HP Mashele, SE Terblanche and JH Venter

Pairs trading on the Johannesburg Stock Exchange

ABSTRACT

Pairs trading strategies aim to profit from temporary deviations in some underlying relationship between the prices of two stocks. The trader takes appropriate long and short positions in the two stocks and waits for their prices to revert back to the underlying relationship or even to deviate in the opposite direction from the current deviation, at which time the trader may exit at a profit. We formulate formal trading rules that implement pairs trading strategies and discuss their profitability and risk by means of back-testing on stocks listed on the Johannesburg Stock Exchange (JSE).

1. INTRODUCTION

Pairs trading can be described as follows. Two different stocks may be exposed to similar underlying market conditions which cause a tendency in their prices to move together. From time to time short term disturbances may influence the individual stocks differently and cause temporary deviations from the longer term underlying co-movement tendency. If this happens consistently enough profitable trading in the two stocks may be done by entering a long position in the stock whose price is lower and a short position in the stock whose price is higher during a temporary deviation from the underlying long term tendency. When the short term disturbance ends and the prices revert to the underlying tendency, the previously higher stock may be priced lower, the previously lower stock may be priced higher and the entered positions may be exited at a profit.

Of course there are no certainties in these events materializing exactly as described. What appears to the trader as a short term disturbance may eventually turn out to cause a more permanent shift in the underlying relation between the stocks in which case the envisaged profit on exit may not arrive. Pairs trading strategies therefore deal with ‘statistical arbitrage’ opportunities in the sense that the rewards from these opportunities carry considerable risk. Pairs trading strategies can also be described as ‘contrarian’ in the sense that the trader believes the disturbance that caused the relative prices to deviate from the underlying norm will not continue and therefore takes a position contrary to the current market movements. They are ‘reversal’ strategies in the sense that the trader expects a reversal back to the underlying norm and enters positions accordingly.

The influential paper of Gatev, Goetzmann and Rouwenhorst (2006) reviewed the background and history of pairs trading and studied its profitability in depth. Their implementation involved a pairs formation period of twelve months followed by a trading period of six months. From their universe of US stocks, pairs were formed that had the best co-movement tendency over the formation period in terms of the sum of squared deviations between their normalized price series. During the trading period positions were taken when price deviations exceed more than two historical standard deviations. Among many results they report that “the profits typically exceed conservative transaction-cost estimates”. Their results and the further work reported by Andrade, Di Pietro and Seasholes (2005) and Engelberg, Gao and Jagannathan (2009) were reviewed by Do and Faff (2010) who then updated and refined them. Do and Faff (2010) particularly emphasized that returns can be enhanced by not relying solely on pairs of shares obtained by matching statistical properties of their past price series, but by using pairs formed from “fundamentally homogeneous asset groups”. Do and Faff (2012) take commissions, market impact and short selling fees into account and still report profitability when using “well matched pairs that are formed within refined industry groups”. Perlin (2007) finds profitability of pairs trading strategies on Brazilian markets, Broussard and Vaihekoski (2010) find the same for Finland, as do Hong and Susmel (2004) for the Asian ADR market and Bolgun, Kurun and Guven (2007) for the Istanbul stock exchange. Statistical models that may be used to describe the co-movement aspect of the price series of suitable pairs can be based on Ornstein-Uhlenbeck processes (Elliot, van der Hoek and Malcolm (2005) and Mudchanatongsuk, Primbs and Wong (2008)) and cointegration principles (Lin, McCrae and Gulati (2006) and Vidyamurthy (2004)). More extensive modeling along similar lines is discussed in Do, Faff and Hamza (2006).

Our paper contributes to the pairs trading literature in several ways. Firstly, we do not follow the traditional approach using nearest co-movement pair selection from a twelve months formation period followed by a six months trading period. As suggested by Do and Faff (2010, 2011) we simply select pairs from stocks in economically meaningful groups but then develop trading rules that dynamically decide to trade if suitable, suspend trading when no longer suitable and resume when deemed suitable again. Secondly, when evaluating pairs trading rules we do not follow the returns approach popularized by Gatev, Goetzmann

[*]Centre for Business Mathematics and Informatics, North West University, Private Bag X6001, Potchefstroom 2520, Republic of South Africa.
Email: hennie.venter@nwu.ac.za

and Rouwenhorst (2006) because the concept of 'return' is problematical when dealing with self-financing portfolios which require zero capital investment (see e.g. footnote 6 of Broussard and Vaihekoski (2010)). We introduce the notions of 'average cash flow per day' as a measure of reward and base a measure of risk on the 'average clean value per day'. These measures are used to judge the quality of trading rules and to choose the parameter values needed for practical implementation. Thirdly, our treatment of the effects of costs on both profitability and risk differs substantially from the existing literature. Rather than trying to estimate the returns using rules which ignore costs and then judging whether the returns ultimately exceed costs, we incorporate costs from the beginning into all equations. Fourthly, the returns evaluation approach in the existing literature has the side effect that the implicit trade pricing assumptions used in back-testing may be unrealistic but is not subject to scrutiny. In our approach the pricing assumptions used for trading evaluation are specified explicitly and this makes it possible to quantify the effects of trader execution skill which turns out to be very important for successful pairs trading. Fifthly, our contribution demonstrates that pairs trading can be successful for stocks listed on the JSE.

The rest of the paper is organized as follows. In Section 2 we derive formulas for the numbers of shares involved when entering self-financing long/short positions in two stocks taking proportional trading cost into account. We point out that the events under which such positions can be exited profitably depend on the development over time of the log-price difference of the two stocks and indicate a trading strategy that may be used to take advantage of such events. In Section 3 we formulate a quantitative trading rule that implements this strategy in the case where the log-price difference process forms a mean reverting AR(1)-process. For practical application the parameters of such a process must be estimated from a sliding window of past data. Price data of the large retail banks listed on the JSE are used to illustrate and motivate the development of our approach. Section 4 extends the rule of Section 3 by limiting the clean value of the pair positions, which enables management of a pairs trading account to control risk exposure. Section 5 is devoted to the important role of pricing assumptions in back-testing of trading rules. Ten possible pricing scenarios are introduced and they illustrate that the results are strongly affected by trader skill and timing. Trading cost is another important determinant of success with pairs trading and this is dealt with in Section 6. The rules of Sections 3 – 6 can be modified in many intuitively reasonable ways and Section 7 discusses and evaluates some of these possibilities. Section 8 gives a brief discussion on the application of the rules to other data sets, the time series properties of cash flows generated by the rules and issues arising from the efficient market hypothesis. Section 9 concludes with a summary and an indication of items requiring further research.

2. SELF-FINANCING LONG/SHORT PAIRED POSITIONS

Consider two stocks A and B and suppose their prices on day t are a_t and b_t respectively. To create a self-financing 'short B long A' (SBLA) position we sell n_B shares of B and we buy n_A shares of A with the numbers n_B and n_A chosen such that no cash investment is needed. Assume that the sizes of the two trades are S rand each and that proportional costs of c_s and c_b for selling and buying respectively are involved in the two trades. Then selling n_B shares of B yields an inflow of $S = n_B b_t(1-c_s)$ and buying n_A shares of A causes an outflow of $S = n_A a_t(1+c_b)$. Thus

$$n_B = \frac{S}{b_t(1-c_s)} \quad \text{and} \quad n_A = \frac{S}{a_t(1+c_b)} \quad . \qquad \ldots (1)$$

