and agent's relative risk aversion. However, the model reproduces some realistic patterns of traders' behaviour. We use the Ornstein-Uhlenbeck process to model the price process and consider a finite horizon power utility agent. The dynamic programming approach yields a non-linear PDE. It is solved explicitly, and simple formulas for the value function and the optimal trading strategy are obtained. We use Monte-Carlo simulation to check for the effects of parameter misspecification.

Contents

[*]elena@boguslavsky.net
[†]michael@boguslavsky.net

1 Introduction

1.1 Motivation

Many academic papers about optimal trading rules and portfolio selection assume that the assets follow geometric Brownian motions, or, more generally, random walks. These papers are typically concerned with portfolio selection problems faced by long-term investors. In this paper, we consider a problem where the asset price is driven by a mean-reverting process. With some exceptions (e.g. [Lo]), this kind of processes is not widely used to model stock or bond price dynamics. However similar portfolio selection problems arise naturally in many "relative value" strategies assuming some kind of mean reversion in a tradable asset.

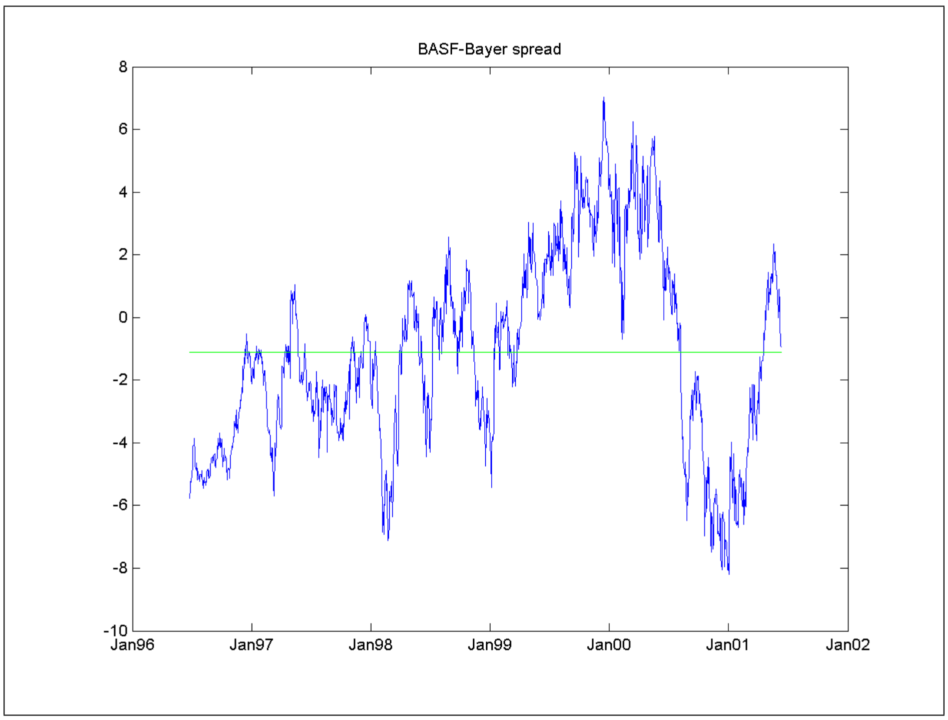

Figure 1: Difference of ordinary share close price for BASF AG and Bayer AG, 1997-2001.

Consider, for example, a limited capital speculator trading the spread (i.e. the difference) between two cointegrated assets or, more generally, an arbitrageur with a limited capital trading a mean-reverting asset. The trader knows the "correct" (long-term average) price of the asset, and he knows that the price will sooner or later revert to the correct level, but the risk is that the position losses may become unbearable for the trader before the reversion happens. The finite horizon assumption is quite realistic because the bonuses to traders and fees to hedge fund managers are usually paid yearly. Just to give an example, Fig. 1 shows the spread of the once famous BASF-Bayer stock pair.

Faced with a mean-reverting process, a trader would typically take a long (i.e. positive) position in the asset when the asset is below its long-term mean and a short (i.e. negative) position when the asset is below the long-term mean. He would then either liquidate the position when the price reverts closer to the mean and take the profit or he might have to close the position before the reversion happens and face the losses. The question is in the size of the position and how the position should be optimally managed as the price and the trader's wealth change and time passes by. An often rule of thumb is that one opens the short position as soon as the spread is above one standard deviation from its mean and a long position as soon as the spread is below one standard deviation.

It is well known that capital and risk-bearing constraints may seriously limit arbitrage activities. Shleifer [Shl] built an equilibrium model for a market with limited capital arbitrageurs.

We solve the optimal problem assuming an Ornstein-Uhlenbeck process for the price and power utility over the final wealth for a finite horizon agent. This model was first formulated for the power utility case and solved for the log-utility case by Mendez-Vivez, Morton, and Naik, [M], [M-VMN].

