Value-add in Technical Analysis on the JSE Bond Market

Zavier Haddad

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

May 16, 2017

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy in the University of the Cape Town. It has not been submitted before for any degree or examination in any other University.

Signed by candidate

May 16, 2017

Abstract

Trading on the JSE Bond Market is still done in an archaic fashion when compared to the highly digitalised trading done within the equities markets in South Africa, indicating there is less market efficiency within bond trading. Technical analysis relies on market inefficiencies to achieve an informational advantage and so there could be technical analysis based trading opportunities within bond trading. Bollinger Bands are one of the more prominent technical analysis methods. In this dissertation they are used in trading simulations to generate buy and sell signals in order to test if there is any value-add in their implementation. The dissertation attempts improve Bollinger Band based trading in two ways. The first involves attempts to more accurately estimate the underlying distribution of the time series, that is assumed to be normal in the standard methodology. It is shown that no additional benefit is derived from the alternative distribution estimation methods. Bollinger Bands make an assumption of stationarity on the time series on which they are implimented and so the second attempt at improved accuracy addresses this notion. Cointegration is used to generate linear combinations of bonds that are stationary, leading to more accurate application of the Bollinger Bands. The stationary combination of bonds produces positive results from the trading simulations, primarily within the combinations that are generated from a linear combination of less bonds and that posses larger variation. Not considering the liquidity assumtions, the positive results show that there is value-add within specific technical analysis based trading strategies.

Acknowledgements

A large debt of gratitude is owed to Obeid Mahomed for the many hours spent helping in the completion of the dissertation.

The extrenal supervisors, D. Polakow and G. Van de Linde of Precient Wealth, are also thanked for their input and proposing the topic.

The last thanks goes out to the Masters in Mathematical Finance class of 2015 for making the year thoroughly enjoyable and a special thanks to a specific few who provided me with help on countless occasions.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Problem Contextualisation

The topic of this dissertation was proposed and supervised by an industry professional and, as a result, the approach and contents are more practical than theoretical in nature. The topics of Bollinger Bands and cointegration within the JSE Bond Market are explored.

The two ideas are often referred to within the field of technical analysis, which makes use of statistical techniques in an attempt to grant a practitioner an informational advantage as to the future movements of certain financial instruments. The advantage will hopefully allow him to beat the market and make a profit over and above other low risk investments.

It is harder to achieve and informational advantage in more efficient markets as information disseminates between participants more effectively. South Africa's equity market has seen massive growth in the volume of trades year on year (Jse.co.za, 2016b) indicating the market is becoming more efficient. The technically based trader is left with less opportunities within the equities market thus increasing the incentive to look into other markets to apply his/her techniques.

One of the potential markets is the JSE Bond Market. Unlike the JSE Equity Market, studies on its efficiency are few and far between as High and Honikman (1995) state in their study of capital markets. It, together with the study by Liu (2013), conclude that the South African bond market is likely to posses weak form market efficiency and hence there should be no value-add found in technical analysis based trading strategies. The more recognised work of High and Honikman is however, very dated and neither sources are widely referenced. The mentioned limitations detract from their suggestions that no value-add can be achieved.

An encouraging factor is the manner in which instruments are traded within the bond market. It is relatively archaic when compared to the digitalised and on screen trading that is the norm within equities markets in South Africa. For market participnts to buy and sell bonds they would usually have to pick up the telephone and call a fellow market participant. This lack of technological advancement in the bond trading channels indicates inefficiencies exist as shown by Kwon and Kish (2002) in their study on market efficiency on the New York Stock Exchange. The implementation of new systems will inevitably change this in future, heightening the urgency of capitalising on potential opportunities technical analysis offers.

The techniques used in technical analysis make bold assumptions about certain price processes to justify their use. One of the stronger assumptions made is that certain time series are stationary. An example of this is in the famed Bollinger Bands, which Polakow (2010) critiques on their statistical inaccuracy and suggests improvements to their construction.

In this dissertation the performance of Bollinger Bands and cointegration based trading strategies are examined and it is show that better results are achieved by being more accurate and relaxing some of the assumptions involved in their implementation.

1.2 Aims

The topic of the dissertation is broad and so to provide greater clarity, the main aims of the dissertation are summarised below:

1. Investigate the level of value-add within the technical analysis fields of Bollinger Bands and cointegration within the JSE Bond Market.

2. Quantify the performance of trading strategies that use the methods and isolate the best performers.

3. Investigate if greater value-add is achieved by being more mathematically accurate and making less assumptions within their application.

With the primary aims articulated above, the structure by which the the problem is tackled can now be described.

1.3 Structure

The dissertation begins with a brief description of the data set used throughout the study. Details on the instruments and the justification of their use is then presented. The price series, named the corrected all-in price (CAIP), is generated from historical bond yields and is introduced as the primary historical time series for testing the trading strategies. Some of the subtleties, limitations and slight data alterations are also highlighted.

With the data set better understood some of the theory involved within the dissertation is presented. First the theory behind the Bolliger Bands and some adaptations to their construction is examined. Some aspects of cointegration are then looked into, as these ideas are used in generating the cointegrated time series that the relative trading strategies make use of.

The methodology behind the testing of the trading strategies is then described. Methodologies vary between absolute trading strategies (trading of one bond at a time) and relative trading strategies (trading of multiple bonds at the same time). The differences are explained to facilitate greater understanding of the results. The three primary metrics of performance; average excess returns, standard deviation of excess returns and average trades per year are then described.

The results that each of the trading strategies achieves within the performance metrics are graphically presented and used to determine their relative performances. First the poor results, achieved by the absolute trading strategies, are discussed and it is show that the different methods of Bollinger Band construction offer little in the way of improved performance. The results of the cointegration analysis are then presented, showing a high degree of cointegration within the data set. Positive relative trading results, which involve trading in linear cointegrated combinations of the 10 bonds, show the significance of relaxing the stationarity assumption.

The dissertation concludes by highlighting the areas in which value-add is achieved, namely within the relative trading strategies using cointegration to determine the combinations of bonds. It is shown that higher variation historical price series, made up of a combination of fewer bonds, produce the best results and suffer the least from liquidity and trading cost assumptions and thus offer the best potential trading opportunities. In implementing Bollinger Bands more accurately, by applying them to stationary linear combinations of cointegrated time series, it is shown that better results are achieved.

With the structure of the dissertation now better understood some aspects of the data used are presented.

Chapter 2

Data

In this chapter the data used within the dissertation is examined, all of which is obtained from Bloomberg. The chapter aims to better prepare the reader in understanding the relevance and limitations of the results to come.

What follows is a brief description of the instruments, creation of the corrected all-in price and the justification for its use. The chapter concludes with a description of the slight data alterations made to the data set.

2.1 The Instruments

The instruments that are the subject of this dissertation are the top ten liquidly traded South African government bonds within the JSE Bond Market. A complete list of the bonds appears in the key of Figure 2.1.

The selection of the bonds is based on their liquidity. Government bonds account for 90% of the liquidity in the JSE Bond Market (Jse.co.za, 2016a) and profiting off technical analysis requires adequate levels of liquidity. Multiple trading simulations are run, which involve many instances of buying and selling, so it is important that in reality the bonds are easily traded. Using the liquid bonds ensures the validity of the findings is not too severely compromised due to the liquidity assumptions that are made.

The technical analysis methods this dissertation is concerned with are more prominent within the equities trading space and so it is important to understand some of the basic differences between bonds and stocks before the analysis starts.

The bonds all pay a fixed coupon twice a year and the price of the bond is determined by the JSE-Bond Pricing formula, which backs out a price from a yield together with other fixed, bond specific variables. The nature of the instruments leads to their all-in price ceasing to fluctuate and liquidity significantly dry up as the bond approaches maturity where a fixed pay-off is guaranteed. No bond data close to maturity is used within the trading simulations to avoid these problems.

Bonds in South Africa trade on yield, with prices being quoted as spreads over the benchmark government bond, the R186. It could appear that the yield is thus the best time series to apply the technical analysis based trading rules to however, the corrected all-in price is preferred as the following section explains.

2.2 Corrected All-In Price Verse Yield as an Indicator

The reason the bonds yields are not used directly for the trading rules is that changes in the yield do not linearly translate into changes in the monetary value of the traded instruments. Traders ultimately care most about the monetary change and so any trading rule should seek to base its buy or sell signals around a monetary change and not a change in yield.

If one examines the non-linear relationship between all-in price and bond yields, this is easily understood. It is clear that the monetary change associated with a change in yield is a function of the initial value of the yield and hence less favourable trade signals could be presented should changes in yield be used as the basis of making trading decisions.

Another point that corroborates this sentiment is that the technical analysis methods this dissertation is concerned with are more prominent in equities trading which uses the share price as the indicator for trading decisions. A share price represents a direct monetary value and thus changes in it are what an actual investor cares most about.

The aforementioned points justify the construction of the historical corrected all-in price time series for the 10 bonds in question, and its use within the trading simulations to come.

2.3 The Back Testing Dataset

From the time series of historical yields a theoretical historical corrected all-in price time series is constructed, with the assumption that coupon payments are re-invested into the bonds. The new time series is named the ”corrected all-in price (CAIP)” and is the price series used in the trading simulations to produce the results within the dissertation. The re-investment of coupons leads to the upward trend of the time series, as seen in Figure 2.1 and the large discrepancy between prices at similar times. Results are unaffected by this as they are calculated on a returns basis. The relative changes in the CAIP during a specific trade’s life is what determines the performance of a specific trade.

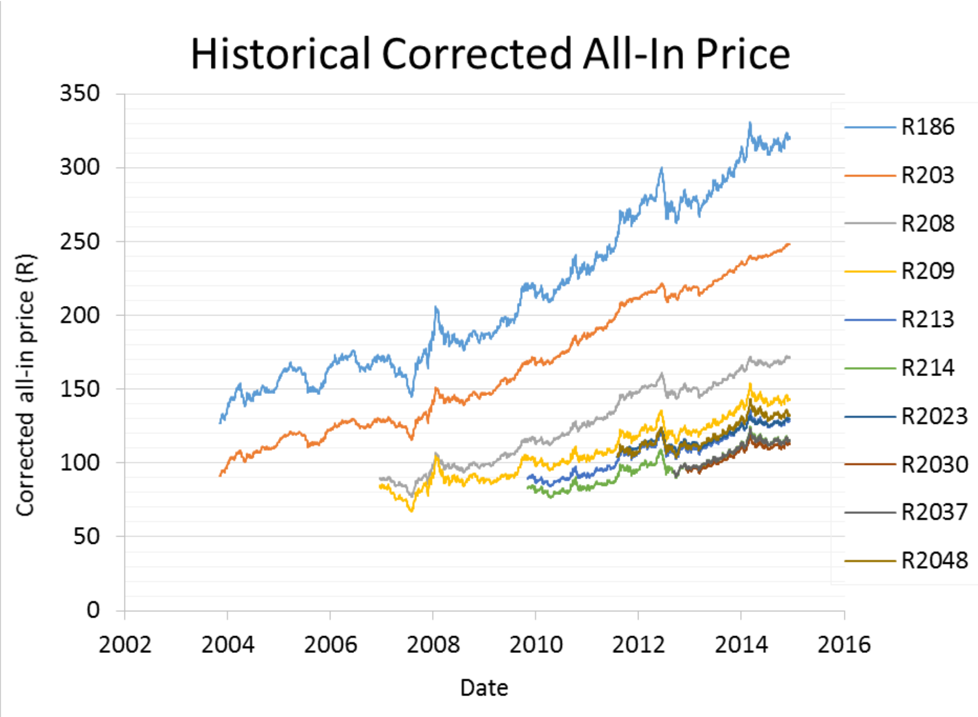

Fig. 2.1: Historical time series of the corrected all-in price for the 10 bonds

An attribute of the dataset Figure 2.1 makes apparent is the different time periods the data for each specific bond spans, due to their different issue dates. In the trading simulation results this has sample size implication as described in greater detail when the results are presented.

