An improved pairs trading strategy based on switching regime volatility

Marco Bee

Department of Economics and Management, University of Trento

Giulio Gatti

Department of Economics and Management, University of Trento

July 27, 2015

Abstract

A pairs trading strategy on energy, agricultural and index futures is developed. The strategy uses different parameters according to a volatility regime detected using a threshold evaluated in two ways, namely by means of a mixture of two Gaussian densities and a Markov switching model. The performance is assessed using different time frames and filters. When associated to cointegration, this investment algorithm gives a larger Sharpe ratio with respect to classical methods; on the other hand, the correlation filter does not work well with the regime switching algorithm.

Keywords: Pairs trading; Gaussian mixture; Hidden Markov Model; cointegration.

1 Introduction

Relative value trading, developed in the mid 80s of the 20-th century by a team led by Nunzio Tartaglia (Gatev et al., 2006), is currently one of the most popular market-neutral strategies. The basic version uses a mean-reverting technique[1] that prescribes to short one asset and buy the other one when they diverge, counting on a long-run equilibrium and on a future convergence of the two assets. By so doing, the market value of the resulting portfolio is expected to be stationary.

There are hundreds of pairs with these characteristics: for example, two companies operating within the same sector, with similar fundamental values (such as PepsiCo and Coca-Cola), the price of a commodity and the firm that extracts or produces refined products (oil and Gazprom), futures with the same underlying asset but with different expiration (calendar spread), bonds with different maturities, and so on.

Since the introduction of relative value strategies, many traders have started using them. Nowadays, with the development of technology and the increasing computing power, many hedge funds and proprietary trading companies have adopted trading algorithms based on this idea.

Various authors have given formal definitions of statistical arbitrage (see, in particular, Avellaneda and Lee, 2010 and Hogan et al., 2004). In general, these models may be seen as a combination of a pure arbitrage strategy and a mean-reverting strategy. The difference between pure and statistical arbitrage is that the latter has a mostly empirical support, while the first is based on theoretical considerations. In

[1] This technique is based on a volatility break-out that employs Bollinger bands (Bollinger, 2001). Each band is defined as a moving standard deviation summed (and subtracted) to a moving average of prices. The standard deviation is usually associated to a multiplier factor.

other words, while pure arbitrage relies on mathematical models that may then be used in practice, statistical arbitrage methods have their starting points in economic and empirical evidence based on the behavior of asset prices, and try to model it according to an investment perspective.

The distinction can also be summarized by saying that the probability of losing money in a pure arbitrage strategy should be zero, whereas with a statistical arbitrage model the probability is zero only in an infinite time horizon. In other words, there is a positive probability of losses in a finite time period (Hogan et al., 2004).

Relative value trading is an investment scheme that focuses on relative pricing instead of absolute pricing. While the latter tries to develop methods based on discounted cash flow models, relative pricing concentrates on an asset price with respect to another. The main rule is that two substitute goods or assets have to be sold at the same price, which is not known in advance (Gatev et al., 2006).

A pairs trading strategy may also be related to equilibrium models or relative value methods of asset pricing, such as the Law of One Price, the Arbitrage Pricing Theory (APT) and the Capital Asset Pricing Model. Since these are some of the most important models of the mechanics of financial markets, we may say that they justify pairs trading strategies from a theoretical point of view.

The market-neutral idea supporting the strategy implies that the trader has to consider three crucial decisions:

(i) select the assets to trade within the set of all tradable instruments;

(ii) given that a spread (i.e., a dynamic weighted difference between two assets) is traded, define the concept of spread appropriately in econometric terms;

(iii) choose the trading algorithm.

When a strategy operates on just one asset, the first problem is almost trivial, the second disappears and only the third is actually important. On the other hand, in a trading pairs setup, all of them have to be considered.

A statistical arbitrage mechanism is deeply connected to the concept of equilibrium between two prices. When such an equilibrium is broken, there are profit opportunities, under the condition that the equilibrium relation will sooner or later be restored. Such models suffer from three kinds of market behaviors, that may be described as follows:

(i) the spread no longer converges: the system will never be profitable;

(ii) the equilibrium relation is restored slowly: some strategies may have to close positions before the spread converges;

(iii) the spread volatility increases: a strategy may become unprofitable, because the trigger for an entry signal is given too early and the difference between two prices will diverge in the near future.

The aim of the model proposed in this paper is to avoid issues related to the raise of volatility, trying instead to benefit from it. The basic version prescribes to buy (sell) the spread when the price breaks the lower (upper) Bollinger band; we suggest to build different entry signals according to the current volatility regime, detected through the identification of a threshold that determines the volatility regime. We develop two approaches to the determination of the threshold: the first one uses a two-population Gaussian mixture, and the second one an Hidden Markov Model (HMM), in order to take into account the possible presence of some temporal dependence.

The performance of these techniques is measured by a backtesting exercise whose outcomes are quite good, as the strategy produces larger returns and Sharpe ratios with respect to other commonly used strategies.