Besides the quantitative result, there is a number of interesting qualitative questions to answer about the optimal strategy.

- When and how aggressively should one open the position?
- When should one cut a loosing position?
- Can a trader ever be happy when the spread widens against his position?
- What is the effect of process parameters on optimal strategy?
- How does the trading strategy and the value function change as the time horizon approaches?
- What is the effect of risk aversion on position term dependence?
- How the process parameters uncertainty affects the optimal strategy?

We answer these questions in section 3.

1.2 Choice of the price process

Without loss of generality we can assume that the long-term mean of the price process is zero. WE stick to the simplest example of a mean-reverting process, namely, the *Ornstein-Uhlenbeck process* given by

$$dX_t = -kX_t dt + \sigma dB_t, \tag{1}$$

where B_t is a Brownian motion, k and σ are positive constants. This process will revert to its long-term mean zero. More exactly, given X_t, the distribution of X_{t+s}, $s > 0$ is normal with parameters

$$E(X_{t+s}|X_t) = X_t e^{-ks}; \qquad Var(X_{t+s}|X_t) = \left(\frac{1-e^{-2ks}}{2k}\right)\sigma^2. \tag{2}$$

Informally, the constant k measures the speed of the mean-reversion and σ measures the strength of the noise component.

1.3 Choice of the utility function

For $-\infty < \gamma < 1$ we consider power utility

$$U = U(W_T) = \frac{1}{\gamma} W_T^\gamma \tag{3}$$

over the terminal wealth U_T. This is a simple but rich enough family of utility functions. Utility functions are defined up to an additive constant. To include the log-utility as a special case, it is sometimes more convenient to consider the family of utility functions $U(W_T) = \frac{1}{\gamma}(W_T^\gamma - 1)$. Taking the limit $\gamma \to 0$ we obtain the log-utility function $U(W_T) = \log(W_T)$. The log-utility version of our problem was solved by A. Morton [M].

The relative risk aversion is measured by $1-\gamma$, so the bigger γ is, the less risk averse is the agent. In the limit $\gamma \to 1$ we have a linear utility function. In section 3 we study the effect of γ on trading strategy.

1.4 The model

The problem can be treated in the general portfolio optimization framework of [Mer]. Suppose a traded asset follows an Ornstein-Uhlenbeck process (1). It is convenient to think about X_t as a "spread" between the price of an asset and its "fair value". Let α_t be a trader's position at time t, i.e. the number of units of the asset held. This parameter is the control in our optimization problem. Assuming zero interest rates and no market frictions, the wealth dynamics for a given control α_t is given by

$$dW_t = \alpha_t dX_t = -k\alpha_t X_t dt + \alpha_t \sigma dB_t. \tag{4}$$

We assume that there are no restrictions on α, so short selling is allowed and there are no marginal requirements on wealth W.

We solve the expected terminal utility maximization problem for an agent with a prespecified time horizon T and initial wealth W_0. The utility function (3) is defined over the terminal wealth W_T. The value function $J(W_t, X_t, t)$ is the expectation of the terminal utility conditional on the information available at time t:

$$J(W_t, X_t, t) = \sup_{\alpha_t} \mathbf{E}_t \frac{1}{\gamma} W_T^\gamma. \tag{5}$$

1.5 Normalization

It is more convenient to work with dimensionless time and money. Let \$ be the dimension of X; we denote it by $[X] = \$$. By T we denote the dimension of time. From Eq. (1) it is clear that $[\sigma] = \$T^{-1/2}$ and $[k] = T^{-1}$. Renormalizing price X_t, position size α_t, and time t

$$\begin{aligned} X &\to \frac{X}{\sigma}\sqrt{k}, && (6)\\ \alpha &\to \frac{\alpha}{\sqrt{k}}\sigma, && (7)\\ t &\to kt, \end{aligned}$$

we can assume that $k = 1$ and $\sigma = 1$. The wealth W does not change under this normalization. Note that this normalization is slightly different from one used in [M-VMN].

1.6 Overview

In [M-VMN] it is proven that for $\gamma = 0$ (log-utility case) the optimal control is given simply by

$$\alpha_t = -W_t X_t.$$

The case $\gamma = 0$ is simpler than the general case because a log-utility agent does not hedge intertemporally (see [Mer]) and the equations are much simpler. The same paper also derives an approximate solution for the case $\gamma < 0$. The approximation does not behave particularly well.

In Section 2, we obtain an exact solution to the problem defined by Eqs. (1) – (5) for the general case $\gamma < 1$, $\gamma \neq 0$. The answer is given by Eqs. (16) and (17). In section 3, we analyze this solution. We are looking at how J and α change as the spread X changes and how the risk aversion affects trader's strategy.

