An Investigation into The Predictive Power of Overnight Gaps on The Johannesburg Stock Exchange

By: Ian de Jager

Dissertation presented in partial fulfilment of the requirements for the degree of Master of Commerce in the Department of Finance and Tax at the University of Cape Town.

January 2023

Supervisor: Professor Paul van Rensburg

Date: 2023/01/24 Signature: **Paul van Rensburg** Approval: Accepted

Plagiarism Declaration

Finance Masters Thesis
Department of Finance & Tax
University of Cape Town

STUDENT NAME **STUDENT NUMBER**

Ian de Jager DJGJOH005

Declaration

1. I know that plagiarism is wrong. Plagiarism is to use another's work and pretend that it is one's own.

2. I have used the Harvard-UCT convention for citation and referencing. Each contribution to, and quotation in, this submission from the work(s) of other people has been attributed, and has been cited and referenced.

3. This paper is my/our own work.

4. I have not allowed, and will not allow, anyone to copy my/our work with the intention of passing it off as his or her / their own work.

5. I acknowledge that copying someone else's assignment, essay or paper, or part of it, is wrong, and declare that this is my/our own work.

Student's Name: Ian de Jager Signature(s): **Ian de Jager** Date: 2023/01/24

Acknowledgements

My heartfelt thanks go out to all of my family and friends that have supported me in my academic and research endeavours.

A special thanks goes to Jessica Garlick, my rock and biggest supporter. You never stopped believing in me an encouraging me to push carry on and push through when the going got tough. Thanks for all your support and love, and always being my side.

Many thanks also go to Ita Schulze, without whom this research would not be possible. Your belief in my and commitment to my education will always be remembered and appreciated.

I would also like to express my sincere appreciation for the support of my parents, who have proudly supported me throughout my academic career. Your encouragement and support will always be appreciated.

My thanks go out to my supervisor, Professor Paul van Rensburg, for his guidance and expertise that were always available. Your experience and valuable input are greatly appreciated.

Additional thanks must also be given to everyone that has directly or indirectly influenced the creation of this thesis and the development of my academic career. A list to long to write, your influences are no less important and just as much appreciated.

Abstract

Stock market gaps occur nearly every day, yet very little is known about their influence on subsequent pricing behaviour, particularly in developing economies like South Africa. The aim of this research is to comprehensively identify and analyse the relationship between overnight price gaps and subsequent intraday returns for publicly traded South African companies. These theoretical findings will also be applied practically, with a simulated trading strategy created and tested based on the theoretical and statistical findings. The primary method of identifying the underlying relationships at play is a collection of multiple linear and multiple logistic regression models, created using data spanning 20 years and 371 companies and split into training and testing sections to ensure accurate and bias-free results.

The robust set of statistical tests and analyses performed indicates a persistent and highly significant inverse relationship that exists between large overnight gaps and subsequent intraday returns. This significant relationship was also applied to a very successful mean-reversion based trading strategy, with a sustained average annual outperformance of the JSE AllShare Index observed under even the highest transaction costs of 1.3% per trade. At the lower transaction cost level associated with CFD trading, an average annual outperformance of 166% was recorded.

The theoretical and practical implications of these findings are in stark contrast to the widely accepted efficient market hypothesis and provide compelling new evidence of the exploitable nature of a relatively under researched market anomaly created by the inefficiencies associated with overnight price gaps. These results also pave the way for further analyses of gap behaviour, and collectively these findings add new and meaningful results to the body of knowledge on market anomalies.

Table of Contents

List of Abbreviations

AIC: Akaike Information Criterion

AMH: Adaptive Market Hypothesis

CFD: Contract for Difference

EMH: Efficient Market Hypothesis

EST: Eastern Standard Time

FINDI: The FTSE/JSE Financial & Industrial 30 (JFNDI)

JSE: Johannesburg Stock Exchange

NYSE: New York Stock Exchange

OLS: Ordinary Least Squares

SEC: Securities and Exchange Commission

SMA: Simple Moving Average

STT: Share Transfer Tax

USA: United States of America

VAT: Value Added Tax

VIF: Variance Inflation Factor

List of Tables

List of Figures

Chapter 1: Introduction & Overview

1.1. Overnight Stock Market Price Gaps

Stock market gaps, also referred to as overnight gaps, occur when an exchange listed security opens at a different price today than the prior trading days closing price. Stock markets in particular have limited trading hours, with trading of instruments listed on the exchange only possible during the specified market hours. Standard market hours are the period of time during which any and all authorised market participants can engage in transactions without restriction, subject to the exchange's ordinary usage terms.

In addition to market hours, many exchanges also offer extended trading hours on many listed securities (Nasdaq, n.d.). Extended trading hours consist of pre-market and after-hours trading periods, which, as the names imply, are a window of time before and after ordinary market hours during which trading may still take place. In contrast to normal market hours, trading in extended hours is usually subject to additional limitations and restrictions (Charles Schwab, 2021). These vary from exchange to exchange and security to security, but in general usually consist of limits on the types of orders that can be placed, number of tradeable securities and margin requirements, for example. Extended hours are also not available to many market participants, with a specialised stockbroker required to access extended hours trading in the case of equities. Partly as a result of these restriction and barriers to entry, both trading volume and liquidity are much lower during extended hours compared to ordinary trading hours.

Even with extended hours taken into account, individual stocks on exchanges around the world are still subject to a period of time during which they cannot be traded at all, and the market is closed to all participants. In the United States, major exchanges offer trading (including extended hours) between 04:00 and 20:00 EST (United States Securities and Exchange Commission [SEC], 2011). This leaves an 8-hour window during which no shares may be traded. In contrast, the Johannesburg Stock Exchange (JSE) does not offer extended hours trading, with full market hours consisting of an opening time of 09:00 and a closing time of 17:00 (MarketHours.com, n.d.). As a result of both extended trading hours and the full market closures that take place daily, price gaps can occur between one day's closing price and the following day's opening price.

Such gaps can be both positive and negative in relation to the prior closing price and can at times be very large in size. In the event of an earnings announcement or a significant news release overnight on a particular company, the stock may gap up or down several times the average daily trading range. Particularly on USA exchanges, low float, low market capitalisation stocks can often have extreme gaps driven by news, with gaps over 100% occurring several times per week, on average. Large float, mega capitalisation stocks also develop gaps from time to time, however, they are not nearly as large in relative terms on average, and the market reaction is often smaller.

The structure of nearly all stock markets allows for overnight gaps to occur on a large number of securities from time to time, and to varying degrees. In the case of extreme gaps, a substantial market reaction can occur. A large gap in either direction can be equivalent to a sudden move in price equal to multiple times the average daily range of a stock. Such gaps essentially result in a spike in volatility, and following these gaps, numerous resting market orders would get triggered once the market opens. The effect of overnight returns and trading activity has been found to have the largest impact on intraday returns during the first thirty minutes after market open (Chan, Chockalingam & Lai, 2000). In the authors experience, when a large overnight gap has occurred, at market open there is often a volatile reaction to this jump in price, with above average volume, volatility and bid-ask spreads recorded, as market participants digest the sudden and large change in prices.

When considering and investigating the predictive power of overnight gaps it is important to note that, in extension, such an analysis also tests market efficiency. The efficient market hypothesis (EMH) states that prices should adjust rapidly and unbiasedly to new and relevant information (Fama, 1970). The weak form of the EMH in particular deals only with historical prices as "relevant information". This means that any form of technical analysis or price pattern (such as gaps) should not have any significant predictive power, or at least not enough to earn consistent market-beating returns.

Market efficiency, both in the South African markets as well as internationally has been a hotly debated topic for many years, with some compelling research presented for both cases, and no widely accepted and conclusive answers existing so far. It can be argued that market efficiency, like many other areas of financial theory, is something that can be discredited by the process of falsification. That is to say, no amount of confirming evidence of market efficiency can allow us to conclusively state that the market is efficient, but just one case presenting clear evidence to the contrary is enough to discredit the theory of market efficiency. To date, a large body of evidence exists that contradicts the traditional theory of market efficiency, such as the momentum and value anomalies, day of the week effects, January effects and many more (Eloff (2014), Page, McClelland & Auret (2020), Van Heerden & Van Rensburg (2015), Lockhart-Ross (2016), Oldfield and Rogalski (1980), Liu et al. (2021), Cooper, Cliff, & Gulen (2008), Dorador (2017)). Out of the observed anomalies, many have proven to be highly pervasive and persistent over time, presenting a compelling case for further research into what other patterns and conditions may exist that could predict market returns.

Given the extreme nature of certain gaps, and the large market reaction that can occur following their development, the question arises: do they exhibit any predictive power or influence on subsequent returns?

Chapter 2: Literature Review

2.1. Introduction

When considering potential research into overnight gaps, it is important and beneficial that such a study builds on the findings of authors investigating similar topics, and, if possible, contributes to existing literature by adding to this body of knowledge in a meaningful way.

To date, a substantial amount of research has been conducted on market efficiency and market anomalies, in many international markets and across a long and well covered time horizon in totality. Given that market gaps and their potential influence on market characteristics constitutes a subset of market efficiency and market anomaly tests, one would expect comparatively less research to have been done on this particular topic. This is indeed the case. However, literature on stock market gaps and their influence on prices and other market characteristics has been covered relatively well to date, with numerous studies being done on this topic in international markets, and several overlaps of market choice and research time period occurring.

When considering the South African market and the JSE in particular, literature on stock market gaps and their predictive power is sparse at best. Thus, a comprehensive review of the literature is conducted, starting with gaps at a broader level and extending to the specific case of returns on the JSE, laying the foundation for such research to be conducted in a relevant and meaningful way in the context of existing literature.

2.2. Stock Market Gaps & Overnight Returns – International Research

Over a historical period of more than fifty years, Liu et al. (2021) look at stock market gap behaviour on major United States of America (USA) and Chinese indices and attempt to determine if gaps and resulting trading activity contain any non-random, predictable behaviour. They find that USA exchanges contain more gaps, as well as gaps of a larger magnitude than on Chinese exchanges. By detrending and using random shuffle techniques, seasonality and the impacts of trends can also be analysed, with the results indicating that, overall, the trend does play a role in gap-fill behaviour, with a gap in the same direction as the trend being less likely to be filled than a gap in the opposite direction to the trend. Such findings are indicative of the potential presence of momentum; however, the authors did not analyse these results further. In terms of predictability, both detrended and unadjusted data showed that gaps did not exhibit random behaviour, and that some level of pattern in terms of gap-filling behaviour exists on both Chinese and USA exchanges and may be widespread at the individual security level in these countries (Liu et al., 2021).

In contrast to the findings of Liu et al. (2021), Plastun et al. (2019) analyse the Ukrainian stock market's UX index and do not find any pattern or seasonality in the gaps observed from 2009 – 2018. Results from an array of statistical tests performed in this analysis indicate that the assumption of market efficiency holds in the Ukrainian stock exchange with regards to price gaps. Despite this result, Plastun et al. (2019) do find strong evidence of the existence and persistence of the momentum anomaly on negative gap days specifically. Unlike Liu et al. (2021) no trend analysis was done in this analysis, and thus it is possible that trend behaviour may help explain why momentum behaviour is only present in gap down days.

Aboody et al. (2018), analysed overnight return behaviour as a means to measure investor sentiment, and the authors find evidence of overnight volatility and returns persisting on a weekly time frame, as well as being higher for harder-to-value firms. These findings are consistent with prior literature that found harder-to-value firms to be relatively more sensitive to market sentiment than more easily valued firms (Baker & Wurgler (2006), Berkman et al. (2009), Hribar & McInnis (2012), Seybert & Yang (2012)). Zhong (2008) finds similar results when looking at a comparison between liquid and illiquid, thinly traded shares, hypothesising that illiquid shares contain more information than more thickly traded stocks. The author also finds that overall, overnight gaps are characterised by significant returns, particularly in comparison to intraday returns. This effect was also observed to be more significant during extended hours trading, with illiquid shares having a slower reaction to news released overnight, and these effects spilling over into the following trading day's activity. This result in particular giving further credence to the significance of the findings of Aboody et al.

(2018) with regards to overnight returns having a bigger impact for harder-to-value firms, a characteristic which is closely linked to illiquidity.

Overnight returns were also found to be a significant predictor of intraday return characteristics by Oldfield and Rogalski (1980), who employed a multiple jump process model in order to explain differences in returns between overnight and intraday trading. Higher volatility during intraday trading in comparison to overnight trading was also observed by Hong & Wang (2000), with new information revelations during periods of market closure being presented as one key influential factor. Overnight returns have thus far been found to have the largest impact on intraday returns during the first thirty minutes of intraday trading after market open (Chan, Chockalingam & Lai, 2000). These results are in line with the findings of Berkman et al. (2012), who also find overnight trading to have the biggest impact on intraday returns around market open.

Overnight gaps have also been found to be non-random and contain statistically significant predictive power by Cooper, Cliff, & Gulen (2008) as well as Dorador (2017). In their analysis of USA indices, Dorador (2017) finds that the average gap is small, at less than 0.05%. However, nearly 78% of the time gaps were greater than 0.1%. Such a gap size is significant, particularly at an index level and during periods of lower market volatility. It was also found that both the frequency and size of the gaps increases significantly during periods of above average market volatility, with the strongest effects on price observed the day following the gap, suggesting again that overnight returns may play a significant role in shorter-term price predictions. Cooper, Cliff & Gulen (2008) theorise that the reason for larger overnight returns is related to the illiquidity of extended hours trading, with far less trade volume required to move markets significantly more than during intraday trading, thus leading to greater volatility in overnight returns in comparison to intraday returns.

Prior research suggests that retail investors in particular are likely to act less rationally than other market participants in their trading behaviour, with numerous studies showing that retail investors tend to gravitate towards attention-grabbing stocks, with over-optimistic trading leading to significant deviations in firm value/price from core fundamentals (Delong et al. (1990), Shleifer & Vishny (1997), Barber & Odean (2008)).

Similar findings on retail investor sentiment were also presented by Berkman et al. (2012), who found that stocks that have attracted increased levels of retail investor attention tend to have higher overnight returns, with intraday reversals occurring following these above average overnight returns. The theory put forward to explain this phenomenon revolves around attention-based overpricing. The authors assert that stocks that have attracted high levels of retail attention have a tendency to result in strong buying pressure at the open, driven by sentiment, with higher overnight returns and a higher opening price the following day observed, followed by reversal-like behaviour in subsequent intraday returns as the sentiment changes.

Partly based on these findings on investor sentiment and retail investor overreactions, Huang, Hu & Truong (2021) looked at average overnight returns as a predictor of longer-term share price movements, with their findings indicating that a significant relationship does exist between overnight returns and share price performance across a 1-to-12-month subsequent period. Correlations between overnight returns and future stock returns were found to be varying substantially depending on the level of momentum effects observed in the stock under analysis.

Similar to the findings of Cooper, Cliff & Gulen (2008), Barclay & Hendershott (2003) also find that overnight gaps and overnight trading activity contains more information than trading occurring during normal market hours. The authors of this research hypothesise that overnight trading contains more information because of the substantially lower levels of liquidity during extended hours trading, along the same lines as the research presented by Zhong (2008). As a result, any trades that do occur have a much larger proportional impact on price, and less shares can be traded without moving the market. It is thus argued that investors willing to put up with the additional challenges of extended hours trading must value their information highly enough to compensate them for these challenges, and thus extended hours trading is likely to be less efficient and contain more information than standard market hours trading (Barclay & Hendershott, 2003).

