Optimal Portfolio Selection in Nonlinear Arbitrage Spreads

Hamad Alsayed[*] and Frank McGroarty[†]

School of Management, University of Southampton. United Kingdom.

March 2010

Abstract

This paper derives, in closed-form, an optimal portfolio strategy for a finite-horizon arbitrageur, whose investment set consists of a relative mispricing between two statistically similar assets (e.g. long/short equity pairs). We model the evolution of the mispricing as a nonlinear stochastic process exhibiting nonlinear mean-reversion which weakens in the mispricing. This approach extends the Ornstein-Uhlenbeck model prevalent in the literature, and jointly captures several important features inherent in arbitrage trading. Though arbitrageurs typically act to exploit pricing inefficiencies, we show that under certain conditions, arbitrageurs unwind positions even when investible opportunities are at their greatest. This provides a theoretical explanation of the increased volatility, exacerbated pricing inefficiencies, and lower liquidity following fund liquidations.

Empirically, we test our strategy on both synthetically-generated and FTSE100 daily data, against the strategy implied in an Ornstein-Uhlenbeck setting and a threshold strategy employed in the empirical literature. The historical data is chosen to test our model's robustness across a wide variety of stock market behaviors. The synthetic data test shows that our nonlinear strategy is more robust to parameter misspecification, while the historical data test shows that our nonlinear strategy is more robust to regime-breaks. These results have novel implications on profitability and risk-management in statistical arbitrage trading.

Keywords: Pairs trading, Hamilton-Jacobi-Bellman equation, Statistical Arbitrage, Stochastic optimal control, Stability bounds.

[*]Email: h.alsayed@soton.ac.uk
[†]Email: f.j.mcgroarty@soton.ac.uk

1. Introduction

In the pricing of related securities (for example shares with either identical or similar characteristics), market efficiency is enforced by the presence of rational arbitrageurs (Shleifer and Vishny (1997), and Mitchell, Pulvino, and Stafford (2002)). Xiong (2001) shows that if the actions of noise traders cause such a price relation to be violated, arbitrageurs seek to profit from this by betting on the elimination of this violation. By demanding the cheap security and supplying the expensive security, the actions of arbitrageurs exert forces on these prices to revert back to their natural level, which in turn generates a profit for the arbitrageurs and enforces price efficiency. However, unlike "textbook arbitrage" opportunities which lock in a riskless profit and require no capital commitments (Dybvig and Ross (1992)), these types of investments present risks to those who seek to engage in them. To see this, note that if arbitrageurs position themselves to benefit from the convergence of a mispricing, it is possible for this mispricing to diverge further before converging, or not converge at all, resulting in substantial losses for arbitrageurs. In 1998, the near-collapse of the hedge fund Long Term Capital Management (LTCM)[1] is frequently cited as an example of this phenomenon (see e.g. Kondor (2009), Xiong (2001), and Liu and Longstaff (2004)). In theory, arbitrageurs can earn large profits by exploiting a relative mispricing, but because of limited capital, they may be forced to unwind positions at a loss upon further divergences (DeLong et al. (1990), Shleifer and Summers (1990), and Shleifer and Vishny (1997)). What then, is the optimal portfolio policy for an arbitrageur to adopt? When is the best time to exploit a mispricing? At what point does it become necessary to unwind a losing position?

In this paper, we develop a theoretical model to provide answers to these questions in an analytically tractable form. By considering the problem under a partial equilibrium setting, and in light of the central questions this paper aims to address, we are able to gain novel insights into the investment behavior of a representative arbitrageur, which are otherwise obtained qualitatively or through numerical solutions in general equilibrium models (e.g. Xiong (2001)). This ties our work closely to the portfolio optimization literature pioneered by Merton (1969, 1971), and in particular Boguslavsky and Boguslavskaya (2004) and Jurek and Yang (2007), who both consider arbitrage models in partial equilibrium and model the mispricing exogenously as an Ornstein-Uhlenbeck[2] (henceforth "OU") process. We depart from this assumption by introducing a novel stochastic process[3] capable of richer dynamics

---

[1]See Edwards (1999) and MacKenzie (2003) for analysis of the LTCM crisis.

[2]It is well-known that the Ornstein-Uhlenbeck process has widespread applications in finance, e.g. Wachter's (2002) and Merton's (1971) application to mean-reverting returns in stock portfolios, applications to stochastic interest rates (Korn and Kraft (2002), Vasicek (1977)), and applications to arbitrage opportunities in Xiong (2001), Boguslavsky and Boguslavskaya (2004), Jurek and Yang (2007), and Lv and Meister (2009).

[3]Our stochastic process is based on a modification to Shaw's (2009) "hyperbolic OU" process, and Wong's (1964) *repulsive Wong process* which occurs in physics.

than those implied by OU. This is motivated by the need for more robust quantitative risk management models in the wake of the recent financial crisis (see Shaw (2009)).

To this end, we begin by specifying a process to describe the evolution of the mispricing. Our main assumption is that the mispricing is exogenously defined,[4] and driven by a nonlinear mean-reverting stochastic process whose strength of mean-reversion weakens with the level of the mispricing. This process is perturbed by Gaussian noise. The nonlinearity is in the mean-reversion component, and can be explained as follows: Intuitively, when the mispricing is close to its *natural level*,[5] reversion is strong relative to the random perturbations, gradually pushing the mispricing towards its natural level. If however, the mispricing is relatively large, the strength of the reversion is weaker relative to the random perturbations, potentially allowing the mispricing to "linger" away from its natural level. This artifact of our model allows us to parsimoniously capture three broad phenomena inherent in arbitrage trading: Horizon risk and divergence risk imply, respectively, uncertainty whether a mispricing will converge prior to the reporting period, and a worsening in the mispricing after an arbitrageur has opened a position.[6] The third is the risk of regime-break: When a mispricing becomes large, the force in our nonlinear model pushing it towards its natural level is weaker, thus its evolution appears more random. If the mispricing is then perturbed in the direction toward its natural level, it will experience an increasing mean-reversion toward its natural level.[7] This is a novel feature captured by our model, absent from OU.

An artifact of the (linear) OU process is a constant half-life of the mispricing. This artifact implies that the larger the mispricing, the more dramatically it is corrected. But this notion seems inconsistent with economic arguments that mispricings can exhibit large and long-lasting departures from fundamentals (see e.g. Abreu and Brunnermeier (2002, 2003)). Our model, on the other hand, implies shorter half-lives for small divergences and longer half-lives for large divergences. If we interpret these departures from fundamentals as persistent divergences from a mispricing's natural level, then we believe our model is therefore more suited to capture this phenomenon.

To derive a trading strategy, we employ stochastic optimal control theory. The result is a closed-form optimal portfolio strategy which is nonlinear in the magnitude of the mispricing.

---

[4]This ties our work to Liu and Longstaff (2004), Jurek and Yang (2007), Boguslavsky and Boguslavskaya (2004), Xiong (2001), and more generally Merton (1969, 1971) and Wachter (2002).

[5]If a mispricing is between two identical shares quoted on different markets, then economic arguments suggest a natural level of zero (Froot and Dabora (1999)). A mispricing between any two cointegrated assets (Engle and Granger (1987)), can have its natural level take any *constant* value.

[6]These two broad measures of risk are defined in Jurek and Yang (2007). In particular, divergence risk is equivalent to Lamont and Thaler's (2003b) *noise trader risk*. See also Shleifer and Vishny (1997) and Abreu and Brunnermeier (2002).

[7]An intuitive analogy of this feature is imagining a particle falling to earth from space. Unlike what OU would imply, the gravitational pull is weak at first, but strengthens as the particle falls towards the earth.

As the mispricing diverges from its natural level, our model implies a diminishing marginal exploitation of the mispricing for any given level of wealth. This novel feature of our model addresses the potential of regime break, which is mechanically absent from assuming OU. Specifically, an arbitrageur faithful to our model grows increasingly sceptical towards larger divergences in the mispricing from its natural level, whereas an arbitrageur faithful to OU interprets these as increasingly attractive investment opportunities. The latter notion is also maintained in works which do not exogenously specify a process describing the evolution of the mispricing (e.g. Kondor (2009), Gatev, Goetzmann, and Rouwenhorst (2006)), and alternative specifications for the mispricing process, such as Liu and Longstaff's (2004) Brownian Bridge process. In Section 3 we illustrate that only within a small neighborhood around the natural level of mispricing do the portfolio strategies implied by both our model and OU approximately coincide. Furthermore, we suggest that conclusions regarding inter-temporal hedging demands[8] assuming OU persist in a more complex dynamic model such as ours, since the dynamics of our model are flexible enough include OU as a special case.

With the exception of Jurek and Yang (2007) and Boguslavsky and Boguslavskaya (2004), the works we mention in this section do not analytically characterize situations in which arbitrageurs unwind losing positions to preserve capital.[9] Exploring this phenomenon is essential to understanding the role and capacity arbitrageurs have in enforcing price efficiency. Our model addresses this point by providing closed-form bounds in the mispricing beyond which arbitrageurs unwind losing positions. We refer to these as *stability bounds*. When a divergence of the mispricing from its natural level causes an arbitrageur to suffer losses on his current position, he will initially continue to allocate capital aimed at exploiting the mispricing, because the improvement in the investment set outweighs the negative wealth effect caused by the current loss. However, once a mispricing diverges beyond the stability bound, the wealth effect dominates, and an arbitrageur begins cutting a losing position to preserve capital. Furthermore, we show that for a common set of model parameters, the stability bounds in our model are tighter than those implied by assuming OU. This has novel implications regarding when an arbitrageur should cut losses. Because our model calls for losses to be cut sooner, it is reasonable to infer that arbitrage trading faithful to our model will result in less market volatility and higher liquidity relative to OU in the event of fund liquidation.[10]

The rest of this paper is organized as follows: Section 2 provides a review of the related literature, placing our model in the context of other models of arbitrage. Section 3 formalizes the model, and derives the optimal portfolio strategy and stability bounds for our model,

---

[8]For a discussion of inter-temporal hedging demands, see Liu (2007), Campbell and Viceira (2002), Wachter (2002), and Kim and Omberg (1996).

[9]Xiong (2001) provides a numerical solution to this in general equilibrium.