We will see that although our model is very simple, it reproduces some of the typical trader behavior patterns. For example, if a trader is more risk-averse than a log-utility one, then he will cut his position as the time horizon approaches. This behavior is similar to the anecdotical evidence on real position management practice.

Section 4 concludes with suggestions for possible generalizations.

2 Main result

2.1 The Hamilton-Jacobi-Bellman equation

We need to find the optimal control $\alpha^*(W_t, X_t, t)$ and the value function $J(W_t, X_t, t)$ as explicit functions of wealth W_t, price X_t, and time t.

The *Hamilton-Jacobi-Bellman equation*[1] is

$$\sup_\alpha \left(J_t - xJ_x - \alpha x J_w + \frac{1}{2}J_{xx} + \frac{1}{2}\alpha^2 J_{ww} + \alpha J_{xw} \right) = 0 \tag{8}$$

The first order optimality condition on control α^* is

$$\alpha^*(w, x, t) = x\frac{J_w}{J_{ww}} - \frac{J_{xw}}{J_{ww}}. \tag{9}$$

Note that the first summand in the right-hand side of Eq.(9) is the myopic demand term corresponding to a static optimization problem while the second term hedges from changes in the investment opportunity set. For a log-utility investor ($\gamma = 0$) the second term vanishes (see [Mer].)

Substituting this condition into the Hamilton-Jacobi-Bellman equation for the value function, we obtain the non-linear PDE

$$J_t + \frac{1}{2}J_{xx} - xJ_x - \frac{1}{2}J_{ww}\left(\frac{J_{xw}}{J_{ww}} - x\frac{J_w}{J_{ww}}\right)^2 = 0. \tag{10}$$

2.2 Main theorem

Let

$$\tau = T - t \tag{11}$$

be the time left for trading and define the constant ν and time functions $C(\tau)$, $C'(\tau)$, and $D(\tau)$ by

$$\nu = \frac{1}{\sqrt{1-\gamma}} \tag{12}$$

$$C(\tau) = \cosh \nu\tau + \nu \sinh \nu\tau \tag{13}$$

$$C'(\tau) = \frac{dC(\tau)}{d\tau} = \nu \sinh \nu\tau + \nu^2 \cosh \nu\tau \tag{14}$$

$$D(\tau) = \frac{C'(\tau)}{C(\tau)}. \tag{15}$$

As we shall see, the function $D(\tau)$ plays a crucial role in determining the optimal strategy.

[1] see e.g. [Fle].

Theorem 1 *Suppose that $\gamma < 0$ or $0 < \gamma < 1$. Then the optimal strategy for the problem (1) – (5) is given by*

$$\alpha_t^* = -wxD(\tau). \tag{16}$$

The value function is given by

$$J(w,x,t) = \frac{1}{\gamma} w^\gamma \sqrt{e^\tau C(\tau)^{\gamma-1}} \exp\left(\frac{x^2}{2}\left(1 + (\gamma-1)D(\tau)\right)\right), \tag{17}$$

where τ, $C(\tau)$, and $D(\tau)$ are defined by Eqs. (11) – (15) and $X_t = x$, $W_t = w$.

We will prove the theorem in the Appendix.

Note that the optimal position is linear in both wealth W_t and spread X_t. The term under the last exponent in (17) measures the expected utility of the immediate trading opportunity. If $X_t = 0$ i.e. there are no immediate trading opportunities, the value function (17) simplifies to

$$J(w,0,t) = \frac{1}{\gamma} w^\gamma \sqrt{e^\tau C(\tau)^{\gamma-1}}.$$

The $\frac{1}{\gamma} w^\gamma$ term is just the expected utility generated by the present wealth. The square root term can be thought of as the value of the time. We will analyze Eqs. (16) and (17) in more detail in section 3.

3 Analysis

In this section, we analyze the behavior of solution (16) - (17). Unless specified otherwise, the parameters used for illustrations were $k = 2$, $\sigma = 1$, and $\gamma = -2$. From Eq. (2), it follows that the long-term standard deviation of the price process value is $\sqrt{2}/2$, so, roughly, an absolute value of X greater than 0.7 presents a reasonable trading opportunity.

3.1 Position management

Let us look at how the value function and trading position change as X_t changes. Using Ito's lemma, we see from (16) that the diffusion term of $d\alpha_t$ is

$$-D(\tau)\left(W_t + \alpha_t X_t\right).$$

Thus, the covariance of $d\alpha$ and dX is

$$\operatorname{Cov}(d\alpha, dX) = -D(\tau)(W_t + \alpha_t X_t) = W_t D(\tau)\left(-1 + X_t^2 D(\tau)\right). \tag{18}$$

This is negative whenever

$$|X| \le \sqrt{1/D(\tau)}.$$

Consequently, as X_t diverges from 0 either way, we start slowly building up the position α_t of the opposite sign than X_t. If X_t diverges further from 0, our position is making a loss, but we are still increasing the position until the squared spread X_t^2 reaches $1/D(\tau)$. If the spread widens beyond that value, we start cutting a loss-making position. Another interpretation of Eq. (18) is that we start cutting a loss-making position as soon as the *position spread* $-\alpha X_t$ exceeds total wealth W_t. Fig. 2 shows how $D(\tau)$ depends on remaining time τ for different values of γ.