The final part of the data set used is a collection of historical yield curves. Discounting and capitalising within the trading simulations is done using this data. Where maturities do not line up, direct linear interpolation is used to obtain an approximate rate for discounting or capitalising.

2.4 Slight Data Alterations

To achieve accurate results, the data from the 10 bonds has to be consistent before testing starts. If trading strategies are to be run with combinations of many bonds, as is done in the relative trading simulations, there has to be a price for each bond at each potential trade date or certain trades on problematic days would lead to errors in the trading model. For this reason slight alterations are required in many of the historical price series.

The methodology behind the changes revolves around keeping the dates consistent with the benchmark bond, the R186, due to it being the most actively traded bond in the set of instruments (Pitsillis and Taylor, 2015). In the case that a day of data is missing a simple linear interpolation is used to approximate a value for the yield of the bond using the two adjacent days. In the case where the R186 has a day of data missing, the corresponding data points are removed from the other instruments.

The changes are performed on very few discrete dates, for each time series of CAIPs and so the structure of the data is not notably compromised. The outcome being a set of completely time consistent traded CAIPs for all ten bonds in question, allowing the testing of trading strategies to proceed with less potential of data based errors.

For some of the bonds initial trade data from the date of issuance is sparse, leading to data only being taken from a later date when more frequent data is present. By doing this extensively interpolated data is avoided, leaving less room for inaccuracy in results.

Having described the data in more detail, focus is now moved to some of the theory involved in the implementation of the trading strategies.

Chapter 3

Theory

Important aspects of the theory involved in the technical analysis used in the dissertation are looked at in further detail within this chapter.

The section starts with a brief overview of how Bollinger Bands work. Two different methods of their construction are then examined, before concluding with some theory behind cointegration.

3.1 Bollinger Bands

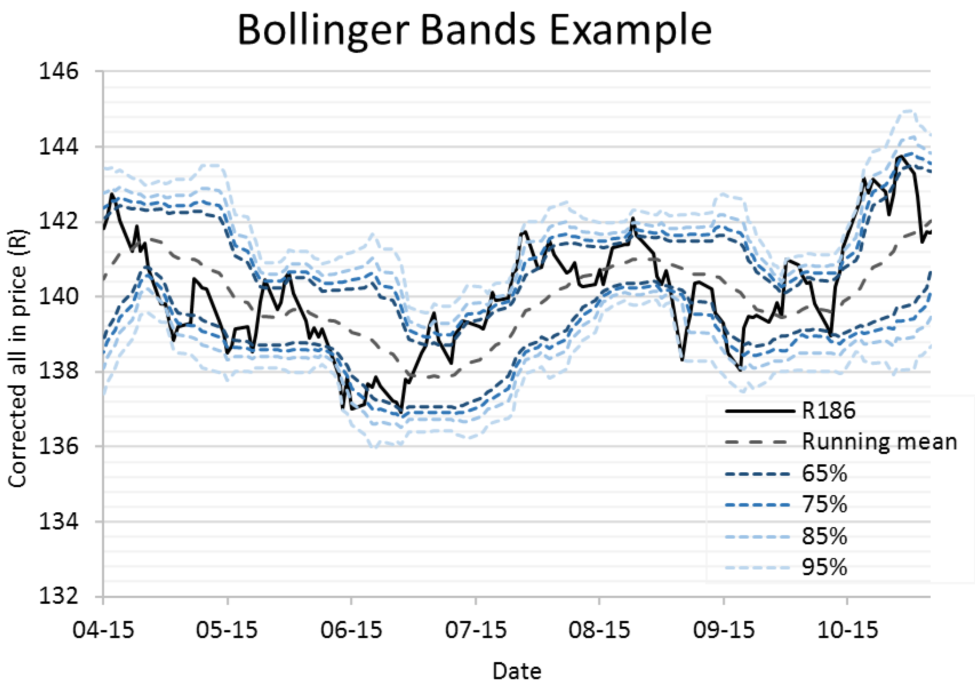

Fig. 3.1: An example of Bollinger Bands plotted around the corrected all-in price of the R186 bond.

Standard Bollinger Bands are constructed by taking a fixed number of the previous values of the times series to generate a probability distribution of the next value, the normal distribution in the case of Bollinger Bands. The distribution is used to determine upper and lower confidence bound for the following value in the time series, based on a specific confidence level. As the series progresses forward through time the rolling sample that is used to construct the distribution changes and thus so do the values of our upper and lower bound that tend to follow, and for the most part, envelope the actual price series from which the sample came.

An example of this is presented in Figure 3.1 where Bollinger Bands are plotted together with the 20 period running mean around the R186 benchmark government bond. Different shades of blue represent Bollinger Bands constructed using various confidence bounds.

Crossing the bands serves to inform a trader of large unexpected moves in the price series that should indicate that the current level is too high or too low and that it should return to normal levels in the short term future. The crossings would be either a buy or a sell signal in the context of this dissertation.

The bands are constructed at various confidence bounds depending on risk preference. Figure 3.1 shows 4 of these confidence bounds. A larger confidence bound means that fewer trades are entered into, at levels that are further away from the running mean. The reason for this, as Figure 3.1 shows, is that the larger the confidence bounds the less times the indicator is expected to breach the bands and trigger a buy or sell signal.

Many assumptions are made when implementing Bollinger Bands. Two of the more notably incorrect ones are the assumptions on the distribution of the time series and the assumption that the time series is stationary. The dissertation looks to improve upon upon these two assumptions in order to increase the performance of the Bollinger Band based trading strategies.

Despite the large assumptions, Bollinger Bands still enjoy a large amount of use in industry. The following sections explore the standard method of determining the distribution of the price series, before looking at some of the theory around potential improvements to the distribution estimation.

3.2 Standard Method

The first, and standard method of constructing the Bollinger Bands tested is that developed by Bollinger (1992) in which it is assumed that the time series is stationary and sampled from a normal distribution. The method constructs bands based around the price series by constantly using a fixed number of previous values of the time series to estimate the mean and variance of the assumed normal distribution. The results of this are then used to plot confidence bounds around the running mean of the time series. Breaches in either the upper or lower bands indicate an unexpected move and generate a trade signal.

3.3 Empirical Distribution Method

In this method of constructing the Bollinger Bands an empirical distribution is generated from a constant running number of previous values in the time series. The empirical distribution is used instead of the normal distribution used within the standard Bollinger Bands construction. The resultant new distribution is used to construct confidence intervals around the time series, similarly to the method described in the previous section.

The Matlab function ”ecdf” is used, which takes as an input, the vector of historical price realisations, $x_t, x_{t-1}, ..., x_{t-n}$. The assumptions that all the realisations are equally probable is made and the points used to generate an empirical cumulative distribution function, $P(z)$. It is this distribution that is then used to determine and plot the respective upper and lower bands at each point in time, for the specified confidence interval.

3.4 Kernel Soothing Method

The primary aim of kernel fitting methods is to construct a smooth curve from a sample of values from a random variable that best approximates the probability density function. The smooth approximation is then used instead of the normal distribution from the standard Bollinger Band methodology.

The method makes use of the Matlab function ”ksdensity”. The inputs to this function at time t are a set of previous values of the time series, $x_t, x_{t-1}, ..., x_{t-n}$ where n is the length of the running sample.

As is shown by Alexander (2008b),the kernel approximation to the density of a random variable X can be defined as:

$$f_h(x) = (nh)^{-1} \sum_{i=1}^{n} K(u), \text{ where } u = \frac{x - x_i}{h} \tag{3.1}$$

Here K and h are the kernel function and the bandwidth respectively. In the dissertation and within the Matlab function the kernel function is a normal density function and hence a Gaussian kernel is being used. The bandwidth, h , is chosen such that it minimises the errors between the empirical and and newly fitted densities.

3.5 Cointegration

Cointegration is concerned with finding combinations of time series that, when combinated linearly in specific proportions, result in a net time series that is stationary. Generally there is some form of common trend within the data that is removed though the weightings of the comprising time series. Cointegration is different from correlation in that cointegration suggests long term asset prices are unlikely to diverge. Correlation is concerned with short term movements in one indicating an immediate movement in the other (Chan, 2009).

Stationarity of the price time series is an assumption the Bollinger Bands construction procedures assumes in order for the statistical inferences made about the distributions to hold. It is a desirable attribute to posses when considering candidates for the Bollinger Band based trading as it would lead to more accuracy within the method. It is for this reason that, when considering relative trading, the metric for determining the weightings of the various assets is their level of cointegration. If cointegration is present then a net stationary price series is theoretically achievable.

When dealing with multiple assets there could exist levels of cointegration among certain combinations of them. The challenge is finding the weighting factor, if any, that lead to this desirable property.

There are two predominant methods for determining the factors and testing for cointegrated relations. The method outlined by Engle *et al.* (1993) and the other by Johansen and Juselius (1990). The EngleGranger methodology is more suited to pairs of cointegrated variables as when dealing with 3 or more time series the factors found do not necessarily correspond to the combination that is most stationary. In contrast, the Johansen methodology seeks to find the most stationary combination given multiple cointegrated variables. Seeing as the dissertation deals with combinations of up to 10 bonds, Johansen's methodology is superior.

An Overview of the method followed by the Johansen test is presented below, closly following the description outlined by Alexander (2008a).