Note that (2.1) may not yield integral numbers of stocks and some rounding may be required in practice. By taking the trade size S large the effects of rounding can be made small and it is convenient to ignore this aspect in the developments below. At some later day T when the prices of the stocks are a_T and b_T respectively, exiting of the SBLA position entails buying back the n_B shares of B causing an outflow of $n_B b_T(1+c_b)$ and selling the n_A shares of A yielding an inflow of $n_A a_T(1-c_s)$. Thus the 'clean value on day T' (mark-to-market but taking into account costs of termination) of this SBLA position is

$$\begin{aligned} CV_T &= -n_B b_T(1+c_b) + n_A a_T(1-c_s) \\ &= -\frac{S b_T(1+c_b)}{b_t(1-c_s)} + \frac{S a_T(1-c_s)}{a_t(1+c_b)} \qquad \ldots (2) \\ &= \frac{S b_T(1-c_S)}{b_t(1+c_b)}\left[\frac{r_T}{r_t} - \frac{(1+c_b)^2}{(1-c_s)^2}\right] \end{aligned}$$

where $r_t = a_t/b_t$ is the price ratio of the two stocks on day t. This clean value (CV) is positive if we have

$$\frac{r_T}{r_t} > \frac{(1+c_b)^2}{(1-c_s)^2} \quad . \qquad \ldots (3)$$

Typically c_s and c_b are small numbers in which case $\frac{(1+c_b)^2}{(1-c_s)^2} \cong 1 + 2(c_s + c_b)$ and it follows that the SBLA position initiated on day t can be exited at a profit on day T only if the price ratio of the two stocks increased by a factor of about $2(c_s + c_b)$. The condition for a profit on exit depends only on the price ratio of the two stocks and it is natural to expect that the factor $2(c_s + c_b)$ should enter the equation since two sells and two buys are involved in the complete round trip trade. Introduce the notation

$l_t = \log(r_t) = \log(a_t) - \log(b_t)$ ... (4)

for the log-price difference (LPD) between shares A and B on day t. Then the condition for a profit becomes $l_T > l_t + \log\left[\frac{(1+c_b)^2}{(1-c_s)^2}\right] \cong l_t + 2(c_s + c_b)$, i.e. the LPD should increase by about $2(c_s + c_b)$. Clearly our strategy should be to enter a SBLA position if we think that the LPD between A and B is currently **too low** and expect it to increase sufficiently in future to enable us to exit the SBLA position at a profit.

By interchanging symbols similar expressions and results can be written down for a self-financing short A long B (SALB) position. In this case termination is profitable if $l_T < l_t - \log\left[\frac{(1+c_b)^2}{(1-c_s)^2}\right] \cong l_t - 2(c_s + c_b)$ and our strategy should be to enter such a position on day t if we think that the LPD between A and B is **too high** and will decrease sufficiently so that we can exit it profitably on some future day.

3. AR(1)-MODELS AND PAIRS TRADING RULE 1

The pairs trading strategy should be exploitable if the LPDs l_t form a mean reverting process that tends to fluctuate around some mean level μ over time. When the LPD is far below (above) μ at some point in time we could enter a SBLA (SALB) position and then the mean reversion property will cause the subsequent LPD values to increase (decrease) making it profitable to exit when this happens.

The simplest mean reverting process is provided by the AR(1)-model. This model is discussed in most textbooks on time series (see e.g. Tsay (2010) p38) and we mention briefly a few facts needed for our purposes. If the LPD process l_t follows an AR(1)-model then

$l_t = \mu + \varphi(l_{t-1} - \mu) + \sigma e_t$ … (5)

where μ is the mean reversion level, φ is the AR-coefficient, the e_t's are the innovations, assumed to have expectation 0 and variance 1, and σ is the innovation standard deviation. If the LPDs of a given pair of shares follow such a model then intuitive terms such as 'far below (or above) μ' can be quantified more exactly, making it possible to formulate the pairs strategy into operational trading rules.

Assuming that $|\varphi| < 1$, the AR(1)–process has a stationary distribution with expectation μ and standard deviation $\sigma' = \sigma/\sqrt{(1-\varphi^2)}$. This standard deviation gives us a measure in terms of which distances of l_t from μ can be judged. The event that l_t is far below μ can be taken to mean that $l_t < \mu - d\sigma'$ with $d \geq 0$ a suitably chosen multiple. Alternatively, defining the B-factor

$B_t = 100(l_t - \mu + 2\sigma')/4\sigma'$, … (6)

the event $l_t < \mu - d\sigma'$ is equivalent to $B_t < 50 - 25d = B^*$ with B^* a suitably chosen percentage threshold factor. Similarly the event that l_t is far above μ is taken to mean that $B_t > 50 + 25d = 100 - B^*$, assuming that the same multiple d is used for the down and up sides. This B-factor is obviously inspired by the so-called 'percent B' indicator of technical analysis associated with Bollinger bands (see e.g. http://en.wikipedia.org/wiki/Bollinger_Bands) but here it is defined in terms of the parameters of the AR(1)-model.

Estimates of the parameters involved in these expressions will be needed for practical application. To do this on any given day t we fit the AR(1)-model to the LPD values of the last w days with w a suitably chosen window length. We used PROC AUTOREG of SAS for this purpose in the results reported in this paper. Assuming that the model fits well and will continue to describe the future development of the LPD process the trading strategy can now be quantified into the following basic trading rule.

Rule 1

1. Choose a window length w and a threshold $B^* < 50$, assume that we have no position to start with and track the LPD series of the two stocks A and B over time for at least w days.

2. On any day t fit an AR(1)-model to the LPDs of the last w days and compute B_t by (3.2).

3. If $B_t < B^*$ then we have a 'low signal' and enter a SBLA position. If $B_t > 100 - B^*$ then we have a high signal and enter a SALB position. If $B^* \leq B_t \leq 100 - B^*$ then we do nothing but continue to the next day and apply the same rules.

4. Once we do have a SBLA (or SALB) position we continue tracking the LPDs doing nothing until the first day on which we get a subsequent opposite high (low) signal. On that day we terminate the SBLA (SALB) position and enter into a new SALB (SBLA) position.

More will be said on the choice of the parameters w and B^* below. Note that we employ a sliding window of past LPDs to fit the AR(1)-model and estimate the parameters. Therefore we are not assuming that one fixed AR(1)-model describes the development of the LPD series but allow the estimates to change over time. We illustrate some features of the rule in Figure 1. In this paper we mostly use data of the four banks Absa (ASA), First Rand (FSR), Nedbank (NED) and Standard Bank (SBK) for this purpose. They all operate in the same economic sector and one could reasonably expect their prices to move together at least to some extent and this makes them candidates for pairs trading strategies. Their price data were extracted from the intraday database (IDDB) of the BMI-center of North-West University and covers the 1764 trading days from 11 January 2005 (the first day for which the IDDB provides usable intraday price data for our purpose) to 31 January 2012.

Figure 1 illustrates Rule 1 in action for the pair ASA/FSR using **daily closing prices**. We chose the window length $w = 20$, the threshold $B^* = 35$, the trade size $S = 10000$ rand and both trading costs $c_s = c_b = 0{,}002$ (i.e. 20 basis points (bps)). More will be said about the choice of parameters as we proceed below. The top panel of Figure 1 shows an extract of the B_t's and the thresholds over days 1000 to 1100 (about January to June 2009). A SALB (here meaning 'short ASA, long FSR') position was entered on day 1003 and exited and replaced by a SBLA (here meaning 'short FSR, long ASA') position on day 1012. The middle panel of Figure 1 shows that the LPD value declined by more than 0,1 from day 1003 to day 1012. This is more than the round trip cost of about 4x0,002=0,008 so that a profit occurred on day 1012 and this is shown by the positive cash flow on day 1012 in the bottom panel of Figure 1. The SBLA position of day 1012 was exited on day 1027; since the LPD rose by about 0,08 this again yielded a profit shown as a positive cash flow in the bottom panel on day 1027. A SALB position was entered on day 1027 but this was exited much later on day 1073. The top panel shows that the LPD increased over this period so that this exit was done at a loss shown as a negative cash flow in the bottom panel. Three further trades were done and they all went according to plan and ended up with positive cash flows.

The reason for the one failure is clear from the middle panel of Figure 1. The mean reversion level μ was quite stable up to day 1040 and also onwards from about day 1060 but between these two dates it shifted upwards. This shift in reversion level prevented the anticipated decrease in LPD from occurring when entering the SALB position on day 1027 and resulted in the exit on day 1073 yielding a loss. Clearly sudden shifts in the mean reversion level may adversely affect the operation of Rule 1.