The findings of Barclay and Hendershott (2003) give further credence to the research of Kelly & Clark (2011), who looked at return differences of daily and overnight returns and the Sharpe Ratio of these returns as well, with compelling evidence presented that overnight returns in USA exchanges over the period 1999 - 2006 were substantially higher than daily returns, while also having lower volatility levels. Collectively, the literature on the significance of overnight gaps in international markets builds a clear picture of the stark contrast between intraday and overnight risk/return characteristics.

One of the most comprehensive studies of price gap anomalies belongs to Plastun et al. (2020), who looked at three major USA indices over the period 1928 to 2018 with the aim of investigating whether price gaps create exploitable market inefficiencies. Through the results of various statistical tests (Student's t-test, ANOVA, Mann-Whitney test & regression analysis), the authors find strong evidence of non-random price gap behaviour. Specifically, on the day of a price gap, prices tend to change in the direction of the gap. A trading strategy built around these observations was found to be effective, with results showing that the indices under analysis were not efficient from a weak-form EMH perspective. Similar to the research presented by Berkman et al. (2012), Plastun et al. (2020) also find that the strongest effect of price gaps to be on the day that they occur, and not the day after, with the authors concluding that one trading day is enough time to profit from price gap behaviour and patterns. As with many other market anomalies, both the occurrence and predictive power of price gaps was found to evolve, with a decrease in the quantity and follow-through of these gaps observed in more recent data. The authors also find that the consideration of what constitutes a significant gap is crucial in such an analysis. Too small a gap size and there are too many gaps observed to be considered anomalous behaviour, while too big of a gap size may lead to insufficiently large sample sizes for any meaningful statistical significance.

2.3. Price Gap Predictive Power & Proposed Trading Strategies

Based on the research of the significance of overnight returns and the link between overnight and intraday returns, De Gooijer, Diks & Gatarek (2009) employ linear regression models as well as non-linear models in an attempt to explain and predict overnight returns in several international indices and exchanges. The results of this analysis indicate that the prediction of both the sign and size of overnight returns was possible when using non-linear models. Given the link between overnight returns and intraday returns, the ability to predict overnight returns using information revealed from other international exchanges leading up to the domestic market open is a valuable finding for both academics and practitioners.

Despite the relatively large sample of international exchanges used, the timeframe of analysis used by De Gooijer, Diks & Gatarek (2009) was short, between 2007 and 2008 only and less than one year in total. The study was also performed around the start of the 2008 global financial crisis, a key factor that could skew results. Stock market characteristics, predictability and overall behaviour have been found to differ significantly during periods of economic crisis, or heightened fear and volatility (Ellington, 2008). In addition to overnight returns being statistically significant, De Gooijer, Diks & Gatarek (2009) also found returns over weekend market closures to be significantly higher than weekday returns.

Serletis and Rosenberg (2009) and Leung (2009) find evidence of mean-reverting behaviour in stock markets at a broader level, and based partly on this body of research, at least two studies of strategies to exploit overnight gaps and mean-reverting behaviour have also been developed (Dunis, Laws & Rudy (2011), Stübinger & Schneider (2019)). Looking at the S&P500 from 1998 to 2015, Stübinger & Schneider (2019) find that a profitable mean-reversion strategy built around overnight gaps can be formulated. The strategy explored centres around the most extreme gaps, with a position in the opposite direction to the gap taken at market open with the aim of capturing the mean-reverting behaviour exhibited in the markets and created by overnight gaps. Over the period of analysis (1998 – 2006), the strategy developed by the authors was found to outperform the industry standard “buy and hold” benchmark severalfold, earning annualised returns of 51%. These returns were generated after taking associated transaction costs into account, specifically as they were over the period of analysis.

Over the period 2000 to 2015, Dunis, Laws & Rudy (2011) conducted similar research to that of Stübinger & Schneider (2019). The authors analysed the performance of a multi-factor model built with a focus on overnight gaps in major USA indices, with results also indicating that a profitable trading strategy could be built with such a model. The largest and most significant results were reported for two of the three indices under analysis, namely the S&P 600 SmallCap Index and the S&P 400 MidCap

Index, with returns from gap trading on the S&P500 being significantly lower. Such results largely fall in line with research explored so far in this thesis, particularly as it relates to market efficiency differing across market segments, with the market generally tending to be more efficient the more market participants and fund flows there are. Dunis, Laws & Rudy (2011) also found their strategy to be insensitive to transaction costs and bid-ask spreads, and even after these associated trading costs the gap trading strategy was found to outperform all major benchmark performance metrics, with annualised returns as high as 215%. The authors conclude that there is strong evidence showing that overnight trading is information rich, and not all information revealed in overnight trading may be fully incorporated at market open, leading to exploitable asset mispricing behaviour.

2.4. Market Efficiency – South African Research

At this point, it is important to note that the results and findings of the literature presented so far pose some major challenges to the traditional efficient market hypothesis (EMH). Additional findings of market efficiency in South African markets overall can shed some light on whether there is merit for additional research on pricing and market anomalies.

Under the EMH, securities should always trade at their intrinsic values, with any exploitable deviations presenting a contradiction of this theory. In their in-depth analysis of the JSE and various well-documented market anomalies, Alexandroi (2019) finds compelling evidence that the JSE was not efficient in their analysis. Significant and exploitable differences between trading prices and intrinsic company values were also observed by Atsin & Ocran (2015). If the EMH holds, then active fund managers should not be able to consistently outperform, however, in an analysis of active management performance, specifically as it relates to trading perceived overactions on the JSE, Eloff (2014) found that consistent market-beating returns could be generated, posing a further challenge to the EMH concept.

The existence and persistence of market anomalies, such as the momentum effect, would also pose a challenge to the notion of the EMH. The momentum anomaly in particular can become something of a self-fulfilling prophecy. That is to say, the more market participants trade based on momentum, the greater the momentum observed in the market will be, and the bigger the effect will be, including the mean-reverting correction that occurs after the short-term momentum (Eloff, 2014). Such market swings and deviations from fundamentals were observed by Page and Way (1993) in their analysis of JSE listed companies from 1974 to 1989. Although the research is now dated, the momentum anomaly has been found to persist in numerous subsequent studies (Eloff (2014), Page, McClelland & Auret (2020), Van Heerden & Van Rensburg (2015), Lockhart-Ross (2016)). Although not based on the momentum anomaly, Frisch, Kolaric & Schiereck (2014) find that extreme moves (both upwards and downwards) on JSE listed companies were likely to be followed by positive returns.

Exploitable market patterns are also likely to exist under the behavioural finance school of thought. Theories put forward by behaviouralists revolve around the notion of market participants largely being human, and humans being irrational in many aspects, with human decision-making flaws, biases and heuristics showing up in investment and trading decisions taken by market participants as well, leading to repeatable and exploitable patterns in markets (Heymans and Santana, 2018).

On the other hand, it is argued that the fact that many market anomalies that have been observed and documented over the years but have since disappeared is due to the efficiency of markets, such that any predictable pattern that develops will quickly and efficiently be exploited until it no longer exists, and the market is driven back to equilibrium (Fama, 1998).

Although proponents of the EMH and behaviouralist view can be regarded as two extremes of the possible state of markets, something of a middle ground does exist in other theories of market activity, such as the adaptive market hypothesis (AMH), which states that markets tend to undergo cycles of efficiency and inefficiency, rather than being permanently in one state or the other (Lo, 2004). Such a theory conforms to the findings of Heymans & Santana (2018) in South African context, with the authors analysing the efficiency of JSE indices and sub-indices. Levels of efficiency were found to differ not only across indices, but over time within the same indices as well, presenting evidence of market “cycles” and varying levels of efficiency on the JSE.

An additional market anomaly that has been observed on the JSE is the weekend effect. This phenomenon refers to the fact that stocks and stock markets tend to have differing average returns on different days of the week. Specifically, returns on Mondays tend to be lower than returns on other weekdays. It is theorised that this occurs due to the tendency of companies to release more negative news towards the end of the week, mostly on a Friday night or over the weekend. This information can then only be acted on by traders and investors on Monday, with the result being significantly lower returns on Mondays compared to other days (Atsin & Ocran, 2015). Looking at the JSE, results are mixed as to whether the weekend effect exists and/or is prominent on the JSE. Atsin & Ocran (2015) analyse various indices and test whether the weekend effect is present or not, and over their timeframe of analysis (2002 – 2013) the authors observe that the weekend effect is present on some indices, but not all. Specifically, the All-Share Index and Top40 Index do not appear to contain the weekend effect, however the MidCap and SmallCap indices do shower statistically significant differences in returns on Mondays in comparison to other days of the week (Atsin & Ocran, 2015). These findings present further evidence in support of the weekend effect anomaly overall, in particular as it relates to price gaps.

In addition to the weekend effect, Atsin & Ocran (2015) also looked at the whether the January effect exists in South African markets. The January effect is a market anomaly based on the observation that stock returns tend to be significantly lower in December than they are in January for stocks with a weak year-to-date performance, with respective returns in both cases being below and above the monthly average as well. It is thought that this occurs due to multitude of factors that lead both institutional and retail investors to sell their underperforming holdings in December. Especially for institutional investors, on behalf of their clients, this is thought to be done for tax purposes, and once the losses have been realised, the shares are bought back again the following month in January, leading to a share price appreciation again (Atsin & Ocran, 2015). The authors did not find conclusive evidence of the January effect in the indices under analysis and conclude that the January effect does not appear to be present on the JSE at an index level. These results are in contrast to other research on African markets, with Enowbi, Guidi, & Mlambo (2009) finding the January effect to be present and persistent in the Moroccan and Egyptian stock markets.

Building on the body of research on market anomalies in South Africa, Page, McClelland & Auret (2020) find that both the momentum and value effects have existed and persisted on the JSE over their period of analysis, 1992 – 2014, similar to the findings of Atsin & Ocran (2015) with regards to the momentum and size effects. The focus of this research also extends to the effect of direct and indirect transaction costs on well documented market anomalies. Both the momentum and value factors were found to be sensitive to transaction costs, with the momentum factor being more sensitive to direct transaction costs and the value effect being more sensitive to indirect transaction costs. It is argued that, since both are sensitive to transaction costs, this may be the component of arbitrage costs that has led to the existence and persistence of these anomalies both in South African and international markets over the years.

2.5. Literature Overview

Overall, literature to date provides compelling evidence of the persistence and prevalence of several market anomalies, such as the momentum and value effects (Eloff (2014), Page, McClelland & Auret (2020), Van Heerden & Van Rensburg (2015), Lockhart-Ross (2016)). These anomalies pose a challenge to the efficient market hypothesis and lays the foundation for the investigation into potential additional market anomalies that may exist.

With regards to overnight gaps, research shows that overnight gaps are significant in predicting certain stock and market related characteristics. Most notably, research from international markets on the effect of overnight gaps on intraday and daily returns following gaps have been found to be significant by numerous researchers to date (Oldfield and Rogalski (1980), Liu et al. (2021), Cooper, Cliff, & Gulen (2008), Dorador (2017)). In extension, overnight returns have also been found to contain more information and a higher degree of information asymmetry than intraday returns (Zhong (2008), Berkman et al. (2012)).

In terms of using overnight gaps for price prediction, Stübinger & Schneider (2019) as well as De Gooijer, Diks & Gatarek (2009) find that such predictions are both possible and profitable. These findings are consistent with research that shows overnight gaps to be significant and presents a real-world potential use of such a market anomaly.

Based on these findings collectively, a comprehensive study of overnight gap behaviour in the South African market is proposed, which will add to the body of knowledge on this topic in a significant and meaningful way.

Chapter 3: Data & Descriptive Statistics

3.1. Data Overview

Given the comparatively small size of the JSE in comparison to other international exchanges, the number of listed securities is low enough that an extensive analysis over a long time period can be conducted. The period of analysis for this research covers 2000 – 2021 and includes all JSE primary-listed securities that were listed over this time period. Data for pricing and additional fundamental variables was collected on a daily aggregation basis, with all relevant pricing and fundamental information sourced from Bloomberg.

The timeframe of 2000 – 2021 covers two of the most recent major financial recessions, and these will be analysed further in this study. The 2008 global financial crisis sub-period as well as the 2020 Covid-19 recession will be analysed in terms of structure and influence in relation to the overall behaviour of gaps and their predictive power. The 20-year total time horizon is long enough to cover a substantial amount of data and facilitate an in-depth analysis, in a similar fashion to the 17-year period of analysis for Stübinger & Schneider (2019) in their research on the USA market.

In totality, there are 371 companies in this dataset. Over the roughly 20-year analysis period, this constitutes more than 720,000 unique observations. Variables included in the dataset are:

Price and volume data

- Open, high, low & close.
- Volume traded and value traded.

Company information

- Full company name and ticker symbol.
- Primary operating sector and head office location.
- Market capitalisation.

- A sl indicating whether the company is a FTSE/JSE Financial & Industrial 30 Index (FINDI) constituent.

3.2. Additional Variables

In order to perform an in-depth analysis of gaps and their behaviour, certain additional variables were created in order to facilitate this, based on the authors prior research on gaps at an index level:

Gap: Given the pricing information contained in the dataset, aggregated on a daily basis, the overnight returns can be calculated for all the stocks in the dataset. These overnight returns are calculated as the difference between the closing price of a security the prior trading day, and the opening price of that security the following trading day. Given that the JSE does not facilitate extended hours trading in the same way some USA exchanges do, these overnight returns constitute only the key variable of interest in this study: the overnight gap. The overnight gap is calculated as follows:

$$\mathrm{Gap}_{\mathrm{i,t}}(\%) = \frac{(\text{Opening Price}_{\mathrm{i,t}} - \text{Closing Price}_{\mathrm{i,t-1}})}{\text{Closing Price}_{\mathrm{i,t-1}}} \times 100 \qquad (3.1)$$

Intraday returns at time t and t-1: These returns are calculated with the opening and closing intraday prices of a given security. This calculation is done for both the current trading day (t) as well as the prior trading day (t-1) in order to analyse the effect of an overnight gap without hindsight bias. If an overnight gap were to form, it would be between the periods t – 1 and t. Hence, the prior days return (t – 1) in combination with the overnight gap will constitute the two parts of the total return of a security in the preceding 24-hour period. The subsequent days intraday return (t) will thus allow for an examination of the ex-post effect of the gap on intraday returns. This return is calculated as follows:

$$\text{Intraday Return}_{\mathrm{i,t}}\,(\%) = \frac{(\text{Closing Price}_{\mathrm{i,t}} - \text{Opening Price}_{\mathrm{i,t}})}{\text{Opening Price}_{\mathrm{i,t}}} \times 100 \qquad (3.2)$$

Cumulative preceding returns, from t-5 to t-1: These variables deal with the total cumulative return of the security in question, leading up to, but not including the occurrence of an extreme overnight gap. This total return thus includes both the overnight and intraday portion of the prior returns, extending as far back as 5 trading days. The cumulative preceding return at t-5 thus incorporates the return of the security from market open 5 days ago until the close of trading 1 day ago. The same returns are calculated at intervals of t-3 and t-2, and lastly, t-1, as described above. The 5-day prior cumulative return is calculated as follows:

$$(\text{Cumulative Return})_{\text{i,t}-5}(\%) = \frac{(\text{Closing Price}_{\text{i,t}-1} - \text{Opening Price}_{\text{i,t}-5})}{\text{Opening Price}_{\text{i,t}-5}} \times 100 \quad (3.3)$$