Not surprisingly, for the log-utility case $\gamma = 0$, the threshold $1/D(\tau)$ equals identically one and we are getting the same result as in [M] and [M-VMN].

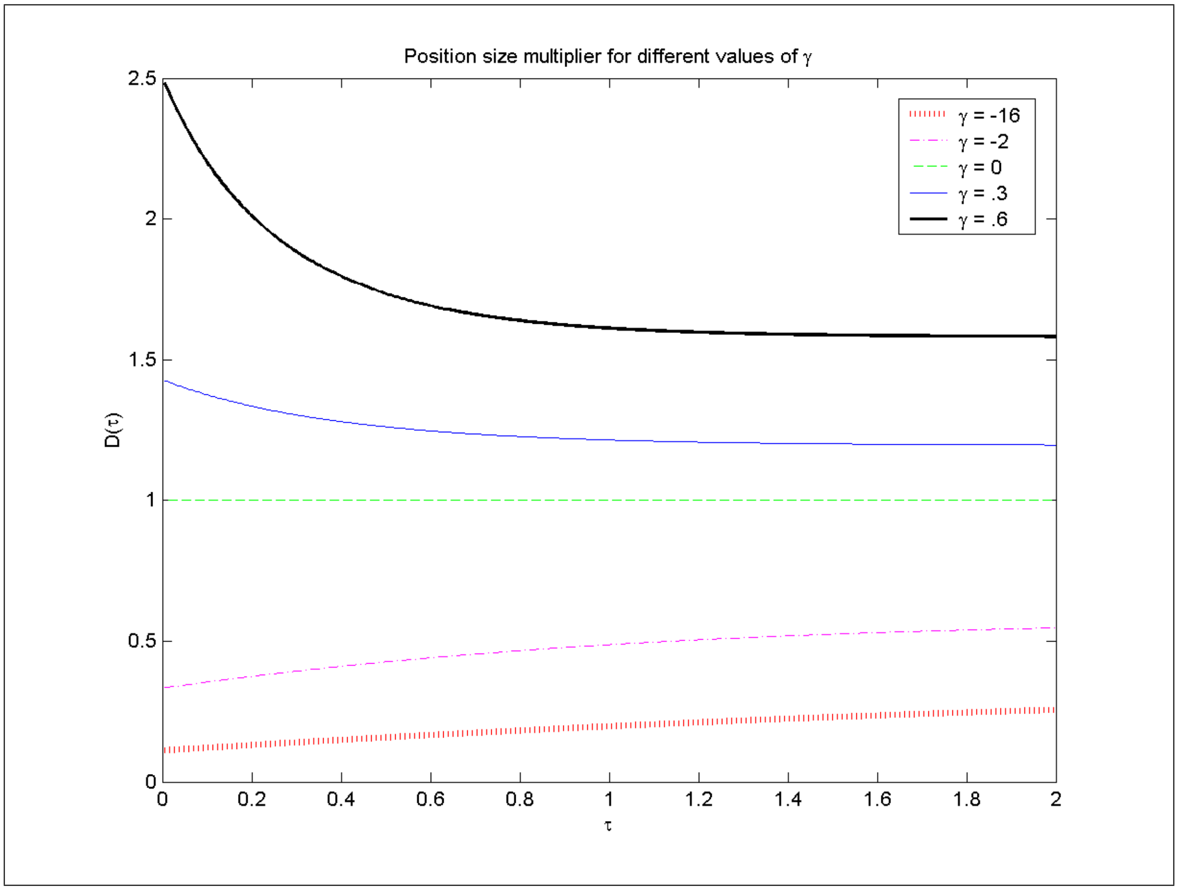

Figure 2: $D(\tau)$ as a function of the remaining time τ for five different values of γ.

3.2 Value function dynamics

Let us check now how the value function $J(W_t, X_t, t)$ evolves with X_t. In [M-VMN], it is shown that a log-utility agent's value function always decreases as the spread moves against his position. It might be the case that a more aggressive agent's value function sometimes increases as the spread X_t moves against his position because the investment opportunity set improves. Let us check whether this ever happens to a power utility agent.