In addition to the cash flows at the times of entering and exiting positions, the bottom panel of Figure 1 also shows the CV of the position on each day just **before** trading if any trading is done on that day. By definition this CV will equal the cash flow when exiting a position. Entering a new position requires no cash flow but the CV immediately afterwards will be negative since the additional trading cost required to exit the new position is taken into account. Notice that during the period over which the mean reversion level shifted the CV tended to become large negative. Large negative values of the CV signal that the current position is not developing according to plan and should possibly be terminated. An extension of Rule 1 in the next section will do this but first some further results for Rule 1 are reported.

With four shares a portfolio of six pairs can be formed to trade in concurrently. Table 1 shows summary statistics for the individual pairs and the total portfolio when trading over the whole period by Rule 1. Again daily closing prices and the same choices of parameters as above were made and all existing positions on the last day of trading were exited.

The first column of Table 1 names the pair for each row with the last row listing the results for the portfolio of all six pairs. The second column shows the average cash flow per day (ACFPD, i.e. the total of all cash flows divided by the total number of days available for trading (i.e. $1764 - w$)). The ACFPD is the main **measure of reward** to be associated with a trading rule from here onwards. All pairs contributed positively with ASA/FSR highest and ASA/NED lowest. The third column shows the main measure of risk used in this paper, namely the average negative clean value per day (ANCVPD). If we were forced to close our positions on a particular day we would lose the clean value of the positions on that day if it was negative. Therefore the ANCVPD reflects the amount we would expect to lose if forced to close our positions on a randomly chosen day. The fourth column shows a more extreme measure of risk, namely the minimum (most negative) clean value (MCV) experienced over the period. Note that the total over pairs of these risks will not add up to the risk shown for all pairs due to off-setting between pairs, i.e. there is some risk diversification in this respect. The fifth and sixth columns show the numbers of positive and negative cash flows obtained while the last two columns show their average sizes. The average sizes of the negative cash flows tend to be larger than those of the positive cash flows but the greater abundance of the positive cash flows amply compensates and the overall impression is that a trader who consistently managed to trade according to Rule 1 would have done quite well over the period. Note that choosing the trade size $S = 10000$ rand is only a matter of scale. With a 10 time larger (or smaller) value all the cash flows and clean values in Table 1 would change by the same factor.

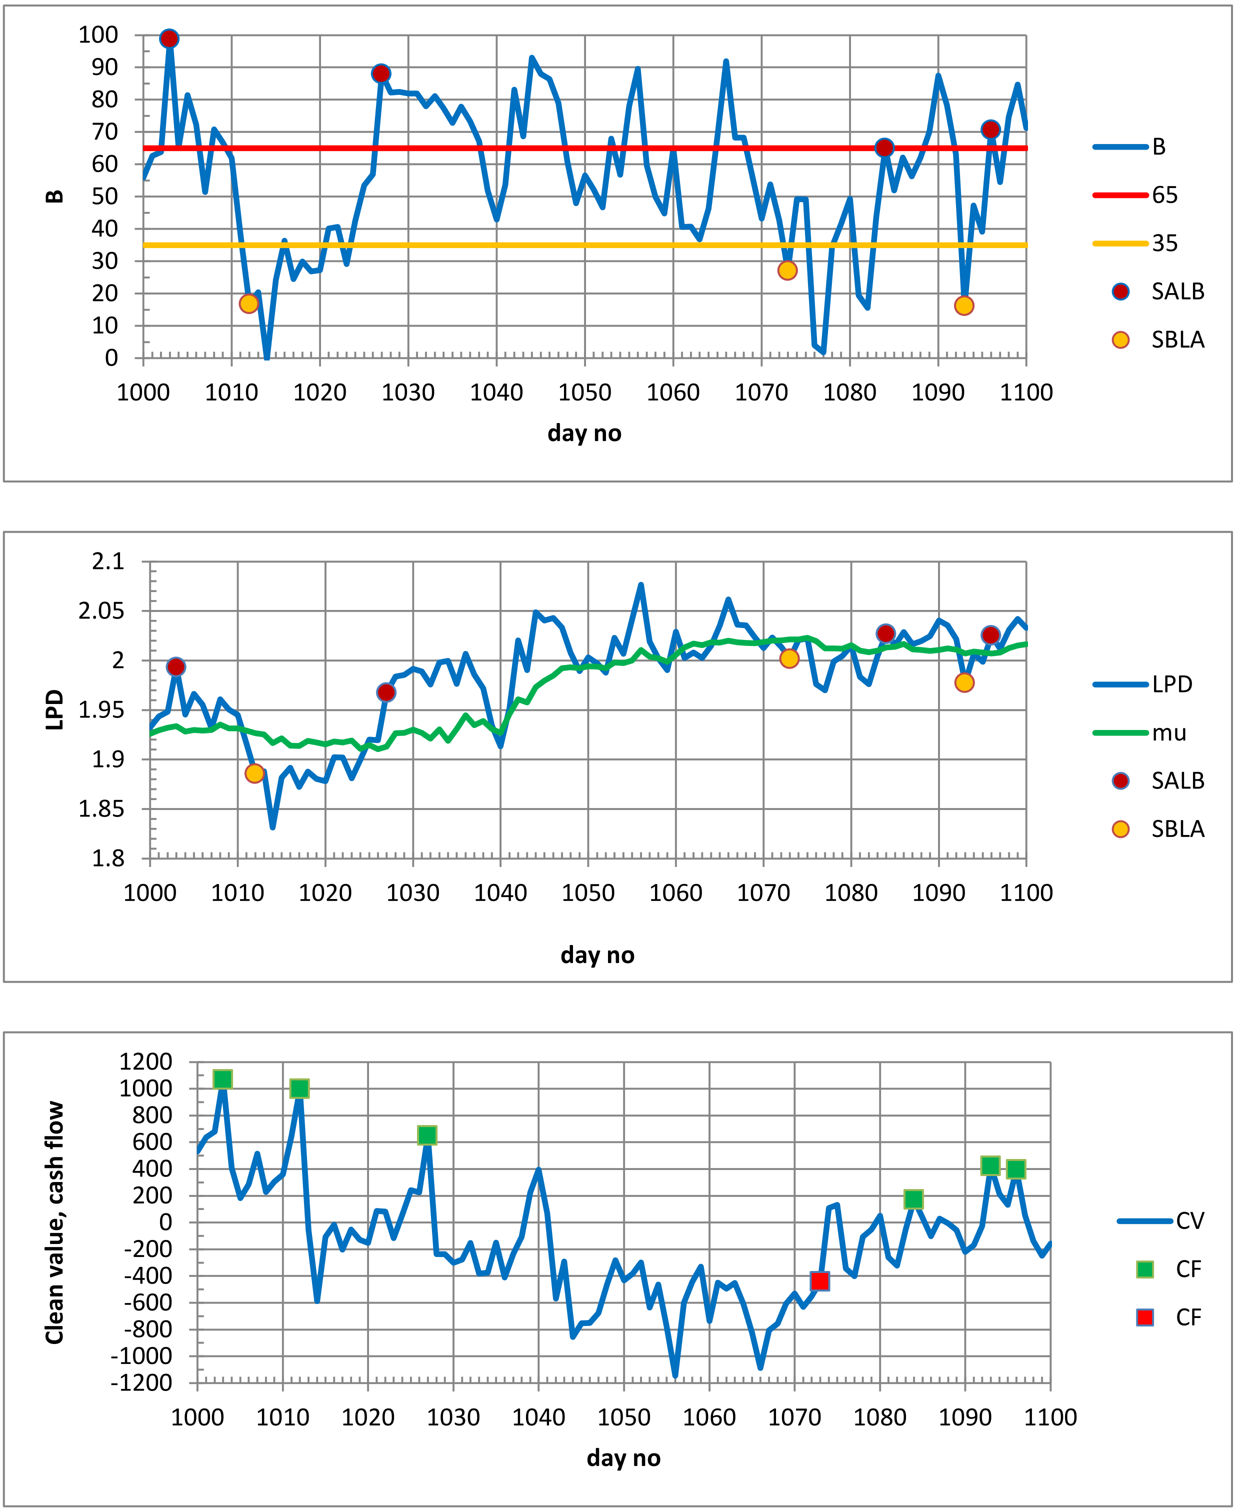

Figure 1: Top panel: B factor and positions entered by Rule 1 over days 1000-1100 for pairs trading in ASA/FSR using closing prices. Middle panel: LPD values, mean reversion estimates and times of positions entered and exited. Bottom panel: Clean values and cash flows resulting from positions.