Total return index change: Similar to the intraday return variables, this variable measures the change of the JSE AllShare Total Return Index (ALSI TRI) over a 24-hour period. Unlike the variables above, this total return is not broken down into intraday and overnight returns and stated only in total return terms. The calculation of these returns remains useful as a measure of comparison, particularly for looking at the overall size or volatility of intraday or overnight returns on an individual security in comparison to the ALSI TRI, which is one of the most commonly used equity benchmarks in South Africa. The Index total return is calculated as follows:

$$\text{Total Index Return}_{i,t}\ (\%) = \frac{(\text{Closing Price}_{\text{i,t}} - \text{Closing Price}_{\text{i,t}-1})}{\text{Closing Price}_{\text{i,t}-1}} \times 100 \quad (3.4)$$

9 & 20 Period simple moving averages (SMAs): The creation of a simple moving average of the closing prices of a security was necessary in order to create an additional variable that can represent the overall trend of a security, either bullish, bearish or non-trending. The two values chosen for the number of samples used in the calculation of the SMAs was based on some of the most commonly used moving average lengths in technical analysis (TradeCiety, n.d). The commonly used long-term 200-day SMA was omitted in this case due to the method of constructing the portfolio of tradable assets for this analysis, whereby securities can fall in and out of this portfolio, as discussed in more depth later in Chapter 3.4 – Portfolio Construction. The 200-day SMA is thus too long to allow for enough meaningful observations. Rather, the 9SMA was used to capture an estimate of the short-term trend and the 20SMA for the medium/longer-term trend. The price of a security in relation to these SMA's, as well as the relative positions of the SMAs in relation to each other can greatly assist in determining the overall context of a security's prior price movements up to the current date of analysis. The simple period moving averages of length n can be calculated as follows:

$$\mathrm{SMA_n} = \frac{\mathrm{A_1 + A_2 + \cdots + A_n}}{\mathrm{n}} \tag{3.5}$$

The total daily value traded: A key characteristic of the JSE is the large variance in liquidity for the listed companies. Notably, the ALSI has been shown to have comparatively low levels of diversification compared to indices of other countries and exchanges (Raubenheimer, 2010). One of the primary reasons for this is the heavily skewed distribution of size and liquidity of JSE listed companies. A few large companies dominate the market and correspond to a disproportionately large share of the ALSI. At the other end, the majority of companies listed on the JSE have been shown to have substantial liquidity issues. This is of particular concern to the analysis of this study, as the analysis and potential exploitation of gap behaviour will require accurate and frequently updated intraday data, something which is not present for thinly traded (illiquid) securities. As such, certain measures need to be put in place to exclude shares that are too illiquid and will skew results if they are included in the gap analysis. The starting point for this adjustment is the calculation of the daily value traded of the JSE-listed companies. Crucially, in comparison to the commonly quoted volume (of shares traded per day) variable, value traded takes both volume and price into account. In this case, data in the sample is aggregated daily, and intraday information on the value traded is not available for most observations in the dataset. As such, the most accurate measure of value traded is impossible to obtain. A close estimation of the true daily value traded is possible, however. The estimation of daily value traded is crucial in selecting the portfolio of assets suitable for analysis, as discussed further under Chapter 3.4 – Portfolio Construction. The estimate of value traded in this case was provided through the following formula:

$$\text{Value traded (R)} = \frac{\text{(Open+High+Low+Close Prices)}}{4} \times \text{Volume (in shares) traded} \tag{3.6}$$

Intraday range and the 1-week average intraday range: A crucial component of this research involves identifying extreme gaps and differentiating them from insignificant gaps. Overnight gaps in the South African stock market are not uncommon. In the data sample for this analysis, there was a total of 284,873 gaps. However, the majority of these gaps are small in size and cannot be considered meaningful. Considering an average daily range of 1.6% for the equities under analyses here, an overnight gap of 0.1% is not of a significant size, and prior research, such as that by Plastun et al. (2020) indicates that the relative size of the gap is a key influential factor in determining the subsequent market/price reaction. As such, for an analysis of gaps for the JSE-listed companies to be performed, a set of adaptive criteria needs to be implemented to aid in the differentiation of extreme and non-extreme gaps. One such option is to use a subset of the data, and based on this distribution implement a fixed absolute cut-off value. For example, based on the distribution of the data, the positive and negative absolute values may be determined such that a certain percentage (1% or 5% potentially) of gaps are larger, in absolute terms, than this value. There are numerous flaws to this approach, however. Firstly, in order to avoid hindsight bias, only a sample of the total data can be used. This is not ideal, as it has been shown with extensive research and evidence that market characteristics (such as volatility, volume and return tendencies) can change drastically during times of market crisis as opposed to more normal conditions (González-Hermosillo & Hesse, 2009). Secondly, as discussed earlier in Chapter 2, the characteristics of JSE-listed securities can change substantially over time. Notably, the average intraday and overnight volatility of stocks shows a wide range of values. In this case, using a fixed value to determine extreme gaps will skew the data in favour of stocks that naturally trade with higher levels of volatility, and less volatile stocks will erroneously be excluded.

It is proposed that the criteria for an extreme gap be adaptive to both company and market characteristics. In this case, the use of daily ranges becomes a valuable measure to use. The daily range is calculated as the difference between the intraday high and low, divided by the opening price. The 5-day average intraday range is then calculated on a rolling basis for each security. By using the average of the intraday range as a baseline for normal volatility, both stock characteristics as well as overall market conditions are taken into account. The average intraday range is thus a far better measure to use when identifying extreme gaps. In this case, an extreme gap is classified as an overnight gap which is larger, in absolute terms, than the 5-day average intraday range of the stock. The formula for the intraday range and average intraday range are:

$$\text{Intraday Range}_{\text{i,t}-1}(\%) = \frac{(\text{High Price}_{\text{i,t}-1} - \text{Low Price}_{\text{i,t}-1})}{\text{Open Price}_{\text{i,t}-1}} \times 100 \tag{3.7}$$

$$5 - \text{Day Average Intraday Range}_{\text{i,t}-1} = \frac{\sum_{\text{t}-5}^{\text{t}-1} \text{Intraday Range}_{\text{i,t}}}{5} \tag{3.8}$$

Although the dataset for this analysis is extensive and comprehensive, there are certain structural and data-based limitations that need to be considered before an accurate analysis can be done. Data cleaning and processing is outlined hereafter.

3.3. Data Cleaning & Outlier Detection

3.3.1. Missing Data

The starting point for data processing is the identification of corrupt and missing data. In this case, the data collected from Bloomberg is of a high quality, and as such minimal data loss and corruption was found. Notably, five companies (tickers ACP, AIB, BCX, FPT and TMG) did contain missing pricing information between 2000 and 2015. The source of the data was checked and found to be unrecoverable. As such, data for these companies was removed from 2000 to 2015.

Additionally, missing pricing information on individual days was also found in other companies, seemingly at random. The gaps in pricing information necessitated the removal of the affected days, as the analysis of gap behaviour would have been misleading or incomplete had they been included. These missing values accounted for less than 1% of the data.

3.3.2. Outlier Detection

With missing data removed, an analysis of corrupt data was undertaken. The starting point for the identification of erroneous or corrupt data was a plot of each of the key variables. Along with an analysis of summary statistics, errors in the data were revealed. Although there are potential explanations to some of the large data disparities observed (for example: stock splits/reverse splits for a disproportionately large jump in prices), due to the extended analysis timeframe, such fundamental information is no longer available. As such, extreme outlier values were identified and treated as erroneous values and removed from the data. One method that can be used to identify outliers is the Tukey 2.2 rule, this is what was applied to the identification and removal of outlier values. In terms of outlier detection techniques, the Tukey 2.2 method is preferred as it is comparatively simple to apply and has been found to be robust to extreme values and recursion (Hoaglin & Iglewitz, 1987).

Figure 3.1 below illustrates the distribution of gaps over the full period of analysis and before the removal of any outliers. As can be seen, there is a slight positive skew, with a marginally higher proportion of positive gaps. This can also be seen from the interquartile range given in Table 3.1 on page 22, with a median value of 0.1%.

Figure 3.1: Histogram of Overnight Gap Returns From 2000 To 2021

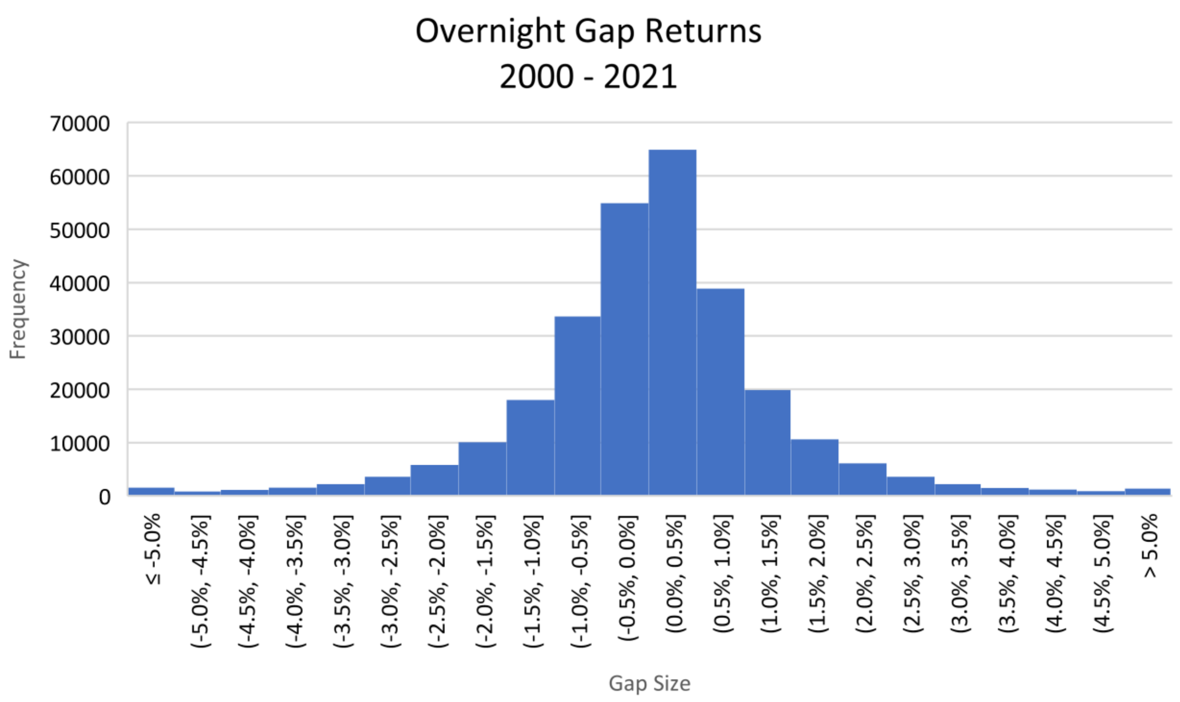

Table 3.1: Minimum, Maximum & Interquartile Range for All Gaps From 2000 To 2021

Minimum	Lower Quartile	Median	Upper Quartile	Maximum
-98.3%	-0.6%	0.1%	0.7%	8400.0%

The minimum and maximum values for gaps of -98% and 8400% respectively clearly indicate that there is either corrupt data, or there are extreme values in the dataset that have a technical/fundamental explanation that can no longer be accessed for clarification. In both cases, these extreme outlier values do not constitute anything close to normal trading ranges and values, and as such have the potential to falsely skew the results of further analysis.

The standard measures of outlier detection that revolve around the use of the interquartile range were all found to be too strict in trimming large values. This is likely due to the use of the entire data sample for calculating the interquartile range, wherein the majority of observations lie around the mean and median of 0.1%. As discussed in more depth earlier in Chapter 3 on page 19, these small values are far less valuable in an analysis of gap behaviour, and it is specifically at the extremes where the largest effects and predictive power is expected to exist.

A slightly revised approach to outlier detection is proposed in this case. The same Tukey 2.2 methodology is used to detect outliers, however the dataset to be used in determining these outlier values is the one that contains only the extreme gaps, for which the methodology of identification has been outlined earlier in Chapter 3 on page 19. Given that this dataset of extreme gaps already contains a much-reduced sample of gaps, in particular those that are meaningfully large in the context of the underlying volatility of the asset, the interquartile range calculated on these values gives a far larger threshold in terms of allowing for extreme values, while still eliminating erroneously large values. This adjusted interquartile range is provided in Table 3.2 on page 23, and clearly shows the larger range covered as a result of using the dataset of extreme gaps only.

Table 3.2: Tukey 2.2 Modified Minimum, Maximum & Interquartile Range for All Gaps From 2000 To 2021

Minimum	Lower Quartile	Median	Upper Quartile	Maximum
-98.3%	-3.4%	-1.4%	3.0%	8400.0%

Based on this interquartile range, the Tukey 2.2 cut-off values can be calculated according to the formulas for the lower and upper-bound cut-off values:

$$\mathrm{Lowerbound} = \mathrm{UQ} + \mathrm{k}(\mathrm{UQ} - \mathrm{LQ}) \qquad \text{where k is the constant of 2.2} \qquad (3.9)$$

$$\mathrm{Upperbound} = \mathrm{LQ} - \mathrm{k}(\mathrm{UQ} - \mathrm{LQ}) \qquad \text{where k is the constant of 2.2} \qquad (3.10)$$

The lower and upper-bound cut-off values using this technique are -17.5% and 17.1%, respectively. In comparison, the lower and upper-bound values calculated using the dataset of all gaps with the same methodology yielded cut-off values of -3.3% and 3.4% respectively. Had these values been used to trim outliers, more than half the observations of extreme gaps would have been removed. In this case, using the modified approach, only 77 values were removed from the dataset as extreme/erroneous outliers, corresponding to only 0.77% of the extreme gap observations, and 0.03% of the dataset of all gaps.

Finally, we are left with the post-trimming interquartile range in Table 3.3 below. As can be seen, the Tukey 2.2. method of outlier removal applied to extreme gaps only is successful in removing erroneously large data values, without significantly changing the core distribution of the data.

Table 3.3: Post Outlier Removal Minimum, Maximum & Interquartile Range for All Gaps From 2000 To 2021

Minimum	Lower Quartile	Median	Upper Quartile	Maximum
-17.5%	-3.4%	-1.4%	3.0%	16.8%

3.4. Portfolio Construction

With the data fully cleaned and prepared for further analysis, the portfolio of tradeable assets can be constructed. The method of stock selection involves first calculating the average daily value traded for each company in the first two calendar weeks of each year. With the daily value traded already calculated, it is a simple task to calculate the average of these values for each stock in the first two weeks of each year. The two-week average daily value traded figures give us a clear indication of liquidity levels across stocks. Stocks are then ranked in descending order of average daily value traded, and the bottom 50% of the most illiquid shares are discarded. In this case, shares that are too illiquid to be used in any extreme gap analysis are removed, and only the shares that are liquid enough will remain in the portfolio of applicable stocks. This process is repeated for each of the 21 years under analysis, such that the portfolio of applicable stocks is rebalanced each year. This is necessary as an analysis of market capitalisation and daily value traded over time reveals that a non-negligible portion of shares experience changes in their liquidity ranking year on year, leading to these shares dropping in and out of the portfolio of applicable stocks. By reconstructing the portfolio of applicable stocks annually as outlined above, the sample of stocks that are used for analysis will dynamically adjust to changes in the overall liquidity levels of the JSE as well as the number of listed securities on the exchange, thus ensuring a dynamically adjusted and representative sample throughout the entire period of analysis.

3.5. Descriptive Statistics & Exploratory Data Analysis

With the dataset cleaned and prepared for further analysis, we can gain a better understanding of the nature of the data by looking at some key summary statistics and visual plots. Table 3.4 on page 25 provides these summary statistics in tabular form.