[10]Clearly this claim can only be definitively investigated in a general equilibrium setting with price-feedback and wealth effects.

then compares these with models which assume OU. We also present a synthetic data test to determine the effect of parameter misspecification on the optimal portfolio strategy. Section 4 empirically applies our optimal strategy to historical FTSE100 data, and shows that our strategy is more robust to regime-break than assuming OU. Section 5 concludes.

2. Related Literature

There is considerable evidence that prices can diverge from fundamental values and similarly, that relative prices can diverge from their natural levels. Shleifer and Summers (1990) and Black (1986) highlight the effect of noise traders' irrational trading behavior on the formation of prices. Noise traders react disproportionately to news in the belief that they have insider information regarding the future direction of prices, creating a mispricing. On the one hand, Friedman (1953) argues that such mispricings cannot exist for long, as rational arbitrageurs will trade against noise traders, hence pushing prices toward fundamental values. On the other hand, Lamont and Thaler (2003a) find examples of persistent deviations from *the law of one price*, Froot and Dabora (1999) find mispricings between *Siamese twin* shares, Malkiel (1977) finds mispricings in the valuation of closed-end funds, and Lamont and Thaler (2003b) find mispricings in tech stock carve-outs. Consistent with the latter view, our model is based on the notion that mispricings can persist for significant periods of time. But why in general does capital not materialize to eliminate such mispricings?

The persistence of mispricings can be attributed to many factors. Abreu and Brunnermeier (2002, 2003) show that when individual arbitrageurs attempt to "time the market", informational asymmetries can cause coordination problems that consequently allow mispricings to persist. Brunnermeier and Pedersen (2005) highlight the role of predatory trading behavior in the persistence of mispricings. Shleifer and Vishny (1997), Xiong (2001), and De Long et al. (1990) show that impediments to arbitrage can arise endogenously. When arbitrageurs allocate capital aimed at exploiting a mispricing, further divergences may cause them to unwind positions at a loss to preserve capital, a phenomenon described by Friedman (1953) as "buying high and selling low". This creates a price feedback mechanism which exacerbates the mispricing. Lamont and Thaler (2003b), Liu and Longstaff (2004), Gromb and Vayanos (2002), and Basak and Croitoru (2000) suggest that portfolio constraints can allow mispricings to persist. Moreover, Kondor (2009) shows that arbitrageurs do not necessarily act to eliminate a mispricing even when portfolio constraints do not bind, because "opportunities might get better tomorrow".

Our partial equilibrium model is aimed at capturing these intuitions into an exogenously-defined mispricing process. We develop an optimal portfolio policy for an individual rational

arbitrageur to adopt, in light of these mechanisms and the risks associated with arbitrage trading. This relates our model to Boguslavsky and Boguslavskaya (2004), Jurek and Yang (2007), and Liu and Longstaff (2004).[11] The common theme in these works is that the evolution of the mispricing is modelled exogenously in a partial equilibrium framework. Liu and Longstaff (2004) model the arbitrage opportunity using a Brownian Bridge process, which has two fixed values *a priori*, typically at the start and end of the investment horizon. Under this setup, arbitrageurs have perfect *ex-ante* knowledge regarding the date of the elimination of a mispricing. The authors incorporate portfolio constraints to preclude arbitrageurs taking infinite positions and attaining nirvana (Kim and Omberg (1996)). Boguslavsky and Boguslavskaya (2004) and Jurek and Yang (2007) observe that, apart from finitely-lived opportunities such as deviations from option put-call parity, arbitrageurs typically would not have perfect knowledge regarding the elimination of a mispricing. Therefore, these works model the arbitrage opportunity using an OU process, which implies uncertainty regarding the mispricing at all future dates. This uncertainty naturally precludes arbitrageurs taking infinite positions even in the absence of portfolio constraints, and this feature is present in our model.

Another common theme across the partial equilibrium models of arbitrage is that arbitrageurs have perfect knowledge regarding the natural level of a mispricing[12]. This assumption ignores Lamont and Thaler's (2003b) *fundamental risk*, which highlights changes in economic fundamentals that can affect the natural mispricing level itself. Fundamental risk is particularly important when the natural level is ascertained statistically, and not hinged on economic intuition (e.g. spread between two different bank shares). Our model considers this risk. Even though the natural level is assumed fixed, the fact that an arbitrageur faithful to our model exhibits a diminishing marginal exploitation of the mispricing can be interpreted in terms of his decreased confidence that the mispricing will revert to its natural level. On the other hand, the fact that models based on OU or Brownian Bridge call for increasing the allocation as divergences get large implies that arbitrageurs faithful to these models have complete confidence in the estimated natural level, which completely ignores fundamental risk.

The dynamics of our model are simple enough to yield a closed-form solution to the arbitrageur's allocation problem and to include OU as a special case. Thus our results can be readily compared to Jurek and Yang (2007) and Boguslavsky and Boguslavskaya (2004). Our work also relates to Xiong (2001), Basak and Croitoru (2000), and Liu and Longstaff (2004) in the sense that we assume log-utility over terminal wealth.

---

[11]Similar works exploring optimal portfolio selection in arbitrage under a partial equilibrium framework include Mudchanatongsuk, Primbs, and Wong (2008), Kim, Primbs, and Boyd (2008), Lv and Meister (2009), and Elliott et al. (2005).

[12]One could use economic arguments or Likelihood Estimation methods (see the appendix of Jurek and Yang (2007)) to establish the natural level, but either way this level is assumed to be *ex ante* true and constant.

More generally, considering the arbitrageur's portfolio allocation problem in partial equilibrium gives us two main advantages over the general equilibrium models discussed in this section. Firstly, we are able to construct an analytically tractable and empirically implementable optimal portfolio policy for a representative arbitrageur, which yields novel insights into the behavior of rational arbitrageurs. This is particularly important in illustrating one of the central novelties in our model, namely the skepticism toward exploiting large mispricings. Secondly, we ascertain precise conditions under which arbitrageurs cease exploiting a mispricing and unwind losing positions, complementing works on the impediments to arbitrage.

3. The Model

This section considers the behavior of a rational arbitrageur who has access to an arbitrage opportunity and a riskless asset.[13] The riskless asset yields a continuously compounded return. We model the arbitrage opportunity as a mean-reverting asset S_t. Intuitively, this asset can be thought of as a long/short position in an equity pair. In this context, holding long one unit of S_t is equivalent to holding one unit long in some undervalued security and one unit short in an overvalued security at time $t \in [0, T]$.

3.1 The Investment Opportunity:

Consistent with Merton (1969, 1971), Wachter (2002), and Kim and Omberg (1996), we assume the investible universe consists of two assets. The first is a risky asset S_t which describes the evolution of the arbitrage opportunity (i.e. mispricing). The second is the risk-free asset B_t. The dynamics of these, respectively, are given by the following equations:

$$
(1) \qquad dS_t \;=\; -\frac{k}{c} \tanh\left(c(S_t - \overline{S})\right) dt + \sigma dZ_t,
$$

$$
(2) \qquad dB_t \;=\; rB_t dt.
$$

where $\overline{S}$ is the natural level of the mispricing, r is the risk-free rate, Z_t is a Brownian Motion with respect to the real-world probability measure, $k > 0$ is the parameter of mean reversion, and $\sigma > 0$ is the volatility parameter. The investment horizon is continuous and finite from $0 \leqslant t \leqslant T < \infty$. The parameter $c > 0$ measures the nonlinearity in the mean-reversion component of (1). We explore this further below.

---

[13]Because Jurek and Yang's (2007) model occurs as a special case of ours, we endeavor to use the same notation, such that both models can be readily compared.

3.2 Contrasting our dynamics with OU:

The parameter $c > 0$ is a novel feature in our model. This parameter specifies the nonlinearity of mean-reversion in our stochastic process (1). To make clear the connection between our model and OU, we observe that in the special case as $c \to 0$, the dynamics of the mispricing (1) reduce to:

$$(3) \qquad \lim_{c \to 0} \left[-\frac{k}{c} \tanh \left(c(S_t - \overline{S}) \right) dt \right] = -k(S_t - \overline{S}) dt.$$

The right hand side of (3) is precisely the mean-reversion component of an OU process with parameters k and σ. Indeed, this would reduce (1) to:

$$dS_t \to -k(S_t - \overline{S}) dt + \sigma dZ_t,$$

which implies that our model is reduced to precisely the specification in Jurek and Yang (2007). Further setting $r = 0$ and $\overline{S} = 0$ reduces the model to precisely the specification in Boguslavsky and Boguslavskaya (2004).

As long as c remains strictly positive, the stochastic process in (1) is nonlinear in the spread. Specifically, the strength of reversion in S_t weakens and attains a low magnitude in the limit as the mispricing diverges significantly from its natural level, i.e. $|S_t| \gg \overline{S}$. To see this, we force S_t infinitely far above its natural level, and obtain:

$$(4) \qquad \lim_{S_t \to \infty} \left[-\frac{k}{c} \tanh \left(c(S_t - \overline{S}) \right) dt \right] \to -\frac{k}{c} dt.$$

The right hand side of (4) implies that the strength of mean-reversion pulling S_t to its natural level has weakened to its limit. This reduces the dynamics of the mispricing (1) to:

$$dS_t \to -\frac{k}{c} dt + \sigma dZ_t.$$

Depending on the values of c and k, the dynamics now are arguably close to a pure diffusion process. But what are the economic implications of properties (3) and (4)?

We believe that, to adequately model the evolution of a mispricing exogenously using a stochastic process, it is desirable for the process to command dynamics simple enough to be practically implementable, yet rich enough to capture realistic market characteristics. We have shown that our model addresses an important phenomenon absent from the dynamics of OU and the Brownian Bridge, namely regime break. Consistent with Shaw (2009), the "hyperbolic OU" process (1) is approximately OU for a small mispricing (i.e. when S_t is close

to $\overline{S}$), but the reversion term weakens as the mispricing becomes arbitrarily large, and tends to a constant. The reduced reversion strength as mispricings diverge captures the intuition that investors may have diminished confidence in that the mispricing will dramatically revert to its natural level. Conversely, OU implies mean-reversion which strengthens with the spread.