The system is backtested on the 10 and 20 minute closing prices of 7 futures pairs from October 2012 to September 2014. There are 2 pairs within the energy futures market, an inter-market spread between Light Sweet Crude Oil futures (CL) and Brent Crude futures (BRN) and an intra commodity pair formed by Heating Oil (HO) and RBOB Gasoline (RB). As for agricultural futures, the pairs are Soybean Meal (ZM) versus Soybean Oil (ZL) and Soybean (ZS) versus Corn (ZC). The last three spreads are created with the mini futures of Dow Jones Index (YM), Nasdaq (NQ) and S&P 500 (ES).

The paper is organized as follows. In Sect. 2 we give an overview of the relative value trading literature. Sect. 3 summarizes the standard model and gives the details of the methodology employed for describing the volatility shifts. Sect. 4 proposes a switching regime algorithm where the regime leads to different parameter choices, according to the standard deviation of the spread. A comparison of the two models to each other and to more classical approaches is carried out in Sect. 5. Finally, Sect. 6 discusses the results.

2 Literature overview

We can summarize the most important contributions in the literature according to the three choices needed to implement the market-neutral approach (see the discussion in Sect. 1 and Delledonne, 2012).

2.1 Choosing the assets

The minimum distance criterion is a simple and intuitive method proposed by Gatev et al. (2006). Correlation filters allow one to find a threshold and trade pairs whose correlation exceeds it (Modi et al., 2011). Cointegration filters (Burgess, 1999) have been developed as alternatives to correlation with the aim of finding a long term relationship between two time series. Beta's significance is another simple tool based on the significance of the beta of the regression between two time series. ETF methods, first proposed by Avellaneda and Lee (2010), derive from the well-known concept of APT's common risk factors. They use ETFs, namely baskets of stocks with common risk factors, as hedging instruments. The same paper also proposes a procedure based on Principal Component Analysis to select the best components for trading purposes.

2.2 Modeling the spread

The minimum variance ratio is a well-known hedging technique (see, for example, Hull, 2012), where the beta used to weight the two assets is such that the variance of the resulting portfolio is minimized.

The Total Least Squares (TLS) approach (Teetor, 2011) takes care of the fact that, for pairs trading purposes, buying an asset and shorting the other one or vice versa should make no difference. However, the classical OLS method only accounts for the variance of one leg; hence, it implicitly treats the market in an asymmetric way. On the other hand, the hedge ratio found by the TLS algorithm is consistent and the variance is treated symmetrically.

Another possibility consists in modeling the spread as a first order autoregressive process. Burgess (1999) uses a discrete process, while Bertram (2010) proposes a continuous setup.

More recently, the spread has also been described by HMMs. A first version was presented by Bock and Mestel (2008): according to the authors, the time series of the spread contains structural breaks and the assumption of constant parameters is not reasonable. Structural breaks in economic and financial time series are caused by financial crises, wars, political changes or bubbles in commodity prices, and may occur in autoregressive parameters, regression coefficients, means, variances or GARCH parameters. Bock and Mestel (2008) develop a model where regime shifts are governed by a Markov chain, the current regime is determined by an unobservable variable and inference on regimes is based on state probabilities. Formally, the model is a two-state first-order Markov Switching process with switching mean and variance.

Bucca and Cummins (2012) analyze instead structural breaks on the cointegration relationships among different time series. They describe how a cointegration model may have non-constant parameters and may be seen as a switching mean process or switching mean trending process or two-regime process.

The spread is also sometimes modeled via a latent variable dynamic process, whose parameters are estimated using the Kalman filter. Specifically, Chan (2013) proposes a model that recursively estimates the hedging ratio avoiding discretionary decisions or arbitrary choices.

Finally, a spread can be described according to physical relations originating from empirical evidence based on economic reasons rather than on mathematical models. Typical examples are crush and crack spreads (Mitchell, 2010; Bucca and Cummins, 2012; CME, 2006; CME, 2013).

2.3 Choosing the trading algorithm

As for the trading algorithm, the simplest solution is the white noise spread hypothesis, which just assumes that the spread is white noise. Some implementations optimize the volatility multiplier of Bollinger bands; in other cases (e.g. Gatev et al., 2006) a multiplier of the volatility equal to 2 is set ex ante, according to the 95% critical value of a Gaussian distribution.

Similarly to returns computed from asset prices, in most cases empirical evidence suggests that spreads do not follow a normal distribution; accordingly, a more sophisticated approach assumes a non-Gaussian spread density. Different, possibly non-parametric, techniques may be developed in order to take into account these features; for example, Vidyamurthy (2004) uses an alternative method based on mixtures of Gaussian densities.

3 Modeling volatility shifts

In order to define whether the volatility regime is high or low, various tools can be employed. A common assumption is that the spread is stationary with expected value equal to zero. As for the problem of modeling the whole distribution, in this paper we use two different methods: first, we assume that the spread is well described by a mixture of two Gaussian densities; second, the data-generating process is supposed to be a stochastic process with regime switches in volatility. The latter is a generalization of the former, in that the sequence of hidden variables determining the population that generated each observation is not independent, but rather follows a Markov process; thus, possible temporal dependencies are included in the model.

In the next sections, after properly defining the notion of spread, we provide an in-depth analysis of the two methods. When moving to real data, in both the normal mixture and the Markov switching approach, we expect (i) to find two regimes with identical expected value but different standard deviation and (ii) the density with lower variance to be associated to the larger probability parameter.