It starts by considering a set of n integrated variables $X_1, ..., X_n$. In the context of the dissertation these will be the corrected all-in prices of the 10 bonds. The data can be represented in vector autoregressive representation form as follows:

$$
\begin{aligned}
&X_{1t} = \alpha_1 + \beta_{11}X_{1,t-1} + ... + \beta_{1n}X_{n,t-1} + \epsilon_{1t} \\
&\vdots \\
&X_{nt} = \alpha_1 + \beta_{n1}X_{1,t-1} + ... + \beta_{nn}X_{n,t-1} + \epsilon_{nt}
\end{aligned}
$$

Defining the following matrices the above system of equations is expressed in matrix form as follows.

$$
X_t = \begin{pmatrix} X_{1t} \\ \vdots \\ X_{nt} \end{pmatrix}, \alpha = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{pmatrix}, B = \begin{pmatrix} \beta_{11} & \cdots & \beta_{1n} \\ \vdots & \ddots & \vdots \\ \beta_{n1} & \cdots & \beta_{nn} \end{pmatrix}, \epsilon_t = \begin{pmatrix} \epsilon_{1t} \\ \vdots \\ \epsilon_{nt} \end{pmatrix}
$$

$$
X_t = \alpha + BX_{t-1} + \epsilon \tag{3.2}
$$

Expressing this in the form of 1 period changes in X_t and defining the new variable Π , where I represents the $n * n$ identity matrix results in the following:

$$
\Pi = B - I
$$

$$
\Delta X_t = \alpha + BX_{t-1} - X_{t-1} + \epsilon_t
$$

$$
\Delta X_t = \alpha + \Pi X_{t-1} + \epsilon_t \tag{3.3}
$$

The next step is to augment the above equation with sufficient lagged dependant variables so at to allow the residuals to be free of auto correlation.

$$
\Delta X_t = \alpha + \Pi X_{t-1} + \Gamma_1 \Delta X_{t-1} + ...\Gamma_q \Delta X_{t-q} + \epsilon_t \tag{3.4}
$$

It now follows that ΠX_{t-1} is stationary due to the fact that each of the variables $X_1, ..., X_n$ is integrated and hence each equation in the above set of equations has a stationary dependant variable. The result is that the right hand side of the equations must also be stationary.

The next step in the procedure is to test the rank of matrix Π as if the rank is 0, no insight into the relationships between $X_1, ..., X_n$ is gained. If on the other hand the rank is $r > 0$ then there will be r independent linear relations that will be stationary, implying cointegration is present.

If it is found that the matrix Π has rank r then it can be expressed in row reduced Echelon form as follows, where there will be r non zero rows.

$$\Pi = \begin{pmatrix} 1 & \theta_{12} & \cdots & \cdots & \theta_{1n} \\ 0 & 1 & \theta_{23} & \cdots & \theta_{2n} \\ 0 & 0 & 1 & \cdots & \theta_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \tag{3.5}$$

Each non zero row now defines a co-integrating relation by containing the coefficients of the stationary processes as follows.

$$\begin{aligned} Z_1 &= X_1 + \theta_{12}X_2 + ... + \theta_{1n}X_n \\ Z_2 &= X_2 + \theta_{23}X_2 + ... + \theta_{2n}X_n \\ &\vdots \end{aligned}$$

The dissertation is concerned with the result if it returns a cointegrating relation containing all of the fed in components as all possible combinations are tested one by one and if one of the factors is zero then the result represents a different unique combination. In the above system of equations this is denoted as Z_1 and as it can be seen to be comprised of a linear combination of all the input components $X_1, ..., X_n$.

For greater clarity an example is presented below in Figure 3.2 where the resultant time series is plotted along with the three bonds used to generate the net stationary time series. The figure uses different y-axes for the combination and single time series and shows the general upward trend of the individual bonds has been removed within the cointegrated combination of them.

The equation and factors used to generate the stationary linear combination are presented in equation 3.6 where X_1, X_2 and X_3 represent the historical corrected all-in prices of the $R209$, $R213$ and $R214$ bonds respectively.

$$Z = -2.89X_1 + 2.07X_2 + 1.32X_3 \tag{3.6}$$

An important aspect of the cointegrated series, when compared to the single bond series, is the differences in their relative changes. As a percentage of its initial value, the change in the conintegrated time series value is far greater than the changes in the single bond time series. As a result of this it is expected that the standard deviations of the returns on trades executed on the cointegrated series will be much larger than those executed on the individual bond price series.

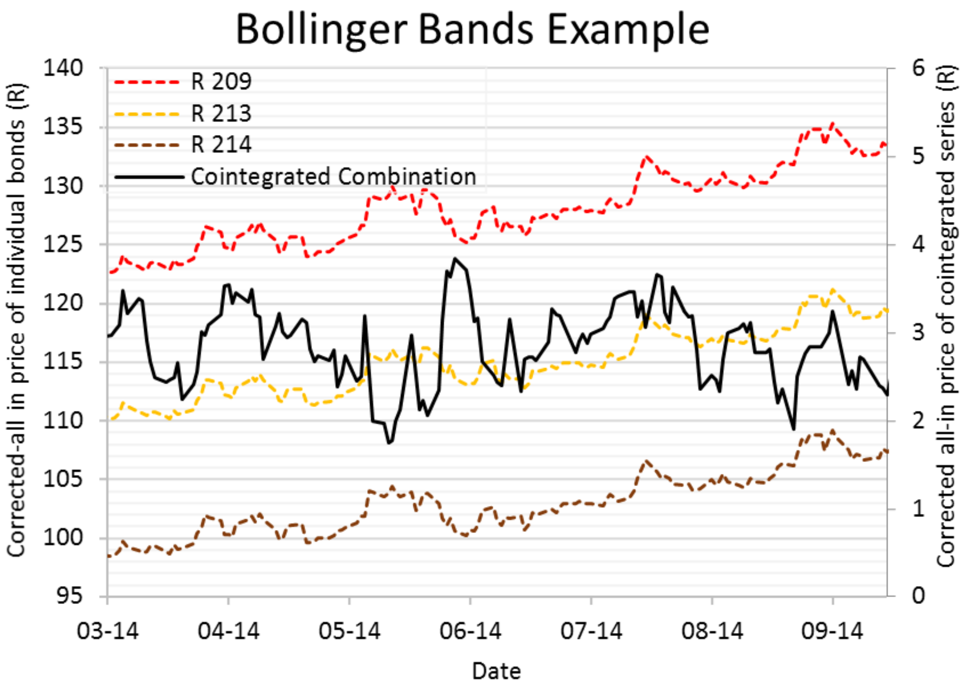

Fig. 3.2: An example of the resultant cointegrated linear combination of bonds.

In the results section the difference in the relative changes is apparent in many of the generated cointegrated combinations as it is the reason the standard deviation risk metrics see significant increases when dealing with cointegrated series.

With a better understanding of the theory at hand the dissertation progresses onto the methodology section.

Chapter 4

Methodology

The methodology undertaken to produce the results of the dissertation is described in this chapter. An overview of the 4 main steps in the testing procedure are first unpacked. The section then leads into the results chapter by defining the three performance metrics, displayed in the results section, and briefly summarising the testing methodology.

The path followed to generate the performance metrics for absolute and relative trading strategies differs in a few key areas, as is explained at a later stage.

4.1 Testing Methodology

The four main steps required to generate the 3 performance metrics, necessary to assess the value-add in the methods considered, are outlined below.

4.1.1 Indicator Selection

The indicators serve as the base time series around which the Bollinger Band trading simulations are run. Different indicators are used for absolute trading and relative trading strategies.

Absolute Trading Indicators

For the absolute trading simulations the corrected all-in price for the individual bonds is the indicator series and used to test the trading strategies on the bonds one by one. This corresponds to trading in one of the considered bonds at a time, and thus participating in an absolute trading strategy.

Relative Trading Indicators

The relative trading indicators are chosen through their cointegration properties. This involves trading in long and short positions of multiple bonds at the same time and is thus termed relative trading.

The cointegration analysis finds only 1 set of factors (associated with the most stationary linear combination) for each unique combination of bonds so the testing is confined to this combination when running trading simulations. One unique possible net price series for each unique combination is thus tested.

The results are grouped into categories depending on how many bonds are used in the linear combination. Using formula 4.1 below the number of possible results are obtained and listed in Table 4.1 below.

For $n = 10$ and $k = 2, 3..., 10$

$$\frac{n!}{k!(n-k)!} = \binom{n}{k} \tag{4.1}$$

Tab. 4.1: Table containing the number of possible unique combinations

Number of containing bonds (k)	2	3	4	5	6	7	8	9	10
Number of unique combinations	45	120	210	252	210	120	45	10	1

Table 4.1 shows that the number of possible combinations for each value of k (number of unique component bonds in the linear combination) vary significantly and thus it may seem unwise to group the results in this way.

The above opinion is false as the dissertation seeks to produce realistic results, and the liquidity assumptions and trading costs associated with trading in a combination of 2 bonds is very different form those pertaining to trading in a combination of 9 bonds. For this reason it is pragmatic to group series with common values of k together as they are more likely to share the aforementioned costs and limitations. Large liquidity and trading cost assumptions in the higher number of component bond categories distorting the more realistic assumptions in the lower categories is hence avoided.

The matlab function ”nchoosek” is used to output all the possible unique combinations of the bonds in the 9 categories. Price series associated with these combinations are then tested using the Matlab function ”jcitest” to determine in which combinations cointegration is present. Bond combinations that pass the test for cointegration are used in back testing relative trading strategies, further explained in the back testing section of this chapter.

4.1.2 Trading Band Construction

The methodology behind the Bollinger Band construction is now looked into.

The theory behind the bands construction is found in the theory chapter. The ideas from that chapter are used to generate symmetric (around the running mean) trading bands for a range of confidence bounds from 65% up to 95% in intervals of 2.5% for each tested indicator. Three unique time series bands are the result of this. The upper, lower and running mean bands associated with the upper confidence bound, lower confidence bound and the running mean respectively. All of these being computed using the running sample of the nearest 20 adjacent historical time values as this is the recommended look-back length (BCM, 2016).

The crossing of the bands defines buy and sell signals for the trading strategies to be used within the back testing simulations.

4.1.3 Back Testing the Trading Strategies

In the dissertation the theoretical trader is confined to two actions. He/she may either buy or sell an asset, with short selling permitted, triggered by the indicator time series crossing one of the constructed Bollinger Bands.