3.3 Modelling Risk-Preferences:

A central part of establishing an optimal portfolio strategy is to obtain an idea of the risk preferences of arbitrageurs. In our partial equilibrium framework, the representative arbitrageur is assumed to maximize his expected log utility over terminal wealth.[14] Specifically, the arbitrageur seeks to maximize the *value function*:

$$(5) \qquad\qquad V(S_t, W_t, t) = \max_{N_t} \left[\mathbb{E}_t \left(\ln W_T \right) \right],$$

Here N_t is the number of units held in the mispricing S_t, and W_t denotes wealth at time t. Equation (5) is an economic statement that, through optimally exploiting the mispricing, an arbitrageur maximizes log utility over terminal wealth. In the next section we derive an explicit expression for N_t.

Log utility occurs as a special case of the general CRRA family of power utility.[15] In this paper, we assume log-utility for two main reasons. Firstly, in discrete time, Breiman (1961) showed that investors who assume log-utility will, with probability one, outperform any other strategy asymptotically, without bankruptcy. The notion of maximizing log-utility has well been established in the sports betting market, and is known therein as the *Kelly Criterion* (Kelly (1956)). It has also been established in the portfolio optimization literature. Kyle and Xiong (2001) and Lv and Meister (2009) suggest that Brieman's (1961) theorem holds in continuous time: as wealth approaches zero, log utility arbitrageurs become infinitely risk-averse. This condition guarantees non-bankruptcy. The second reason we assume log-utility relates to the mathematical tractability of the resulting stochastic optimization problem (Kim and Omberg (1996), Boguslavsky and Boguslavskaya (2004), and Kargin (2003)). This is particularly relevant to our optimization problem, given the nonlinearity present in (1).

3.4 The Optimal Portfolio Strategy:

Here we derive the optimal portfolio policy for the representative arbitrageur. First, we impose the self-financing condition. This requires that changes in wealth are directly attributable to

---

[14]Translating this to the investment actions of a typical hedge fund, for example, terminal wealth can be defined as the fund's performance at the end of a reporting period.

[15]See e.g. Wachter (2002)

investment in the risky and risk-free asset, and no external source. If N_t is the number of units held in the mispricing, and M_t is the number of units held in the riskless asset, then the self-financing condition implies that our budget constraint is given by:

$$
\begin{aligned}
dW_t &= N_t dS_t + M_t dB_t \\
&= N_t dS_t + \frac{W_t - N_t(S_t - \overline{S})}{B_t} dB_t,
\end{aligned}
$$

(6)

where the second equality in (6) imposes that whatever remains after investment in the mispricing is invested in the riskless asset.[16] To facilitate the analysis, we define a new time variable $\tau = T - t$ as the time left for investing, and rewrite (5) as:

(7)
$$
V(S_t, W_t, \tau) = \max_{N_t} \left[\mathbb{E}_t \left(\ln W_T \right) \right].
$$

Maximizing (7) subject to (6) yields the optimal portfolio strategy. The optimal number of units held in the mispricing depends on the current mispricing level, current wealth, and time left for trading.[17] Therefore, it is denoted by $N_t(S_t, W_t, \tau)$. It is derived as follows:

Theorem 1. (OPTIMAL PORTFOLIO STRATEGY) *The optimal portfolio strategy for a log-utility maximizing arbitrageur facing investment opportunities (1) and (2) is given by:*

(8)
$$
N_t(S_t, W_t, \tau) = \left(\frac{-\frac{k}{c} \tanh\left[c(S_t - \overline{S}) \right] - r(S_t - \overline{S})}{\sigma^2} \right) W_t.
$$

Proof. The proof is given in the Appendix.

$\square$

Remark 1. It is interesting to contrast our optimal portfolio strategy with the equivalent OU strategy. Firstly, taking the limit $c \to 0$ in (8) yields the corresponding OU optimal strategy, which we label N_{OU}:

$$
\begin{aligned}
N_{OU} &= \lim_{c \to 0} \left(\frac{-\frac{k}{c} \tanh\left[c(S_t - \overline{S}) \right] - r(S_t - \overline{S})}{\sigma^2} \right) W_t \\
&= \left(\frac{-(k + r)(S_t - \overline{S})}{\sigma^2} \right) W_t.
\end{aligned}
$$

---

[16]Similarly, an arbitrageur willing to make an investment in the mispricing larger than current wealth can do so by shorting the riskless asset.

[17]An artifact of log-utility is that the strategy is independent of the time horizon (Kim and Omberg (1996) and Wachter (2002)). We include the notation for generality.

For a common set of parameters and $\overline{S} = 0$, N_{OU} yields the specification in Jurek and Yang (2007). Setting $r = 0$ yields the specification in Boguslavsky and Boguslavskaya (2004).

While c is positive, our optimal holding is nonlinear in the mispricing, unlike OU. We illustrate this point by comparing our strategy (8), labeled N_{TANH}, to the equivalent OU strategy N_{OU} using common parameter values.
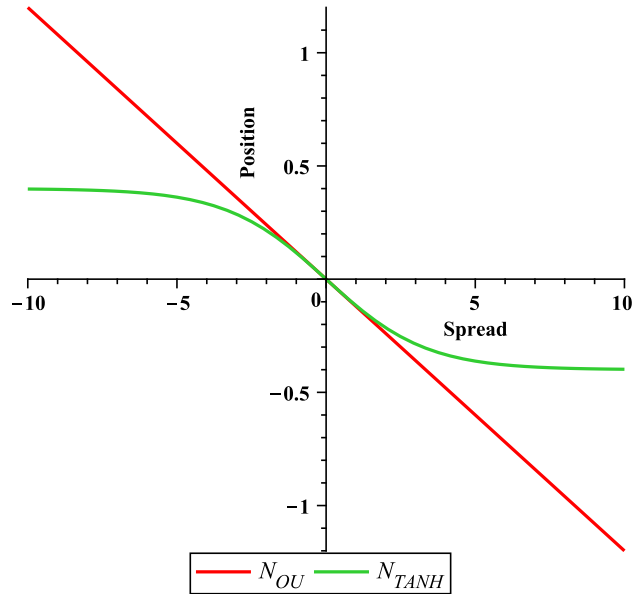

FIGURE 1. **Illustration of position size in the mispricing (spread)**. The OU (red) and TANH (green) optimal portfolio strategies with parameters for nonlinearity $c = 0.3$, mean-reversion $k = 3$, volatility $\sigma = 5$, wealth $W = 1$, risk-free rate $r = 0$, and natural mispricing level $\overline{S} = 0$. The mispricing range is $S \in [-10, 10]$, and the position in the mispricing taken by both strategies is on the vertical axis.

Figure 1 illustrates an interesting property: Unlike OU, our strategy implies that, as the mispricing diverges away from its natural level, an arbitrageur becomes less interested in exploiting it. This idea represents an important difference between our model and OU. Our model interprets large divergences as potential regime breaks, consistent with the notion of fundamental risk (Lamont and Thaler (2003b)). On the other hand, OU interprets large divergences as attractive investment opportunities.

3.5 Do Arbitrageurs Always Exploit a Mispricing?

It seems intuitive to think that an arbitrageur will always exploit a mispricing in pursuit of a profit. Here, we investigate this claim. Consequently, we scrutinize our optimal portfolio

policy (8) to determine the direction in which an arbitrageur trades as a response to the evolution of the mispricing process. We show that there is a critical level in the mispricing beyond which arbitrageurs unwind positions at a loss to preserve capital.[18] These are our *stability bounds*, so-called because an arbitrageur only seeks to enforce price efficiency as long as the mispricing lies within these bounds. As long as the magnitude of the mispricing relative to its natural level lies within the stability bounds, further divergence causes the arbitrageur to continue investing against the direction of the spread asset (i.e. exploiting the mispricing). However, once this stability bound is breached, further divergence causes the arbitrageur to unwind a losing position.[19] The stability bounds are analytically characterized in the following theorem. For simplicity and without loss of generality, we set the riskfree rate $r = 0$.

Theorem 2. (STABILITY BOUNDS) *An arbitrageur will only trade against a mispricing if the magnitude of the mispricing beyond its natural level,* $\left| S_t - \overline{S} \right|$*, lies within a fixed bound. Beyond this bound, an arbitrageur will unwind a losing position. The bound is as follows*:

$$(9) \qquad \left| (S_t - \overline{S}) \right| < \frac{1}{c} \operatorname{arcsinh} \left(\frac{\sigma c}{\sqrt{k}} \right).$$

Proof. We begin by expanding our optimal strategy (8) using Ito's lemma, suppressing the time-subscripts for neatness:

$$(10) \qquad dN = \frac{\partial N}{\partial S} dS + \frac{\partial N}{\partial W} dW + \frac{\partial^2 N}{\partial S^2} (dS)^2 + \frac{\partial^2 N}{\partial W^2} (dW)^2 + \frac{\partial^2 N}{\partial S \partial W} (dS \cdot dW) - \frac{\partial N}{\partial \tau} dt.$$

Next, substituting from (1), (6), and (8), along with the relevant partial derivatives of N, and using the fact that $(dS)^2 = \sigma^2 dt$, $(dW)^2 = \sigma^2 N^2 dt$, and $(dS \cdot dW) = \sigma^2 N dt$, we rewrite (10) as:

$$(11) \qquad \begin{aligned} dN &= \left[\frac{-wk \, \operatorname{sech}^2(c(S - \overline{S}))}{\sigma^2} + \frac{wk^2 \tanh^2(c(S - \overline{S}))}{\sigma^4 c^2} \right] dS \\ &+ \left[\frac{wk(2c^2\sigma^2 + k) \, \operatorname{sech}^2(c(S - \overline{S})) \tanh(c(S - \overline{S}))}{c\sigma^2} \right] dt. \end{aligned}$$

---

[18]Jurek and Yang (2007) and Boguslavsky and Boguslavskaya (2004) explore this notion using OU. The economic mechanism behind this notion is the relation between the investment effect and wealth effect. See Xiong (2001).

[19]In partial equilibrium, it is not possible to analyze the magnitude of the effect arbitrageurs' actions would have on the mispricing itself; we do however state the direction the effect of the representative arbitrageur's action is likely to have on the mispricing. Xiong's (2001) general equilibrium model analyzes this effect assuming OU.