Using Ito's lemma, we see from (17) that the diffusion term of dJ_t is

$$J_t X_t \left(1 - D(\tau)\right).$$

Thus,

$$\operatorname{Cov}(dJ, dX) = J_t X_t \left(1 - D(\tau)\right). \tag{19}$$

For $\gamma < 0$, the utility function is always negative, so the value function is also always negative. Similarly, for $\gamma > 0$ the value function J is always positive. It is easy to check that the sign of $1 - D(\tau)$ is opposite to the sign of γ for all τ. Thus, $\operatorname{Cov}(dJ, dX)$ is positive for $X_t < 0$ and negative for $X_t > 0$. This means that any power utility agent suffers decrease in his value function J as the spread moves against his position. This is true even for an agent with an almost linear utility $\gamma \to 1$.

For $0 < \gamma < 1$ there is a non-zero bankruptcy probability.

3.3 Time value

Let us look once more at how the value function depends on the time left for trading. Recall that

$$J(w,x,t) = \underbrace{\frac{1}{\gamma} w^\gamma}_{A} \underbrace{\sqrt{e^\tau C(\tau)^{\gamma-1}}}_{B} \underbrace{\exp\left(\frac{x^2}{2}\left(1+(\gamma-1)D(\tau)\right)\right)}_{C}, \tag{20}$$

where τ, $C(\tau)$, and $D(\tau)$ are defined by Eqs. (11) – (15) and $X_t = x$, $W_t = w$. Thus, the value function J can be split into three multiplicative terms. Term A is the value derived from the present wealth, term B is the time value, and term C is the value of the *immediate* investment opportunity. Fig. 3 shows[2] dependence of the value function J on time τ assuming that there is no immediate opportunity, i.e. $X = 0$.

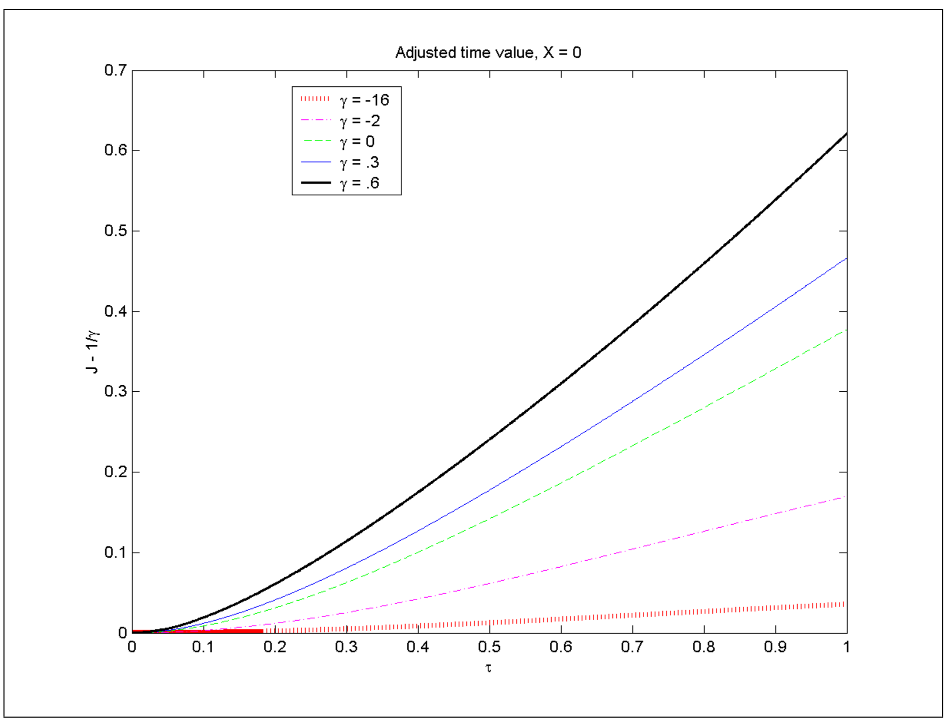

Figure 3: Adjusted time value, $X = 0$, $W = 1$.

Since log-utility agent's strategy does not depend on time, his value function J grows linearly with time (the green line on Fig. 3.) Extension of the trading period beyond a certain minimal

[2]The figure shows the graphs of the function $J(w,x,\tau) - \frac{1}{\gamma}$ for $w = 1$, $x = 0$, $0 \le \tau \le 1$ and several different values of γ. Substraction of $\frac{1}{\gamma}$ from the value function makes easier the comparison with the log-utility case $\gamma = 0$.

length does not significantly increase the value function of a sufficiently risk-averse agent (the pink and the red lines on Fig. 3.)

The value of time grows roughly exponentially in X_t^2 if there is an immediate investment opportunity.

3.4 Effect of risk-aversion on time inhomogeneity

The ratio $D(\tau)$ defined by Eq. (15) plays a crucial role in most of our formulas: it determines the position size in Eq. (16), the threshold at which we start unwinding a loosing position (Eq. (18)), and it also enters equations (17) for the value function and (19) for the covariance of J and X. Fig. 2 shows the graphs of $D(\tau)$ for different values of γ. Recall that Eq. (16) implies that for given wealth W_t and spread X_t, position size is proportional to $D(\tau)$.