Table 1: Summary statistics for Rule 1 pairs trading in ASA, FSR, NED and SBK over the period Jan 2005 to Jan 2012 using daily closing prices with $\mathbf{w = 20}$, $\mathbf{B^* = 35}$, $\boldsymbol{S = 10000}$ and costs of 20 bps.

Pair	Average cash flow per day	Average neg clean value per day	Minimum clean value	Number of positive cash flows	Number of negative cash flows	Average positive cash flow	Average negative cash flow
ASA/FSR	15,76	-220,38	-1232,88	129	29	268,37	-245,94
ASA/NED	5,69	-385,48	-3143,21	84	33	288,26	-432,87
ASA/SBK	9,58	-221,72	-1843,80	119	33	204,45	-230,75
FSR/NED	12,35	-299,06	-2715,86	116	33	276,36	-318,84
FSR/SBK	15,07	-259,30	-2370,09	138	26	254,22	-338,60
NED/SBK	8,56	-284,27	-2515,45	108	23	231,03	-435,80
ALL	67,01	-1559,93	-8122,11	517	137	325,56	-375,47

A glance at Figure 2 helps to explain the relative merits of trading in the different pairs. Figure 2 shows the graphs (shifted for better visibility) of the LPDs of the six pairs over the whole time period. It is clear that the pairs can be ranked in terms of near stability of their LPDs over time in the same order as their ranking in terms of ACFPD in Table 1. Evidently ASA/NED was not a good pair to trade in since co-movement of their prices was particularly poor. It is remarkable that Rule 1 did manage to get a positive ACFPD from trading in this pair but it is clear that this was accompanied by the most risk.

4. CONTROLLING RISK: RULE 2

Here we return to the issue of refining Rule 1 in an effort to make it more robust against the effects of possible shifts in the mean reversion level. It was noted above that a level shift may cause building up of a large negative CV, implying growing risk inherent in the currently held position. If the underlying reason for the LPD to have the mean reversion property ends permanently, then the opportunity may never arise to exit the position when relying solely on the stipulations of Rule 1 and the trader may be stuck with a position with ever larger negative value. This suggests that we put a threshold on how negative the CV is allowed to become and to exit the position when this threshold is exceeded. Such an event happening also suggests that trading in the pair should be suspended and that we must wait until a new stable enough reversion level is achieved before resuming trading. Therefore some signal is needed to indicate when trading may be resumed. The following extension of Rule 1 formulates these considerations.

Rule 2

In addition to the stipulations of Rule 1, choose a threshold $\mathrm{CV}^* < 0$. On each day calculate the CV due to the current position and if $\mathrm{CV} < \mathrm{CV}^*$ then exit this position. If the exited position was SBLA (SALB) suspend entering new SBLA (SALB) positions on further low (high) signals occurring until such time as the B-factor first crosses the level of $100 - \mathrm{B}^*$ upwardly (B^* downwardly) at which time trading resumes.

From a management point of view this rule makes good sense: the maximum loss that can be experienced for whatever reason is now limited when trading the pair. The top panel of Figure 1 illustrates the reasoning behind the resumption signal for the case of an upward shift in reversion level: during the transition the B-factors tended to remain high and it was only after the first downward crossing of the threshold B^* that 'normal' up and down fluctuations returned which made further trading feasible.

Table 2 shows the same summary statistics for Rule 2 as was shown for Rule 1 in Table 1. The same parameters were chosen and the CV threshold was set at $\mathrm{CV}^* = -0{,}1S = -1000$. Comparing the third and fourth columns of the two tables, it is clear that there were substantial reductions in risk as measured by both the ANCVPD and the MCV. Note in particular that the MCVs of the individual pairs were controlled to be above $\mathrm{CV}^* = -1000$; that the ANCVPD of the whole portfolio was also close to -1000 was mere coincidence. The number of negative cash flows and their average sizes became larger than those of Table 1 while the number of positive cash flows and their sizes changed little. This decreased the ACFPD which is an indication of the price that was paid for controlling risk as indicated by the lower ANCVPD and MCV.

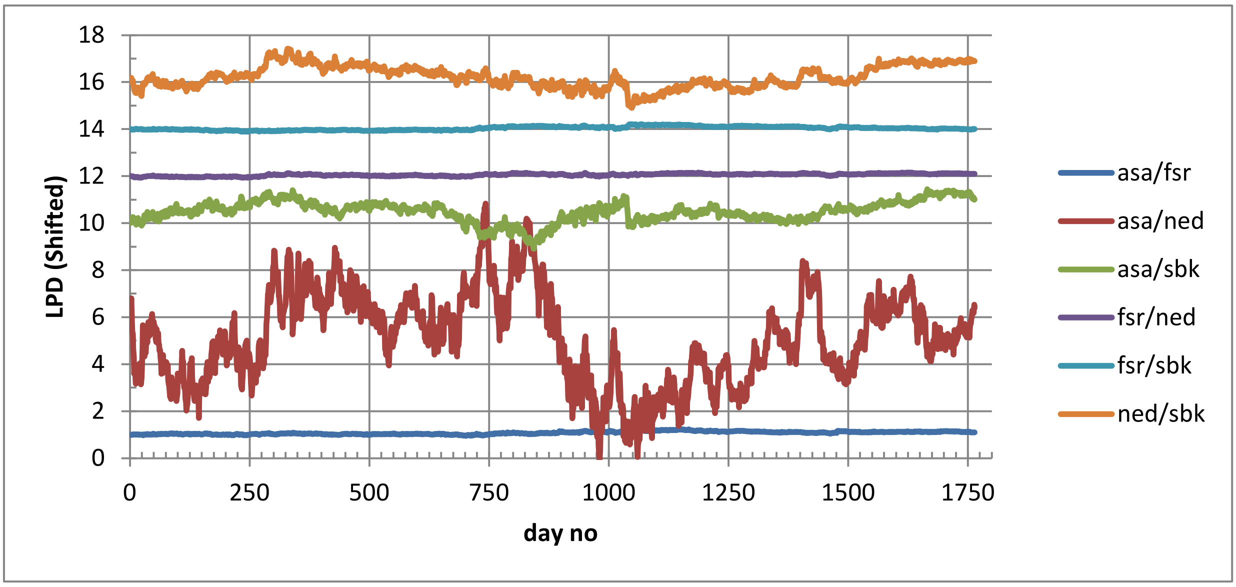

Figure 2. Graphs of shifted log-price differences (LPDs) over time of the six pairs of banks.

Table 2:Summary statistics for Rule 2 pairs trading in ASA, FSR, NED and SBK over the period Jan 2005 to Jan 2012 using daily closing prices with $\mathbf{w = 20, B^* = 35, CV^* = -1000, S = 10000}$ and costs of 20 bps.