Equation (11) separates the instantaneous changes in the optimal portfolio allocation as a response to changes in the mispricing dS and the time variable dt. Therefore, the direction of the arbitrageur's portfolio allocation in response to changes in the mispricing is governed by the sign of the first set of square brackets in (11), namely:

$$\left[\frac{-wk \ \text{sech}^2(c(S - \overline{S}))}{\sigma^2} + \frac{wk^2 \tanh^2(c(S - \overline{S}))}{\sigma^4 c^2} \right],$$

and this is negative when:

$$\frac{wk^2 \tanh^2(c(S - \overline{S}))}{\sigma^4 c^2} - \frac{wk \ \text{sech}^2(c(S - \overline{S}))}{\sigma^2} < 0.$$

Simplifying this expression, we obtain the result:

$$(12) \qquad \left| (S - \overline{S}) \right| < \frac{1}{c} \ \text{arcsinh} \left(\frac{\sigma c}{\sqrt{k}} \right).$$

This completes the proof of Theorem 2.

$\square$

Remark 2. Taking the limit $c \to 0$ shows that our stability bounds (12) tend to those implied by OU (Jurek and Yang (2007), Boguslavsky and Boguslavskaya (2004)):

$$(13) \qquad \lim_{c \to 0} \left[\frac{1}{c} \ \text{arcsinh} \left(\frac{\sigma c}{\sqrt{k}} \right) \right] = \frac{\sigma}{\sqrt{k}}.$$

Direct calculation in (13) shows that for any values of reversion $k > 0$, volatility $\sigma > 0$, and coefficient of nonlinearity $c > 0$, the stability bounds in our model are tighter than those implied by OU, which means that our model calls for smaller losses to be cut sooner, rather than larger losses cut later. Based on this fact, we suggest that the potentially negative impacts of arbitrage trading within our model would be lower relative to OU.[20]

3.6 Robustness to Parameter Uncertainty:

It has been implicit up to this point that the investor has perfect *ex-ante* knowledge regarding the values of the parameters k, c, and σ. This assumption is maintained in Brieman's (1961) and Lv and Meister's (2009) theorems regarding non-bankruptcy of the trading strategy. In reality, parameters are often estimated from historical data, and these theories seem predicated on the assumption that the future perfectly mimics the past, and that the model used describes the time series perfectly. Of course one cannot affirm that the future will

---

[20]For a discussion of the wider effects of arbitrage trading, see Shleifer and Vishny (1997), Xiong (2001), and De Long et al. (1990)

mimic the past perfectly, nor that any diffusion model can describe market evolution perfectly. Thus the risk of bankruptcy persists in reality,[21] which is in part due to model risk or Knightian uncertainty (Knight (1921)). This realization leads us to an interesting empirical question, namely "What is the effect of parameter misspecification on the profit-generating and risk-management capabilities of the optimal strategy?"

To answer this question, we compare our optimal strategy (8) against the OU strategy on synthetically generated data, then forced trading in both models based on deliberately wrong reversion and volatility parameters. The methodology is as follows:

We simulate 50,000 paths from a discretized OU process over 1 year with a time increment of $\frac{1}{252}$ (corresponding to 252 trading days). The OU data is generated by a Euler-Maruyama discretization scheme (see Kloeden and Platen (1999)). We set $c = 0.3$, $k = 3$, $\sigma = 5$, $r = 0$, and $\overline{S} = 0$. The initial value for each simulated path is set at zero, corresponding to no immediate arbitrage opportunity. We then apply our optimal strategy (8), which we refer to now as N_{TANH}, against the optimal strategy implied by OU, which we call N_{OU}. For clarity, these are given explicitly as follows:

$$(14) \qquad N_{TANH} \;=\; \left(\frac{-\frac{k}{c}\tanh\left(c(S-\overline{S})\right) - r(S-\overline{S})}{\sigma^2} \right) W.$$

$$(15) \qquad N_{OU} \;=\; \left(\frac{-k(S-\overline{S}) - r(S-\overline{S})}{\sigma^2} \right) W.$$

To test the effect of parameter misspecification in (14) and (15), we trade using deliberately wrong parameters k and σ. Specifically, we set $k = k(1+K)$ and $\sigma = \sigma(1+\Sigma)$, with $K \in [0,3]$ and $\Sigma \in [-0.5, 0]$. This means that the two strategies each trade with k ranging from its true value 3 up to quadruple its true value 12, and σ ranging from its true value 5 down to half its true value 2.5. We define our parameter grid such that it contains 30 values for each parameter, implying 900 parameter combinations in total.

For each parameter combination and each strategy, we measure two quantities: the proportion of trials which result in bankruptcy, and the mean terminal wealth over all trials. The results are presented in Figure 2.

---

[21]Jurek and Yang (2007) point out that another source of bankruptcy risk comes from the discretization procedure used in calibrating a continuous-time model to historical data.
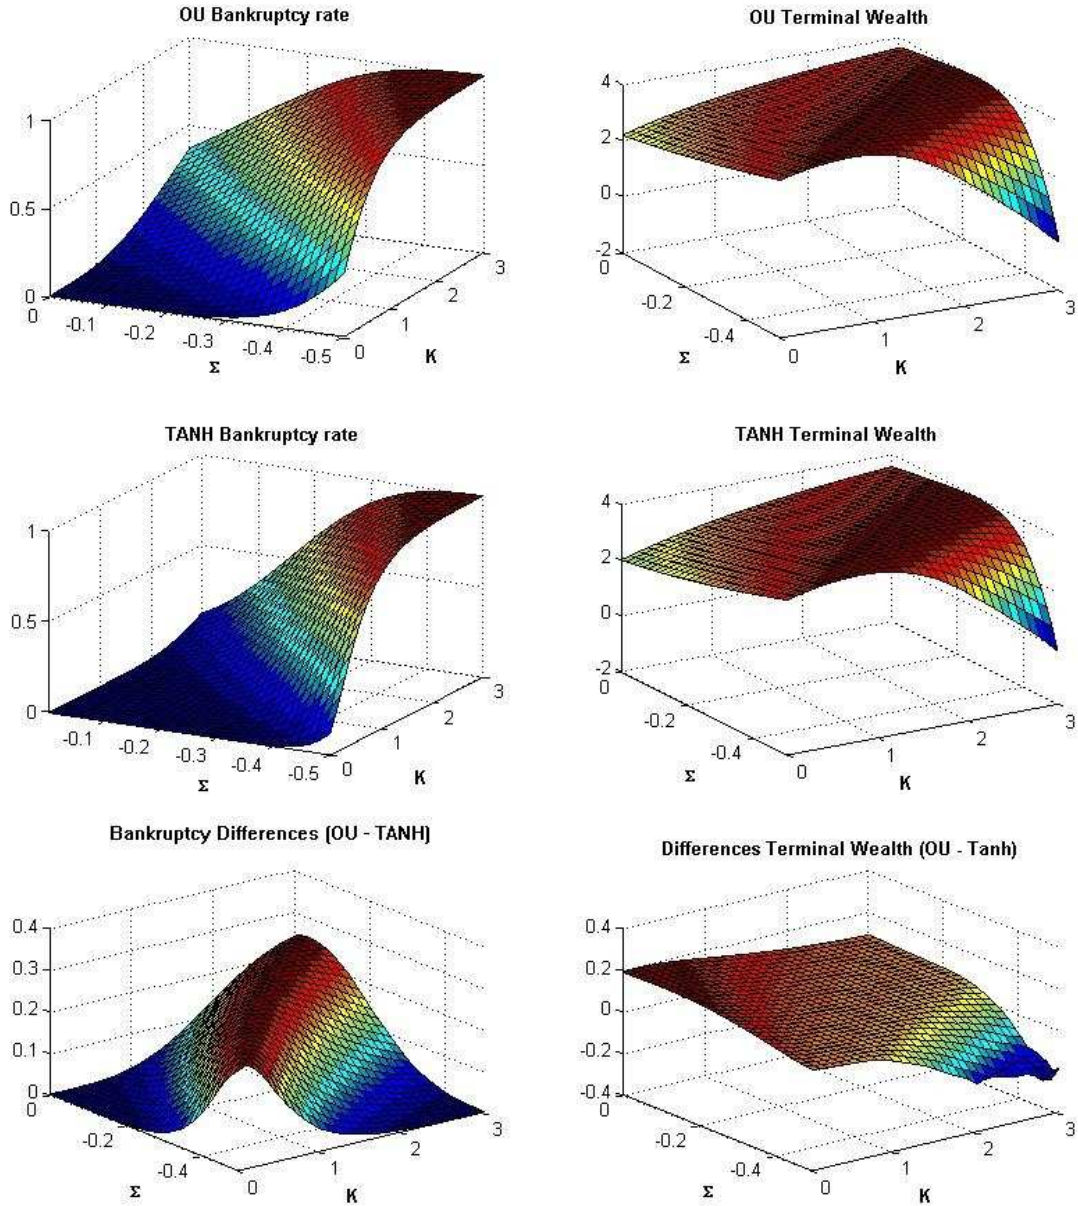
FIGURE 2. **Effect of parameter misspecification on the optimal strategy.** Simulation of 50,000 OU paths over 252 trading days with parameters $k = 3$, $\sigma = 5$, $r = 0$, $c = 0.3$, and $\overline{S} = 0$. Both OU and TANH optimal strategies (14) and (15) trade assuming deliberately wrong parameter values. $K \in [0, 3]$ implies reversion $k \in [3, 12]$, while $\Sigma \in [-0.5, 0]$ implies volatility $\sigma \in [2.5, 5]$. Initial wealth is set to 1. The left panels show the effect of parameter misspecification on the probability of bankruptcy, while the right panels show this effect on mean terminal wealth.

The top-left panel in Figure 2 shows the proportion of trials of the OU strategy that result in bankruptcy. Trading using the correct parameters universally causes none of the trials to result in bankruptcy. However, we see that overestimating the reversion parameter leads to an increase in the proportion of trials that result in bankruptcy. Underestimating the volatility parameter has the same effect. This is intuitive, as the portfolio strategy seeks to exploit a mispricing which it believes will revert faster than it actually would, or exhibit less volatility than it actually would, respectively. Our findings are consistent with Lv and Meister (2009) in the sense that quadrupling the estimated reversion parameter has the identical detrimental effect to halving the volatility estimate. For N_{OU} in the top-left panel, 37.04% of trials ended in bankruptcy as a result of trading using half the volatility or double the reversion. The corresponding figure for N_{TANH} in the center-left panel is 12.24%. This shows that, given a common set of model parameters, our optimal TANH strategy (14) is more robust to parameter misspecification than the corresponding OU strategy in terms of bankruptcy risk. This point is illustrated in the bottom-left panel, which subtracts the proportion of trials resulting in bankruptcy of the TANH strategy from that of the OU strategy for each combination of k and σ. That the bottom-left panel is positive reflects the notion that the TANH strategy is less susceptible to parameter misspecification.