Table 3.4: Summary Statistics Table for Key Variables of Interest

Variables					
	Extreme Gaps (%)	Intraday Returns (%)	Total Daily Return (%)	Daily Value Traded (ZAR)	5-day Average Intraday Range
Observations	10 757	10 757	10 757	10 757	10 757
Minimum	-17.5%	-28.3%	-39.2%	0	0.0%
Maximum	16.8%	28.6%	50.0%	1.33T	16.5%
Range	34.3%	56.8%	89.2%	1.33T	16.5%
Median	-1.4%	0.0%	0.0%	484M	2.4%
Mean	-0.3%	0.2%	0.0%	6.34B	2.6%
Standard Deviation	4.2%	3.4%	3.3%	26.96B	1.4%
Skewness	-	0.32	0.54	18	1.72
Kurtosis	-	*2.92*	13.86	626	7.26

Starting with extreme gaps, we note the decreased range of 34.3%, owing to the data cleaning and outlier removals performed earlier in Chapter 3. We can also identify an interesting shift in the median of the data. Comparing the median of all gaps identified earlier (0.1%) to the median of extreme gaps only (-1.4%), we can observe a significant decrease. This is also reflected to a lesser degree in the mean of all gaps calculated earlier (0.0%) compared to the mean of extreme gaps only (-0.3%).

This disparity in the measures of centrality of data indicate that negative gaps appear to contain a higher concentration of extreme values than positive gaps. This is not entirely unexpected, as research has largely shown that volatility during periods of market crisis, as characterised by falling asset prices, is much higher than under normal conditions, or broadly increasing asset prices (Bala & Takimoto, 2017). We would thus expect the higher levels of volatility during market turmoil (such as the 2008 global financial crisis) to impact gap sizes too, and this does appear to be reflected in the positively skewed distribution of extreme gaps.

Additional observations of note include the differences in range between intraday returns and total daily returns. Given that total daily returns consist of intraday returns and overnight returns, we can deduce the overnight returns range from the data collected. In this case, the total and intraday ranges were 50% and 28.6%, respectively, leaving a 21.4% range for overnight returns. This is proportionally large, particularly considering that trading activity only occurs during market hours on the

JSE. These overnight returns thus constitute large changes in price without any trading activity.

Additionally, we can compare our observations about the positive skew of overnight returns to that of intraday returns. In this case, the differences are marginal, with a mean return of 0.2% for intraday returns and -0.3% for overnight returns. Nonetheless, this constitutes an interesting observation, as the summary statistics indicate the possibility of conflicting return tendencies. Namely, as discussed above, overnight returns tend to have a higher concentration of extreme events/returns, as well as a positively skewed distribution, whereas intraday returns may be more negatively skewed, as evidenced by the mean return of 0.2%. If nothing else, this lays the foundation for an investigation into possible mean-reverting tendencies of intraday returns following overnight gaps, as opposed to a momentum effect.

Lastly, it is worth pointing out that the difference in overnight returns/gap sizes in comparison to the 1-week average intraday range can be quite large. If we compare the absolute minimum and maximum overnight gap returns to the average 1-week intraday range, we observe gaps at the extreme ends of greater than six times the average intraday trading range of a stock, even after trimming outliers. Particularly in conjunction with the conflicting return tendencies described above, these preliminary results and findings certainly provide a useful baseline for conducting further analysis into gaps, their behaviour and potential influence on subsequent returns.

Figure 3.2 on page 27 shows the frequency of positive and negative gaps over the full 21-year period of analysis, while Figure 3.3 on page 27 shows the average size of the extreme gaps each year over the same timeframe. These two graphs together allow us to get a better picture of the structure of extreme gaps in this dataset, by looking at both the size and frequency of these gaps, split by the sign of the gap.

Figure 3.2: Yearly Frequency Distribution of Positive and Negative Gaps

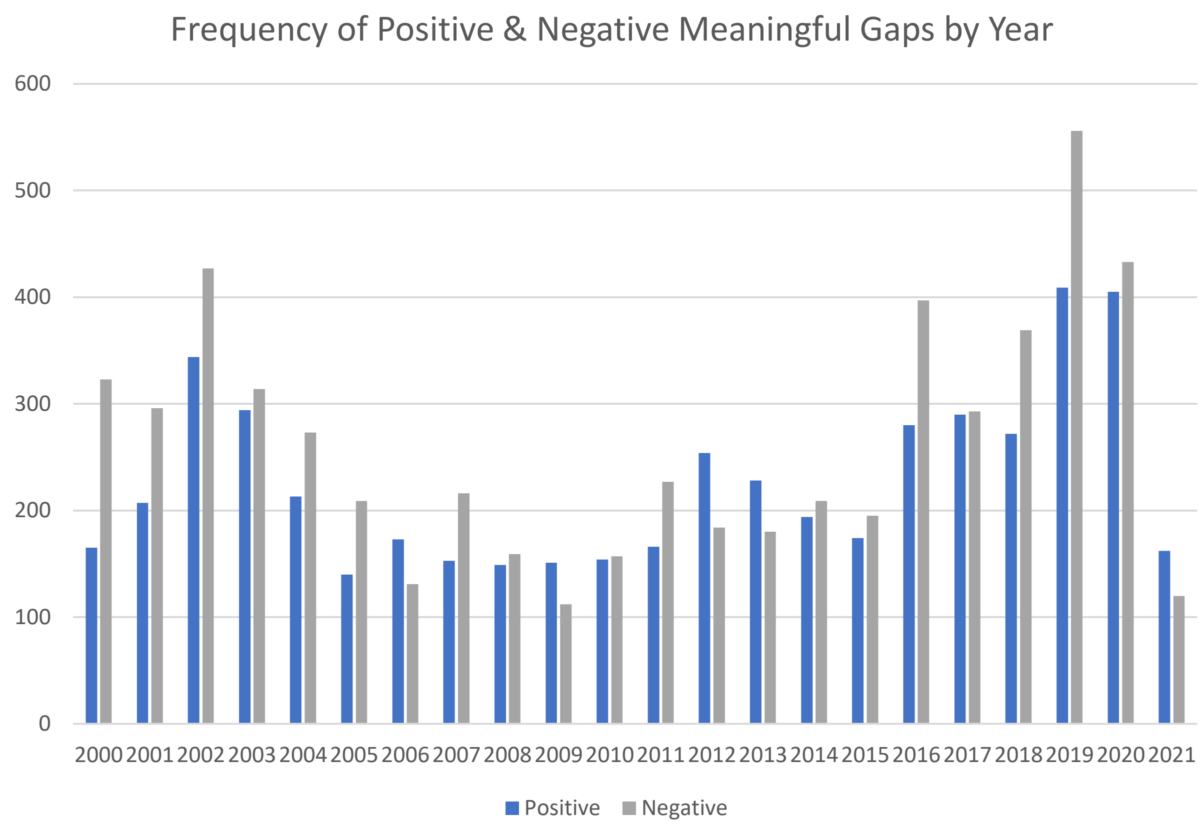

Figure 3.3: Yearly Distribution of Average Positive and Negative Gap Sizes

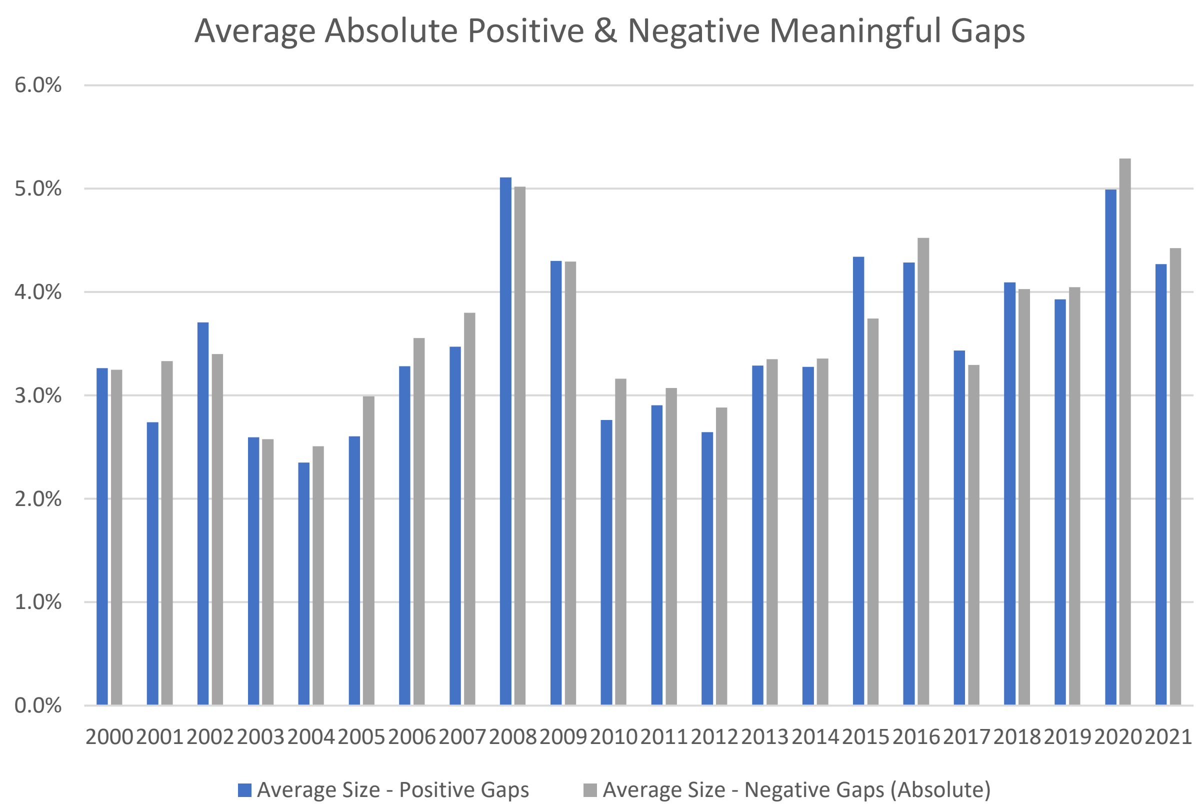

Looking at the frequency of positive and negative gaps each year, there are some noticeable divergences between the number of positive and negative gaps in some years. Overall, there is a definite bias towards a greater number of negative gaps occurring, and these observations match the analysis of descriptive statistics from Table 3.4 on page 25. It is also interesting to note that during the 2008 global financial crisis as well as the 2020 Covid-19 market crash periods, the frequency of positive and negative gaps appears to be matched evenly. The conclusion from this preliminary analysis is thus that there does not appear to be a greater concentration of negative gaps occurring during periods of market turmoil/downturn, which presents an interesting case for further research into whether gap behaviour may change in other aspects during market crash periods.

The graph of the gap sizes presents far less variation, in contrast. Both positive and negative gaps are very similar each year, with far less bias towards negative gaps also observed. It is worth noting the year-on-year variation in the total average gap size, however, as substantial variance does exist in this area. The difference between the lowest and highest average gap sizes for each year is large at a nearly 100% range, going from 2,3% to 5,1%. Looking at periods of market crisis, we can also observe a substantially greater average gap size during these periods, particularly around 2008 and 2020. Thus, while there does not appear to be a meaningful change in the average size of the gap during market crashes, the same cannot be said for the frequency of these gaps, and it appears that the greater frequency of gaps could be the driving force behind variations in returns associated with market crashes.

Figures 3.4 and 3.5 on page 29 illustrate the distribution of positive and negative gaps in percentage form over the entire period of analysis, rather than being grouped by year. While the distributions are closely matched, there is also a visibly thicker tail to the right for negative gaps. Considering the tendency for a greater frequency of negative gaps, we can also identify a slight bias towards larger negative gaps as well. In contrast, positive gaps follow a more clustered distribution with less bias, whereby more values are centred around the mean. Overall, it appears that negative gaps may be more impactful than positive gaps, due to both the increased frequency of these negative gaps as well as the slight skew towards larger values. The findings from the preliminary analysis of this chapter will now be explored in more depth in Chapter 4.

Figure 3.4: Histogram of The Absolute Value Sizes of All Negative Gaps From 2000 To 2021