3.1 Gaussian mixture models

The density of a two-population mixture of normal univariate densities is defined as follows:

$$f(x) = \pi_1 \phi(x; \mu_1, \sigma_1^2) + \pi_2 \phi(x; \mu_2, \sigma_2^2), \tag{1}$$

where $\phi(x; \mu, \sigma^2)$ is the normal density with parameters μ and σ^2, $\pi_j \geq 0$ $(j = 1, 2)$ and $\pi_1 + \pi_2 = 1$, so that we just set $\pi_1 = \pi$ and $\pi_2 = 1 - \pi$. Equation (1) can therefore be rewritten as:

$$f(x) = \pi \phi(x; \mu_1, \sigma_1^2) + (1 - \pi) \phi(x; \mu_2, \sigma_2^2).$$

The vector of unknown parameters to be estimated is $\theta = (\pi, \mu_1, \sigma_1^2, \mu_2, \sigma_2^2)'$, namely the (prior) probability of belonging to the first distribution (π), and the expected value and variance of both probability distributions.

3.2 Markov Switching Models

The second model has originally been proposed for pairs trading purposes by Bock and Mestel (2008). However, Markov switching models have been used long before. As for estimation, MLE based on the EM algorithm has been developed by Hamilton (1990).

Within this framework, regime switches are led by Markov chains. The regime is determined by hidden variables and inference about regimes relies on state probabilities. The process used here has two states, with switching mean and variance. Even though, in principle, it may not be unreasonable to consider only the fixed mean specification, a switching-mean model seems to produce better estimates.

The main difference between our model and the Bock and Mestel (2008) one relies on what is modeled. While their underlying variable is the hedging ratio, our time series is the spread itself and not one of its components. Formally, the process is defined as $spread_t = \mu_{s_t} + \epsilon_t$, where $\mathrm{E}(\epsilon_t) = 0$ and $\sigma^2_{\epsilon_t} = \sigma^2_{s_t}$.

Note that, in the setup of Bock and Mestel (2008), a switching mean was indeed necessary. On the other hand, since the model we use for the spread leads to a stationary process, we do not need it.

4 A systematic model robust to changes in volatility regimes

Building on the approaches by Bock and Mestel (2008), Bucca and Cummins (2012) and Vidyamurthy (2004), in this section we develop a novel framework with the goal of improving the entry signals.

4.1 Some common problems in spread trading

In Sect. 1 we have identified three issues related to market-neutral strategies. Here we refine that classification and consider some difficulties concerning spread trading.

A long lasting divergence in prices is difficult to find with algorithms based only on prices. Usually, fundamental analysis combined with traders' experience may help to understand if the divergence is permanent. The time of mean reversion can be estimated using continuous models; in this case, the position's holding period depends on the model used for the spread and, as a consequence, on the hedging ratio updating frequency.

Many factors may cause jumps in spread volatility, so that jumps themselves cannot be considered signals to stop trading on some pair.

When frequently updating the hedging ratio, it is more likely to find jumps in volatility than in mean. On the other hand, if the ratio is changed less frequently, switches in mean can also be found, since the spread starts to show a strong dependency on the future relationship of prices.

4.2 A three-step testing strategy

4.2.1 Pairs identification

In this section we employ some of the concepts introduced above as an on/off filter for our algorithm. As we use intraday data, the filters are evaluated with a daily frequency, i.e. the length of the training and test set consists of the 11 hours of the previous day (9:20am - 8:15pm).

Seven pairs have been chosen with the aim of avoiding spurious correlations or insane and casual matchings. Since we test strategies on commodities or index futures, we do not focus on fundamental analysis; this would be recommended in equity trading.

We consider the following three simple filters:

- 5% significance of the slope parameter of a linear regression;
- correlation between prices larger than 80%;
- 5% significance of the cointegration relation between two asset prices, using the Kwiatkowsky, Phillips, Schmidt e Shin (KPSS) test (Kwiatkowski et al., 1992).

4.2.2 Spread modeling

The spread is defined as the weighted difference between two prices: $spread_t = P_{1t} - \beta P_{2t}$. The sampling frequency is equal to 10 or 20 minutes, and the hedging ratio β is updated daily using previous day data by means of the minimum variance ratio method (see Sect. 2.2). The updating frequency leads to a strictly intraday trading model.

These choices have the advantage of avoiding overnight margins for futures trading; at the same time, the position cannot be held for more than a trading day, which can be an annoying limitation.

4.2.3 Trading algorithm

The trading model is a straightforward mean-reverting strategy that exploits Bollinger bands. The model proposed in this paper uses four bands, the widest as an entry signal and the narrowest as an exit. As a consequence, the strategy is not in the market all the time. There is only one additional exit signal (which is not a take profit), corresponding to the end of the day.