Pair	Average cash flow per day	Average neg clean value per day	Minimum clean value	Number of positive cash flows	Number of negative cash flows	Average positive cash flow	Average negative cash flow
ASA/FSR	14,44	-192,09	-951,37	129	29	268,37	-325,43
ASA/NED	3,30	-209,32	-986,80	84	33	288,26	-559,26
ASA/SBK	7,63	-178,03	-971,12	119	33	204,45	-334,26
FSR/NED	7,15	-187,65	-993,64	114	35	278,29	-549,97
FSR/SBK	13,79	-174,72	-994,62	137	27	255,02	-403,49
NED/SBK	7,71	-176,60	-971,43	108	23	231,03	-499,89
ALL	54,02	-1013,21	-3265,87	517	142	325,35	-521,08

Next we look at other possible choices of the three parameters w, B^* and CV^* that control the performance of Rule 2. Figure 3 shows a scatter plot in the risk-reward plane with risk taken as ANCVPD and reward as ACFPD when varying the three parameters. The window length w was varied over the values 15 – 60 by increments of 5, the B-threshold B^* was varied from 0 to 50 by increments of 1 and the CV threshold CV^* was varied from -25 to -10000 with growing increments to cover this range in 30 steps. At each combination of parameter values Rule 2 was run over the whole period and the resulting ANCVPD and ACFPD values were calculated for the portfolio of all pairs using closing prices, trading costs of 20 basis points and trade size $S = 10000$ as before. Figure 3 shows that the parameter values used in Table 1 (effectively Rule 2 with $CV^* = -10000$) yield an ACFPD close the overall optimum achievable. Also the parameter values in Table 2 for Rule 2 clearly limits the ANCVPD risk above -1000 while still yielding an ACFPD that is about as high as can be achieved under this risk constraint, i.e. this choice is close to the efficient frontier attainable under the modeling used here. These results justify using the parameter choices in the illustrations above.

The same parameter choices were used for each of the pairs above. Of course we can study optimal parameter choices for each pair separately and this may improve profitability and risk control. Instead of pursuing this possibility we next attend to the practically more important issue of using prices other than closing prices.

5. TRADE PRICING ASSUMPTIONS

All the illustrations so far were based on daily closing prices both for modeling to estimate LPD deviations and for trade implementation of decisions. This is a convenient starting point since closing prices are widely available, but from an actual trading point of view this is not a realistic assumption. It is simply not possible to wait until the closing prices of the day are available, do the required modeling to judge the extent of LPD deviations from the reversion levels and then still be able to enter positions depending on the results of the analysis. In order to make back-testing of trading rules on historical share price data more realistic, we must distinguish between the prices on which modeling of the LPD deviations are based and the prices at which trading is assumed to take place. In this section we study the implications of a variety of more realistic trade pricing assumptions.

The first group of assumptions still do the model fitting on daily closing prices as above but then implement trading decisions on the **next** trading day according to one of the following assumptions:

1. at the opening price of each stock

2. at the closing price of each stock

3. at the median price of all trades on that day for each stock

4. if a stock is bought (sold) then the price is taken as the 40% (60%)-percentile of all trades on that day for each stock

5. if a stock is bought (sold) then the price is taken as the 60% (40%)-percentile of all trades on that day.

Assumptions 1 and 2 make it possible to see what the effects are of time elapse between estimating the sizes of deviations in LPD from the reversion levels and trading implementation. Assumption 3 may be interpreted as reflecting that the trader's skill level is on a par with his peers, while assumption 4 (5) implies that the trader is 10 percentage points more (less) skilled than his peers, both when buying and selling. Assumptions 3, 4 and 5 allow us to judge the effect of the trader's skill level on pairs trading results. Instead of using the median prices in assumption 3, the volume weighted average prices could be used to reflect the 'average trader', but we have found that this makes little difference to the results and therefore continue with the median prices below.

The second group of assumptions do the model fitting at the ruling prices at a particular time during the trading day (referred to as the **revision time**) and the implementation during the rest of the **same** trading day according to one of the following assumptions:

6. at the ruling price of each stock at revision time

7. at the closing price of each stock on that day

8. at the median price of all trades over the rest of the trading day for each stock

9. if a stock is bought (sold) then the price is taken as the 40% (60%)-percentile of all trades over the rest of that day

10. if a stock is bought (sold) then the price is taken as the 60% (40%)-percentile of all trades over the rest of that day.

Assumption 6 may also be unrealistic but is included for comparison purposes, while assumptions 8 - 10 again allow judging the effects of the trader's skill level but now in the intraday trading context.

Of course AR(1)-model fitting still requires a choice of window length w and implementing Rule 2 also requires choices for the parameters B^* and CV^*. To save space we report the next results only for $w = 20$, $B^* = 35$, $CV^* = -1000$ which delivered results close to the efficient frontier in the previous context and was found to continue to do so in the present context as well. The trade size is again at $S = 10000$ and the trading costs at 20 bps as before.

As a first example of the effects of pricing assumptions, Table 3 shows the results for the individual pairs and the total portfolio when trading under assumption 3. Thus the only difference between Tables 2 and 3 is that trade executions are at closing prices for Table 2 and at the next day's median prices for Table 3. Comparing these two tables it is clear that the reward as measured by ACFPD dropped a great deal individually for each pair and also in total. However, the risk as measured by both ANCVPD and MCV remained similar. Although the sizes of the positive and negative cash flows tended to be smaller, the numbers of positive (negative) cash flows decreased (increased) and this caused the large drop in overall profitability. Thus it is evident that pricing assumptions have important implications for pairs trading.

Table 4 shows the results for all the trade pricing assumptions listed above as well as when the revision time is varied over the values 11, 13 and 15 hours in assumptions 6 - 10. Keeping space limitations in mind we show results only for the total portfolio. The results under assumptions 1 (open prices) and 3 (median prices) are quite similar but under assumption 2 (closing prices) the ACFPD is much lower although this is partially compensated for by lower risk. Under the higher trader skill assumption 4 ACFPD is dramatically higher and risk is lower, while under the lower trader skill assumption 5 the ACFPD drops to a negative value and more risk is experienced as well. These effects are also reflected in the behavior of the numbers and sizes of the cash flows. It is now evident that the level of trader skill is a major determinant of successful pairs trading.

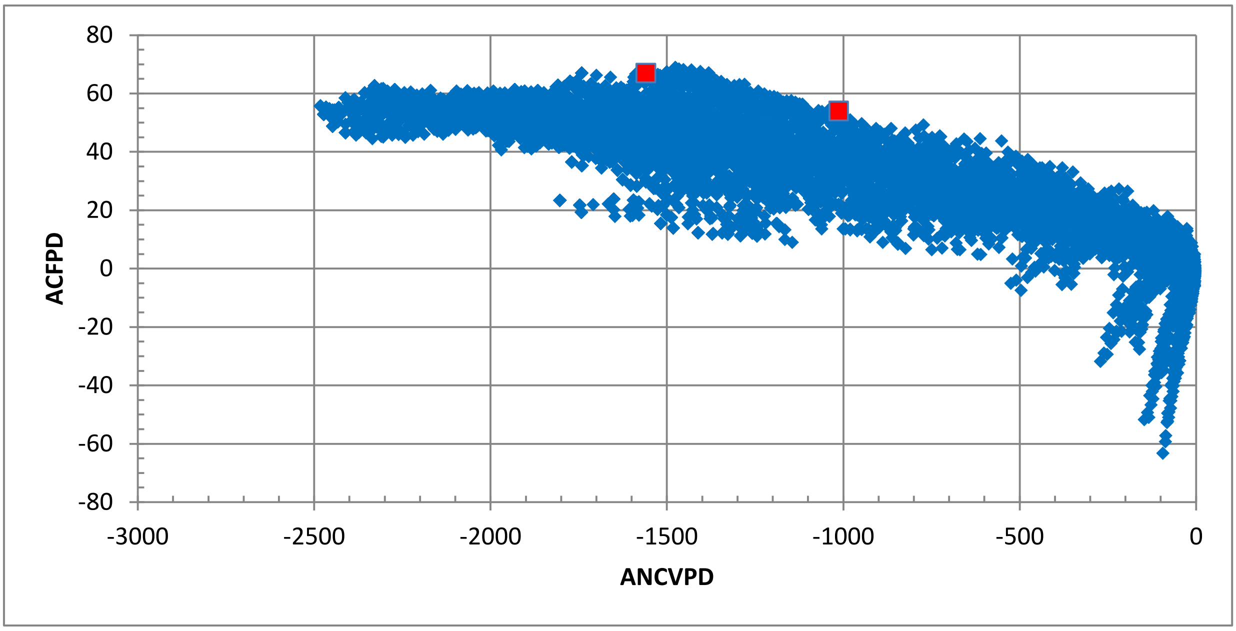

Figure 3: Scatter plot of average cash flow per day (ACFPD) against average negative clean value per day (ANCVPD) for Rule 2 when varying the window length w, B-threshold B^* and CV threshold CV^*. The parameter choices used in Tables 1 and 2 give ANCVPD and ACFPD values located at the red dots.