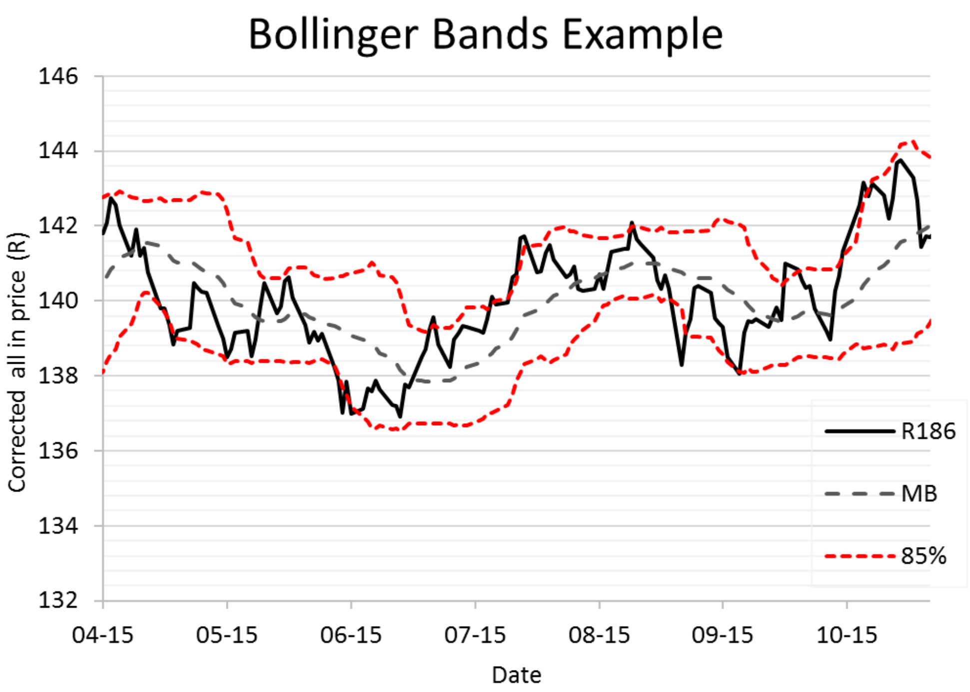

Fig. 4.1: An example of historical Bollinger Bands plotted at the 85% confidence bounds together with the moving average.

The two types of trades this leads to are explained with the aid of Figure 4.1. In the long only trading strategies the long position in the price series in entered into when the bottom band is crossed as opposed to the short only trading strategy, that is entered into when the upper band is crossed. In both cases the position is closed out when the time series crosses the moving average, indicating a return to normal price levels.

Back testing is a vital stage in the dissertation as it directly produces the results from which the conclusions are made. For this reason a large amount of time is spent testing and coding up the back testing framework so as to achieve reliable results.

The back testing methodology for the absolute and relative trading strategies are similar but there are distinct differences.

Absolute Trading Back Testing

Testing for the absolute trading strategies is fairly straight forward. The trading rules are applied and the simulations run over the entire data set for all bonds to achieve a final set of time indexed excess returns on all executed trades. The results of this are used to compute the three metrics by which we evaluate the performance of the trading strategy.

Relative Trading Back Testing

The relative trading strategies back testing methodology is a little more complicated as the actual historical time series is used to generate the factors that are used to produce the linear combinations.

When discovering cointegrating relations it makes sense to use the entire length of the time series concerned but when back testing this would lead to unrealistic results. To find the factors requires a certain length of the time series and as a trader one would not have access to values in the time series at future times. For this reason, to find these factors the historical time series is divided into half and the first half (further back in time) used to compute the factors for each bond in the combination and the remaining half used to then back test the trading strategy. The dividing of the time series into two keeps the testing strategy in line with what is achievable in reality.

Once the time indexed trade results of the relevant sections of the linear combinations have been produced by the trading simulation performance metrics are output.

4.1.4 Evaluate Performance

In evaluating the value-add of the trading strategies their performance is quantified in the form of a series of performance metrics. The three metrics are defined below.

1. **Average Excess Returns**

   The mean of the returns generated by the discounted or capitalised cash flows, from the profit and loss simulation, is named the average excess returns. It is the first of the three performance metrics.

   Long term average returns over the risk free rate are captured, in a standardised and comparable fashion within this metric. Any potential fund manager or trader would value this sort of metric as it gives an indication of the long term performance of the strategy.

   Problems arise when traditional formula are used to compute returns within the results, as the dissertation contains short strategies as well as strategies involving time series that start positive and breach zero, to end up negative through the linear combinations of short and long positions in the cointegrated historical price series. For this reason a new modified returns metric is defined.

   The price of the asset, or linear combination of assets, is defined at t_0 and t_1 as Z_0 and Z_1 respectively. In the case where the trading strategy involves a cointegrated linear combination of assets these will represent the the net price of this linear combination of both positively and negativity weighted assets. The variables Z_0 and Z_1 can thus be negative or positive.

   The discount factor associated with the risk free rate, applicable over the period from t_0 to t_1, we denote $D_{0,1}$. For the purpose of this dissertation $D_{0,1}$ is derived from historical Jibar yield curves as described in the data chapter.

   Two portfolios are used in the argument to generate the metric of returns over the risk free rate. An active portfolio and a standard risk free bank account portfolio.

   The risk free portfolio consists of only a bank account which is denoted B_0 and B_1 at times t_0 and t_1 as respectively.

   The risky portfolio P also contains a bank account component, $\widehat{B}$, as well as positions in the risky assets.

   The argument starts with $|Z_0|$ in the risk free bank account portfolio which is capitalised to time t_1 at the risk free rate.

$$B_1 = \frac{|Z_0|}{D_{0,1}}$$

The active portfolio for the long trading strategies also starts with $|Z_0|$ in the bank account and then takes a long position in Z_0 (can have negative value). The remainder is then stored in the bank account (not the same bank account as the passive risk free portfolio). The bank account component of the active portfolio thus has the following value at t_0.

$$\widehat{B_0} = (|Z_0| - Z_0)$$

The active position is closed out at t_1 with the following net value.

$$P_t = \frac{(|Z_0| - Z_0)}{D_{0,1}} + Z_1$$

Where what is in the bank is capitalised by the by the risk free capitalisation factor. The final value of the active portfolio is divided by the final value of the risk-less portfolio and subtracted by 1 to achieve a measure of returns over the risk free rate for long trading strategies, R_l suitable for a price series that can go both negative and positive.

$$R_l = \frac{P_t}{B_t} = \frac{(|Z_0| - Z_0)/D_{0,1} + Z_1}{|Z_0|/D_{0,1}} \tag{4.2}$$

A similar argument is used to achieve the equivalent result for short trading strategies seen below.

$$R_s = \frac{(|Z_0| + Z_0)/D_{0,1} - Z_1}{|Z_0|/D_{0,1}} \tag{4.3}$$

The metrics above measure the returns for one instance of the trading strategy. The complete set of results are averaged over the life of the trading strategy to achieve a measure of average excess returns over the risk free rate.

The metric is malleable enough to be applied to all the historical price series that are dealt with in this dissertation as well as quantifying a measure of value-add in the sense that it gives a measure of the performance for the trading efforts verse a zero effort strategy of investing in a bank account.

Standardising the performance metrics in this way leads to more comparable results and more hence meaningful conclusions.

2. **Standard Deviation**

   The average excess returns metric gives an indication of the average performance achievable by a trading strategy, but it does not give an indication of the riskiness of the trading strategy. For this purpose the standard deviation of the series of excess returns is used.

   Trading strategies would all not possess the same type of risk. A strategy yielding higher average excess returns but that possessed a lot more risk could be judged inferior depending on risk appetite of the practitioner. For the various trading strategies the standard deviation of the excess returns is presented to give an indication of how the riskiness of the strategy varies across the set of results.

   The standard deviation's value, in terms of how much information it conveys, is a function of the underlying distribution. If the excess returns are normally distributed more value will be taken from this statistic than if the underlying distribution follows some sort of non-symmetric, more strangely shaped distribution. Regardless of these facts there is a still a degree of value that is taken from the statistic as it speaks of how sparsely distributed the excess returns are and hence presents an indication of how risky the respective strategy is.

3. **Average Trades Per Year**

   Regardless of the levels risk and expected excess returns form a strategy, trading opportunities could present themselves too infrequently to warrant a strategy's use. With this in mind the last metric of performance is defined as the average number of trades per year that a specific strategy presents when back tested on the historical dataset.

   A strategy that has superior average excess returns, but only presents itself once a year would be judged inferior to one that yielded slightly lower average excess returns but presented itself far more often. Average trades per year quantifies this difference and is therefore a valuable metric when comparing the various trading strategies.

The three fundamental comparison metrics are now defined and so the evaluation methodology can be described before the results are presented.

4.2 Evaluation Methodology

The three primary performance metrics, defined above, are the foundations on which the trading strategies performances are assessed. A desirable trading strategy has high average excess returns as well as trades per year for the lowest possible amount of risk.

4.2.1 Evaluation of Absolute Trading Strategies

The evaluation of the absolute trading strategies is straight forward. For each method of Bollinger Band construction the strategy is run along all 10 bonds, at various confidence bounds, and the results combined to represent the complete set of results for the absolute trading trading strategies. Long only and short only strategies are tested separately with the results for the three performance metrics presented graphically. Results of the evaluation are found in the first section of the results chapter.

4.2.2 Evaluation of Relative Trading Strategies

Evaluation of the relative trading strategies is done a little differently. The results are grouped by the number of bonds within the linear combinations. This is preferable as the costs and limitations associated with taking positions in the same number of multiple bonds would be similar as previously explained.

The categories into which the results are grouped are labelled $C2, C3, C4...., C10$ which represent the collections of the net price series containing linear combinations of two, three, four all the way up to 10 component bonds respectively.

Care is also taken, within the relative trading back testing, to not use the same historical data to obtain the cointergration factors and perform the back testing as is previously described in the relative trading back testing section, within this chapter.

4.2.3 Summary of Dissertation Testing Methodology

The testing methodology is briefly summarised below, before the final results are presented.

1. **Indicator Selection** - The indicator time series, around which the trading bands are constructed, is selected. Either single or linear combinations of historical corrected all-in prices are used.

2. **Bollinger Band Construction** - The Bollinger Bands are constructed and the rules set up for the buy and sell signals associated with breaches in the bands. In long strategies a long position is entered into when the lower band is breached and the position closed out when the mean band is crossed. For short strategies a short position is entered into when the upper band is breached and the position is closed out when the mean band is crossed.

3. **Back Testing** - The strategies are back tested on the historical time series of the selected indicators. In the case of relative trading, the historical time series is split in half. The first half is used to compute the factors for the linear combinations. The second half is used for back testing to keep the results realistic.

4. **Evaluate Performance** - The three performance metrics are computed and presented in order to adjourn whether value-add is achieved achieved form the respective trading strategies and identify which have shown the greatest potential for further study and possible implementation.

Chapter 5

Results and Discussion

In this chapter the results of the dissertation are presented within their respective categories. First absolute trading strategy results are looked into, followed by the results of the cointegration analysis. The relative trading results are then presented together with results pertaining to the effect the price series variation has on the performance of the cointegration based relative trading strategies.

5.1 Absolute Trading Strategies

The absolute trading strategies involve using the individual corrected all-in prices for the bonds as the indicator time series for the trading signals, as the bonds are simulated to be traded one at a time.