This strategy is compared to a similar one that uses a volatility threshold to identify the Bollinger multiplier that better suits the current volatility state. Thus, in this case

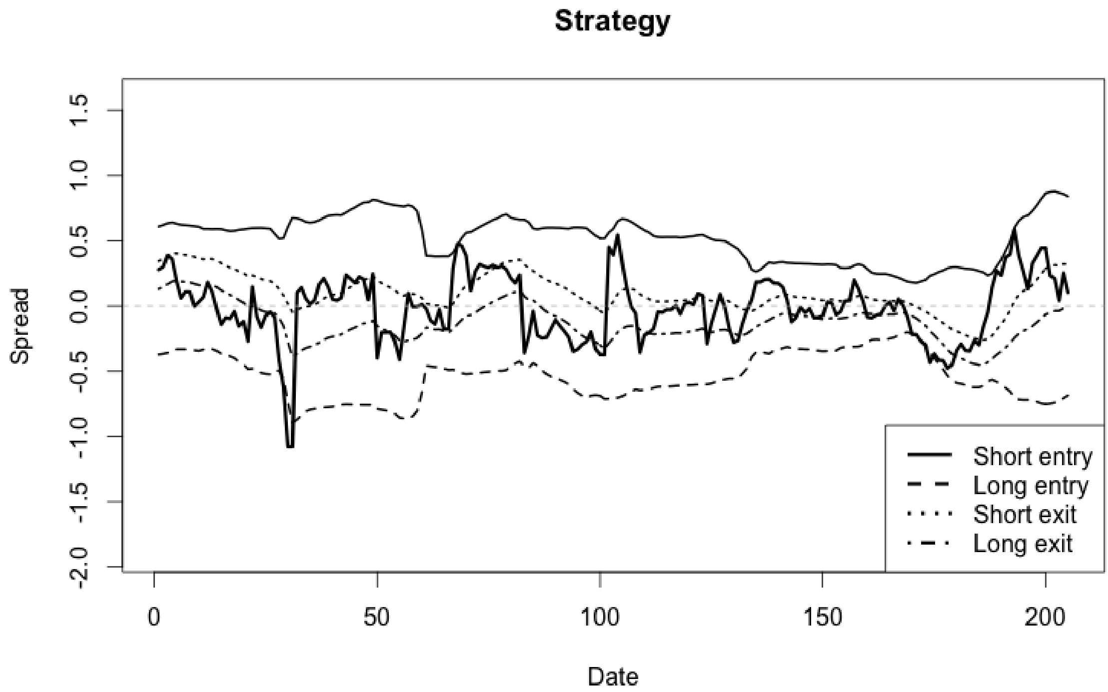

Figure 1: The trading algorithm with one regime

there are eight bands, four when the volatility is low and four when it is high.

An example of the first model (with 4 Bollinger indicators) is shown in Figure 1: when a candle's close bar is above the continuous line, we enter short, and vice versa when the close price is smaller than the lowest line. As concerns the narrowest band, the upper line is the long exit level while the lower one is the trigger of a short exit. Since we expect a countertrend spread behaviour, we are betting on an intraday mean reversion.

The full system is based on two different regimes. The indicators are divided in three groups: entry levels, exit levels and a switching regime threshold. The threshold is the difference between the basic model and our approach, and Gaussian mixtures and HMMs play a key role for its determination.

As the threshold indicates the volatility level corresponding to the shift from one regime to the other, the problem consists in defining it properly. In the mixture setup, we use as threshold the larger of the two mixture populations standard deviations; similarly, in the HMM framework, we employ the larger of the standard deviations of the two switching models.

5 Results

We split the data into four samples, using the first and the third one as training sets. In particular, we set the threshold equal to the average of the values of the two training sets. Then we employ this value to test the strategy using a single threshold for all the samples.

All the computations have been carried out using the R software (R Core Team, 2015). As concerns estimation, we used the `mixtools` and `MSwM` packages respectively for Gaussian mixtures and Markov switching models. In both cases the likelihood is maximized by means of the EM algorithm. As the sample sizes are approximately equal to 4500, the computational cost is quite small: a few seconds are necessary in the Gaussian mixture case, and around a minute, on average, for HMMs.

The three panels of Figure 2 graphically represent how the regimes are detected. The top panel shows that there are periods of high volatility followed by periods of low volatility and that this behavior seems to persist over time. The middle panel displays the rolling standard deviation along with the threshold (the dashed line): high-volatility regimes correspond to periods with rolling standard deviation above the dashed line, low-volatility regimes to the opposite situation. Finally, the bottom panel displays the regimes. Figure 3 shows how entry levels change simultaneously