Table 3: Summary statistics for Rule 2 pairs trading in ASA, FSR, NED and SBK over the period Jan 2005 to Jan 2012 using daily closing prices for modeling and next day median prices (assumption 3) for implementation with $w = 20$, $B^* = 35$, $CV^* = -1000$, $S = 10000$ and costs of 20 bps.

Pair	Average cash flow per day	Average neg clean value per day	Minimum clean value	Number of positive cash flows	Number of negative cash flows	Average positive cash flow	Average negative cash flow
ASA/FSR	8,27	-195,91	-981,85	114	45	230,59	-263,80
ASA/NED	0,57	-194,73	-990,57	68	49	304,95	-402,96
ASA/SBK	2,92	-186,30	-989,18	98	55	177,99	-224,63
FSR/NED	2,78	-185,33	-985,16	106	44	233,97	-453,35
FSR/SBK	6,62	-171,69	-973,46	116	48	218,33	-287,21
NED/SBK	0,96	-177,31	-966,92	87	45	217,64	-383,71
ALL	22,11	-987,21	-3396,26	433	237	291,88	-370,58

The results with revision time at 11 hours and pricing assumptions 6 - 10 largely mirror those of assumptions 1 - 5. In particular the prominent effects of trader skill level are again clear. The same holds for the results of revision times 13 and 15 hours. The effects of the revision time on the results differ under the various pricing assumptions. Under assumption 6 (ruling prices) there are little differences, but this assumption is rather unrealistic. Under assumption 7 (closing prices) the later revision time appears to do better and this may be due to this set-up getting closer to both revision time and pricing approaching the closing time and prices set-up of the previous section. Under assumption 8 (median prices) the results are again similar. Under assumption 9 (higher skill) the ACFPD drops slightly with later revision time while risk increases; it is likely that this effect can be attributed to the trader having less time to exercise his superior skill as the revision time becomes later in the day. Under assumption 10 (lower skill) there is slight improvement as the revision time increases, maybe attributable to the inferiorly skilled trader having less time to make mistakes.

Overall the results tend to conform quantitatively with intuitive expectations.

6. TRADING COST EFFECTS

In all the illustrations above the trading costs were kept at 20 bps for both buying and selling. Since each round trip of entering and exiting pair positions involve four trades it is intuitively clear that trading costs will have important effects on profitability and risk. We computed cost effects in a variety of settings but report only for trading under assumptions 6 - 10 with revision time at 15 hours which demonstrates these effects adequately. Table 5 shows the results when the costs were varied over the values 40, 30, 20 and 10 bps. Both Rules 1 and 2 were used with the choice of parameters $w = 20$, $B^* = 35$ and $CV^* = -1000$. To save space we do not show the numbers and sizes of the cash flows; they generally follow the impressions visible from the ACFPD and risk measures as in the cases of Tables 3 and 4.

Columns 3 - 5 of Table 5 show the rewards and risks for Rule 1. This rule does not control risk and both the ACFPD and the risks improved with diminishing trading costs under all trade pricing assumptions. At the high costs end of 40 bps the ACFPD was low to negative with the superior (inferior) skill assumption 9 (10) producing the best (worst) results as before. At the low costs end of 10 bps the results were excellent in all cases and even the inferior skilled assumption produced relatively good results. Columns 6 - 8 show the rewards and risks for Rule 2 which does control risk. The ACFPDs followed the same patterns as for Rule 1 but were lower everywhere, reflecting the price paid for having lower risk. At costs of 40 bps the ACFPDs were mostly negative and trading was essentially not profitable. At costs of 30 bps only the superior skilled trader produced worthwhile profits, but as costs dropped further to 10 bps profitability improved substantially under all assumptions.

Similar results were obtained in other settings and they confirm the intuitive expectation that low trading costs in addition to trader skill are imperative determinants of successful pairs trading.

7. MODIFICATIONS OF RULES 1 AND 2

Rules 1 and 2 can be modified in many ways and here we report briefly on four of these. In Figure 1 note that the SALB position that was entered on day 1027 and exited at a loss on day 1073 could have been exited at a profit on day 1040 if the low exit signal was changed to downward crossing of a higher B-level such as 50 rather than 35. This suggests that it might be better to use rules that are **asymmetric** as far as entry and exit of positions in terms of B-levels crossings are concerned. We experimented with such asymmetric rules on our data but found that they produce too little improvement over the symmetric rules to justify the burden accompanying the selection of an appropriate value for yet another parameter. Lines 7 – 11 in Table 6 compare the results with the (35,50) asymmetric specification for Rule 1 with the (35,35) symmetric choice given in lines 2 – 6, both for assumptions 6 – 10 with revision time 15 hours. The ACFPD went down in all cases and even severely in some cases, but risks also went down somewhat.

Table 4: Summary statistics for Rule 2 for the total portfolio of six pairs obtained from ASA, FSR, NED and SBK over the period Jan 2005 to Jan 2012 using different revision times and trading pricing assumptions with parameter choices $w = 20$, $B^* = 35$, $CV^* = -1000$, $S = 10000$ and costs of 20 bps.

Revision time and prices	Price as-sump-tion	Average cash flow per day	Average neg clean value per day	Minimum clean value	Number of pos cash flows	Number of neg cash flows	Average positive cash flow	Average negative cash flow
Closing time and prices	1	20,33	-1094,66	-3355,36	441	227	296,83	-420,44
	2	12,39	-931,01	-3208,69	388	280	316,25	-361,05
	3	22,11	-987,21	-3396,26	433	237	291,88	-370,58
	4	59,64	-742,17	-3180,40	501	163	337,02	-397,72
	5	-12,66	-1269,65	-3748,62	376	290	238,41	-385,25
11 hrs ruling prices	6	31,43	-1047,31	-3543,21	460	143	282,58	-525,74
	7	9,00	-968,17	-3269,32	360	247	306,84	-383,69
	8	22,00	-1002,39	-3160,51	408	194	280,15	-391,43
	9	47,60	-791,53	-3448,49	460	142	315,71	-438,10
	10	-6,46	-1244,02	-3436,07	362	238	235,49	-405,53
13 hrs ruling prices	6	26,02	-1098,69	-3494,36	446	159	286,29	-517,65
	7	5,97	-1061,23	-3776,44	371	229	289,54	-423,58
	8	19,97	-1067,76	-3222,69	427	172	273,18	-475,67
	9	42,14	-878,75	-2985,00	460	142	315,11	-503,23
	10	-1,28	-1266,11	-3269,05	385	212	242,35	-450,63
15 hrs ruling prices	6	28,62	-1102,24	-3344,61	455	153	286,98	-527,21
	7	11,42	-1106,12	-3369,65	388	222	296,46	-428,43
	8	21,55	-1099,83	-3662,33	428	179	279,24	-457,74
	9	39,44	-961,63	-3574,10	456	154	311,19	-474,78
	10	1,79	-1257,99	-3361,38	403	205	242,04	-460,62

Table 5: Summary statistics for Rules 1 and 2 for the total portfolio of six pairs obtained from ASA, FSR, NED and SBK over the period Jan 2005 to Jan 2012 at revision time 15 hours and trading pricing assumptions 6 - 10 with $S = 10000$ and costs varying between 10 and 40 bps.