We see that for $\gamma = 0$ (log-utility) optimal position does not depend on time. For $\gamma > 0$ the agent is less risk-averse than a log-utility agent. So, for given price X_t and wealth W_t, his position *increases* as the final time approaches. In practice, traders often tend to become *less aggressive* as the bonus time approaches. This is consistent with the optimal behavior of a power utility agent with $\gamma < 0$. For example, assume that $k = 8$ and $\gamma = -2$ an let us measure the time in years. Then for the same wealth W and spread X, the position just a week before the year-end is a third lower than it is at the beginning of the year.

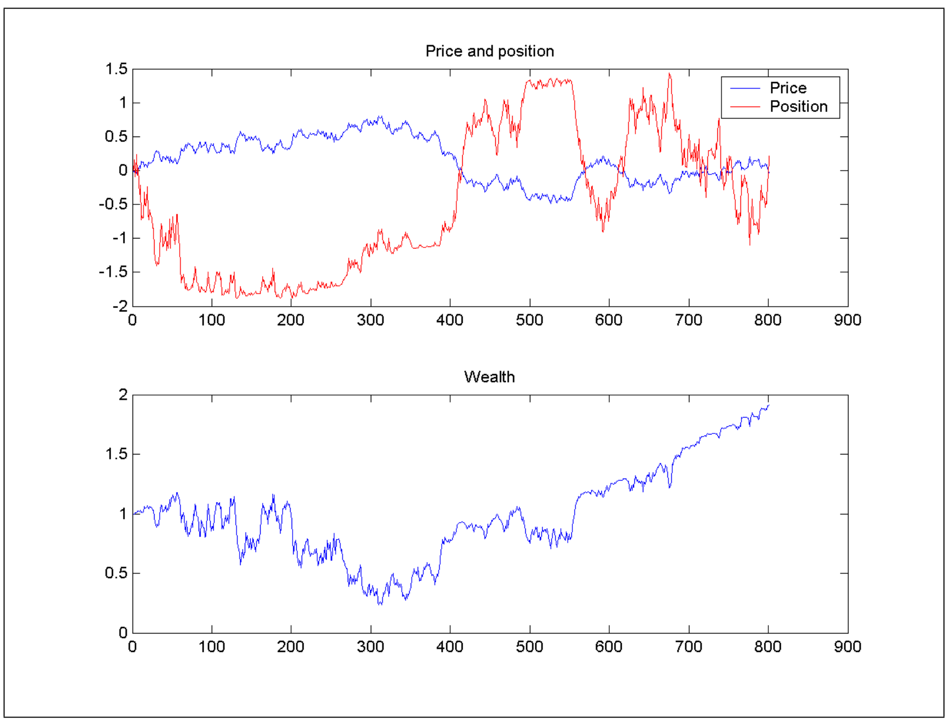

Figure 4: A simulated price sample together with position and wealth dynamics.

3.5 Simulation results

To study the effect of parameter misspecification, we performed a Monte-Carlo simulation. Fig. 4 shows a sample price trajectory with the corresponding optimal position and the wealth trajectories. A simulation without variance reduction also gives a good proxy to the discretisation and sampling errors, i.e. to the deviations of accumulated wealth from the predicted wealth due to the sampling error and non-continuous rehedging.

In reality, it is very hard to predict the mean-reversion parameter k. Even if we assume that the price series is stationary, k has to be estimated from the past data. Figure 5 shows the effect of trading with wrong k.