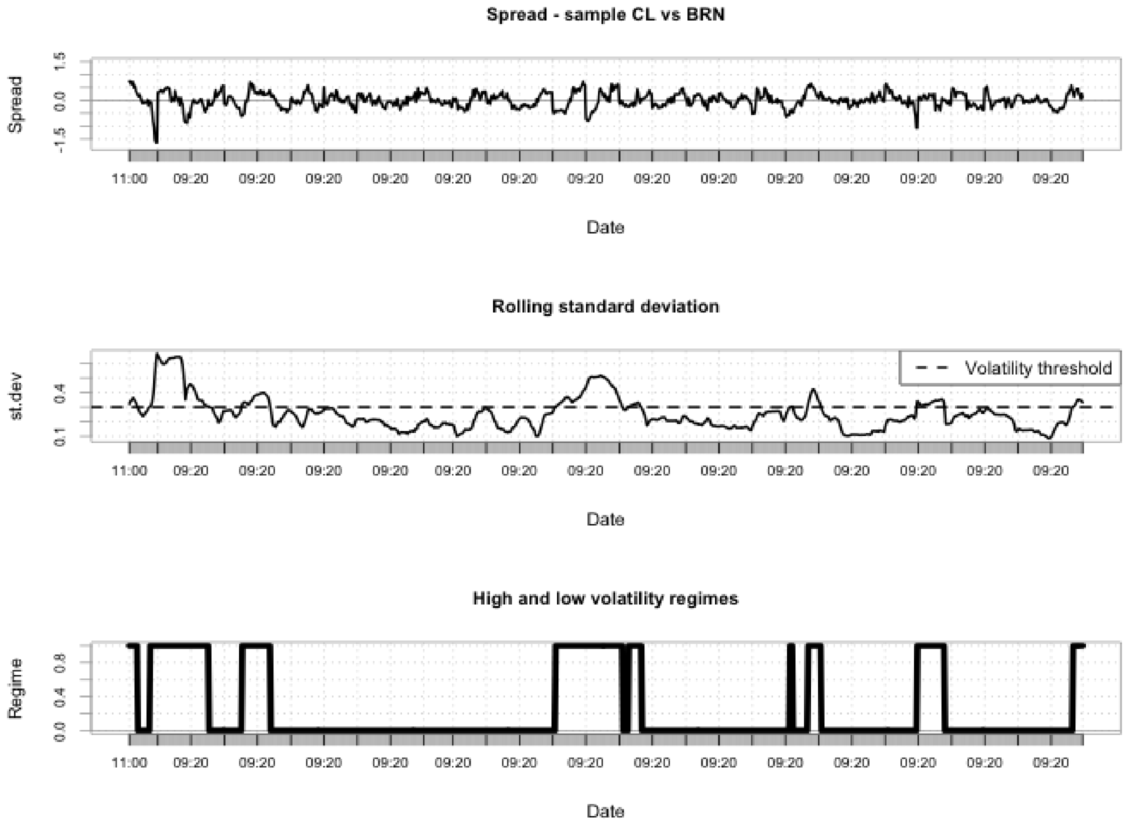

Figure 2: CL - BRN spread subsample, volatility and regimes.

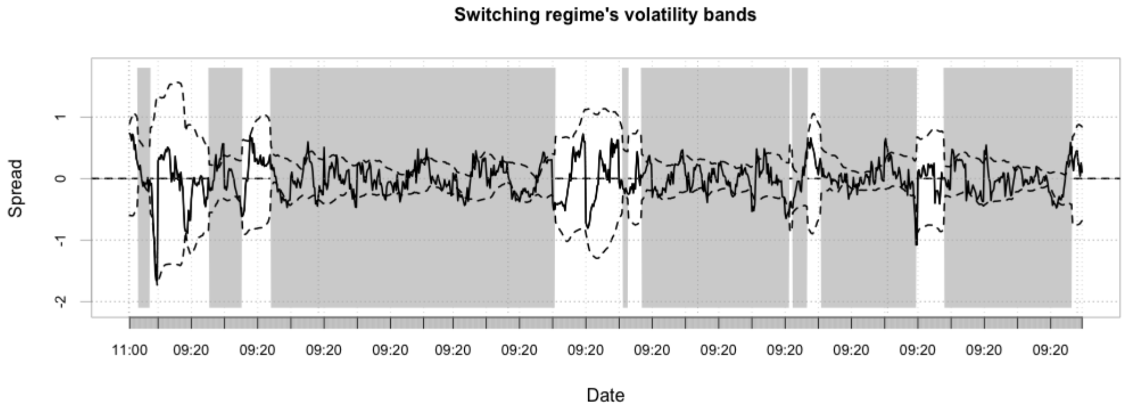

Figure 3: Switching regime entry level.

with the underlying volatility. Note that exit levels behave in the same way.

As of spread modelling, we assume to buy (sell) five contracts for one asset, selling (buying) a quantity equal to $\lfloor \beta \times 5 \rfloor$, where $\lfloor x \rfloor$ is the integer part of x, for the other one, as the quantity must be an integer. When spread futures are traded, initial and maintenance margins are required; these margins should be compensated, and the total requirement might be smaller than the sum of the two margin positions (because we are hedging). However, not every broker considers the hedge portfolio and asks for the whole margins sum.

We assume to trade with 500 000\$ and our margin-to-equity ratio to be approximately 20%. We do not use any money management technique, the margin is 20% only at the beginning of the backtesting period, and decreases when we gain from the strategy.

The overall results are analyzed via some of the most common indicators employed by practitioners: annualized Sharpe ratio, annualized returns, maximum drawdown, percentage of days with positive returns, percentage of days with negative returns, percentage of days without trades, ratio between the mean (in dollars) of positive days returns and negative days returns, correlation with a fully invested portfolio in the S&P 500 Index. In all cases we show the net results, taking into account a 3$ commission fee and one tick of slippage for trade.