5.1.1 Long Only Trading

The long only results, for the absolute trading simulations, are from the trading strategies that trigger a buy signal when the indicator breaks the lower Bollinger Band and then trigger a sell signal when the indicator crosses the running mean.

Initial results for the long only trading strategy show very low average excess returns across all three methods of Bollinger Band construction, as seen in Figure 5.1. This metric is seen starting in the negative and ending off slightly in the positive at the higher confidence bounds. In general the results don't come close to breaking the 0.5% return over the risk free rate. The trading costs are also not included and so from a average returns perspective, this trading strategy performs poorly.

The standard deviation gives an indication of how dispersed the excess returns achieved in the trades are. The relatively low standard deviations of around 2 to 2.5%, as seen in Figure 5.2, across the spectrum of confidence bounds indicates that there is a low chance of making large gains (or losses) on individual trades too. If

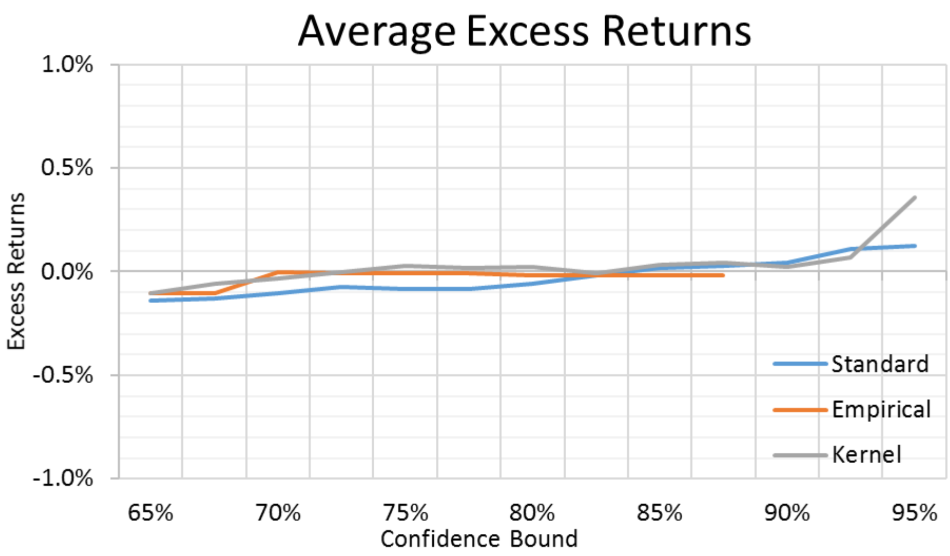

Fig. 5.1: The average excess returns for the long only strategy.

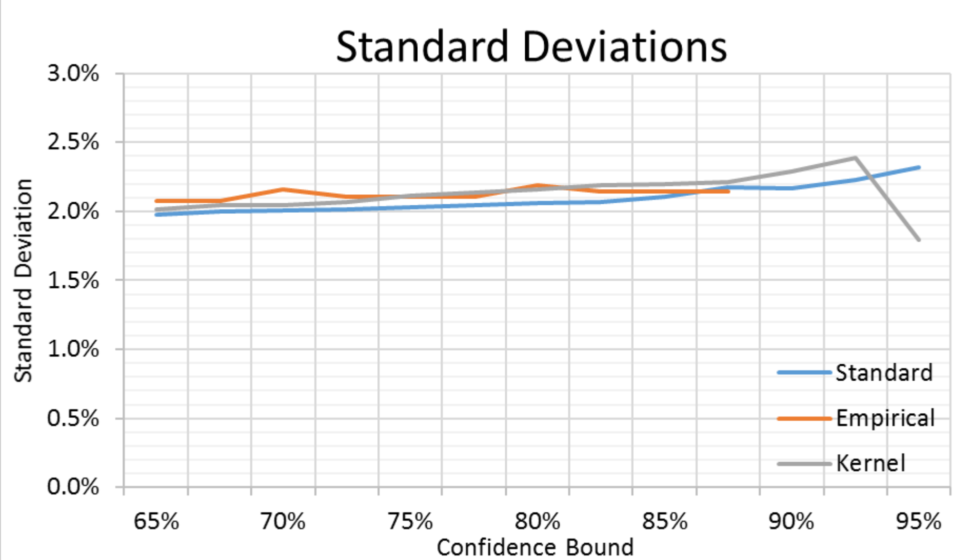

Fig. 5.2: The standard deviation of the excess returns for the long only trading strategy.

one examines Figure 3.2 and observes the small relative changes in the stand alone bond time series the low standard deviation levels achieved make sense.

Further damning evidence that this strategy is not worth pursuing is the low number of trading opportunities presenting themselves per year. Expectedly in Figure 5.3 the number of opportunities taper off slowly as we increase the confi-

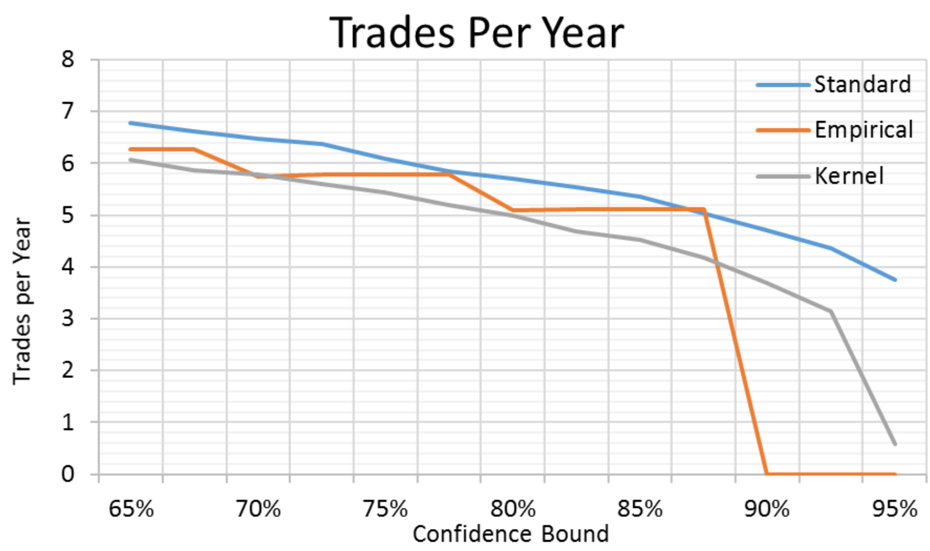

Fig. 5.3: The average trades per year for the long only trading strategy.

dence bound of the Bollinger Bands. In this metric of performance a notable distinction between the methods of band construction is observed, with the standard Bollinger Bands clearly the better performer especially as the higher levels of the confidence bounds are reached.

A breakdown of the empirical method of band construction at the 90% confidence bound is apparent, due to the nature of the pmf method, whereby it does not have continuous long tails. Due to there only being 20 data points in the sample, and the current value being included in the sample, the Bollinger Bands are constructed so that they take on a value that is outside the range of the input sample. As this method uses a pmf generated from this sample the probability of breaching the bands becomes effectively zero and the number of trades goes to zero as seen in Figure 5.3.

In general the two attempts to improve the Bollinger Bands (Empirical and Kernel) have little effect on the results achieved in the average excess returns and standard deviations metric and lead to worse results in the trade frequency metric, with less opportunities presenting themselves due only to the method of band construction.

The set of results advocates the use of the standard Bollinger Bands as the preferred method of constructing the bands as the other provide no improvement. It is also clear form the poor results that the long only absolute trading strategy does not perform well and no value-add is achieved.

5.1.2 Short Only Trading

The short only results for the absolute trading simulations, are from the trading strategies that triggers a short sell signal when the indicator breaks the upper Bollinger Band and then triggers a signal to close out the position when the indicator crosses the running mean.

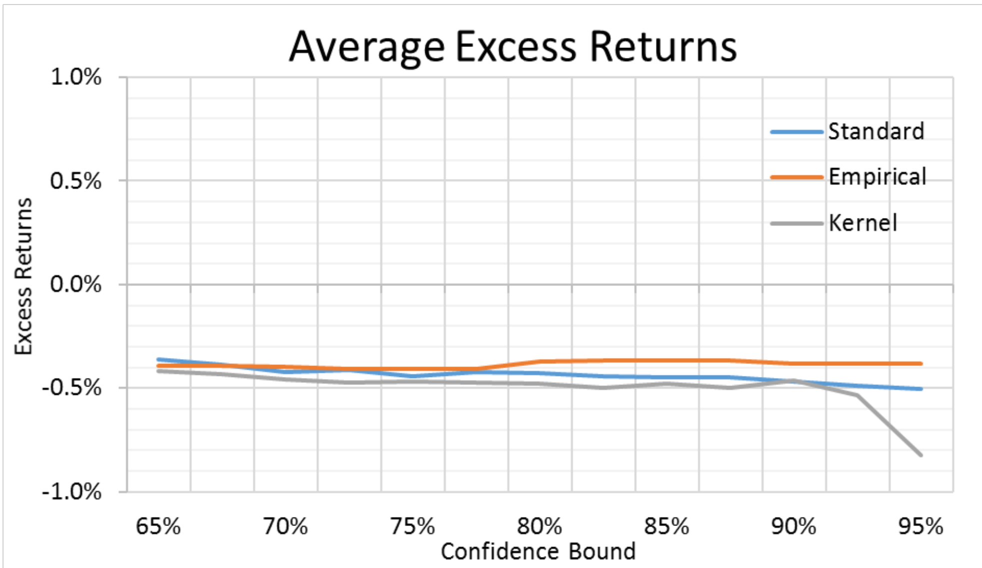

Fig. 5.4: The average excess returns for the short only strategy.

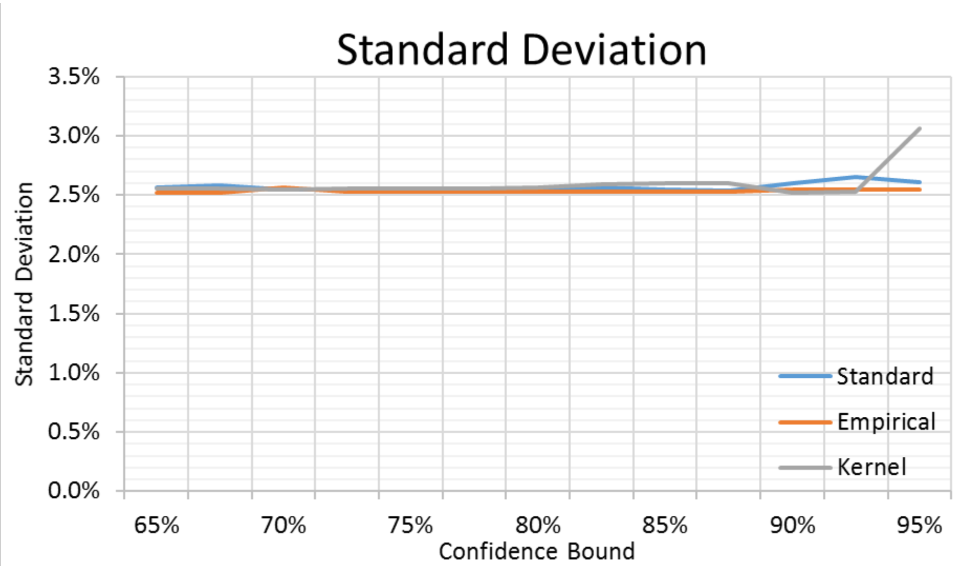

Fig. 5.5: The standard deviation of the excess returns for the long only strategy.