The top-right and center-right panels show the effect of parameter misspecification on terminal wealth in both strategies. These panels show that OU results in higher mean terminal wealth. The reason for this is that the diminishing marginal exploitation of the spread in the N_{TANH} strategy for a common set of parameters implies that the TANH strategy always invests relatively less in the spread than N_{OU} for a given wealth. Therefore, should the larger position assumed by N_{OU} not result in bankruptcy, mean terminal wealth would be greater, reflecting the greater risk OU takes. This fact is illustrated in the bottom-right panel, which subtracts the average terminal wealth of the TANH strategy from the OU strategy for each combination of k and σ. That the bottom-right panel is largely positive implies that, conditional on survival, the OU strategy outperforms the TANH strategy in terms of maximizing terminal wealth.

In our simulation, bankrupted trials were not excluded from the calculation of average terminal wealth. This potentially introduces a bias in the calculation, but excluding bankrupted trials results in a large positive bias in terminal wealth. This is intuitive, as the surviving trials would have taken extremely large positions in the mispricing, which "luckily" paid off, and exaggerate the wealth-maximizing capability of the models.

4. Empirical Application

Now that we have characterized the optimal portfolio strategy (14) and explored its properties, we are ready to implement our results to historical data. We do this by testing the empirical efficacy of the N_{TANH} strategy (14) against two alternative strategies. The first is the N_{OU} optimal strategy (15). The second is a popular "rule of thumb" strategy based on Gatev, Goetzmann, and Rouwenhorst (2006). For consistency we call this strategy N_{GGR}, and it works as follows:

(16) N_{GGR}: Open a position in the mispricing when the mispricing deviates more than 2 standard deviations away from its historical mean. Hold the position constant until the mispricing reverts fully to its natural level. The size of the position opened and held is the equal to what N_{OU} would suggest (i.e. the position of an equivalent OU arbitrageur).

Remark 3. The nature of N_{GGR} is such that a position is not opened in a mispricing immediately as a mispricing diverges from its natural level. Because of this, some tests result in N_{GGR} not trading at all. Also, we do not implement a stop-loss mechanism in N_{GGR}.

4.1 Data, Historical Period, and Methodology:

The data is obtained from DataStream, and chosen for four pairs of FTSE100 companies using closing daily prices over the period Jan 1, 2001 - Jun 30, 2008 (the *training* period), and Jul 1, 2008 - Jun 30, 2009 (the *trading* period). When these days are non-trading (holidays, new years), the previous day's closing value is chosen as a proxy.

The four pairs are chosen based on the highest correlation of daily returns over the training period. Beyond the statistical relationship, they are commercially intuitive choices, and are as follows:

- HSBC and Barclays Plc.
- Legal & General and Prudential.
- Royal Dutch Shell-B and BP.
- Royal Bank of Scotland and Lloyds Plc.

In constructing the mispricing S_t, we take the linear difference (spread) in daily closing prices for each pair. For example, denoting these prices by $P_{i,t}$ for $i = \{1, 2\}$ implies:

$$S_t = P_{1,t} - P_{2,t}.$$

The calibration procedure involves calibrating a OU process for the set of reversion, volatility, and natural level parameters $\{\hat{k}, \hat{\sigma}, \hat{\overline{S}}\}$ respectively, over the training period Jan 1, 2001 - Jun 30, 2008 (a total of 1956 trading days). We use the well-known Maximum Likelihood Estimation method to estimate the parameters of OU.[22] Once we obtain estimates for $\{\hat{k}, \hat{\sigma}, \hat{\overline{S}}\}$, we substitute them into both N_{OU} and N_{TANH}[23] and set the coefficient of nonlinearity $c = 0.02$ into N_{TANH}. For N_{GGR}, we set the historical mean equal to the natural level implied by the OU calibration, i.e. $\mu_{\hat{GGR}} = \hat{\overline{S}}$, and calculate the historical standard deviation $\sigma_{\hat{GGR}}$ based on the 1956 daily data points over the training period.

With this in hand, we proceed by backtesting the three strategies (14), (15), and (16) over the trading period Jul 1, 2008 - Jun 30, 2009, a total of 261 trading days. Initial wealth is set arbitrarily at 1.

4.2 Application to FTSE100 data:

The results of our backtest are shown in Figures 3 through 6. Throughout the backtest, we set $r = 0$ to illustrate clearly the differences between the three strategies without the complication of an additional asset. We further assume no transaction costs, full use of the short proceeds, and no other institutional frictions.

The results show that, when mean reversion is poor, or there are apparent signs of a regime-break, both the N_{OU} and N_{GGR} strategies lead to very large losses, or bankruptcy. Conversely, assuming a coefficient of curvature $c = 0.02$, the N_{TANH} strategy never resulted in bankruptcy, despite poor mean reversion and regime breaks in the data. Figure 6 shows that although N_{OU} returns are volatile, they are very high when reversion is strong about the natural level $\overline{S}$. Given that the nature of OU is such that an arbitrageur has complete confidence in reversion to the natural level, this result is expected.

The backtesting period is chosen specifically to test the models' ability to handle the tumultuous market conditions characteristic of that period. Nevertheless, our four examples in Figures 3 through 6 represent a wide range of market characteristics regarding the evolution of the spread. Overall the results are consistent with the intuition from the works mentioned in sections 1 and 2 of this paper.

---

[22]We refer to Appendix E of Jurek and Yang (2007) for the method, and omit the details here for brevity.