The most remarkable finding is that the cointegration filter performs better than the other two. Figure 4 shows how, during the last semester, at the beginning of the oil price collapse, the only strategies whose results are satisfactory are those based on cointegration. On the other hand, the algorithm leads to poor results if used with correlation or significance regression as filters.

One of the most recognizable problems of intraday spread strategies concerns commission fees. Within our tests the 10 minutes time frame is more affected than the others by these costs.

We only report here (see Table 1) the results for the 20-minute time frame data with cointegration filter, as in the 10-minute frame the outcomes are very similar. Except for the spread between ES and YM contracts, all other strategies exhibit a better performance with two regimes than with one.

Moreover, all methods show positive returns and rather large Sharpe ratios. There is, however, a huge difference between the returns of the commodity and the financial index sector; since index futures contracts are more liquid than commodities, lower returns would have been expected. In both sectors Sharpe ratios larger than 2 have been obtained, so that even the basic methodology appears to behave well.

The overall risk is definitely low, as the magnitude of drawdowns in Table 1 can be surely absorbed without any problem. On average, the percentage of negative return days is low, but a thorough evaluation needs also to take into account the importance of the cointegration filter: the percentage of no-trade days is between 60 and 80%.

		Sharpe Ratio	Annualized Returns	Max Drawdown	% positive days	% negative days	% no trade days	mean W / |mean L|	S&P500 correlation
CL vs BRN	1 regime	5.55	22.00%	-1.11%	22.52%	4.08%	73.40%	2.79	0.01
	2 regimes	6.28	28.07%	-0.86%	25.63%	3.88%	70.49%	2.99	0.01
HO vs RB	1 regime	2.88	16.98%	-2.24%	16.12%	6.02%	77.86%	1.60	0.02
	2 regimes	4.76	41.94%	-3.61%	27.96%	6.60%	65.44%	1.65	0.01
ZS vs ZC	1 regime	3.64	22.92%	-1.80%	17.10%	10.34%	72.56%	2.89	-0.02
	2 regimes	5.45	36.80%	-1.12%	26.24%	7.36%	66.40%	2.59	0.04
ZM vs ZL	1 regime	3.44	23.75%	-1.38%	22.47%	10.54%	67.00%	2.12	-0.01
	2 regimes	4.67	36.38%	-2.89%	28.43%	8.15%	63.42%	1.91	0.01
NQ vs YM	1 regime	3.79	3.85%	-0.30%	12.04%	4.27%	83.69%	2.26	0.03
	2 regimes	5.26	9.06%	-0.15%	20.97%	5.24%	73.79%	2.86	0.01
ES vs YM	1 regime	2.93	1.84%	-0.30%	9.32%	1.94%	88.74%	1.24	0.03
	2 regimes	2.87	2.77%	-0.44%	16.89%	7.77%	75.34%	1.36	0.16
ES vs NQ	1 regime	3.14	3.15%	-0.59%	10.29%	2.52%	87.18%	1.45	-0.03
	2 regimes	4.13	7.28%	-0.79%	17.67%	4.85%	77.48%	1.84	0.03

Table 1: Performance Statistics: 20-minute time frame overall results with cointegration filter. The comparison is between the naive model (one-regime) and the two-regime algorithm.

Another important issue concerns the difference between "no-trade days" with one and with two regimes: the novel strategy operates on average more frequently than the simple one. As a consequence of this increase, the number of positive return days has increased as well, while, in some cases, there is a decrease in the "mean wins on mean losses" ratio. Since the commission fee does not erode the net profit, a strategy with an higher number of operations is preferred, as the overall performance is more stable.

Among all different choices that can be made, the decision concerning the filters or parameters to use mainly depends on the trader risk aversion. As said above, the cointegration filter leads to a smaller number of trades and to less volatile returns, and should therefore be preferred. In particular, the number of trades is always between four and eight per day and the percentage of trade days is, on average, equal to 25%, so that the number of transactions is large enough to lead to robust results.

Finally, the absence of correlation with the S&P500 index is a valuable characteristic to both institutional and retail investors, as it adds diversification effects to a portfolio. More evidence comes from Figure 5, 6 and 7, where the equity line of the model, compared to the naive strategy, confirms the superior performance of our method.

6 Conclusion

In this paper we have developed a new model that explains how pairs trading strategies are suitable to different underlying contracts. We have analyzed two alternative ways to detect the switch of volatility regimes. Since the estimates are similar in both cases, the strategies performances are also approximately the same.

This paper exploits some methods first proposed by Bock and Mestel (2008) and Vidyamurthy (2004), respectively for HMMs and Gaussian mixtures. However, there are some significant differences with respect to these works. The most important one is that, while Bock and Mestel (2008) apply the model to the hedging ratio, we model the spread itself. On the other hand, Bucca and Cummins (2011) exploit structural changes in cointegration in the oil sector, whereas we extend the set of underlying instruments to include agricultural commodity and stock index futures.

The proposed method has several strengths but also some weaknesses. The performances associated with the cointegration filter are robust among all pairs and time frames. The correlation with a classic long only portfolio is essentially zero, so that this investment is interesting for all portfolio managers. Both the basic and the novel models produce a low volatility of returns but the novel approach makes is possible to increase returns proportionally more than volatility.