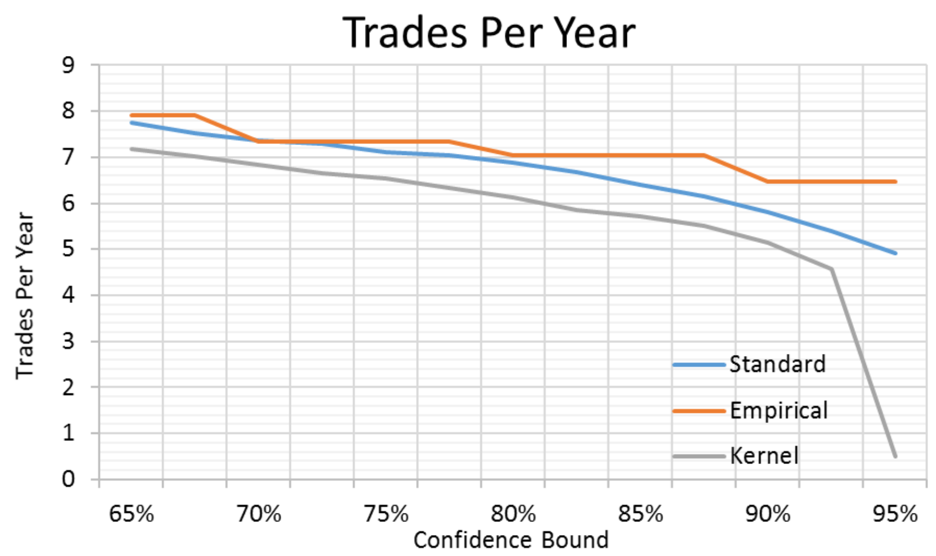

Fig. 5.6: The average trades per year for the long only strategy.

The results for the short only strategy average returns show poor levels are achieved as seen in Figure 5.4. Losses of around the 0.5% level are seen throughout the different confidence bounds.

The lower average excess returns achieved by the short strategy in comparison to the previous long only returns can be explained by the general upward trend of the bonds corrected all-in price series as is clear from Figure 2.1. In the short only strategy positive returns are made on downward movements in the price and thus a general upward trend is detrimental to the averaged effectiveness of the strategy over the long run.

Standard deviations remain at similar levels to the long only strategy as seen in Figure 5.5. Little in the way of deviation is observed as one moves across different confidence bounds again suggesting that even for the more risk prone investor this is a poor trading strategy.

The average trades per year are similarly in line with the previous results. In this metric the different band construction methods deviate the most, however the difference is not large enough to indicate notably better performance.

Overall the short only strategy shows worse performance than the long only strategy, which has already been shown to be poor.

Little value is again derived from the different methods of Bollinger Band construction, as seen by the closeness in the proximity of the results.

In terms of both sets of results pertaining to absolute trading strategies none indicate that there is any value-add achieved when trading in single bonds.

Having shown absolute trading of the bonds to be fruitless attention is turned to relative trading and the field of cointegration to generate a new set of trading strategies. The innovated methods of band construction are shown to not be effective at improving performance and hence are not used in the relative trading strategies sections to come.

5.2 Relative Trading Strategies

The relative trading strategies involve using linear combinations of the corrected all-in prices of the bonds as the indicator time series. Trading positions in multiple bonds is equivalent to this.

Before moving onto the results for the relative trading strategies the results for the analysis of the prevalence of cointegration within the 10 bonds are presented. Greater value will be taken out of the coming results as a consequence of this, as it highlights issues around the sample sizes of the tested data and the number of time series tested within each grouped collection.

5.2.1 Cointegration Analysis

The cointegration analysis across all possible bond combinations returns very positive results. As is shown in Figure 5.7, the high number of co-integration combinations, represented by the blue portion of the bars, indicates that there are many combinations of assets that pass the test and are used within the relative trading simulations.

The prevalence of cointegration as a percentage (indicated by the length of the blue bars in relation to the red bars) increases as the number of component bonds increases. The increase is expected as the likelihood of finding a cointegrated combination given two assets will be less than the likelihood of finding a cointegrated combination of more bonds with the same two assets included.

As the prevalence of cointegration increases as a percentage the net number of positive results obtained starts to decrease after the C5 mark due the net number of possible combinations decreasing. For this reason the most positive results obtained are in the C5 and C6 categories, before the diminishing effect starts to be become significant.

The combinations discovered, represented by the blue regions within the graph, serve as the data set of historical price series with which the relative strategies are tested. As they have passed the test the series have the desirable property of stationarity and thus the application of the Bollinger Bands is more accurate.

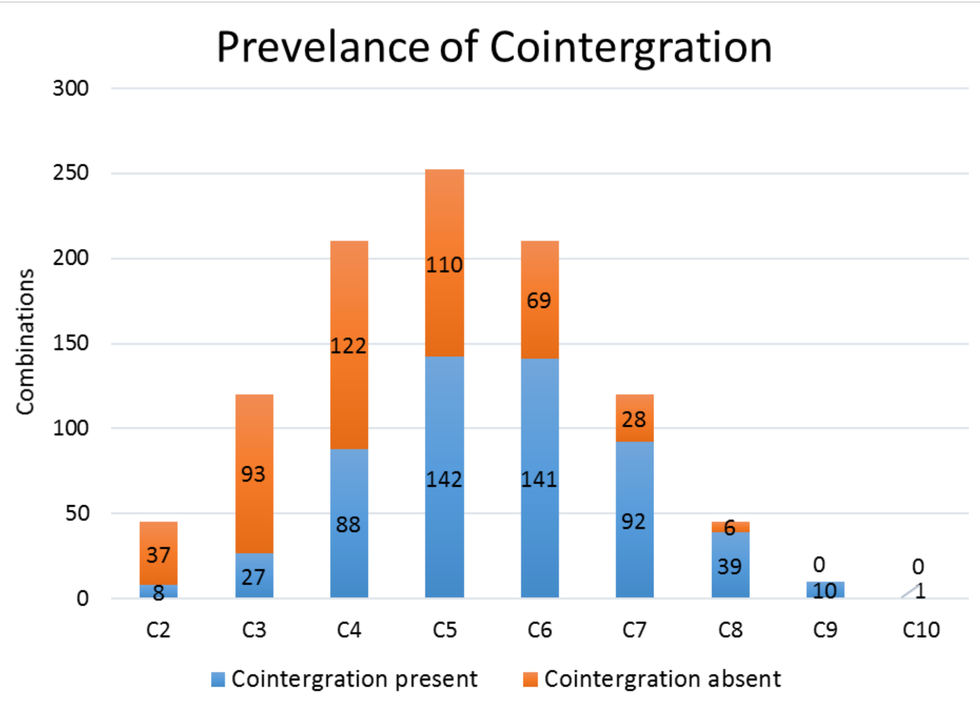

Fig. 5.7: Results of the cointegration analysis across all possible combinations of bonds

It is important to understand the two main factors affecting sample sizes within the collections of bonds that will be tested.

Firstly, as seen in Figure 5.7 the number of testable time series in each category varies and thus the number of trades executed is affected. For example, it is expected that more trades will be entered into within the C3 category than the C2 category because the strategies are run along 27 historical time series in the C3 category as oppose to the far smaller number of 8 historical time series in the C2 category.

The second factor has to do with the actual length of the tested time series. As seen in Figure 2.1 the lengths of the individual bonds historical corrected all-in prices vary significantly. When combining the bonds in linear combinations there needs to be a price for each common day. For this reason the strategies are limited to only run along the time period of the shortest price series within the combinations, greatly affecting the number of trades executed and hence the size of the sample from which the performance metrics are generated. It is expected that the categories with the higher number of component bonds are more likely to be shorter than the categories with a lower number due to the fact that there is a higher probability that the higher number of component bond categories will have a shorter testable historical length.

With the aforementioned limitations better understood, the results of the relative trading strategies applied on the cointegrated time series are now presented.

5.2.2 Long only strategies

The long only strategy for the relative trading results, are from the trading strategy that triggers a signal to go long the net position when the indicator breaks the lower Bollinger Band and then triggers a signal to close out the position when the indicator crosses the running mean.

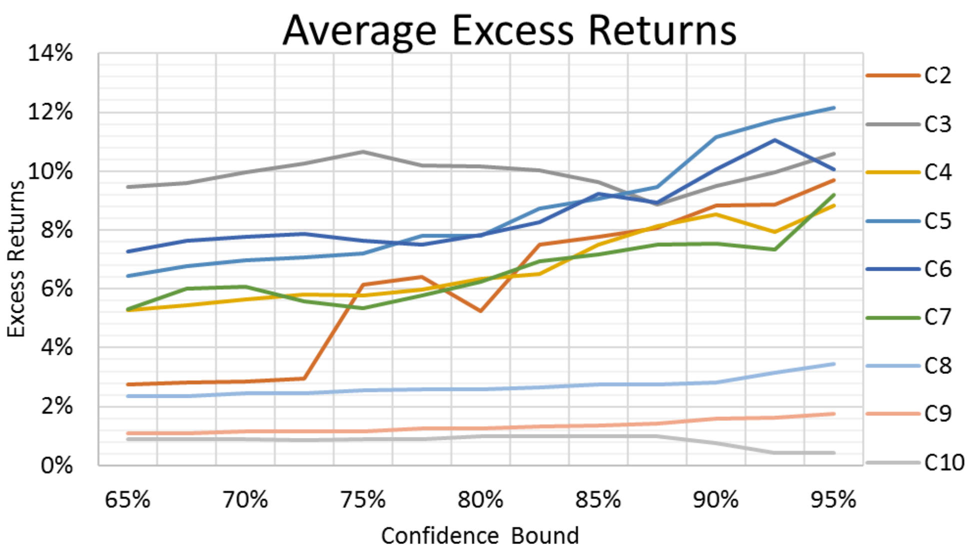

Fig. 5.8: The average excess returns of the long only strategy.