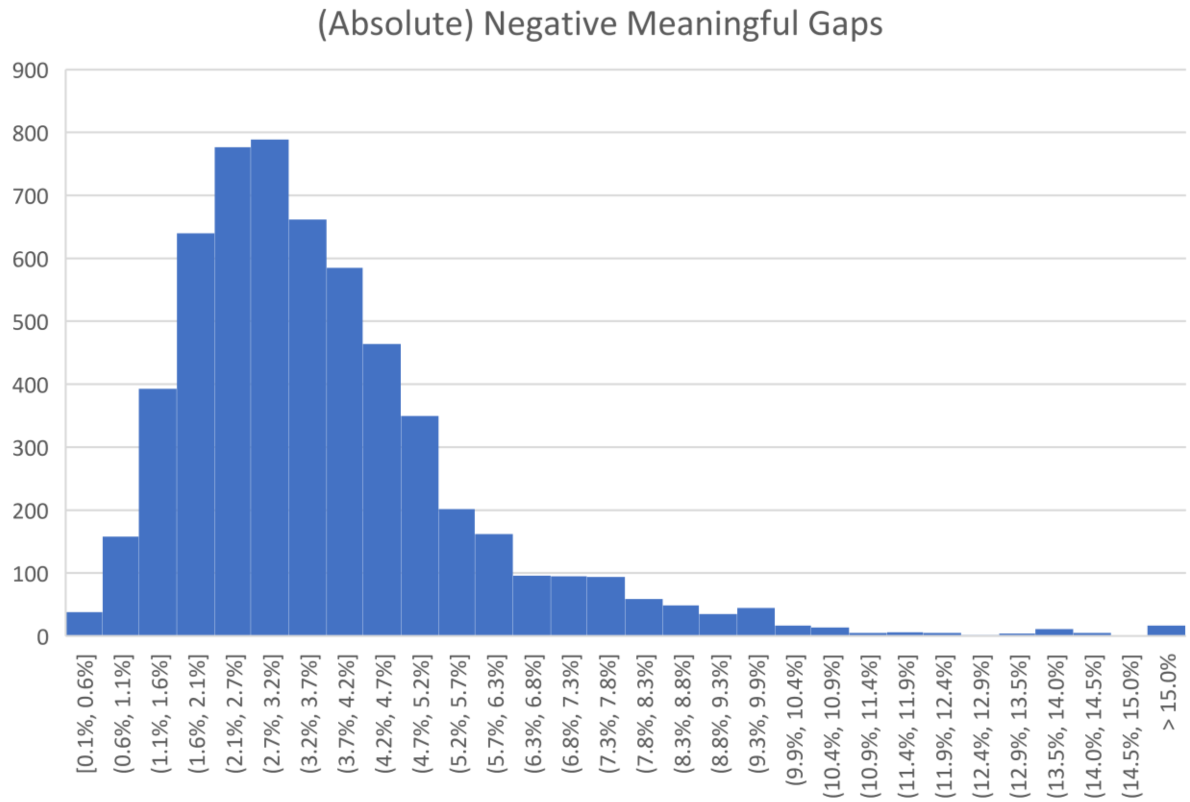

Figure 3.5: Histogram of The Sizes of All Positive Gaps From 2000 To 2021

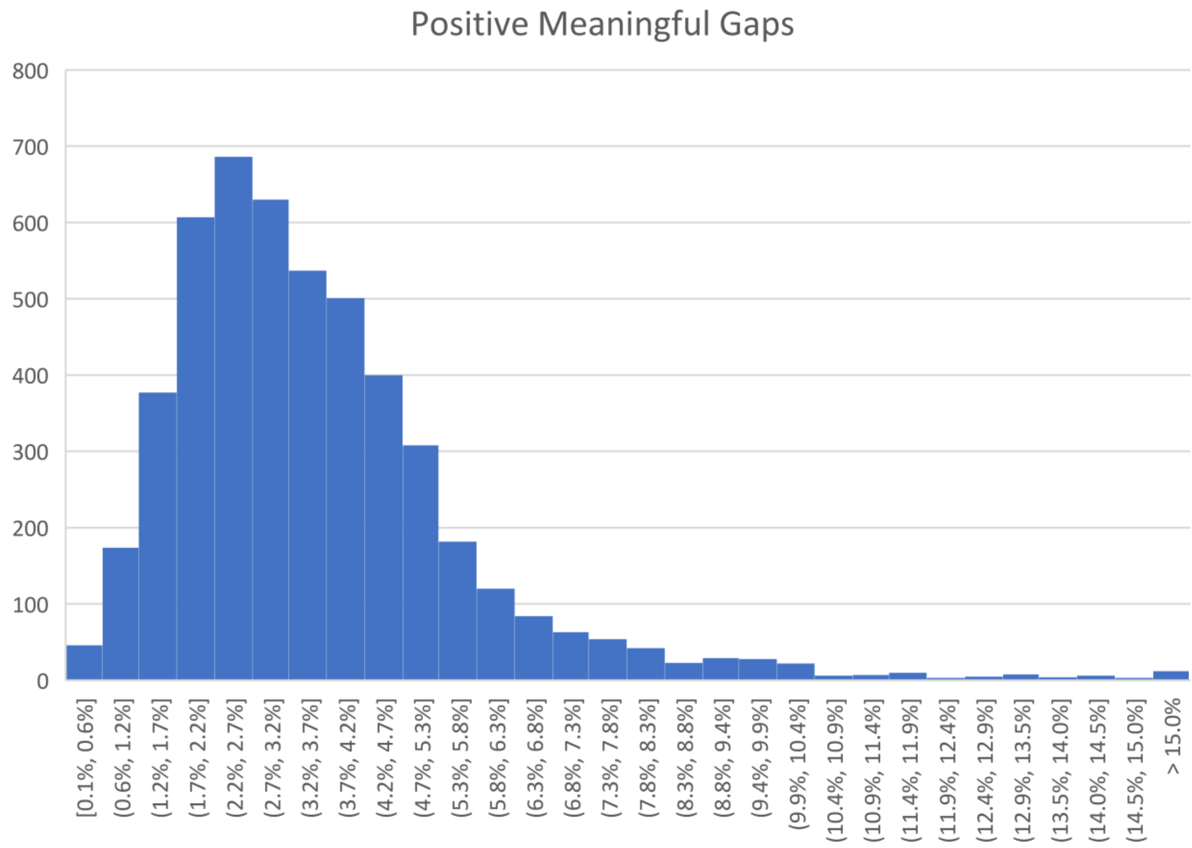

Chapter 4: Research Hypothesis Development & Methodology

Based on the empirical evidence presented to date as well as the preliminary data analysis and summary statistics of Chapter 3, a solid framework exists for establishing primary and secondary research objectives in this chapter.

4.1. Primary Research Objective

The predictive power of extreme overnight gaps on share returns

Do overnight gaps have any predictive power or influence over subsequent returns?

This is the core focus of this research, and the primary question that this thesis will answer. Prior literature from international markets, such as that by Plastun et al. (2020) and Serletis and Rosenberg (2009) as well as the authors own research in the South African market have found that gaps and gap behaviour are significant in predicting subsequent prices and returns, and thus a comprehensive analysis of the JSE is proposed.

4.2. Secondary Research Objectives

4.2.1. The effect of technical and/or fundamental data on gap behaviour and predictive power.

Do market or individual stock characteristics influence the effectiveness of gaps as a predictor of subsequent share returns?

This research question seeks to answer whether or not technical and/or fundamental data has any effect on gap significance or predictive power. Based on the authors own prior research as well as existing literature, such as that by (Liu et al., 2021). and Plastun et al. (2020), there is a case to be made that overall significance could be influenced by some external influences, be that individual stock characteristics, market influences or technical conditions. This research question will allow for the incorporation of additional factors that have been found to be significant predictors of subsequent share returns into the analysis of gaps. This should allow for a more comprehensive and robust model to be built, as well as an analysis of the potential interactions of predictive factors.

4.2.2. The effect of changes in the overall market and economy on predictive power.

Do certain sub-periods of market turmoil and volatility (such as the 2008 global financial crisis) have a significant effect on the overall predictive power and relationship between extreme gaps and subsequent returns?

A large amount of research has been conducted on market characteristics and how they may change during periods of market crises, and the general consensus is that many characteristics can change drastically during such periods, most notably volatility (Ellington, 2008). As such, a natural link exists that extends the research on gap behaviour to also look for potential changes in predictive power during periods of market turmoil.

4.2.3. The potential performance of a hypothetical trading strategy centred around overnight gaps.

Based on the findings of the primary and secondary research objectives above, is it possible to construct a trading strategy based on overnight gaps, and what does the performance of such a strategy look like?

Research conducted by Plastun et al. (2020), Dunis, Laws & Rudy (2011) and Stübinger & Schneider (2019) in international markets indicates the possibility of market beating returns through the creation of a relatively simple trading strategy that seeks to capture mean-reverting behaviour following abnormal overnight gaps. Not only does this research indicate that a strategy likely can be built, but it also provides a baseline model that can be used as a reference. Although most research along these lines focuses on a strategy based on mean-reversion, the details of the hypothetical trading strategy created later will be more heavily influenced by the results of the analysis performed in Chapter 5, as very limited research along these lines exists in the South African market.

4.3. Methodology

4.3.1. Multivariate Regression Model

Based on the precedence set by prior research, such as that of De Gooijer, Diks & Gatarek (2009) and Plastun et al. (2020), the primary method of testing and analysing the research hypotheses will be based on multivariate regression models. The purpose of a regression model is to describe the relationship between one or more independent variables and a single dependent variable. This aligns closely with the core objective of this research, namely, to identify the relationship, if it exists, between overnight gaps and subsequent returns.

4.3.2. Regression Construction

Given a random sample of n individuals and k independent variables, if we let y denote the dependent variable that is linearly related to k independent variables $X_1, X_2, \dots X_k$ through the parameters $\beta_1, \beta_2, \dots \beta_k$, a linear model can be constructed in the form: $y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_k X_k + \epsilon$ where ϵ is the random error component reflecting the difference between the observed and fitted linear relationship, and provided that the assumptions of a multiple linear regression model are met, namely:

(i) The population model is linear in parameters with additive covariates, such that the relationship between the dependent variable and independent variables can be represented in the form: $\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_k X_k + \epsilon$

(ii) The data is a randomly drawn sample from the population, with normally distributed residuals.

(iii) There is no multicollinearity between the independent variables.

(iv) Homoscedasticity of the error terms such that $Var(\epsilon_i | x_{i1}, x_{i2}, \dots, x_{ik}) = \sigma^2$

(v) Zero conditional mean, whereby $E(\epsilon_i | x_{i1}, x_{i2}, \dots, x_{ik}) = E(\epsilon) = 0$ and thus with condition (iv), $\epsilon \sim N(0, \sigma^2)$

If assumptions (i) to (v) are met, then the ordinary least squares (OLS) estimator of $\widehat{\beta_J}$ will be the most efficient estimator, resulting in the lowest variance of all the unbiased estimators of $\widehat{\beta_J}$. The OLS estimators of $\widehat{\beta_J}$ and $\widehat{\beta_0}$ are calculated by minimising the sum of squared residuals and result in the following β estimator equations:

$$\widehat{\beta_J} = \frac{\Sigma_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\Sigma_{i=1}^{n}(x_i - \bar{x})^2} \tag{4.1}$$

and

$$\widehat{\beta_0} = \bar{y} - \hat{\beta}\,\bar{x} \tag{4.2}$$

4.4. Model Construction

The primary method of statistical analysis employed by this research will be that of a series of linear regression models. Various models will be built, ranging from multiple linear regression models to multiple logistic regression models. The timeframe of analysis will be broken up into segments, such that the entire dataset is split in a ratio of 60:40. The first 60% of the dataset (in terms of length, in days) will be used as training data to build multiple linear and multiple logistic regression models, while the remaining 40% of the data will be used as test/unseen data to validate the robustness of the model and analyse the accuracy of the predictions made by the models on unseen data. In addition, following this analysis of accuracy, 2 further subperiods will be analysed independently, specifically from July 2008 to December 2009, and February 2020 to September 2020. The first subperiod corresponds to the 2008 global financial crisis, while the second subperiod corresponds to the 2020 Covid-19 market crash. Although the USA records the Covid-19 recession as officially only lasting for two months (National Bureau of Economic Research, 2021), the effects of this recession on the financial markets arguably extends past these two months to a longer time. As such, an additional six months of data are used to capture a greater portion of the market crash and recovery.

4.5. Model Specification

With initial data analysis and additional variable creation complete, a correlation matrix between the various relevant variables can be analysed to determine potentially significant variables to include in the various regression models that follow. A correlation matrix is a particularly useful measure of summarising data and obtaining an overview of the relationships between variables. Table 4.1 on page 35 provides the full correlation matrix that will be used for variable selection. Given that the first regression models will be built on only the first 60% of data, the correlation matrix is also only constructed with the first 60% of data in order to avoid look-ahead bias. The legend for this correlation matrix is provided below.

Legend

Variable	Description
Intraday Change (t)	The current intraday return following an overnight gap occurring
Gap (t-1)	The gap formed from the prior days close to today's open.
n-Day Total Return (t-n)	The total (overnight + intraday) returns of the previous n days.
Daily Value Traded (t-1)	The prior days total trading activity in Rands (shares * prices).
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Market Value (t-1)	The total company market capitalisation as at yesterdays close.
1-Week Average Intraday Range (t-1)	The rolling 5-day simple moving average of the intraday range of the equity, as at yesterdays close.

Table 4.1: Full Correlation Matrix of All Potential Explanatory Variables

	Intraday Change (t)	Gap (t-1)	1-Day Total Return (t-1)	2-Day Total Return (t-2)	3-Day Total Return (t-3)	5-Day Total Return (t-5)	Daily Total Value Traded (t-1)	Intraday Range (t-1)	Market Value (t-1)	1-Week Average Intraday Range (t-1)
Intraday Change (t)	—									
Gap (t-1)	**-0.608**	—								
1-Day Total Return (t-1)	**0.186**	-0.183	—							
2-Day Total Return (t-2)	0.059	-0.065	0.351	—						
3-Day Total Return (t-3)	0.068	-0.073	0.335	**0.972**	—					
5-Day Total Return (t-5)	0.067	-0.080	0.313	**0.932**	**0.954**	—				
Daily Total Value Traded (t-1)	0.003	0.008	-0.031	-0.021	-0.027	-0.021	—			
Intraday Range (t-1)	0.067	0.015	-0.011	0.005	0.004	-0.005	0.014	—		
Market Value (t-1)	-0.003	0.007	-0.016	-0.008	-0.010	-0.006	**0.496**	-0.041	—	
1-Week Average Intraday Range (t-1)	0.046	0.019	-0.028	0.003	-0.001	-0.003	0.029	**0.772**	-0.028	—

Looking at the dependent variable, Intraday change, the highest correlation is observed with the primary independent variable of this analysis, the overnight gap. A negative correlation of -0.6 is comparatively high and gives solid evidence of a potential inverse relationship between the overnight gap and intraday returns. The total 24-hour return preceding the gap exhibits a positive correlation to intraday returns. Although this correlation is smaller, it is still significant and worth exploring further in a regression model. The correlations between the subsequent independent variables and the dependent variable are comparatively much lower, and insignificant in the case of the market value and daily value traded. As such, the near 0 correlations of these variables to the dependent variable indicates that they are not good candidates to be included in the regression models and would only reduce the degrees of freedom and adjusted R-squared of these models without adding enough explanatory power.

Although correlations between the dependent and independent variables is important in determining which variables to include, it is also critical that correlations between independent variables be considered. Specifically, these correlations should not be too high or the assumption of no perfect multicollinearity in a linear regression model may be violated. Looking at the correlation matrix on page 35, it is clear that the 2-day, 3-day and 5-day total return values are extremely correlated to each other. This is unsurprising as these variables deal with the preceding cumulative return. As such, the 5-day prior return will necessarily include the values of the 1-day return through to the 4-day return as well. As well as having a near perfective positive correlation, the 2-day, 3-day and 5-day preceding returns also have a significantly lower correlation to the dependent variable, indicating that they will likely offer less explanatory power on intraday returns as well. As a result of these factors, the 2-day and 3-day preceding returns will be excluded from the model, and only the 1-day and 5-day prior returns will be included. Although the 3-day prior return does have a marginally more significant correlation to the dependent variable, this is offset by the proportionally much greater degree of collinearity to the one-day prior return. In comparison, the 5-day cumulative prior return has roughly the same degree of correlation to the dependent variable, but with a correlation to the 1-day prior return of 0.31 as opposed to 0.36. In addition, the 5-day prior return does include more pricing and return information than the 3-day prior returns, and thus from a theoretical perspective may also be preferred.

The 1-week average intraday range will also be excluded from the regression models. This is due to the prior days intraday range offering a higher correlation to the dependent variable in comparison to the rolling average. In addition, the correlation between the prior intraday range and the rolling average is concerning at 0.77. As a result of both of these factors, the 1-week rolling intraday range will be dropped from the regression models in favour of the preceding days intraday range only.

Lastly, based on the significant correlations of the prior days total return, this total 24 hour return will be broken down into its two components, the overnight and intraday returns, in much the same fashion that the overnight gap and intraday returns have been separated as primary independent and dependent variables, respectively. This will allow for a more granular look at what the underlying relationship may be between the prior days return and the current intraday return, as well as allow for a comparison between the two components of the total 24-hour returns in both the current and one-day prior periods, as opposed to looking at the prior days return as the full 24-hour return only. Table 4.2 below contains the updated correlation matrix based on the changes described above.

Table 4.2: Trimmed Correlation Matrix of Potential Explanatory Variables

	Intraday Change (t)	**Gap (t-1)**	**Gap (t-2)**	**Intraday Change (t-1)**	**Cumulative Return (t-5)**	**Intraday Range (t-1)**
Intraday Change (t)	—					
Gap (t-1)	**-0.608**	—				
Gap (t-2)	**-0.171**	0.151	—			
Intraday Return (t-1)	**0.303**	-0.287	-0.365	—		
5-Day Total Return (t-5)	0.067	-0.080	0.055	0.272	—	
Intraday Range (t-1)	0.067	0.015	-0.083	0.047	-0.005	—

Legend

Variable	Description
Intraday Change (t)	The current intraday return following an overnight gap occurring
Gap (t-1)	The gap formed from the prior days close to todays open.
Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
5-Day Total Return (t-5)	The total (overnight + intraday) returns of the previous 5 days.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.

The modified correlation table shows an interesting combination of effects between the intraday and overnight return components of the preceding days total return. Notably, the combined positive correlation of 0.186 observed in Table 4.1 can be seen as the summation of a negative correlation between the gap 2 days ago and the intraday return, and a stronger positive correlation between the prior days intraday return and the current intraday return. These results suggest that the inverse relationship between the overnight and intraday returns appears to hold over multiple timeframes, whereas the relationship between the intraday returns appears to be the opposite, with continuation patterns and momentum effects expected to be observed based on the correlations shown.

It is worth pointing out that one of the variables of interest described in Chapter 3, trend, is absent from the correlation matrix. This is due to the fact that the trend variable is categorical with 3 levels. The correlations of a categorical variable do not provide any valid interpretation, and hence this variable was omitted from the correlation matrix. Based on the results of this analysis, there is enough evidence to begin creating the various regression models using the reduced subset of applicable variables identified through the analysis of Chapter 3 as well as the correlation matrices in Chapter 4.