Some difficulties may arise in practical implementation. First, there might be lack of liquidity in commodity markets. Second, the scalability issue is important, and each trader has to take it into account. Finally, we cannot afford the use of different filters (like correlation), since the risk of the strategy would increase exponentially. The cointegration filter results in no trade for the majority of days, so that a large dataset is needed to ensure some robustness of the results.

Not only institutional investors may face liquidity problems; retail investors may have the opposite difficulty, due to the high margins required in futures trading. Spread positions and strictly intraday systems might help to reduce margin requirements. This limit can always be overcome by trading stocks instead of futures, even if the strategy

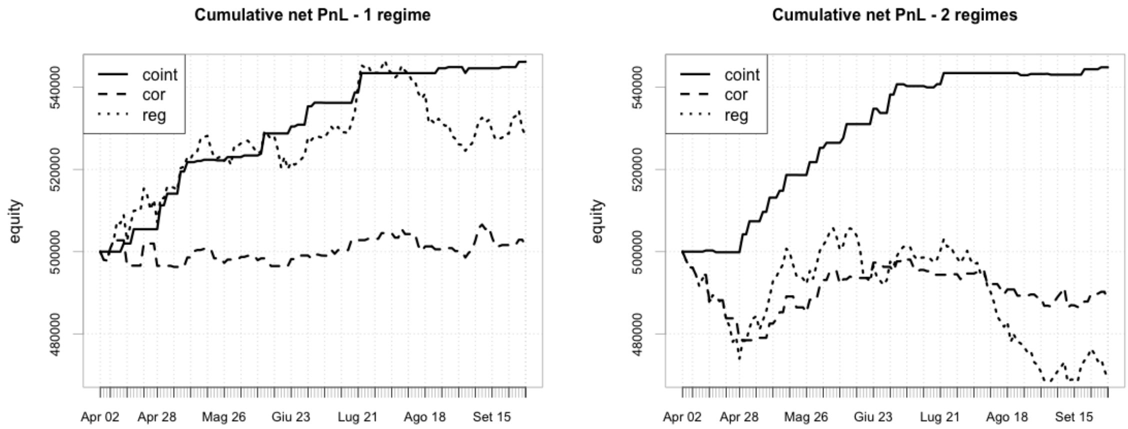

Figure 4: CL vs BRN, 2nd out of sample: cointegration, correlation and regression significance filters. The left anf right panel resepctively show the equity line for one- and two-regime strategies.

may require fundamental analysis to help the systematic model.

Finally, the hedging ratio estimates are based on the OLS method. The use of Total Least Squares may reduce the risk as the estimates are obtained by solving a symmetric optimization problem.

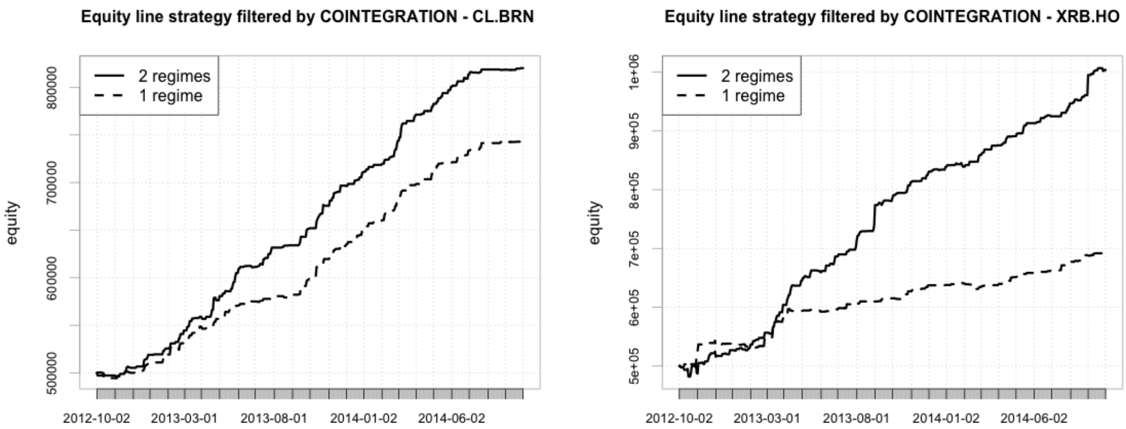

Figure 5: Oil futures: Equity line for one- and two-regime strategies.