In this set of trading strategies the application of the Bollinger Bands is more rigorous. The stationarity assumption is discarded as price series are chosen for the reason that they are stationary through the process described in the previous section. The positive effects of being more mathematically accurate are immediately evident in the results that follow.

Promising results are achieved in the average excess returns generated. All the average excess returns remain positive at all tested confidence bounds with many showing very positive values upwards of 6%. One of the more notable features Figure 5.8 shows is the clear distinction between excess returns generated from the series containing a lower number of component bonds (2-7) and those containing a higher number of component bonds(8-10). The former achieving decent average excess returns mostly upward of 6% and the latter remaining statically in the 0-3% region.

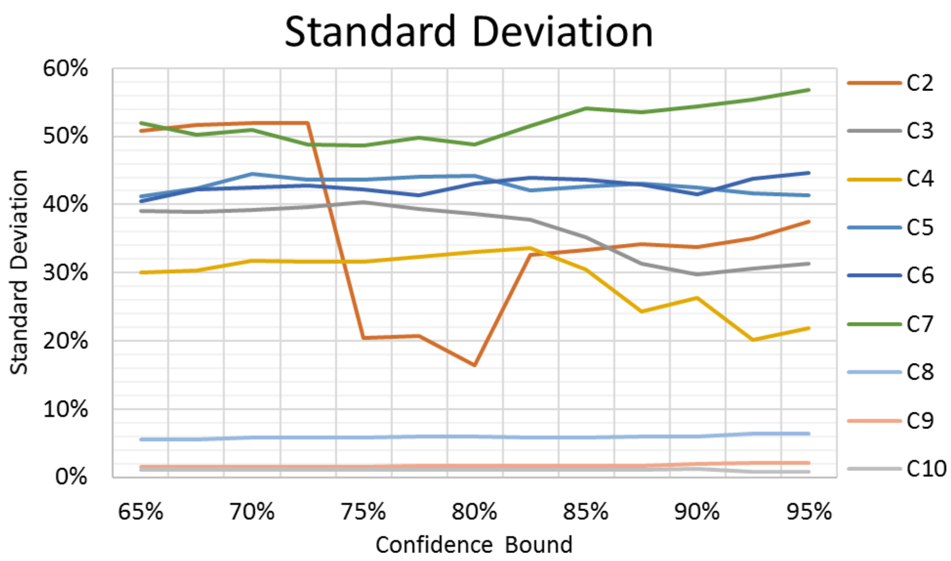

Fig. 5.9: The standard deviation of the long only strategy.

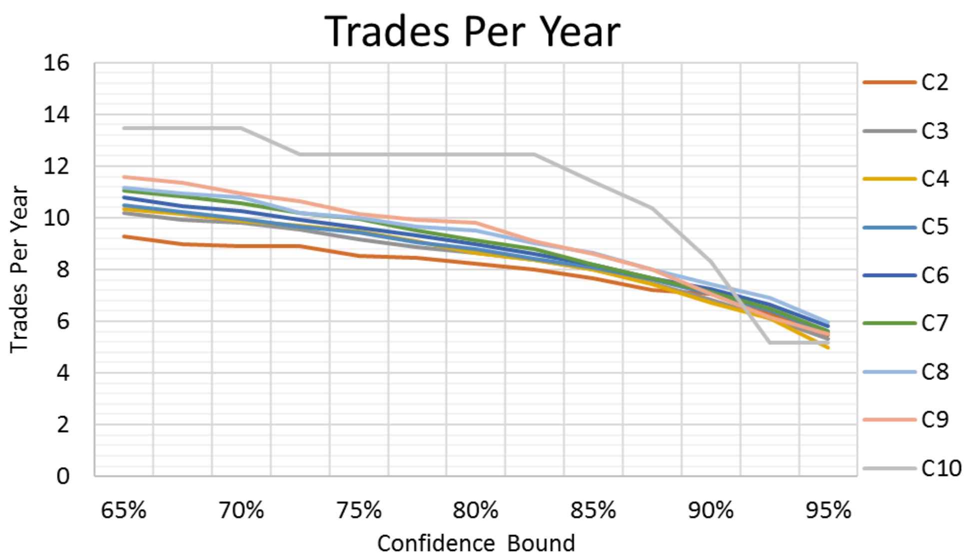

Fig. 5.10: The average trades per year of the long only strategy.

Within the better performing groups there is a notable increase in performance as the confidence interval becomes larger. The positive increase is offset by a decrees in trade frequency as seen in Figure 5.10.

In the markets it is expected that higher returns are accompanied by greater risk and in the results this too is prevalent. The standard deviation of the excess returns achieved on the combinations with higher average excess returns is notably higher and vice versa for the worse performing collections, as shown in Figure 5.9. Within the collections that achieve higher returns there is a great deal of difference within the standard deviations achieved on the tested trades.

The levels reached within the standard deviation metric are far larger than those achieved within the absolute trading strategies. Understanding the significant differences between the sets of results is made easier with the aid of Figure 3.2. In the Figure the relative changes in the corrected all-in price of the cointegrated series are far larger than the relative changes in the equivalent series for the individual bonds. It is this difference that leads to the vast differences between the sets of results.

A more desirable trading strategy is one with higher average excess returns for lower risk (standard deviation). An example of this is the collection of 7 bonds (C7) which generally achieves lower returns of the better performing collection but possess the highest standard deviation. The C7 collection would be adjourned inferior to the collection of three bonds (C3) which achieves some of the higher returns and lower standard deviations of the set of better performing bonds.

A practitioner with a lower risk appetite could even be pleased with the lower positive returns offered by the series containing more bonds for the much lower levels of risk indicated for these results. It must be noted that these combinations are the sets of price series least realistically achieved as they require positions to be held in certain proportions for a large selection of bonds, a hard trade portfolio to get into and potentially also get out of. The attractiveness of the high component strategies is significantly reduced because of the aforementioned points, making it hard to justify the low average excess returns they offer.

Referring specifically to the average trades per year metric, a fairly consistent performance across all grouped collections is apparent from when Figure 5.10 is examined. The exception to this would be the collection of corrected all-in prices comprised of ten bonds which, as shown by Figure 5.7, contains only one price series and short time series length due to the reasons described in the previous sections pertaining to restrictions of the series length based on minimum length of the comprising bonds. The sample size is hence significantly lower and so less reliable results for this specific combination are achieved.

The average trades per year frequency metric indicates that, when compared to the absolute trading results, more trading opportunities present themselves. As seen in Figure 5.10 it is clear that the levels of average trades per year are superior to the results displayed in Figures 5.3 and 5.6. The results are better than the previous however, they still indicate that there are very few opportunities per year, with Figure 5.10 showing that at best it is expected that around 10 or 11 signals to trade present themselves in a year.

The mostly promising results of the long only strategies have been presented and so focus is now moved to the short only trading strategies results.

5.2.3 Short only strategies

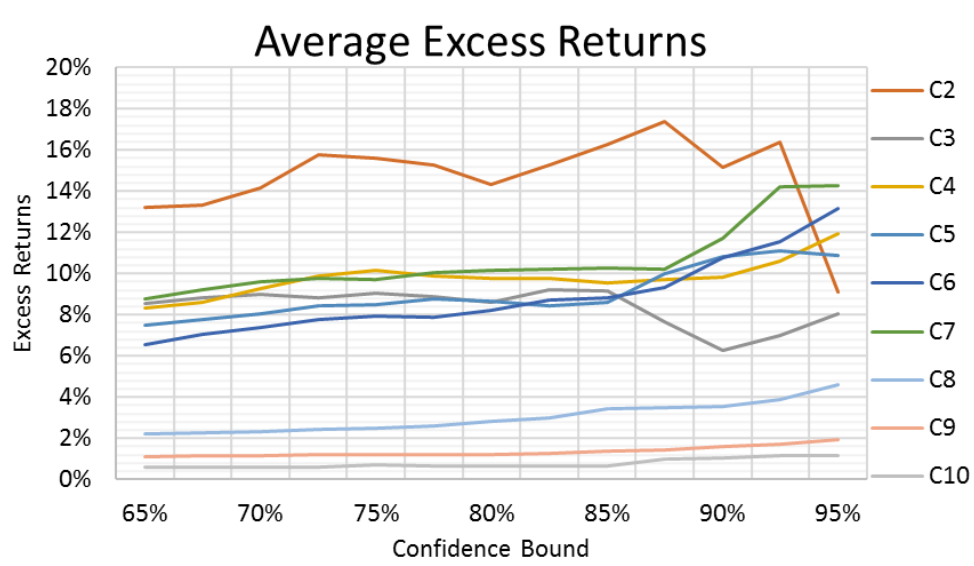

Fig. 5.11: The average excess returns of the short only strategy.

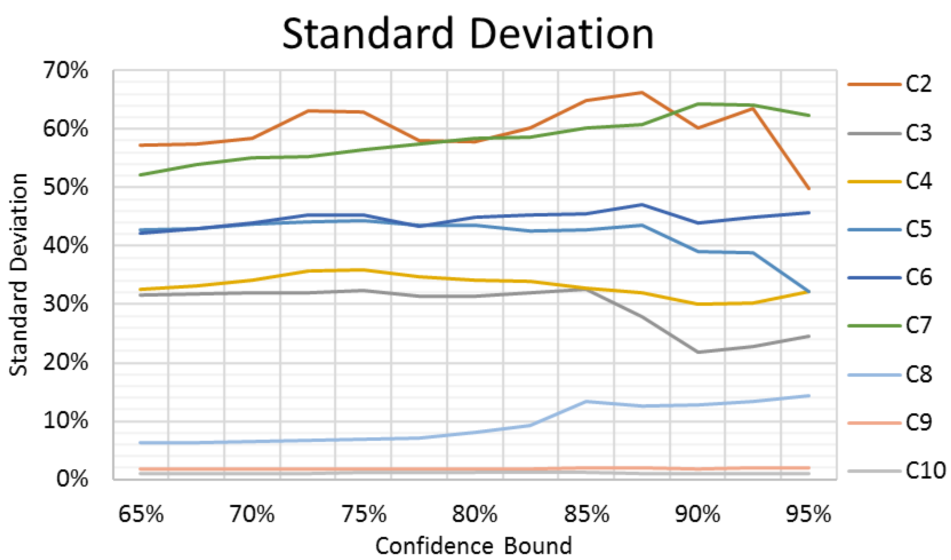

Fig. 5.12: The standard deviation of the short only strategy.