4.6. Regression Models

4.6.1. Model 1A (M1A)

$$Y_{i,t} = \beta_0 + \beta_1\left(Gap_{i,t-1}\right) + \beta_2\left(Intraday\ Return_{i,t-1}\right) + \beta_3\left(Gap_{i,t-2}\right) + \beta_4\left(Trend_{i,t-1}\right) + \beta_5\left(Cumulative\ Return_{i,t-5}\right) + \beta_6(Intraday\ Range_{i,t-1})$$

Where

Yi,t is the dependent variable, in this case the intraday return at t, the day on which a stock has gapped up or down. This return is calculated as follows:

$$Intraday\ Return_{i,t}\ (\%) = \frac{\left(Closing\ Price_{i,t} - Opening\ Price_{i,t}\right)}{Opening\ Price_{i,t}} \times 100 \qquad (3.2)$$

β0 is the intercept term of the regression. This term captures the predicted intraday return (Yi,t) if all the independent variables take a value of 0.

β1(Gapi,t−1) is the regression term that deals with the overnight gap. As discussed in Chapter 3, the overnight gap is calculated as follows:

$$Gap_{i,t}(\%) = \frac{\left(Opening\ Price_{i,t} - Closing\ Price_{i,t-1}\right)}{Closing\ Price_{i,t-1}} \times 100 \qquad (3.1)$$

$\beta_2(Intraday\ Return_{i,t-1})$ is the regression term that deals with the intraday return of the prior trading day only. The total return includes both the overnight return component as well as the intraday returns, and here the coefficient term deals only with the intraday portion of the total 24-hour return. It is calculated as follows:

$$(Intraday\ Return)_{i,t-1}(\%) = \frac{\left(Closing\ Price_{t-1} - Opening\ Price_{t-1}\right)}{Opening\ Price_{i,t-1}} \times 100 \qquad (3.2)$$

$\beta_3(\mathrm{Gap}_{i,t-2})$ constitutes the other portion of the prior 24-hour return, and thus in conjunction with the β_2 term breaks down the prior trading days returns into its two respective components. It is calculated as follows:

$$\mathrm{Gap}_{i,t-2}(\%) = \frac{(\text{Opening Price}_{i,t-1} - \text{Closing Price}_{i,t-2})}{\text{Closing Price}_{i,t-2}} \times 100 \qquad (3.1)$$

$\beta_4(\mathrm{Trend}_{i,t-1})$ refers to the categorical variable, outlined in Chapter 3, that indicates the approximate short-term trend of the stock based on the positions of the 9 & 20 SMA's as well as the delta, or change, of these SMA's. The trend variable is based on three categories, namely: Up, Down & Flat. The trend is classified as up/bullish if:

$$\begin{aligned} &9\mathrm{SMA}_{i,t-1} > 20\mathrm{SMA}_{i,t-1}\ \& \\ &\Delta 9\mathrm{SMA}_{i,t-1} > 0\ \& \\ &\Delta 20\mathrm{SMA}_{i,t-1} > 0 \end{aligned} \qquad (4.1)$$

Similarly, the trend is classified as down/bearish if:

$$\begin{aligned} &9\mathrm{SMA}_{i,t-1} < 20\mathrm{SMA}_{i,t-1}\ \& \\ &\Delta 9\mathrm{SMA}_{i,t-1} < 0\ \& \\ &\Delta 20\mathrm{SMA}_{i,t-1} < 0 \end{aligned} \qquad (4.2)$$

The trend is classified as flat/non-trending in all other cases.

$\beta_5(\text{Cumulative Return}_{i,t-5})$ is the total cumulative return of the previous 5 trading days, excluding the overnight gap at t-1. It is calculated as follows:

$$(\text{Cumulative Return})_{i,t-5}(\%) = \frac{(\text{Closing Price}_{i,t-1} - \text{Opening Price}_{i,t-5})}{\text{Opening Price}_{i,t-5}} \times 100 \qquad (3.3)$$

$\beta_6(\text{Intraday Range}_{i,t-1})$ is the regression term that deals with the prior trading days intraday range, it is calculated as follows:

$$\text{Intraday Range}_{i,t-1}(\%) = \frac{(\text{High Price}_{i,t-1} - \text{Low Price}_{i,t-1})}{\text{Open Price}_{i,t-1}} \times 100 \quad (3.7)$$

Timeframe: the timeframe over which this regression will be run is the first 12 years and 2 months of the dataset. This corresponds to 60% of all the data, as outlined earlier under Chapter 4.4 – Model Construction. The regression for M1A will thus contain data from 1 January 2000 to the 29 February 2012, with 5257 observations in this period.

4.6.2. Model 1B (M1B)

$$Y_{i,t} = \beta_0 + \beta_1(\text{Gap}_{i,t-1}) + \beta_2(\text{Intraday Return}_{i,t-1}) + \beta_3(\text{Gap}_{i,t-2}) +$$
$$\beta_4(\text{Trend}_{i,t-1}) + \beta_5(\text{Cumulative Return}_{i,t-5}) + \beta_6(\text{Intraday Range}_{i,t-1})$$

Model M1B is identical to M1A in all aspects bar the dependent variable. For this model, the dependent variable, $Y_{i,t}$ is categorical, as opposed to continuous as in model M1A. The categorical dependent variable still refers to intraday returns, but can now only take on one of two possible values:

Unchanged/closed lower, or
closed higher.

The choice of a categorical dependent variable necessitates that a logistic regression be used for M1B. A logistic regression, run in conjunction with a multiple linear regression can serve as a useful robustness check to the linear model. Although the interpretation of the coefficients will differ between the models, the sign of these coefficients should be the same. The binary outcome of the dependent variable is a simplification of the continuous variable case, and as such the logistic model is expected to offer a lower level of explanatory power. Nonetheless, it can still serve as a valuable tool for comparison to the multiple linear model M1A. Furthermore, from both an academic standpoint as well as for practitioners, the interpretation and usage of a logistic model also makes sense. If predictions on a continuous scale are not significant or accurate enough, simply being able to predict whether a stock will close higher or lower may be more useful.

Timeframe: the timeframe over which this regression will be run is the same as M1A, and the regression for M1B will contain data from 1 January 2000 to the 29 February 2012.

4.6.3. Model 2A (M2A) & Model 2B (M2B)

Following on from M1A and M1B, these regression models are identical except for the timeframe of analysis. M2A and M2B are the multiple linear and multiple logistic regressions run with the same independent variables that have been specified but will be constructed with the remaining 40% of the dataset, which will serve as the unseen/test data from which comparisons can be drawn between model sets 1 and 2.

Timeframe: the timeframe for model set 2 will span the period of 1 March 2012 to 30 April 2021, constituting 5500 unique observations.

4.6.4. Model 3 (M3)

Provided that the results of the robustness and accuracy checks of models M2A and M2B offer are satisfactory and valid, the next regression model built will be M3. This model will again be identical in variable choice and structure to previous models and differ only in the timeframe of analysis. M3 will constitute the complete, full-scale model. The regression will be run over the full timeframe of the dataset, from January 2000 to April 2021. This complete model will be used as a baseline for comparisons to all other subsequent models as well, most notably models M4 and M5 (financial crises sub-periods).

It is also worth noting that for this full period of analysis, the logistic regression models will no longer be included.

4.6.5. Model 4 (M4) & Model 5 (M5)

The last variations of the regression analysis will relate to the secondary objectives of this research, namely, to identify and analyse potential differences in the level of predictive power and accuracy gaps may have during periods of market crisis/instability. As before, the makeup of the regression models is unchanged in terms of structure and variables employed. M4 will deal with the 2008 global financial crisis, while M5 will deal with the 2020 Covid-19 recession. As discussed earlier in Chapter 4.4 – Model Construction, M4 will deal with the sub-period of July 2008 to December 2009. M5 will be regressed on data from February 2020 to September 2020.

4.7. Regression Assumption Checks

In order to conduct a valid and unbiased regression analysis, it is critical that the assumptions of a regression analysis are met. These assumptions have been briefly outlined earlier in Chapter 4.3.2 – Regression Construction but will be discussed in further detail below. The regression assumption checks will be analysed and discussed in depth in for Model 1A, and subsequently the results of the regression checks will be summarised for the remaining models.

4.7.1. Linear relationship and the additive property of covariates

The regression assumption of linearity necessitates that for the OLS estimators to be the most efficient, the population relationship between the dependent and independent variable/s must be linear in nature. Furthermore, covariates are also assumed to be additive. The influence of an indepenent variable on the dependent variable is required to be independent of any other variables and influences.

The assumption of a linear relationship is particularly important in validating the results of the regression. This is because if the assumption of linearity is violated, then not only will the ordinary least squares (OLS) estimators no longer be the most efficient, but bias will also be introduced into the model and its predictions. This can also manifest through the test statistics for the coefficient estimates being overinflated, and in extension lead to an incorrect conclusion on the significance of a predictor. In extreme cases, a non-linear relationship can also lead to completely inaccurate relationship being captured by the linear model.

The primary means to test for this assumption is by analysing the visual plots of the residuals to the fitted values of the model. This residual plot visually illustrates the difference between the actual observed values and the predicted response values. Given that the residuals are calculated as the difference between observed and expected values, plotting these residuals against the fitted values of the model should yield a random distribution of points. Ideally, the plot should have no discernible pattern, be relatively symmetrically distributed and with a central clustering tendency. If a clear pattern is present, this can indicate an unbalanced model, and a violation of the assumption of linearity and additive covariate, ultimately allowing for potential biased predictions and inefficiency of the OLS predictors.

Figure 4.1 on page 45 provides the plot of fitted and residual values. Although there do appear to be two distinct clusters of datapoints on each side of the mean of 0, this is not considered to a discernible pattern that may be problematic. These groupings can directly be attributed to the data cleaning and processing that was performed in Chapter 3, which has essentially led to the creation of a dataset with two distinct distributions on either side of the mean of approximately 0. A key part of this data processing was the removal of non-meaningful gaps, such that only extreme gaps that are significantly above or below 0% are included. This removal of insignificant gaps is clearly also reflected in the fitted values against residuals plot. Apart from this explainable and unproblematic clustering of groups, there does not seem to be any discernible pattern, and overall, the plot does not present any problematic patterns that may indicate a violation of the assumption of linearity.

Figure 4.1: Plot of Residual Values Against Fitted Values for M1A

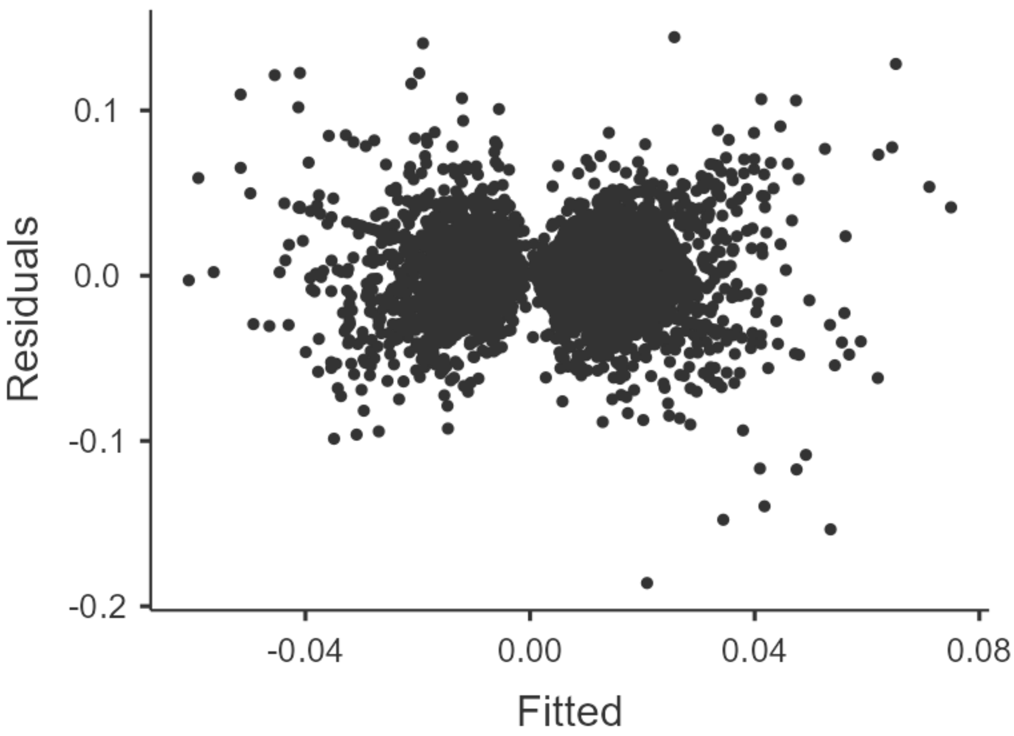

4.7.2. No multicollinearity between the independent variables

Collinearity occurs when independent variables are correlated with one another. This is problematic, as independent variables are required to be independent and uncorrelated to each other for an efficient regression. When a regression coefficient is being analysed and interpreted, this is done under the assumption that it is being done in isolation. That is to say, we are looking at the effect of an independent variable on a dependent variable, holding all other variables constant. If two or more of the independent variables covary, then it is impossible to look at the isolated effect of only one independent variable. As such, the interpretation can become inaccurate, as the standard errors of the regression coefficients will be affected in this case. The possibility exists for one or more of the covarying independent variables to become redundant.

The identification of multicollinearity is usually done through an analysis of a correlation matrix, as well as an accompanying level of knowledge on the topic at hand and variables involved in any interpretation of this matrix. The correlation matrix of variables along with an analysis thereof was performed earlier in Chapter 4 based on Table 4.1 on page 35. An additional statistical test that can provide insight on multicollinearity is the variance inflation factor (VIF). The VIF values provide the decimal odds of what percentage of the total variance in observed values is inflated for each of the coefficients. This necessitates that the VIF values have a minimum value of 1, corresponding to a 0% inflation increase in variance for a given coefficient. Similar to the correlation matrix, the interpretation of the VIF values can be subjective, and there are no universally agreed fixed values that indicate an existence, or lack of, multicollinearity. In general, however, certain ranges of values can provide a good indication of whether or not there may be multicollinearity. Values greater than 5 are considered to be problematic and indicate a high degree of multicollinearity, with values close to 1 indicating very low levels of multicollinearity that will not violate the regression assumption of no perfect multicollinearity.

The VIF values can be calculated by regressing an individual predictor against every other predictor in the model. This provides the R-squared value of an individual coefficient, and can be used to calculate the VIF value of that predictor according to the following formula:

$$VIF_i = \frac{1}{(1-R_i^2)} \tag{4.3}$$

The table of VIF values for each of the predictors in M1A is presented in Table 4.4 on page 47. We can clearly identify that the VIF values for all variables are close to 1, with the largest VIF value being that of the cumulative 5-day prior return at 1.32. This is still considered to be close to 1, and well below the threshold of 5 for problematic multicollinearity.

Table 4.3: Collinearity Statistics for Predictor Variables in M1A

Collinearity Statistics

	VIF	Tolerance
Gap (t-1)	1.01	0.990
Gap (t-2)	1.09	0.918
Intraday Change (t-1)	1.19	0.837
Trend (t-1)	1.06	0.947
Cumulative Return (t-5)	1.32	0.760
Intraday Range (t-1)	1.00	0.996

4.7.3. Homoscedasticity of the error term

Another assumption critical to the successful interpretation of a linear regression is that of homoscedasticity. Homoscedasticity in this case refers to the variance of the error term being constant across all independent variable values, while the complementary notion is called heteroscedasticity. Homoscedasticity is required by nature of how the OLS method of regression coefficients are calculated. The OLS method gives equal weighting to all observations, and thus assumes that errors have a constant variance. If they do not, then errors that have either a greater or lesser variance than the assumed constant level will be disproportionately weighted, leading to the OLS estimators no longer being the most efficient. If heteroscedasticity is present, the weighted least squares method may be preferred, or alternatively a data transformation.