		Rule 1 $(\mathbf{w} = 20\ ,\ \mathbf{B}^* = 35)$			Rule 2 $(\mathbf{w} = 20\ ,\ \mathbf{B}^* = 35, \mathbf{CV}^* = -1000)$		
Costs in basis points	**Price as-sump-tion**	**Average cash flow per day**	**Average neg clean value per day**	**Minimum clean value**	**Average cash flow per day**	**Average clean value per day**	**Minimum clean value**
40	6	4,61	-2124,78	-8883,00	-9,20	-1475,73	-3653,69
	7	-12,03	-2129,89	-8998,32	-27,91	-1432,48	-3593,43
	8	-3,69	-2133,96	-8860,31	-17,57	-1456,58	-3415,66
	9	13,77	-1914,91	-8727,66	1,06	-1302,95	-3395,46
	10	-21,86	-2358,18	-9035,13	-37,52	-1605,73	-3545,30
30	6	23,14	-1887,86	-8608,36	9,31	-1279,75	-3497,48
	7	6,50	-1895,42	-8725,02	-7,99	-1264,31	-3630,37
	8	14,84	-1896,80	-8586,06	1,90	-1280,10	-3345,09
	9	32,30	-1682,22	-8453,43	20,71	-1123,72	-3700,09
	10	-3,33	-2119,17	-8760,73	-17,41	-1446,63	-3602,85
20	6	41,67	-1656,35	-8333,75	28,62	-1102,24	-3344,61
	7	25,03	-1666,48	-8451,74	11,42	-1106,12	-3369,65
	8	33,37	-1663,27	-8311,84	21,55	-1099,83	-3662,33
	9	50,83	-1457,07	-8179,23	39,44	-961,63	-3574,10
	10	15,20	-1881,90	-8486,37	1,79	-1257,99	-3361,38
10	6	60,20	-1433,69	-8059,17	48,42	-934,94	-3114,87
	7	43,56	-1446,20	-8178,51	31,36	-929,75	-3108,93
	8	51,90	-1437,83	-8037,65	39,71	-929,69	-3428,19
	9	69,36	-1244,28	-7905,06	59,19	-784,86	-3339,98
	10	33,73	-1648,51	-8212,04	20,64	-1079,04	-3570,11

Table 6: Summary statistics for Rule 1 and modifications for the total portfolio of six pairs obtained from ASA, FSR, NED and SBK over the period Jan 2005 to Jan 2012 using revision time 15 hrs and trading pricing assumptions 6 - 10 with parameter choices $\mathbf{w} = 20$, $\mathbf{B}^* = 35$ (and 50 for the asymmetric modification), $S = 10000$ and costs of 20 bps.

Rule 1 variation	**Price assump-tion**	**Average cash flow per day**	**Average neg clean value per day**	**Minimum clean value**	**Number of pos cash flows**	**Number of neg cash flows**	**Average positive cash flow**	**Average negative cash flow**
Symmetric entry 35 exit 35	6	41,67	-1656,35	-8333,75	450	147	286,49	-382,65
	7	25,03	-1666,48	-8451,74	384	213	293,68	-324,54
	8	33,37	-1663,27	-8311,84	423	174	277,72	-340,70
	9	50,83	-1457,07	-8179,23	453	144	309,38	-357,62
	10	15,20	-1881,90	-8486,37	396	201	241,15	-343,25
Asymmetric entry 35 exit 50	6	35,64	-1368,30	-7236,82	587	202	204,73	-287,24
	7	15,05	-1362,34	-7583,76	452	337	241,37	-245,84
	8	23,62	-1372,07	-7496,23	528	261	201,31	-249,42
	9	48,13	-1198,80	-7074,41	590	199	232,71	-268,19
	10	-0,93	-1549,99	-7808,23	447	342	176,28	-235,16
Wait for next signal	6	35,91	-1481,21	-7836,75	377	140	304,27	-372,07
	7	22,68	-1494,21	-7952,69	338	179	304,46	-353,97
	8	29,92	-1487,43	-7893,76	374	143	285,15	-380,82
	9	44,73	-1303,15	-7770,19	391	126	321,49	-378,57
	10	14,72	-1677,55	-8325,63	347	170	254,37	-368,24
Multiple max 2 positions	6	39,08	-1540,88	-8085,25	457	140	264,33	-375,96
	7	24,53	-1560,03	-8202,22	404	193	263,27	-329,39
	8	32,09	-1548,51	-8102,80	438	159	253,75	-347,00
	9	47,78	-1356,01	-7974,71	462	135	284,42	-356,15
	10	15,75	-1751,24	-8303,22	413	184	220,53	-345,74
Dividend modifica-tion	6	28,47	-1184,37	-7855,37	419	135	289,30	-530,15
	7	12,94	-1221,83	-7681,64	354	200	295,90	-410,93
	8	20,10	-1199,65	-7522,73	389	165	280,32	-448,40
	9	37,97	-1032,41	-7200,74	423	131	308,97	-492,15
	10	1,99	-1381,05	-7861,06	365	189	243,10	-451,09

As one would expect from earlier exiting of positions the numbers of cash flows increased in all cases while their average sizes decreased. While this is good for the positive cash flows it is bad for the negative ones and on balance these modifications do not seem worthwhile. It is notable that the superior trader (assumption 9) did comparatively well and the inferior trader (assumption 10) particularly poorly in this form of Rule 1.

To motivate the second modification, note in Figure 1 that a SBLA position was entered on day 1012; however, it would have been better to enter this position on day 1014 since the LPD was lower then and exiting the position on day 1027 would have produced more profits. A similar remark applies to the SBLA position entered on day 1072 in which case better profits would have followed if we entered it somewhat later on day 1076 when the first subsequent low signal occurred. This thinking suggests that first crossings of B-thresholds be used for exiting positions after which we wait some days for a (hopefully) larger deviation before entering a new position which may then eventually yield greater profits when exiting. The question is how to formulate this additional waiting time before entering. Possibly the simplest choice is to wait for the first next high or low signal following the one that caused termination of the previously held position. We experimented with this **waiting** variation of the rules but again did not obtain convincing additional value from it. Lines 12 – 16 of Table 6 compare the results for this modification with those of the original Rule 1 in lines 2 - 6. Again the ACFPDs went down but the risks also went down somewhat. The sizes of the positive cash flows increased in line with the motivation for this modification, but unfortunately their numbers went down also and this canceled the hoped for advantage.

As formulated, Rules 1 and 2 are **single position** rules in the sense that at most one SBLA or SALB position is held on any day. For example, if a SBLA is held all further low signals are ignored. An obvious modification is to allow **multiple positions**. This means that if a SBLA position is held an additional SBLA position is entered into on each day giving a further low signal until the first subsequent high signal or crossing of the CV threshold at which moment the total position is exited. Similarly for multiple SALB positions. Lines 17 - 21 of Table 6 also illustrate this modification of Rule 1 with the stipulation that at most two positions of the same type can be entered before exiting them. Compared to the unmodified version of Rule 1 in lines 2 - 6 the results changed only slightly. We calculated the results for three and more multiple positions as well but obtained no material benefit from that either.

So far we did not take dividends and their effects on cash flows into account. In practice a trader will know dividend (and other relevant announcements) relating to the stocks traded in and would take these into account when applying (or temporarily overriding) his formal rules. It is difficult to bring traders' reactions to such announcements into back-testing exercises. In an attempt to do this for dividends, we tested the following **dividend modification** to the rules: on the last day to trade prior to going ex-dividend of a stock do not take any positions in a pair involving that stock and if a position is already held in that stock, exit the positions both in that stock and the other stock involved in the pair. Doing this implies that the trader will neither receive dividends nor be obliged to pay away dividends and that dividends can be ignored in the cash flow calculations. Lines 22 - 26 of Table 6 illustrate the effects of this modification of Rule 1. Compared to the unmodified version of Rule 1 in lines 2 - 6 the ACFPD went down in all cases but so did the risk measures. The number of positive and negative cash flows diminished, the average size of the positive cash flows stayed about the same, but the average size of the negative cash flows increased, most likely due to exiting positions at times which were inconvenient in terms of the benefits expected when the positions were taken. Note that increasing the trade size by about 40% to $S = 1400$ for the modified rule would bring its ANCVPD risk to about the same level as that of the unmodified rule and simultaneously increase the ACFPD by 40% also. Under the price median assumption 8, for example, this would take the ACFPD up to about 28 which compares better with the value of 33 for the unmodified rule. This suggests that the conclusions drawn when ignoring dividends may be somewhat optimistic but not excessively so. Moreover, this dividend modification is likely rather drastic; in practice the trader will have more time between the dividend announcement and the last day to trade to either postpone taking positions or to exit positions more opportunely.