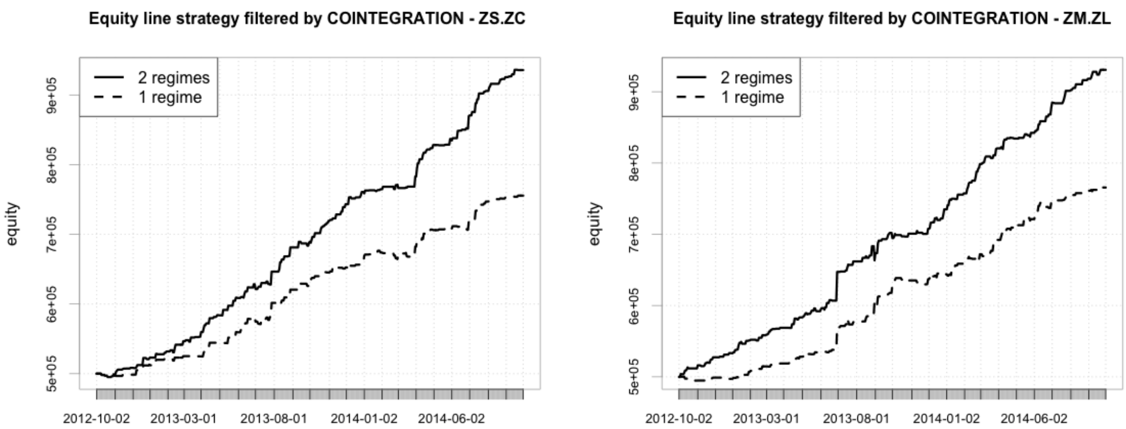

Figure 6: Agricultural commodity futures: Equity line for one- and two-regime strategies.

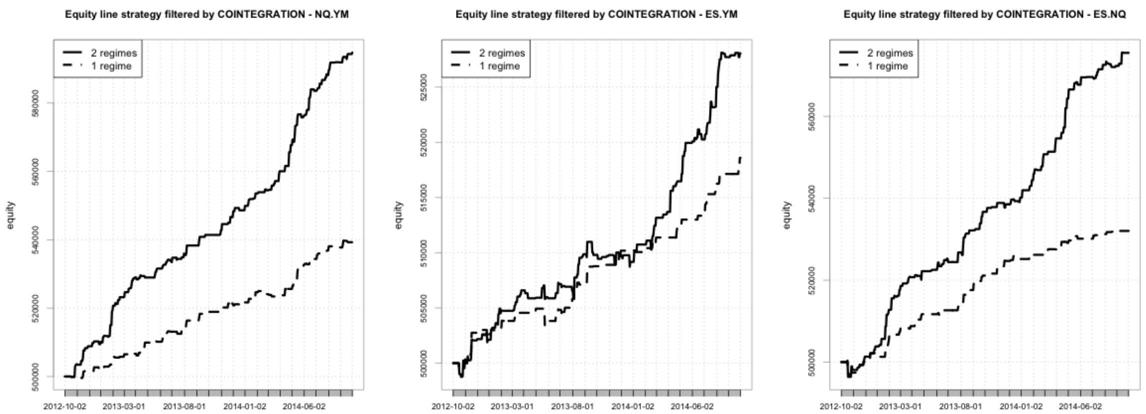

Figure 7: US Index futures: Equity line for one- and two-regime strategies.

References

M. Avellaneda and J. H. Lee. Statistical arbitrage in the US equities market. *Quantitative Finance*, 2010.

W. K. Bertram. A non-stationary model for statistical arbitrage trading. *SSRN, http://ssrn.com/abstract=1632718*, 2010.

R. Bock and M. Mestel. A regime-switching relative value arbitrage rule. *Operations Research Proceedings*, 2008.

J. A. Bollinger. *Bollinger on Bollinger Bands.* McGraw-Hill, 2001.

A. Bucca and M. Cummins. Synthetic floating crude oil storage and optimal statistical arbitrage: A model specification analysis. *SSRN, http://ssrn.com/abstract=1932432*, 2011.

A. Bucca and M. Cummins. Quantitative spread trading on crude oil and refined products markets. *Quantitative Finance*, 12:1857–1875, 2012.

A. N. Burgess. *A Computational Methodology for Modelling the Dynamics of Statistical Arbitrage*. PhD thesis, London Business School, 1999.

E. P. Chan. *Algorithmic Trading - Winning Strategies and Their Rationale*. Wiley, 2013.

CME. Cbot soybean crush. Reference Guide, 2006.

CME. Introduction to crack spreads. Reference Guide, 2013.

F. Delledonne. Pairs trading: Un modello a variabili latenti basato sul filtro di Kalman. Master's thesis, Politecnico di Milano, 2012.

E. Gatev, W. N. Goetzmann, and K. G. Rouwenhorst. Pairs trading: Performance of a relative-value arbitrage rule. *Review of Financial Studies*, 19:797–827, 2006.

J. D. Hamilton. Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45:39–70, 1990.

S. Hogan, R. Jarrow, M. Teoc, and M. Warachka. Testing market efficiency using statistical arbitrage with applications to momentum and value strategies. *Journal of Financial Economics*, 73:525–565, 2004.

J. C. Hull. *Options, Futures, And Other Derivatives*. Pearson, 8th edition, 2012.

D. Kwiatkowski, P. C. B. Phillips, P. Schmidt, and Y. Shin. Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54: 159–178, 1992.

J. B. Mitchell. Soybean futures crush spread arbitrage: Trading strategies and market efficiency. *Journal of Risk and Financial Management*, 2010.

A. G. Modi, J. Desai, Dave A. R., and K. J. Desai. Finding the identical twin. *Econometrics: Multiple Equation Models eJournal*, 2011.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015. URL `http://www.R-project.org/`.

P. Teetor. Better hedge ratios for spread trading. Unpublished note, 2011. URL `http://quanttrader.info/public/`.

G. Vidyamurthy. *Pairs Trading - Quantitative Methods and Analysis*. Wiley, 2004.