The primary method of identifying homoscedasticity is a combination of analysing the plot of the fitted values of the model plotted against the residuals. This analysis can be supplemented with a statistical test as well, such as the Harrison-McCabe test for heteroscedasticity. The residual/fitted plot is contained in Figure 4.1 on page 45, as it was already provided and analysed as part of the assumption of linearity, but this figure is also repeated on the following page.

Figure 4.1: Plot of Residual Values Against Fitted Values for M1A

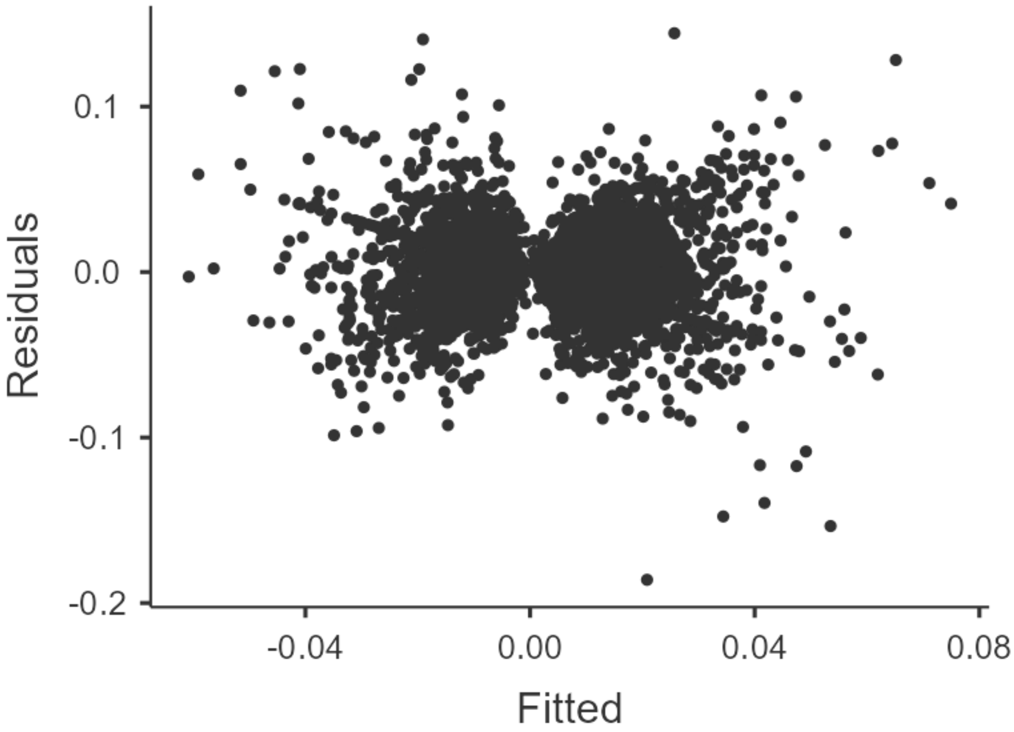

For the case of homoscedasticity, the plot should again contain no discernible pattern, and in particular should not contain a “fanning” pattern of residuals, identifiable by the variance of the residuals increasing as the fitted values increase. Figure 4.1 does illustrate some potential signs of increased variance for higher fitted values, however this appears to be marginal and should not violate the assumption of homoscedasticity, particularly given the large sample size and data pruning undertaken to focus exclusively on large/extreme cases of overnight gaps. In this case, the plot of residuals against fitted values does not provide a clear enough indication of a violation of the homoscedasticity assumption, and a quantifiable statistical test may be needed. The Harrison-McCabe is presently one of the most powerful tests of heteroscedasticity (Uyanto, 2019). The Harrison-McCabe test quantifies the proportion of the residual sum of squares that corresponds to the proportion of data before the chosen breakpoint (Harrison & McCabe, 1979). In the default case, the test statistic should thus be close to 0.5, and if it deviates significantly from the breakpoint value then the null hypothesis of homoscedasticity is rejected.

The test statistic of model M1A was calculated to be 0.483 with a corresponding p-value of 0.05. As was the case with the visual analysis of the residuals/fitted plot, there is insufficiently strong evidence to conclude decisively that the errors are heteroscedastic, but simultaneously the evidence of homoscedastic errors is also weak. Both the visual plot and accompanying test statistic are thus marginal, however, based on the aforementioned data processing and its effect on the distribution of the residuals, the results of this regression are expected to remain valid and efficient. Nonetheless, the case of potential heteroscedasticity of the error terms may justify the use of a modified regression method, such as the weighted least squares method, and this is identified as an area of potential future research.

4.7.4. No serial autocorrelation

This assumption also relates to the error terms, and states that the error terms of a regression must be independent for each observation. Serial autocorrelation refers to the case where an error term of one or more of the observations is influenced by the error term of another observation. The violation of this assumption is particularly prevalent in time-series data, and most commonly at lag 1, whereby an error term is influenced by the immediately preceding error term. For an OLS regression, serial autocorrelation is problematic because of the potential effect that the lack of error term independence has on the standard errors of the estimates, with a tendency of underestimating the standard error. This in turn affects the test statistic of an individual coefficient as well as the confidence interval, and in turn can lead to an incorrect conclusion on the overall significance of a predictor.

One of the most commonly used methods of identifying serial autocorrelation is the Durbin-Watson test. The test statistic is calculated as a function of the error terms of the regression, and tests for the existence of first order correlations in these error terms. The test statistic ranges between values of 0 and 4. A value of 2 indicates no first order autocorrelation exists. Values from 0 up to 2 indicate positive autocorrelation, and correspondingly values between 2 and 4 indicate negative autocorrelation. In general, a test statistic between 1.5 and 2.5 is considered to be normal, with a test statistic smaller or greater than these values considered to be potentially problematic with regards to the assumption of no serial autocorrelation.

The Durbin-Watson test statistic of model M1A is 1.79, with the interpretation that assumption of no serial autocorrelation is not violated.

4.7.5. Normally distributed residuals

The error terms of an OLS linear regression are assumed to be normally distributed, thus $\epsilon \sim N(0, \sigma^2)$. If they are not normally distributed, then this implies that part of the model is left unexplained, and this spills over into the residuals. It is worth noting however that the assumption of normally distributed residuals should be met given a sufficiently large sample size, as well as a correctly specified model (as evidenced by the other regression assumptions being met). In this case, the central limit theorem should dictate that the residuals tend towards a normal distribution. Furthermore, the assumption of normally distributed residuals ranks lower on the list of relative importance. Even if this assumption is violated, the coefficient estimates will not become biased, nor inefficient. A violation of this assumption will most likely show up in the significance of individual coefficient estimates, however the greater concern when the residuals are not normally distributed is that the model has been misspecified, or a key predictor variable omitted. If this is the case, then the model as a whole may be significantly less powerful than it should be.

The assumption of normally distributed residuals is often identified through an analysis of the Q-Q plot of residuals. Although statistical tests do exist that aim to provide a quantifiable and absolute interpretation of the normality of the distribution, it is more common to identify an approximate interpretation from the visual plots, particularly in conjunction with an analysis of the underlying distribution of the data in the regression. The Q-Q plot is created by plotting the theoretical quantiles of the normal distribution on the x-axis, and the ordered observed values of the variable we wish to compare to the normal distribution (the error terms of the regression model, in this case) on the y-axis.

Figure 4.2 on page 51 provides the Q-Q plot of the residuals. Under a normal distribution, the majority of data point should lie along the 45-degree line, with minimal variation in the upper and lower extremities. The Q-Q plot of model M1A illustrates a deviation from the normal distribution, with some variation. In particular, the plot shows that the distribution of residuals contains more outlier values than would be expected under a normal distribution. This can be seen by the plot of points deviating from the standard line at extreme positive and negative values. While the majority of data points do lie closely to the normal line, particularly around the centre, the outlier values indicate the residuals follow a distribution that may be too fat in its tails. It is also worth noting that the distribution of deviations from the normal line is approximately symmetrical, indicative of an error term distribution that is not meaningfully skewed, either positively or negatively.

Figure 4.2: Q-Q Plot of Standardized Residuals for M1A

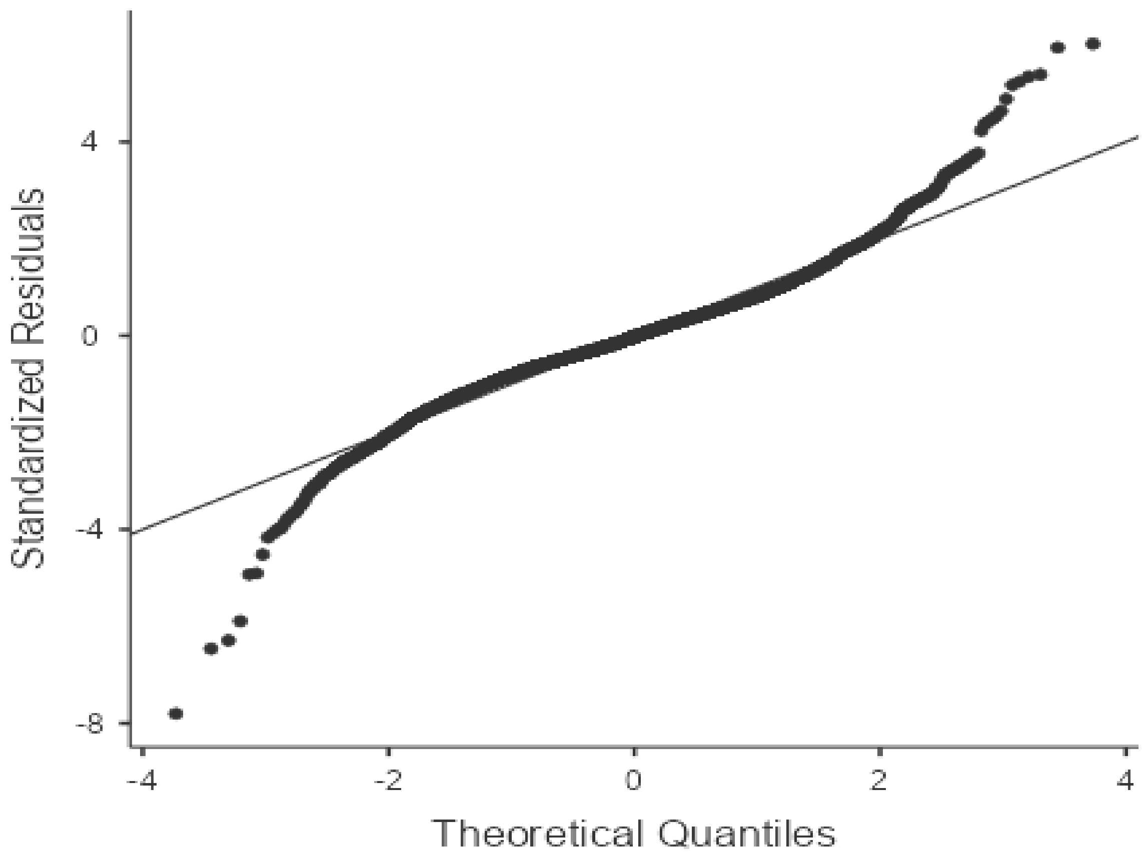

The conclusion from the analysis of the Q-Q plot is that the distribution of residuals is approximately normal in structure, but potentially with an excess of outlier values, identifiable by a distribution with fatter tails than the normal distribution. This potential violation is estimated to be relatively marginal, and once again likely stems from the data processing that was performed to specifically focus on extreme values in this analysis. As such, it is unsurprising that the area where the assumption of normality may be violated is related to outlier values and is not estimated to be problematic. In the event that the distribution is too dissimilar to that of the normal distribution, the net effect on the validity and interpretation of the regression results is also minimal, as estimates are neither biased nor inefficient.

4.8. Models M2 – M5 Assumption Checks

Appendix 1 provide the same statistical tests, graphs and assumption checks just analysed for Models M2A, M3, M4 and M5. Namely, this includes the plot of residual values against fitted values, table of VIF values, Harrison-McCabe homoscedasticity test statistic, Durbin-Watson autocorrelation statistic and Q-Q plot. Results from these regression checks are largely similar to those of model M1A. Interestingly, in the case of Models M4 and M5, the test statistics and graphical analysis indicates a greater margin of error in terms of acceptance of the regression assumptions. These two models are notable for their smaller sample size. In contrast, the largest model in terms of sample size, M3, contains the test statistics and visual plots indicative of a higher number of outliers. Considering the data pruning that was performed in Chapter 3 as well as the primary focus on extreme gaps in this analysis, these results are not too unexpected. Despite these outlier values, the regression assumption checks for M2 – M5 are not considered to be problematic, nor necessitating a data transformation.

Having fully specified the regression models as well as verify that the assumptions of these linear regressions have been met, we can now analyse the output of these models.

Chapter 5: Regression Results & Empirical Analysis

5.1. Introduction

Chapter 3 and Chapter 4 have dealt extensively with the collection and preparation of the data to be used in this study. Given the large sample size and specific focus of this research, these steps are crucial. With data correctly pruned and prepared, as well as the regression assumption checks met to a satisfactory level, we can now safely employ and analyse the results of the regression models built in Chapter 4. Model coefficients and descriptive statistics are provided for each of the models, from M1A to M5, with the results of these models discussed extensively as well. A set of pooled regression is also built in order to analyse differences between models in a more quantifiable manner, whereby differences in coefficients can be compared with statistical significance. The results of these comparisons are also extensively analysed in terms of the primary and secondary research objectives of this research.

5.2. Model 1A

Table 5.1: Regression Coefficient Estimates and Model Summary Statistics for M1A

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
(Intercept)	-5.74e−4	6.50E-04	-0.8826	0.377
Overnight Gap (t-1)	-0.3978	0.00874	-45.5139	**<0.001**
Overnight Gap (t-2)	-0.0621	0.02386	-2.6027	**0.009**
Intraday Return (t-1)	7.90E-04	0.0188	0.042	0.966
Trend: DOWN (t-1)	-1.20e−4	8.35E-04	-0.1439	0.886
Trend: UP (t-1)	8.97E-04	8.08E-04	1.1107	0.267
5-Day Total Return (t-1)	0.0129	0.00951	1.3556	0.175
Intraday Range (t-1)	0.0484	0.01601	3.0262	**0.002**

R-Squared	Adjusted R-squared	Akaike Information Criterion (AIC)	Model Test Statistic	Test Statistic p-value	Degrees of Freedom	Observations
29.3%	29.3%	-24264	310	<0.001	7	5257

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Linear	
Regression Timeframe: 2000/01/01 – 2012/02/29	

Looking at the output from the multiple linear regression on the training data in Table 5.1, it is immediately evident that the main independent variable of interest, the overnight gap, is highly significant, with a similarly large estimated effect size. The coefficient of -0.4 indicates that, all else equal, a 1% increase in the overnight gap return/size is predicted to lead to an intraday return of -0.4%. This constitutes a strong inverse relationship and confirms the expected results that were theorised from the preliminary data analysis performed in Chapters 3 & 4, such as the correlation matrix analysis.