[23]Implementing N_{TANH} as a stand-alone model requires its parameters to be estimated independently through numerical methods (Hurn et al. (2003), Picchini (2007)). In our application, we use the parameter c as a proxy to Knightian uncertainty, and hence use a common set of parameters for N_{TANH} and N_{OU}. Consequently, an interpretation of our methodology is a replacement of the risk aversion parameter inherent in general CRRA utility, in favor of a parameter of model-risk aversion: The higher the value of c set by the arbitrageur, the less confidence he has in the ability of OU to capture the dynamics of the mispricing.
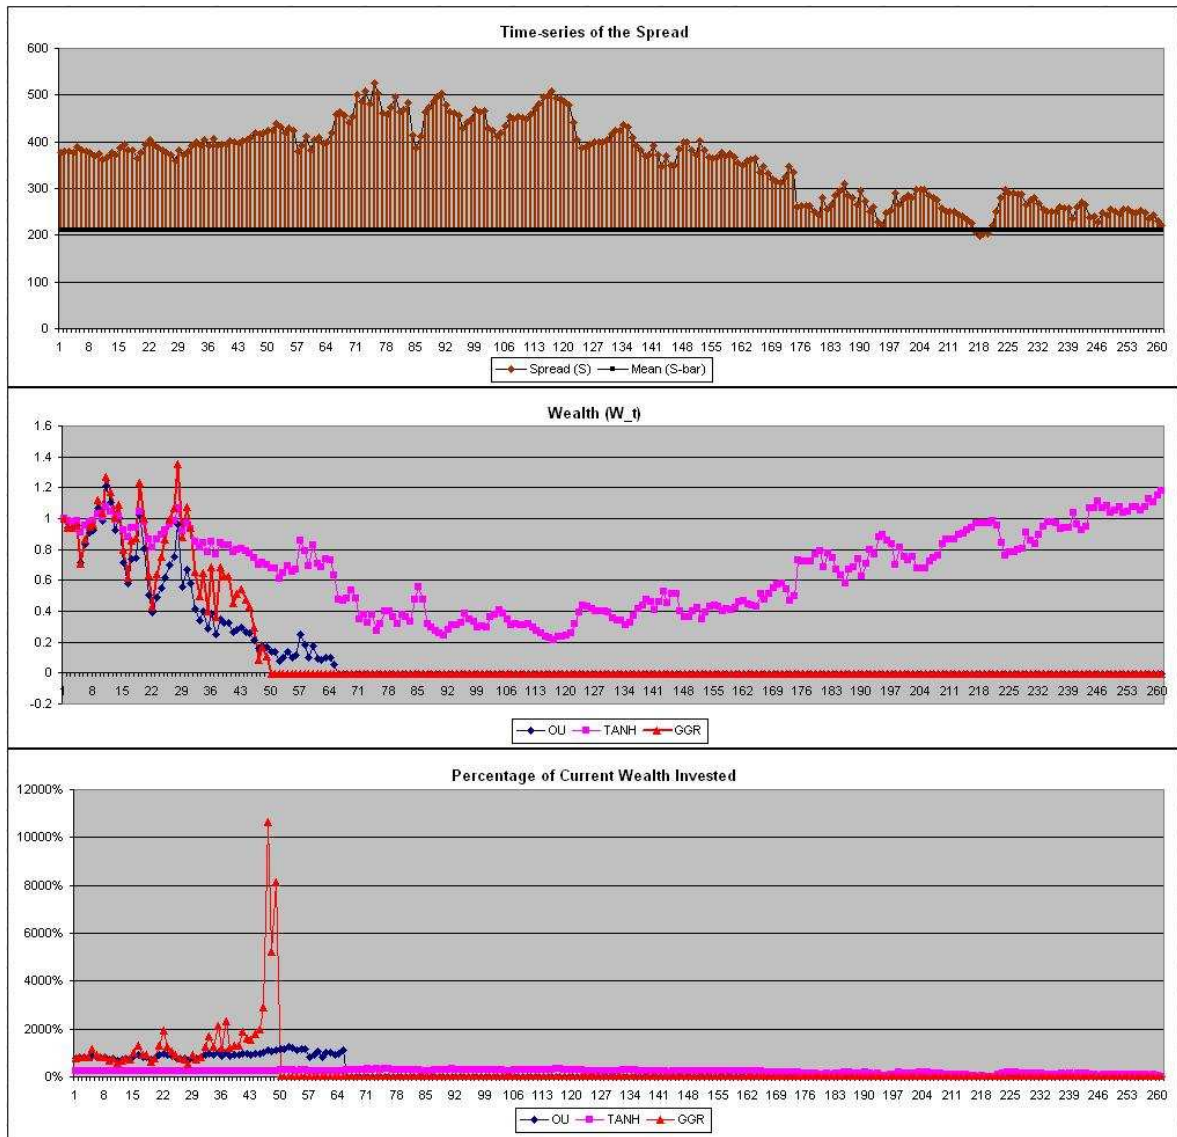
FIGURE 3. **HSBC - Barclays: Poor mean reversion.** The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.
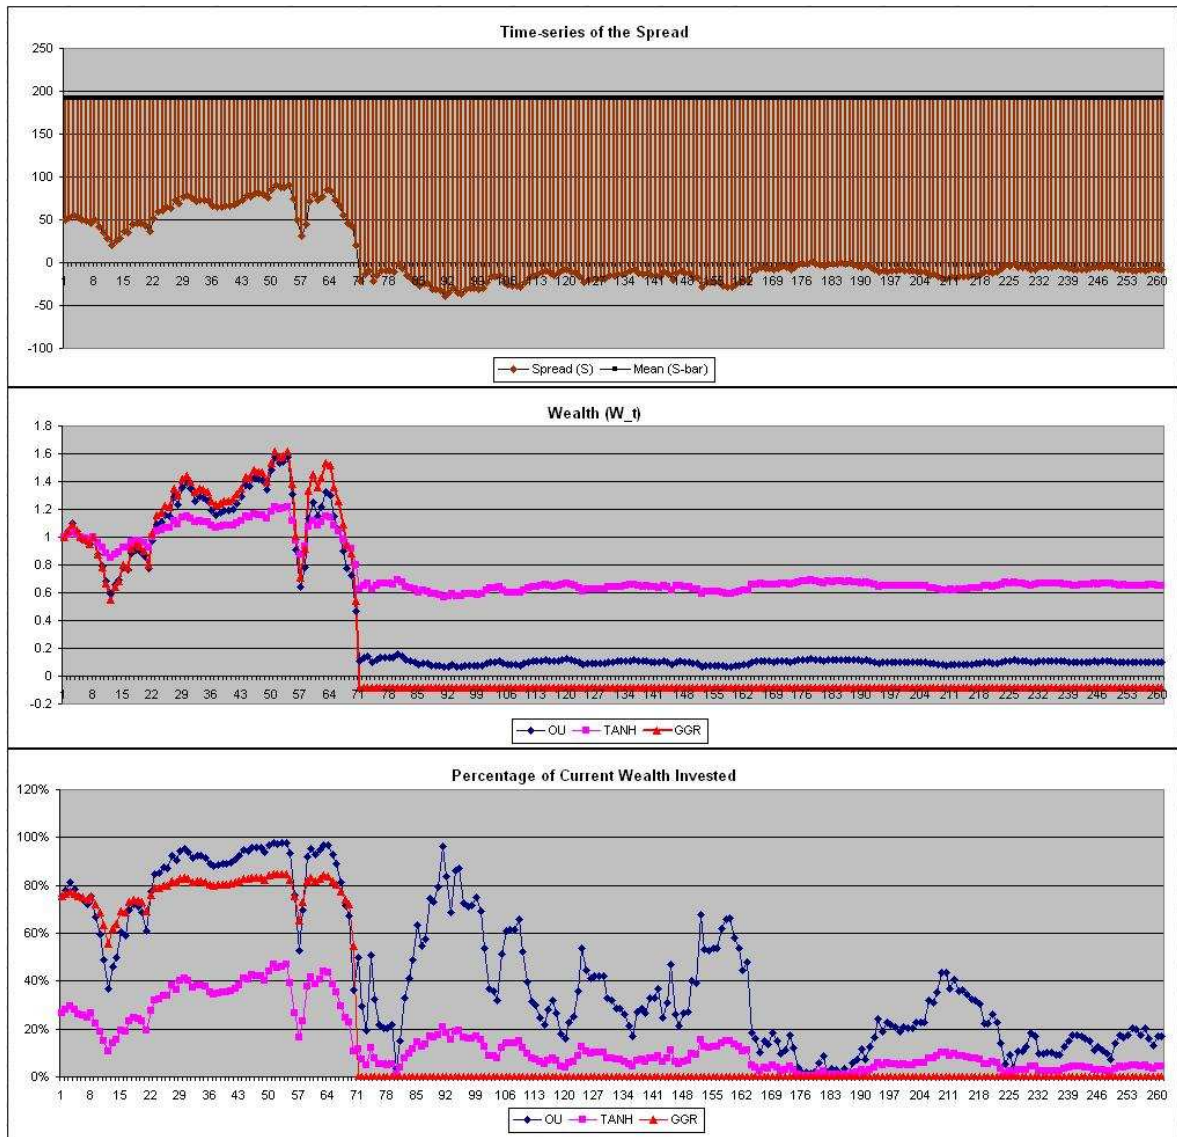
FIGURE 4. **RBS - Lloyds: Regime break.** The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.
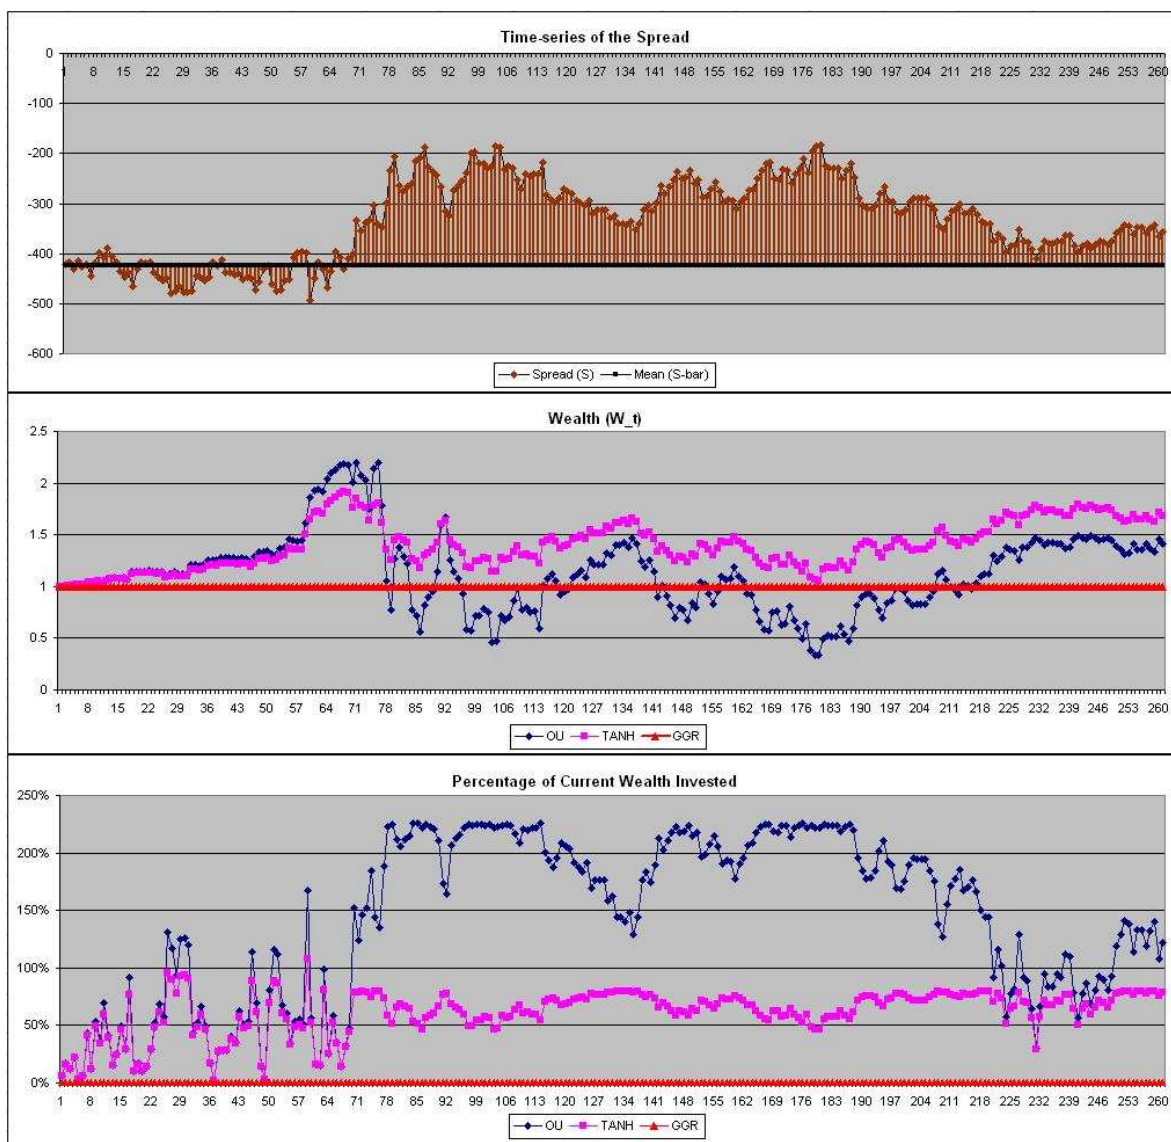
FIGURE 5. **Legal&General - Prudential: Poor mean reversion.** The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.
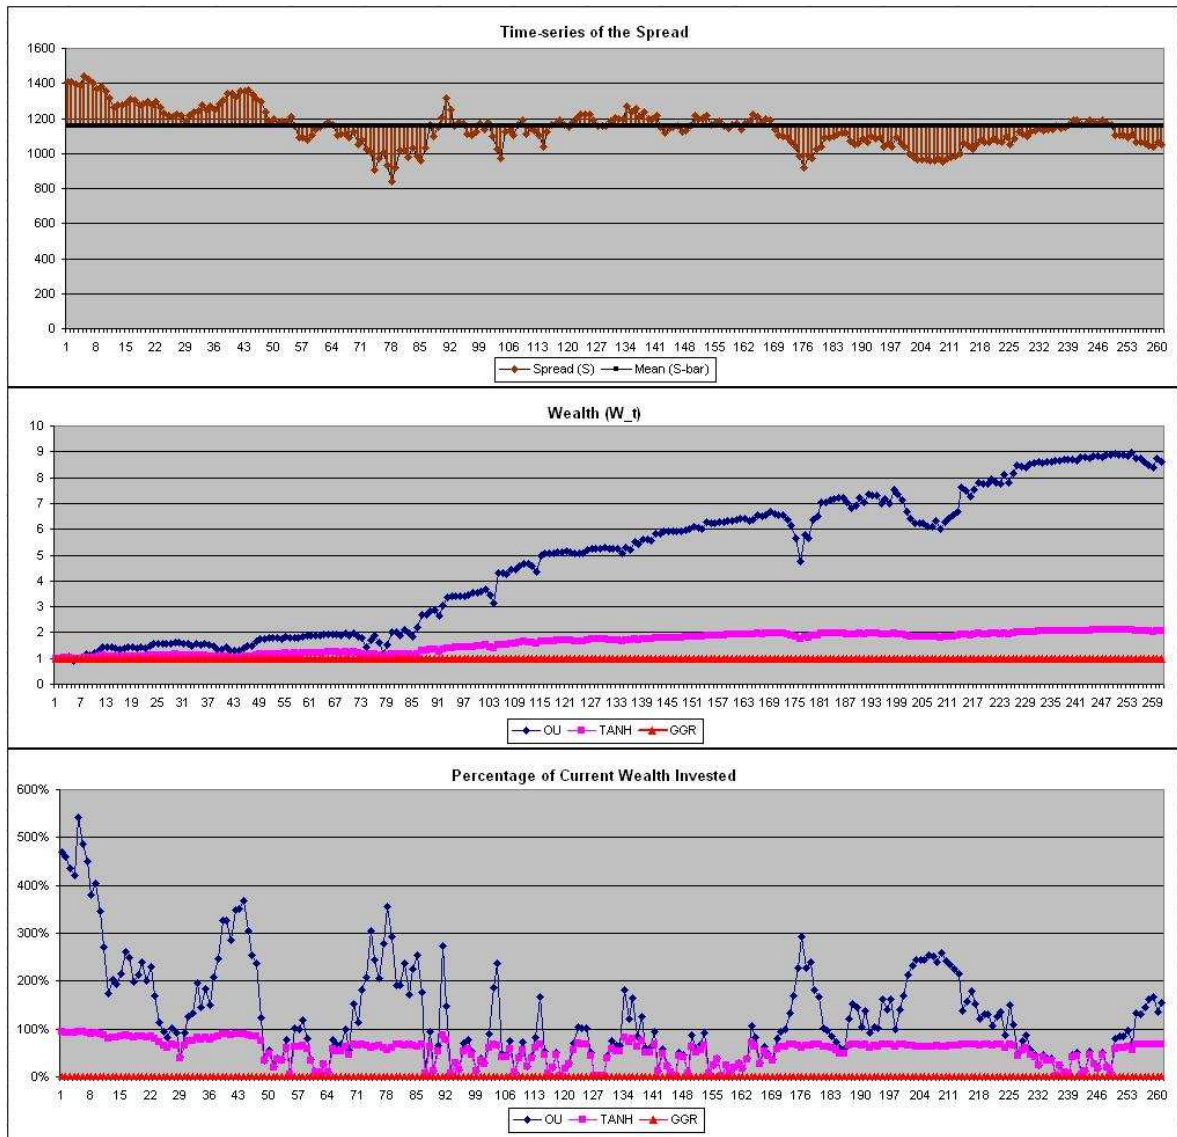
FIGURE 6. **Royal Dutch Shell B - BP: Good mean reversion.** The top panel shows the evolution of the daily spread during Jul 1, 2008 - Jun 30, 2009 (Brown) and the natural spread level (Black). The middle panel shows the corresponding evolution of wealth resulting from trading using N_{TANH} (Pink), N_{OU} (Blue), and N_{GGR} (Red). The bottom panel shows the percentage of wealth invested in the spread.