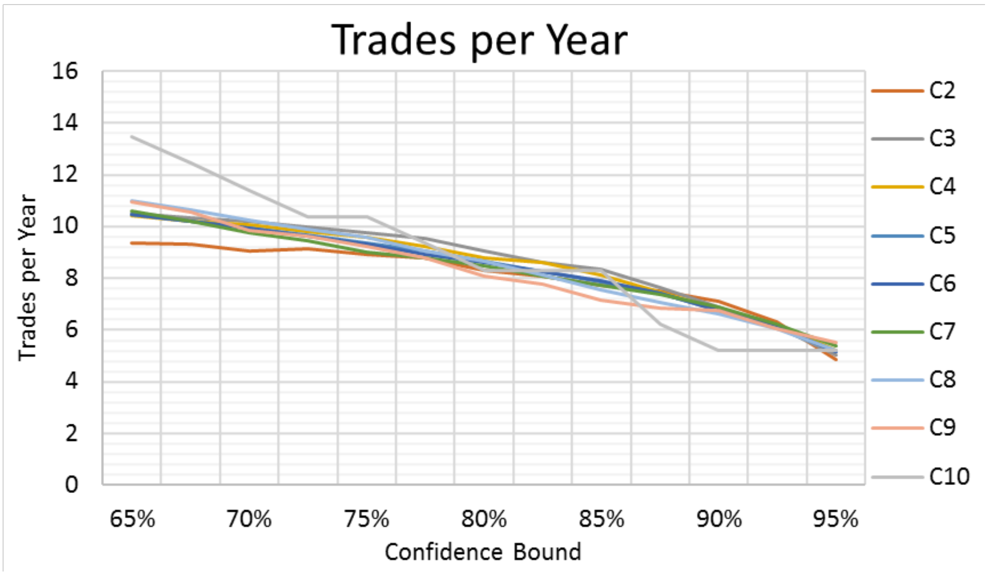

Fig. 5.13: The average trades per year of the short only strategy.

Before the results of the short only trading simulations are unpacked it should be understood that, in the context of cointegratated linear combinations, long only strategies and short only strategies are technically the same thing as the combination of bonds that comprise the cointegrated series can contain long and short positions. A short only strategy is equivalent to a long only strategy at the same confidence bound, with the weighting factors in the cointegrated series being multiplied by -1. The reflected cointegrated time series that this results in is also stationary. For this reason it is expect that very similar results are achieved and this is indeed what the results show.

The average excess returns follow the same trends and achieve roughly the same positive results. An exception of this is the combination of two bond category which sees a notable improvement over the long only strategies. The difference in results in this category could be explained by the smaller sample size associated with this category leaving more room for variation in achieved results.

It is seen again that the results for the series containing the higher number of component bonds achieves far lower average excess returns than the rest, coupled again with a lower level of risk.

Moving focus to the risk metric of standard deviation, the trend of higher risk for greater average excess returns is again present. The levels are also similar, hovering predominately in the 30 - 60% range. Factoring risk into the performance the results show again that C3 shows very desirable levels by obtaining high average excess returns for a lower level of standard deviation when compared to the other high return series. A similar line of reasoning can be followed when determining the relative performance of the other categories and the results are fairly consistent with the ones achieved in the previous section.

The average trades per year's effect in the comparison is again marginal as all the tested time series again achieve very similar levels at all confidence bounds, barring the C10 series, which is plagued by a much lower sample size for previously stated reasons. The levels achieved are almost identical to the long only relative strategies as expected, showing that a lot more opportunities present themselves with the constructed time series when compared to the absolute trading strategies relying on the individual bond time series.

The mostly positive results indicate that there could be value-add within in the bond market within the cointegration based relative trading strategies. Relatively large average returns are the primary reason for this. The positive results in this metric are however dampened by the low number of potential opportunities the average trades per year metric suggests would arise.

In both sets of results for the relative trading strategies it is apparent that the combinations containing less bonds do better. The reason for this could be due to the better performing historical price series generally have higher variation when compared to price series comprised of larger numbers of bonds. Higher variation would lead to larger swings in the position values causing higher average excess returns. The relationship between higher excess returns and time series variation is explored next.

5.2.4 Variation analysis

As the trading strategies within the cointegrated time series centre around large moves in the time series, either up or down, that subsequently are followed by a return to previous levels it is expected that a stationary time series with greater standard deviation will produce, on average, higher returns. The higher standard deviation indicates there are more correcting large moves, or that they are of a greater magnitude. Figure 5.14 corroborates this sentiment below.

It is important to note the distinction between the standard deviation being plotted in Figure 5.14 and the ones plotted in the results of the trading strategy. Plotted on the y-axis of Figure 5.14 is the standard deviation of the actual time series instead of the standard deviation of the excess returns of the trades on the time series, as was previously seen in figures.

The first notable trend that is apparent is the general decrease in variation going from categories with less component bonds to those with higher component bonds with the exception of the C2 and C10 categories. The two mentioned categories

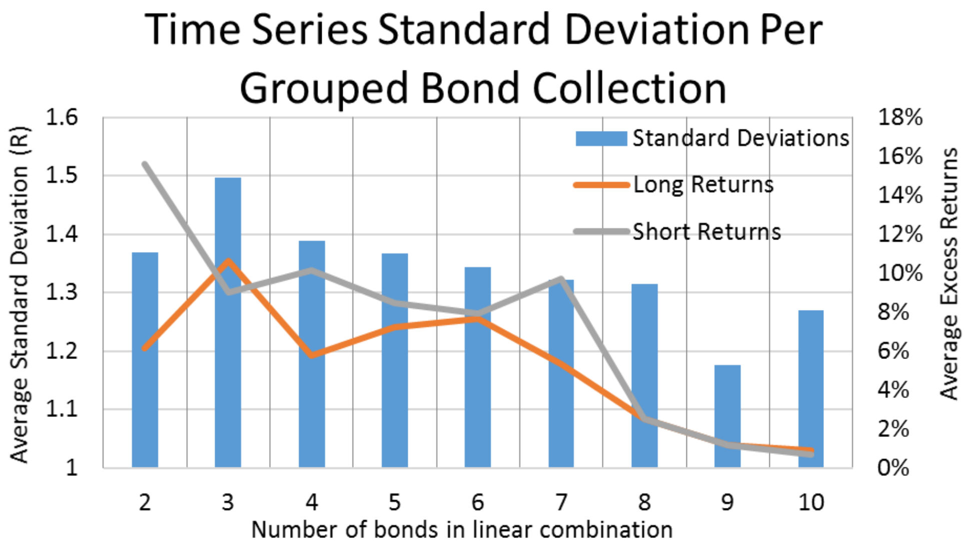

Fig. 5.14: The average standard deviation of the tested time series. The solid lines represent the average excess returns results for the 85% confidence interval.

however suffer from far lower sample sizes and hence the results pertaining to them are less accurate as has been previously explained.

The decreasing trend is also apparent within the average excess returns displayed in the previous sections. A sample cross section of the average returns at the 85% confidence level is overlayed to greater highlight the trend that persists throughout the results at the different confidence levels.

In the previous sections certain collections were highlighted within the higher performing set. The collection of three bonds (C3) was an example of one of the time series collections which has some of the best results. Figure 5.14 shows that this particular this collection has the highest standard deviation.

The series that generated the lower returns comprising 8 to 10 bonds are seen to possess the least variation of the results, further corroborating the notion that stationary series with higher variation produce better results.

The results have been discussed and the dissertation can be concluded in the following chapter.

Chapter 6

Conclusion

Through the back testing of various trading strategies the dissertation has shown that under strong liquidity assumptions there are technical analysis based trading strategies that show enough promise to be tested within the markets. The promising strategies centre around generated cointegrated linear combinations of bonds, specifically combinations comprising of a lower net number of unique bonds.

The two attempts to apply more accuracy to the technical analysis based Bollinger Band trading strategies have vastly different results.

In terms of the Bollinger Band construction methods, through the various results presented in the figures it is clear that of the three explored methods, none seemed to be constantly better. The pmf method however has flaws with the selected sample size of 20 at higher confidence bounds. The kernel and Bollinger methods seem to be superior in this regard however, overall in band construction there is no value-add in the attempts to construct them more accurately.

The attempts to apply the Bollinger Bands to more appropriate stationary time series through cointegration did yield more positive results. In generating new price series, by taking linear combinations of the bonds, significantly improved results are found within the three performance metrics. Average excess returns of over 10% are consistently achieved showing that this area of technical analysis has potential within the JSE bond market and that being more mathematically accurate, through enforcing stationarity in the price series, yields better results.

The positive results were further analysed and a link between high variation stationary series and better trading simulation results was explored. The results, shown in Figure 5.14 indicate that, when implementing the cointergration based Bollinger Band trading strategies, the best results are achieved when dealing with the higher variation price series. The result shows that when searching market data for potential cointergrated series to trade on, variation should be of particular importance.

The liquidity assumptions within the bond market are particularly strong. Having spent some time on the trading floor in a bank it is evident that during many periods liquidity completely dries up in many of the government bonds within the data set. It is also not often possible to trade in the exact proportions of each bond in order to achieve the stationary combinations the promising trading strategies require.

With the aforementioned considerations in mind there is still room for practical testing of the promising strategies in the market with small nationals, as well as the possibility of exploring other more liquidly traded asset classes for cointegration based trading opportunities using the insights gained in this dissertation.

Bibliography

Alexander, C. (2008a). *Practical Financial Econometrics-Suppressed-Added to 1265420.*

Alexander, C. (2008b). *Quantitative methods in finance*, John Wiley & Sons.

BCM, B. C. M. (2016). Classic charts.
URL: *http://www.bollingeronbollingerbands.com/support/?s=ovindi*

Bollinger, J. (1992). Using bollinger bands, *Stocks & Commodities* **10**(2): 47–51.

Chan, E. (2009). Quantitative trading, pp. 116–154.

Engle, R. F., Granger, C. W. J., Hylleberg, S. and Lee, H. S. (1993). The japanese consumption function, *Journal of Econometrics* **55**(1): 275–298.

High, S. and Honikman, A. (1995). The efficiency of the south african capital market, *Investment Analysts Journal* **24**(42): 37–43.

Johansen, S. and Juselius, K. (1990). Maximum likelihood estimation and inference on cointegration with applications to the demand for money, *Oxford Bulletin of Economics and statistics* **52**(2): 169–210.

Jse.co.za (2016a). Debt market.
URL: *https://www.jse.co.za/trade/debt-market*

Jse.co.za (2016b). Market statistics.
URL: *https://www.jse.co.za/services/market-data/market-statistics*

Kwon, K.-Y. and Kish, R. J. (2002). Technical trading strategies and return predictability: Nyse, *Applied Financial Economics* **12**(9): 639–653.

Liu, J. (2013). The bond market in south africa: efficiency and investment issues.
URL: *http://wiredspace.wits.ac.za/handle/10539/13116*

Pitsillis, Z. and Taylor, D. (2015). *The IOSCO Transparency Principle and Modelling the Bid-Ask Spread*, AIFMRM Technical Report.

Polakow, D. (2010). If a portfolio manager who cannot count finds a four-leaf clover, is he still lucky?, *Investment Analysts Journal* **39**(72): 53–71.