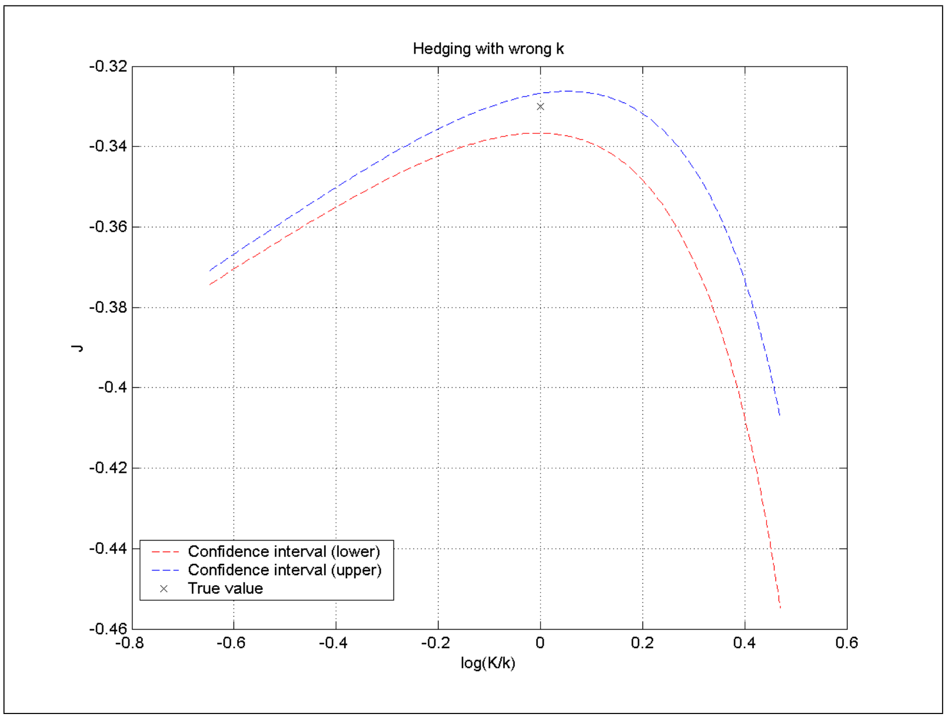

Figure 5: A simulated price sample together with position and wealth dynamics.

In a Monte carlo simulation, we generated a set of Ornstein-Uhlenbeck process trajectories with $k = 2$, $\sigma = 1$ and then simulated trading with a wrong value of k. To look at the dependence of optimal position α^* on mean reversion coefficient K, it is convenient to invert transforms (7) and to rewrite Eq. (16) as

$$\alpha = \frac{k}{\sigma^2} wxD(\tau/k). \tag{21}$$

Thus, we took K in the interval $(1, \ldots, 3.2)$ and simulated trading with position determined by (21), but with K substituted for k. On the horizontal axis of the graph we have $\log(K/k)$. Blue and red dashed lines show the two standard deviations confidence interval bounds for the mean terminal utility when trading with a given value of K. The black cross shows the value function from Eq. (17) for $K = k$.

We can see that the influence of mean reversion coefficient misspecification is asymmetric. Trading with a conservatively estimated k reduces greatly the utility uncertainty. Not surprisingly, overestimation of the mean reversion leads to excessively aggressive positions and big discretisation errors. It is much safer to underestimate k than to overestimate it.

4 Conclusions and possible generalizations

We solved the optimal portfolio selection problem assuming that there is a single risky asset following an Ornstein-Uhlenbeck process with known parameters and there is a representative agent with given wealth, investment horizon and power utility function. The other assumptions used were the absence of market frictions and perfect liquidity of the asset traded.

Most of these assumptions are similar to ones made in the Black-Scholes model. Each of these assumptions is not quite realistic. Even when one manages to find a mean-reverting trading asset, one will need to estimate the parameters of the process. The prices usually seem to follow non-stationary processes, with periodic regime switches and jumps. Market frictions make continuous trading unviable, while the presence of other traders competing for the same trading opportunity and the feedback between trades and prices affect the optimal strategy. A trader usually does not commit all of his capital to trade a single asset, so the real-world problems involve multiple risky asset portfolio selection.

The model considered can be extended to include many of these more realistic features. The resulting PDE is not very likely to have an explicit solution, but singular perturbation theory may be used to obtain approximations by expansions around our solution. A similar problem in discrete setting is considered in [V]. The discrete framework allows to introduce easily transaction costs but, in most cases, lacks explicit solutions. The attraction level of the mean reverting process X_t may be assumed not known *a priori* and to be inferred from observations of X_t. This problem can be treated in the Bayesian framework similarly to [L].

On the other hand, our simple model can serve as a benchmark in practical situations. Quite often, practitioners prefer to introduce *ad hoc* corrections to a simple model than using a more involved model with a large number of parameters.

5 Appendix A.

5.1 A technical lemma

To prove the theorem we need the following lemma.

Lemma 1 *The functions $\alpha^* = \alpha^*(w,x,t)$ and $J = J(w,x,t)$ defined by Eqs. (16) and (17) have the following properties*

1. *$J = J(w,x,t)$ is a solution to Eq. (10);*
2. *boundary condition at T:*
$$J(w,x,T) = \frac{1}{\gamma}(w^\gamma - 1);$$
3. *concavity in current wealth:*
$$J_{ww} \leq 0, \text{ for all } w \geq 0,\, x \in \mathbb{R},\, 0 \leq t \leq T;$$
4. *α^* satisfies the first order optimality condition (9).*

PROOF OF THE LEMMA. All properties can be checked by direct calculations. △

5.2 Proof of the theorem

Let $J(w,x,t,\alpha)$ be the expected terminal utility if the trader follows a particular strategy α. It is enough to show that $J(W_t,X_t,t)$ and $\alpha^*(W_t,X_t,t)$ given by Eqs. (16), (17) satisfy two standard stochastic optimal control conditions:

(A) For any control $\alpha = \alpha(w,x,t)$

$$J(w,x,t,\alpha) \leq J(w,x,t) \text{ for all } x \in \mathbb{R}, w \geq 0, 0 \leq t \leq T$$

(B) The control $\alpha^* = \alpha^*(w,x,t)$ satisfies

$$J(w,x,t,\alpha^*) = J(w,x,t).$$

Condition (A). Applying Ito's formula to $J\left(W_s,X_s,s\right)_{t\leq s\leq T}$, we obtain

$$\begin{aligned} J(W_s,X_s,s) &= J(W_t,X_t,t) + \int_t^s \mathcal{L}(\alpha)J(W_u,X_u,u)du + \\ &+ \int_t^s J_x(W_u,X_u,u)dB_u + \int_t^s \alpha_u J_w(W_u,X_u,u)dB_u, \end{aligned} \tag{22}$$

where

$$\mathcal{L}(\alpha)J = J_t - xJ_x - \alpha x J_w + \frac{1}{2}J_{xx} + \frac{1}{2}\alpha^2 J_{ww} + \alpha J_{xw}. \tag{23}$$

Using the Lemma, we see that

$$\begin{aligned} \mathcal{L}(\alpha)J &= \frac{1}{2}J_{ww}\left(\alpha - \left(x\frac{J_w}{J_{ww}} - \frac{J_{xw}}{J_{ww}}\right)\right)^2 + (J_t + \frac{1}{2}J_{xx} - xJ_x - \frac{1}{2}J_{ww}\left(\frac{J_{xw}}{J_{ww}} - x\frac{J_w}{J_{ww}}\right)^2) = \\ &= \frac{1}{2}J_{ww}\left(\alpha - \left(x\frac{J_w}{J_{ww}} - \frac{J_{xw}}{J_{ww}}\right)\right)^2 \leq 0. \end{aligned} \tag{24}$$

Taking the mathematical expectation $\mathbf{E}_t$ of $J(W_s,X_s,s)$, from (22) we obtain

$$\begin{aligned} \mathbf{E}_t J(W_s,X_s,s) &= J(W_t,X_t,t) + \mathbf{E}_t\int_t^s \mathcal{L}(\alpha)J(X_u,W_u,u)du + \\ &+ \mathbf{E}_t\int_t^s J_x(W_u,X_u,u)dB_u + \mathbf{E}_t\int_t^s \alpha J_w(W_u,X_u,u)dB_u \end{aligned} \tag{25}$$

The stochastic integrals in (25) are martingales, so the mathematical expectation of these integrals is zero. Thus, the last two summands in (25) vanish. Now let $t \to T$. Using (24), we can rewrite (25) as

$$\begin{aligned} J(W_t,X_t,t) &= \mathbf{E}_t J(W_T,X_T,T) - \mathbf{E}_t\int_t^T \mathcal{L}(\alpha)J(X_u,W_u,u)du \geq \\ &\geq \mathbf{E}_t(\frac{1}{\gamma}W_T^\gamma) = J^\alpha(W_t,X_t,t,\alpha), \end{aligned}$$

i.e. condition (A) is satisfied.

Condition (B). It is clear from the Lemma that

$$\mathcal{L}(\alpha^*)J = 0.$$

So for $\alpha = \alpha^*$ we have

$$\begin{aligned} J(W_t, X_t, t) &= \mathbf{E}_t J(W_T, X_T, T) - \mathbf{E}_t \int_t^T \mathcal{L}(\alpha^*) J(X_u, W_u, u) du = \\ &= \mathbf{E}_t(\frac{1}{\gamma} W_T^\gamma) = J^\alpha(W_t, X_t, t, \alpha^*), \end{aligned}$$

i.e. condition (B) is satisfied. This concludes the proof of the Theorem.

References

[Mer] R.C. Merton, *Continuous-Time Finance*, Blackwell Publishers, (1990).

[Lo] J. Campbell, A. Lo, and A. MacKinley. The Econometrics of Financial Markets. Princeton University Press, (1997).

[Shl] A. Schleifer, *Inefficient Markets: An Introduction to Behavioural Finance*, Oxford Univ. Press, Oxford, (2000).

[Fle] W. Fleming and H.M. Soner, *Controlled Markov Processes and Viscosity Solutions*, Springer-Verlag, (1993).

[M] A. Morton, "When in trouble – double", presentation at Global Derivatives 2001 conference, St.Juan - les Pins, May 2001.

[M-VMN] D.Mendez-Vivez, A.Morton, V.Naik, "Trading Mean Reverting Process," unpublished manuscript, August 2000.

[V] Vigodner, A., "Dynamic Programming and Optimal Lookahead Strategies in High Frequency Trading with Transaction Costs," preprint, January 2000.

[L] Lakner, P., "Optimal trading strategy for an investor: the case of partial information," Stochastic Processes and their Applications, **76**, (1998), pp. 77–97.