5. Conclusion

We develop a partial equilibrium, nonlinear model of arbitrage aimed specifically at capturing the fundamental market- and sentiment-based risks inherent in arbitrage trading into an exogenously-driven mispricing process. In so doing, we demonstrate that the incorporation of these risks into our model yields better empirical performance than traditional linear-based and "rule of thumb" models employed elsewhere in the literature. Most importantly, we derive an optimal portfolio strategy in closed form using stochastic optimal control. The analytical simplicity of our framework provides potential both for its practical implementation, and its use in further investigating the role of arbitrageurs.

Our approach draws insight directly from the empirical and general equilibrium perspectives regarding the behavior of rational arbitrageurs facing irrational noise traders. As a result, our portfolio strategy captures three broad phenomena inherent in arbitrage trading: divergence risk, horizon risk, and regime break. Moreover, we are able to provide a proxy to the notion of Knightian uncertainty, or model risk. In this regard, our model delivers empirically superior performance compared to Ornstein-Uhlenbeck or "2-sigma" strategies. This result suggests that our model better captures realistic market conditions, and yields novel insights into the optimal behavior of arbitrageurs.

Though the role of arbitrageurs is typically aimed at exploiting mispricings and enforcing price efficiency, we show that under certain conditions, arbitrageurs unwind losing positions by effectively trading *with* the mispricing, even when investible mispricing opportunities are at their greatest. The closed-form availability of these conditions complements the existing literature on arbitrage in both partial and general equilibrium. Specifically, our model calls for small losses to be cut sooner, in contrast to OU-type models which call for large losses to be cut later. We conjecture that this artifact results in lower price volatility and inefficiency upon the liquidation of funds.

We identify a number of interesting avenues for future research. Mainly, the incorporation of market frictions would allow a serious study of price efficiency in cointegrated-asset and twin-asset markets based on an algorithmic trading approach. Though such an endeavor is unlikely to yield portfolio strategies in closed-form, it remains an area of interest to us.

References

Abreu, Dilip, and Markus K. Brunnermeier (2002) *Synchronization risk and delayed arbitrage*, Journal of Financial Economics 66, 341360.

Abreu, Dilip, and Markus K. Brunnermeier (2003), *Bubbles and crashes*, Econometrica 71, 173204.

Basak, Suleyman and Benjamin Croitoru (2000). *Equilibrium Mispricing in a Capital Market with Portfolio Constraints.* The Review of Financial Studies, 13, 715-748.

Bellman, R.E. (1957). *Dynamic Programming.* Princeton University Press, Princeton, NJ. Republished 2003: Dover, ISBN 0486428095

Bjork, T. (2004). *Arbitrage Theory in Continuous Time.* Oxford University Press.

Black, Fischer (1986). *Noise.* The Journal of Finance, 41, 529-543.

Boguslavsky, M., and E. Boguslavskaya (2004). *Arbitrage under Power.* RISK magazine, pp.69-73.

Breiman, L. (1961) *Optimal gambling system for favorable games.* Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability, 1, 63-8.

Brunnermeier, Markus K., and Lasse H. Pedersen (2005). *Predatory trading*, Journal of Finance 60, 18251863.

Campbell, John Y., and Luis M. Viceira (2002), *Strategic Asset Allocation.* Oxford University Press, New York.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, (1990). *Noise Trader Risk in Financial Markets*, The Journal of Political Economy, 98, 703-738.

Edwards, Franklin R. (1999). *Hedge funds and the collapse of long-term capital management*, Journal of Economic Perspectives 13, 189210.

Elliott, R. J., J. van der Hoek, and W. P. Malcolm (2005). *Pairs trading.* Quantitative Finance, 5(3), 271276.

Engle, Robert F. and Clive W. J. Granger (1987), *Co-Integration and Error Correction: Representation, Estimation and Testing*, Econometrica, 55, p. 149-158.

Friedman, Milton (1953). *The Case For Flexible Exchange Rates.* Essays in Positive Economics, University of Chicago Press.

Froot, Kenneth A. and Emil M. Dabora (1999). *How are Stock Prices Affected by the Location of Trade?*, Journal of Financial Economics, 53, 189-216.

Gatev, Evan, William N. Goetzmann, and K. Geert Rouwenhorst (2006). *Pairs Trading: Performance of a Relative Value Arbitrage Rule*, Review of Financial Studies, 19, 797-827.

Gromb, Denis, and Dimitri Vayanos (2002). *Equilibrium and welfare in markets with financially constrained arbitrageurs.* Journal of Financial Economics 66, 361407.

Hurn, A.S., K.A. Lindsay, and V.L. Martin (2003). *On the efficacy of simulated maximum likelihood for estimating the parameters of stochastic differential equations.* Journal of Time Series Analysis, 24(1):4563.

Jurek, Jakub and Halla Yang (2007). *Dynamic Portfolio Selection in Arbitrage.* Harvard University working paper.

Kargin, V. (2003). *Optimal Convergence Trading.* arXiv: math.OC/0302104

Kelly, Jr, J. L. (1956). *A New Interpretation of Information Rate.* The Bell System Technical Journal, 35(4), 917926.

Kim, S-J, J. A. Primbs, and S. Boyd (2008). *Dynamic Spread Trading.* Submitted.

Kim, Tong Suk and Edward Omberg (1996) *Dynamic Nonmyopic Portfolio Behavior.* The Review of Financial Studies, 9, 141-161.

Kloeden, P.E., and E. Platen (1999). *Numerical Solution of Stochastic Differential Equations.* Springer, Berlin. ISBN 978-3-540-54062-5.

Kondor, Peter (2009). *Risk in Dynamic Arbitrage: Price Effects of Convergence Trading.* Journal of Finance, 64(2), 638-658.

Korn, R., and H. Kraft (2002). *A Stochastic Control Approach to Portfolio Problems with Stochastic Interest Rates.* SIAM Journal on Control and Optimization, 40(4):12501269.

Knight, F.H. (1921) *Risk, Uncertainty, and Profit.* Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company.

Kyle, Albert S., and Wei Xiong (2001). *Contagion as a wealth effect*, Journal of Finance 56, 14011440.

Lamont, Owen and Richard Thaler (2003a). *The Law of One Price in Financial Markets*, Journal of Economic Perspectives, 17(4), 191-202.

Lamont, Owen and Richard Thaler (2003b). *Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs.* Journal of Political Economy, 111, 227-268.

Liu, Jun and Francis A. Longstaff (2004). *Losing Money on Arbitrage: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities*, The Review of Financial Studies, 17, 611-641.

Liu, Jun (2007). *Portfolio Selection in Stochastic Environments*, Review of Financial Studies, 20(1), 1-39.

Lv, Y., Meister, B.K. (2009). *Application of the Kelly Criterion to Ornstein-Uhlenbeck Processes.* arXiv preprint, arXiv:0903.2910.

MacKenzie, Donald (2003). *Long-Term Capital Management and the sociology of arbitrage*, Economy and Society 32, 349380.

Malkiel, Burton (1977). *The Valuation of Closed-End Investment-Company Shares*, The Journal of Finance, 32, 847-859.