Again similar results were obtained in settings other than those reported in Table 6 and our overall impression is that Rules 1 and 2 are simple but efficient pairs trading rules that capture the essential benefits obtainable from pairs trading provided that trading costs are low and the trader has better than average execution skill.

8. OTHER STOCKS AND FEATURES OF RULES 1 AND 2

We applied Rules 1 and 2 to a variety of stocks other than the banks reported above. The main properties of the rules discussed above were found to apply in other stock contexts as well. By way of example and to discuss further features of the rules, we next present some results for the four stocks AngloGold (ANG), Gold Fields (GFI), Amplats (AMS) and Impala Plats (IMP). The pair ANG and GFI operates in similar environments as do the pair AMS and IMP. The two pairs ANG/GFI and AMS/IMP should therefore be suitable for pairs trading, but pairing the gold with the platinum stocks is unlikely to be profitable. Table 7 shows what happens if we do trade in the pairs from these four stocks using Rule 1 with parameter choices $\mathrm{w} = 20$, $\mathrm{B}^* = 35$ and $S = 10000$, revision time 15 hours and median prices over the rest of the day as the trading prices (assumption 8). The last two lines show the total results for all six pairs and only the first two pairs respectively. It is clear that the ACFPD mainly comes from the two 'natural' pairs ANG/GFI and AMS/IMP with the 'unnatural' pairs contributing little or even negatively as anticipated. Considering only the ACFPD it may seem innocuous to trade with our rules in these unnatural pairs: although they do not produce much in terms of profits, they do not cause large losses either. However, taking risk into consideration, another aspect emerges. When the unnatural pairs are included the ANCVPD is -4018 but when they are excluded it improves to -763. There is an even greater reduction in risk as measured by the MCV. This example shows that with the rules used here the advantage of careful pairs selection is better risk control rather than improvement of reward.

Columns 5 - 8 of Table 7 give some additional statistics regarding the time series of cash flows produced by the trading rule for the individual pairs and for the two totals. Column 5 shows the standard deviations of the cash flow per day which is another possible measure of risk. Column 6 gives the corresponding Sharpe ratios (column 2 divided by column 5). It is clear that the two natural pairs and their total have the smallest risk in terms of variability of cash flow over time as well as the largest Sharpe ratios and this is in line with the discussion above based on ANCVPD as risk measure.

In view of the 'efficient market hypothesis' (EMH) statistical significance of profitability found by back-testing of trading rules is an important issue in much of the pairs trading literature. In the present context we examined the various cash flow series for serial dependence and found that the autocorrelations are invariably quite small (below 0.05 in absolute value for lag order up to 10). The ordinary t-statistic can therefore be used to test the statistical significance of the ACFPD and column 7 gives their values. The Newey-West t-statistics which adjust for possible heteroskedasticity and autocorrelation are given in column 8. These two columns are quite similar and imply that the ACFPDs of the natural pairs and their total are statistically significantly positive while those of the others are not significantly different from 0. The latter finding is in line with what the EMH suggests, but at first glance the former finding may appear to go against the EMH. However, keep in mind that all back-testing results are based on certain assumptions (here among others, the assumed costs and pricing assumption 8) and the statistically significantly positive ACFPD may be a reflection of a lack of trading realism in these assumptions rather than an indication that the EMH is false.

9. CONCLUSION

In this paper we presented a fresh approach to pairs trading rules. Compared to the existing literature our approach puts less emphasis on the selection of pairs to trade in and more emphasis on formulation of rules that enter and exit long/short positions appropriately in economically meaningful pairs. Positions are only entered in self-financed form and rules are calibrated and evaluated in terms of their reward/risk implications with reward and risk measured by average cash flow and average negative clean value per day respectively. In the back-testing we differentiate between the prices of the members of a pair used to get high/low signals and the prices at which actual trading takes place. Trading costs are incorporated in each stage of the back-testing and allowance is made for possible trader skill in the evaluation of rules.

Table 7: Summary statistics for Rule 1 pairs trading in ANG, GFI, AMS and IMP over the period Jan 2005 to Jan 2012 using daily revision time 15 hours and trading pricing assumption 8 for implementation with $\mathbf{w} = \mathbf{20}$, $\mathbf{B}^* = \mathbf{35}$, $\boldsymbol{S} = \mathbf{10000}$ and costs of 20 bps.

Pair	Average cash flow per day	Average neg clean value per day	Minimum clean value	Stand-dev of cash flows per day	Sharpe ratio	t-statistic	Newey-West t-statistic
ANG/GFI	7,39	-367,44	-2986,47	98,19	0,075	3,14	3,03
AMS/IMP	13,33	-453,57	-2611,18	154,51	0,086	3,60	3,48
ANG/AMS	-2,91	-912,00	-4694,53	253,37	-0,011	-0,48	-0,49
ANG/IMP	3,57	-764,91	-4953,57	210,80	0,017	0,71	0,73
GFI/AMS	-2,85	-943,49	-5117,61	245,43	-0,012	-0,48	-0,54
GFI/IMP	4,74	-798,86	-5963,94	213,97	0,022	0,92	0,93
All 6 pairs	23,27	-4059,27	-21568,05	570,83	0,041	1,70	1,43
First 2 pairs	20,72	-762,82	-2986,47	183,32	0,113	4,71	4,43

When tested on JSE listed stocks we find that pairs trading can be profitable provided that trading costs are of the order of 20 basis points or lower and the trader has about average skill. With higher skill profitability can be maintained at even higher trading costs.

A number of intuitively reasonable variations on the basic rules were evaluated and found not to be worthwhile. However, more variations are possible among which we briefly mention the following one. As they stand the rules generate much statistical information as they proceed over time and the issue is whether this can be incorporated into the rules to improve them. For example, at any point in time we may compute the average waiting times in past positions that lead to positive and negative cash flows respectively. The waiting times to negative cash flows are typically much longer than those leading to positive cash flows (see the bottom panel in Figure 1 for an example). This suggests that the time in the current position may be compared to past experience in order to predict the likely sign of the cash flow when the current position is eventually exited. If the prediction is negative the position may be exited and trading in the pair suspended immediately and this variation of the rule may improve its performance. We leave the investigation of such further variations for future research.

ACKNOWLEDGEMENT

We thank the reviewers for their comments which helped to improve the paper.

REFERENCES

Andrade S, Di Pietro V and Seasholes M. 2005. Understanding the profitability of pairs trading. Working Paper. UC Berkeley Haas School.

Bolgun KE, Kurun E and Guven S. 2009. Dynamic pairs trading for companies listed in the Istanbul stock exchange. Working Paper.

Broussard JP and Vaihekoski M. 2010. Profitability of pairs trading strategy in Finland. Working Paper.

Do B and Faff R. 2010. Does simple pairs trading still work? *Financial Analysts Journal*, 66:83-95.

Do B and Faff R. 2012. Are pairs trading profits robust to trading costs? *The Journal of Financial Research*, XXXV:261-287.

Do B, Faff R and Hamza K. 2006. A new approach to modelling and estimation for pairs trading. Working Paper.

Elliot RJ, van der Hoek J and Malcolm WP. 2005. Pairs trading. *Quantitative Finance,*3:271-276.

Engelberg J, Gao P and Jagannathan R. 2009. An anatomy of pairs trading: The role of idiosyncratic news, common information and liquidity. Working Paper.

Gatev E, Goetzman WN and Rouwenhorst KG. 2006. Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, 19:797-827.

Hong G and Susmel R. 2004. Pairs-trading in the Asian ADR Market. Working Paper.

Lin YX, McCrae M and Gulati C. 2006. Loss protection in pairs trading through minimum profit bounds: A cointegration approach. *Journal of Applied Mathematics and Decision Sciences*, 2006:1-14.

Mudchanatongsuk S, Primbs JA and Wong W. 2008. Optimal pairs trading: A stochastic control approach. American Control Conference, Seattle, Washington June 11-13, 2008.

Tsay, R.B. 2010. Analysis of financial time series. J Wiley, New York.

Vidyamurthy G. 2004. Pairs trading, quantitative methods and analysis. J Wiley & Sons, Canada.