Merton, R. C. (1969). *Lifetime Portfolio Selection: The Continuous-Time Case*, Review of Economics and Statistics, 51, 247257.

Merton, R. C. (1971). *Optimum Consumption and Portfolio Rules in a Continuous-Time Model*, Journal of Economic Theory, 3, 373413.

Mitchell, Mark, Todd Pulvino and Erik Stafford (2002). *Limited Arbitrage in Equity Markets*, Journal of Finance, 57, 551-584.

Mudchanatongsuk, J. Primbs, and W. Wong (2008). *Optimal Pairs Trading: A Stochastic Control Approach*, Proceedings of the American Control Conference, Seattle, WA, pp. 1035-1039.

Picchini, U. (2007). *SDE Toolbox: Simulation and Estimation of Stochastic Differential Equations with MATLAB*, http://sdetoolbox.sourceforge.net.

Shaw, W.T. (2009). *A model of returns for the post-credit-crunch reality: Hybrid Brownian motion with price feedback*. Submitted.

Shleifer, A., and Lawrence H. Summers (1990). *The Noise Trader Approach to Finance*. The Journal of Economic Perspectives 4(2), pp. 19-33.

Shleifer, A., Vishny, R. (1997). *The limits of arbitrage*. Journal of Finance 52, 3555.

Vasicek, Oldrich (1977). *An Equilibrium Characterization of the Term Structure*. Journal of Financial Economics 5: 177-188.

Wachter, J. (2002), *Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets*, The Journal of Financial and Quantitative Analysis 37(1), 63-91.

Wong, E. (1964), *The Construction of a Class of Stationary Markov Processes*, Proc. Amer. Math. Soc. Symp. AppL Math., 16, 264-276.

Xiong, Wei (2001). *Convergence Trading with Wealth Effects: An Amplification Mechanism in Financial Markets*, Journal of Financial Economics, 62, 247-292.

Zastawniak, T., and Z. Brzezniak (1998). *Basic Stochastic Processes*. Springer.

Appendix

In this section we present the proof of Theorem 1:

Theorem 1. (OPTIMAL PORTFOLIO STRATEGY) *The optimal portfolio strategy for a log-utility maximizing arbitrageur facing investment opportunities (1) and (2) is given by:*

$$(16) \qquad N_t(S_t, W_t, \tau) = \left(\frac{-\frac{k}{c} \tanh\left[c(S_t - \overline{S}) \right] - r(S_t - \overline{S})}{\sigma^2} \right) W_t.$$

Proof. First, we substitute (1) and (2) into budget constraint (6) to obtain:

$$(17) \qquad \begin{aligned} dW_t &= N_t dS_t + \frac{W_t - N_t(S_t - \overline{S})}{B_t} dB_t \\ &= \left(-\frac{k}{c} N_t \tanh\left(c(S_t - \overline{S}) \right) + r(W_t - N_t(S_t - \overline{S})) \right) dt + \sigma N_t dZ_t, \end{aligned}$$

Now we employ stochastic optimal control. First we expand the value function (7) using Ito's Lemma and Ito multiplication.[24] Suppressing the time subscripts for neatness, we have:

$$dV = \frac{\partial V}{\partial S} dS + \frac{\partial V}{\partial W} dW + \frac{\partial^2 V}{\partial S^2}(dS)^2 + \frac{\partial^2 V}{\partial W^2}(dW)^2 + \frac{\partial^2 V}{\partial S \partial W}(dS \cdot dW) - \frac{\partial V}{\partial \tau} dt.$$

Substituting from (1) and (17), we have:

$$\begin{aligned} dV = \quad &\frac{\partial V}{\partial W} \left(\left(-\frac{k}{c} N \tanh\left(c(S - \overline{S}) \right) + r(W - N(S - \overline{S})) \right) dt + \sigma N dZ \right) \\ &+ \frac{\partial V}{\partial S} \left(-\frac{k}{c} \tanh\left(c(S - \overline{S}) \right) dt + \sigma dZ \right) + \frac{1}{2}\sigma^2 \frac{\partial^2 V}{\partial S^2} dt + \frac{1}{2}\sigma^2 N^2 \frac{\partial^2 V}{\partial W^2} dt \\ &+ \frac{1}{2}\sigma^2 N \frac{\partial^2 V}{\partial S \partial W} dt - \frac{\partial V}{\partial \tau} dt. \end{aligned}$$

Next, we employ Bellman's principle of optimality $\mathbb{E}_t(dV) = 0$ to derive the Hamilton-Jacobi-Bellman equation (Bellman (1957)) for this optimization problem:

$$(18) \qquad \begin{aligned} 0 = \, &\frac{\partial V}{\partial W} \left(-\frac{k}{c} N \tanh\left(c(S - \overline{S}) \right) + r(W - N(S - \overline{S})) \right) + \frac{\partial V}{\partial S} \left(-\frac{k}{c} \tanh\left(c(S - \overline{S}) \right) \right) \\ &+ \frac{1}{2}\sigma^2 \frac{\partial^2 V}{\partial S^2} + \frac{1}{2}\sigma^2 N^2 \frac{\partial^2 V}{\partial W^2} + \frac{1}{2}\sigma^2 N \frac{\partial^2 V}{\partial S \partial W} - \frac{\partial V}{\partial \tau}, \end{aligned}$$

along with the terminal condition:

$$(19) \qquad\qquad\qquad V(S, W, 0) = 0.$$

---

[24]for a textbook treatment of Ito's Lemma, see Zastawniak and Brzezniak (1998).

Condition (19) implies that at the end of the investment horizon $\tau = 0$, any arbitrage opportunity is worthless to the arbitrageur, since the trading period is over. In deriving (18), we have used the fact that Z_t is a martingale under the real-world measure. Equation (18) is the Hamilton-Jacobi-Bellman (henceforth "HJB") equation for our optimization problem. In order to derive the optimal policy, we maximize the right hand side of (18) with respect to N. To achieve this, we use the first-order optimality condition with respect to N by differentiating (18) with respect to N then setting this derivative equal to zero and solving for N. Differentiating, we have:

$$0 = \frac{\partial V}{\partial W} \left(-\frac{k}{c} \tanh \left(c(S - \overline{S}) \right) - r(S - \overline{S}) \right) + \sigma^2 N \frac{\partial^2 V}{\partial W^2} + \frac{1}{2} \sigma^2 \frac{\partial^2 V}{\partial S \partial W},$$

and finally solving for N while using subscripts to denote partial derivatives for neatness, we obtain:

$$(20) \qquad N = -\left(\frac{V_W}{V_{WW}} \right) \left(\frac{-\frac{k}{c} \tanh \left(c(S - \overline{S}) \right) - r(S - \overline{S})}{\sigma^2} \right) - \left(\frac{V_{SW}}{V_{WW}} \right).$$

Equation (20) gives us the optimal portfolio policy N. However, in its current form, it is of little use, since it depends explicitly on partial derivatives of the value function $V(S, W, \tau)$. Without obtaining an explicit form for $V(S, W, \tau)$ or safely eliminating the dependence of the optimal portfolio N on the derivatives of the value function, equation (20) remains of little practical use. In order to proceed in finding an explicit form for N, we employ the method of separation of variables. We postulate a trial solution, in terms of a functional form for $V(S, W, \tau)$, then verify whether it solves the HJB equation (18). We assert that it is likely $V(S, W, \tau)$ inherits some structural properties from its components, namely the log-utility function (Bjork (2004)). To this end, we postulate the following solution:

$$(21) \qquad V^*(S, W, \tau) = \ln(W) + f(S, \tau).$$

Here f is an arbitrary function of S and τ only. Condition (19) translates to $f(S, 0) = 0$. We now check whether this solution is successful. The partial derivatives, assuming $V^*(S, W, \tau)$ has the form (21), are:

- $V_W^* = \frac{1}{W}$
- $V_{WW}^* = -\frac{1}{W^2}$
- $V_S^* = f_S$
- $V_{SS}^* = f_{SS}$
- $V_\tau^* = f_\tau$
- $V_{SW}^* = 0$.

A sufficient condition for optimality of the portfolio policy is the concavity of the value function with respect to the state variable W (Boguslavsky and Boguslavskaya (2004)). From the calculations above, we see that the sign of V_{WW} and the fact that wealth is non-negative ensure that this ansatz, if successful in the sense below, would yield the optimal portfolio

strategy. Substituting the relevant partial derivatives into the optimal portfolio policy in (20), we have:

$$(22) \qquad N(S, W, \tau) = \left(\frac{-\frac{k}{c} \tanh\left(c(S - \overline{S})\right) - r(S - \overline{S})}{\sigma^2} \right) W,$$

and substituting (22) together with the relevant partial derivatives into the HJB equation (18) yields a simplification of the HJB equation, namely:

$$
\begin{aligned}
0 =& \left(-\frac{k}{c} \tanh\left(c(S - \overline{S})\right) \right) f_S - f_\tau + \frac{1}{2}\sigma^2 f_{SS} - \frac{\left(\frac{k}{c} \tanh(c(S - \overline{S})) - r(S - \overline{S}) \right)^2}{2\sigma^2} \\
& + r \left(1 - \frac{\left(-\frac{k}{c} \tanh(c(S - \overline{S})) - r(S - \overline{S}) \right)(S - \overline{S})}{\sigma^2} \right) \\
& + \frac{k \tanh\left(c\left(S - \overline{S}\right)\right) \left(-\frac{k}{c} \tanh\left(c\left(S - \overline{S}\right)\right) - r\left(S - \overline{S}\right) \right)}{\sigma^2 c}.
\end{aligned}
$$

(23)

Thus, we have successfully eliminated W from (23), meaning the HJB equation is reduced in dimensionality, and now only depends on S and τ. Moreover, the optimal portfolio policy in (22) is fully characterized and stripped of its dependence on the value function or any of its derivatives. Since we are only interested in the optimal portfolio strategy, the optimization procedure is complete. We need only to make the technical assumption that a smooth function $f(S, \tau)$ twice differentiable in S, once in τ, exists and satisfies (23).[25] This assumption is used in Xiong's (2001) general equilibrium model.

This completes the proof of Theorem 1.

$\square$

---

[25]The explicit solution $f(S, \tau)$ would allow us to characterize an optimal portfolio policy for a general CRRA arbitrageur. See e.g. Jurek and Yang (2007) and Boguslavsky and Boguslavskaya (2004).