The predicted effect of the overnight gap 2 days before is unsurprisingly estimated to be considerably smaller. In part, this is due to the effect of time passing and the various external factors eroding the observable relationship between the gap (t-2) and the intraday returns today. In addition, the dataset was trimmed in order to focus only on large and meaningful preceding gaps (t-1). This same data processing was not performed to remove small gaps preceding the extreme ones (gaps at t-2). As such, a brief analysis reveals that the average gap size at t-2 is indeed several times smaller than the average gap size at t-1. Nonetheless, the coefficient estimate of -0.06 is non-negligible, and is also statistically significant with a p-value of less than 1%. Perhaps the most useful takeaway from including the t-2 gap in this regression is to identify and confirm the predicted inverse relationship that exists between the overnight gap returns an intraday returns. Although the estimated effect is smaller, the same inverse relationship is also observed with the gap at t-2.

We can also observe that the coefficient estimates for the preceding intraday return, trend variable and cumulative prior 5-day returns are all small in size, as well as statistically insignificant. It is interesting to note that the predicted relationship between all three variables and the intraday return have the same estimated relationship, one of continuation. In contrast to the predicted inverse relationship between gaps and intraday returns, these variables are estimated to move in the same direction. That is to say, an increase in the prior days intraday return is predicted to increase the current intraday return as well. It is unsurprising that the same relationship is also observed for the prior 5-day return, considering that the prior intraday return is a component of the 5-day return. Interestingly, the 5-day return has a larger and more statistically significant estimated effect, potentially indicative of a lagged reaction to new market information by market participants. Similarly, an uptrend is predicted to lead to a higher intraday return, and a down trend the opposite. It is worth emphasizing again that the predicted effects in all three cases is very small, and statistically insignificant, and as such the interpretation of the predicted effects is only valid in the context of comparisons to the other regression models, and not necessarily useful for estimating future returns.

Another interesting observation from the coefficient estimates relates to the significance of the prior trading days intraday range. There is an estimated positive relationship between the prior intraday range and the current intraday return. The coefficient estimate of 0.05 indicates that a 1 percent increase in the prior trading days range is predicted to lead to a 0.05% increase in the current intraday return. Although this is not a large effect, the estimate is statistically significant at the 1% level, with a p-value of 0,2%. As a whole, the training data model illustrates a clear divide in influential factors impacting intraday returns. Specifically, the particularly strong relationship between the overnight gap and intraday returns appears to be mean-reverting/inverse in nature. On the other hand, all additional variables in the model are predicted to change positively with the intraday return. This includes the prior return variables, trend variable and intraday return measure, although the statistical significance of these estimates varies, with some being highly insignificant.

Lastly, the model summary statistics can also provide valuable insights into the overall success of the model at capturing the underlying relationships that exist in the data and give a broader indication of how well the model has performed.

The F-statistic and associated p-value provide insight into the significance of the model as a whole. Specifically, a significant F-statistic and p-value indicate that at least one coefficient in the model is a statistically significant predictor of the dependent variable. In this case, the F-statistic is extremely large, with a p-value near 0. This is not unexpected, as the model contains several statistically significant predictors, most notably the overnight gap. The R-squared value is arguably the most significant and often-quoted value indicating the effectiveness of a regression model. The R-squared value indicates the percentage of variation in the dependent variable that the model as a whole is able to capture, with higher R-squared values being indicative of a better model, all else equal. The adjusted R-squared is very similar, however an adjustment is made to counter the effect of simply adding variables to the model. As such, only meaningful independent variables will increase the adjusted R-squared, and insignificant predictor variables that are included will decrease the adjusted R-squared. In this case, both the standard and adjusted R-squared values are 29.3%. This indicates that the training model with its 6 independent variables was able to explain 29% of the variation of intraday returns. For a simple linear regression model with only 6 independent variables, some of which are statistically insignificant in this dataset, this is a comparatively high R-squared value, and indicates that the model as a whole is a good predictor of intraday returns, even if the close to two thirds of the variation in intraday returns is still left unexplained.

Model 1B is the logistic regression model on the same variables and training data. It is expected that the model should offer less explanatory power overall, however its primary purpose is to serve as robustness check, and as such the focus will be on identifying coefficient signs and significance levels. Table 5.2 on page 57 provides the coefficient estimates and model summary statistics for Model 1B.

5.3. Model 1B

Table 5.2: Regression Coefficient Estimates and Model Summary Statistics for M1B

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
(Intercept)	-0.35	0.0611	-5.735	< .001
Overnight Gap (t-1)	-25.355	0.9496	-26.701	**< .001**
Overnight Gap (t-2)	-2.349	2.4247	-0.969	0.333
Intraday Return (t-1)	-0.548	1.9108	-0.287	0.774
Trend: DOWN (t-1)	-0.197	0.078	-2.52	0.012
Trend: UP (t-1)	0.191	0.0742	2.579	**0.01**
5-Day Total Return (t-1)	0.785	0.9551	0.822	0.411
Intraday Range (t-1)	4.033	1.68	2.401	**0.016**

McFadden R-squared	Akaike Information Criterion (AIC)	Model Test Statistic	Test Statistic p-value	Degrees of Freedom	Observations
13.4%	6277	967	<0.001	7	5257

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Logistic	
Regression Timeframe: 2000/01/01 – 2012/02/29	

Looking at the logistic model, we can immediately identify the differences to the multiple linear model through the coefficient estimates. In this case, the dependent variable is binary, while all of the independent variables, bar the trend variable, are continuous. The trend variable is categorical with 3 levels. The coefficients of a logistic regression provide the log odds of the binary dependent variable, in this case, whether the intraday return was positive or negative/unchanged. Largely, the regression coefficients match those of the multiple linear model, with the only exception being the preceding intraday return. The logistic model indicates a negative, inverse relationship, whereas the multiple linear model estimates a positive relationship. This conflict of sign estimates is not particularly problematic or surprising, as in both cases the estimate is very small and highly insignificant. The matching signs of the remaining independent variables as well as the overall model significance provide encouraging results that the multiple linear model has correctly captured part of the variation of intraday returns in the test data.

Table 5.3 and 5.4 on pages 58 and 59 provide the regression models and statistics for M2A and M2B. These models will now be analysed together.

5.4. Model 2A (M2A)

Table 5.3: Regression Coefficient Estimates and Model Summary Statistics for M2A

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
(Intercept)	-0.0046	8.21E-04	-5.608	< .001
Overnight Gap (t-1)	-0.49735	0.00926	-53.7	**< .001**
Overnight Gap (t-2)	-0.00813	0.02443	-0.333	0.739
Intraday Return (t-1)	0.2734	0.01821	15.016	**< .001**
Trend: DOWN (t-1)	-0.00245	9.47E-04	-2.586	**0.01**
Trend: UP (t-1)	0.00139	9.32E-04	1.496	0.135
5-Day Total Return (t-1)	-0.00957	0.00366	-2.614	**0.009**
Intraday Range (t-1)	0.15772	0.0166	9.5	**< .001**

R-Squared	Adjusted R-squared	Akaike Information Criterion (AIC)	Model Test Statistic	Test Statistic p-value	Degrees of Freedom	Observations
46.8%	46.8%	-23543	691	<0.001	7	5500

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Linear	
Regression Timeframe: 2012/03/01 – 2021/04/30	

5.5. Model 2B

Table 5.4: Regression Coefficient Estimates and Model Summary Statistics for M2B

Term	**Estimate**	**Standard Error**	**Test Statistic**	**P-value of Test Statistic**
(Intercept)	-0.245	0.0794	-3.085	0.002
Overnight Gap (t-1)	-33.28	0.9558	-34.818	**< .001**
Overnight Gap (t-2)	-1.541	2.7837	-0.554	0.58
Intraday Return (t-1)	17.306	2.2548	7.675	**< .001**
Trend: DOWN (t-1)	-0.163	0.0869	-1.882	0.06
Trend: UP (t-1)	0.159	0.0835	1.907	0.056
5-Day Total Return (t-1)	-0.294	1.1465	-0.256	0.798
Intraday Range (t-1)	5.849	2.0639	2.834	**0.005**

McFadden R-squared	**Akaike Information Criterion (AIC)**	**Model Test Statistic**	**Test Statistic p-value**	**Degrees of Freedom**	**Observations**
30.8%	5293	2346	<0.001	7	5500

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Logistic	
Regression Timeframe: 2012/03/01 – 2021/04/30	

What is immediately noticeable from the test data model is the even larger predicted effect of the overnight gap on intraday returns, along with higher degree of statistical significance. The coefficient of -0.5 corresponds to a strong inverse relationship, and is significantly larger than the coefficient of -0.4 estimated with the training data model. This pattern of increased significance and predictive power continues with some of the other predictor variables, with a stark contrast observed for predictors like the preceding intraday return, which is estimated to have a much larger effect in the test data. This coefficient was found to be highly insignificant in the training data, but in the test data the opposite is observed, with the estimate now highly significant with a p-value of less than 0.01%. The overall significance of the model is also substantially higher, with R-squared and adjusted R-squared values of 46.8%, indicating that the test model is able to explain nearly half of the variation in intraday returns following a large gap.

A similar trend can also be observed for M2B. Although the primary purpose of this model is to serve as a robustness check for the multiple linear regression model, we can also compare the logistic regression coefficients to each other for the training and test periods of data. The extreme gap coefficient estimate of -33 for M2B is also observed to be larger than the coefficient estimate of -25 for M1B. Another noteworthy difference is the overall model significance of M2B at 30.8% in comparison to the M1B's model significance of 13.4%. Additional results for M2B are largely in line with expectations, and as was the case before, the results of the logistic model confirm those of the multiple linear model, whereby coefficient signs are the same across models. In addition, trends between the two sets of models are also consistent, and the results of these model comparisons does not provide any reason to doubt the accuracy of the multiple linear regression models. Overall, the estimates from the in and out-of-sample models show that the multiple linear regression models have been correctly specified, and as such are accurately capturing the underlying relationships in the data.

While a side-by-side comparison between the two sets of regression models is useful in terms of identifying and analysing differences between coefficient estimates and model significance for in and out-of-sample periods, this approach is more subjective than objective. A more objective comparative analysis can be performed by means of a pooled regression. This involves adding a dummy variable to the regression models previously specified, one which indicates the period that each observation is in (either training or test/unseen data), denoted by a 0 or a 1, respectively. Each term in the regression model is multiplied by this dummy variable, such that an interaction term is added to the previously estimated regression coefficients. In the basic case of a simple linear regression with one independent variable and one interaction dummy variable, estimates of the dependent variable are obtained through the following modified formula, where D_i is the dummy variable:

$$Y_i = \beta_0 + \beta_1 D_{1i} + \beta_2 X_i + \beta_3 (D_{1i} X_i) + u_i \tag{5.1}$$

The addition of this interaction term changes the interpretation of the regression coefficients slightly. β_0 represents the slope estimate when all independent variables are 0 and the dummy variable is also 0, while β_1 represents the additional estimate of the change in slope when all independent variables are 0 but the dummy variable = 1. The interpretation is similar for the β_2 and β_3 coefficients. β_2 is the estimated effect of the independent variable X_i on the dependent variable, Y_i when the dummy variable = 0 (the base case). The new interaction term of β_3 thus provides the estimated additional influence that the independent variable X_i has on Y_i where the dummy variable = 1.

In the case of a dummy variable indicating what period an observation is in, training or test, the β_3 term thus refers to the estimated additional influence that the predictor in question has on the dependent variable in the test/unseen data. By nature of how this pooled regression is constructed, the significance, or effectiveness of the interaction term showing the additional influence also be analysed by looking at the coefficient test statistic and p-value. A large test statistic and small associated p-value indicates that the estimated additive effect of the predictor in the unseen data is statistically significant. Table 5.5 provides the pooled regression between M1 and M2.

5.6. M1 & M2 Comparison

Table 5.5: Regression Coefficient Estimates for Combined Comparison Regression of M1A and M2A

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
Intercept (Training)	0.0019	2.61E-04	7.280	< .001
Intercept (Test)	-0.00105	5.21E-04	-2.015	0.045
Overnight Gap (t-1) (Training)	-0.3978	0.00645	-61.674	< .001
Overnight Gap (t-1) (Test)	-0.09956	0.0129	-7.718	**< .001**
Overnight Gap (t-2) (Training)	-0.0621	0.01735	-3.579	< .001
Overnight Gap (t-2) (Test)	0.05396	0.0347	1.555	0.12
Intraday Return (t-1) (Training)	7.90E-04	0.01336	0.059	0.953
Intraday Return (t-1) (Test)	0.2726	0.02673	10.198	**< .001**
Trend: DOWN (t-1) (Training)	-1.20E-04	6.36E-04	-0.189	0.85
Trend: DOWN (t-1) (Test)	-0.00233	0.00127	-1.835	0.067
Trend: UP (t-1) (Training)	8.97E-04	6.20E-04	1.447	0.148
Trend: UP (t-1) (Test)	4.97E-04	0.00124	0.401	0.689
5-Day Total Return (t-1) (Training)	0.0129	0.00552	2.337	0.019
5-Day Total Return (t-1) (Test)	-0.02246	0.01104	-2.034	0.042
Intraday Range (t-1) (Training)	0.0484	0.01171	4.133	< .001
Intraday Range (t-1) (Test)	0.10928	0.02341	4.668	**< .001**

Looking at the coefficient estimates side by side in the same model, differences become even more apparent. As was mentioned in the analysis of M2A, the primary explanatory variable of interest is significantly larger in both absolute and relative size in the unseen data. This is also reflected in the pooled regression, whereby the additive influence of this variable in the unseen data is estimated to be -0.1 (or -10%) decrease in intraday returns following an extreme gap. This additive influence is also highly statistically significant, and thus we can conclude with confidence that the inverse relationship between overnight gaps and intraday returns is even stronger in more recent data. Considering the small average daily trading ranges of the large capitalisation stocks in this sample, this constitutes a very strong relationship.

Interestingly, the 2-day preceding gap has lost a significant amount of predictive power and statistical significance. Whereas a 1% increase in the gap from 2 days ago was predicted to decrease intraday returns today by roughly 0.062% in the training data, the estimated additional effect of this variable in the unseen data is a nearly proportional increase in expected intraday returns of 0.054%. In totality, these two influences are nearly cancelling each other out. The change in signs between the two sample periods as well as the statistical insignificance of the coefficient of additional influence in the unseen data casts some doubt over the effectiveness of the gap from 2 days ago as a predictor of current intraday returns. This is not unexpected, as the passage of time would likely erode the ability to isolate the effect of the gap from 2 days ago, and hence the estimated effect in totality is considered to be inconsistent and relatively insignificant as well.

In contrast, a drastic change is observed for the prior intraday return variable. The predicted effect is small in the training data, with a coefficient of only 0.0001, or 0.1%. In the unseen data there is an extreme increase, whereby the additive influence of this predictor on intraday returns is estimated at 0.27%. The estimated signs of these coefficients are consistent however, suggesting that the underlying relationship has been correctly captured, and the large differences in size may be attributable to changes in market conditions. This contrast also reflects in the significance of the estimates, whereby the coefficient estimate in the training data was found to be highly insignificant, and the opposite was observed for this coefficient in the unseen data. It is also interesting to note the contrast between the effects of overnight and intraday returns, whereby a strong momentum effect is observed when looking at intraday returns, whereas strong mean-reversion tendencies exist for overnight/gap returns. Although the effect of the gap is predicted to be larger, the final intraday return can be observed as something of a battle between the continuation effect of the prior days intraday returns and the mean-reverting tendencies of the overnight gap, particularly in more recent data. Overall, there appears to be a large shift in market characteristics, whereby a strong momentum/continuation effect is present in more recent data. While the underlying cause for this shift is unclear, this does present an interesting observation for potential future research.

Looking at the categorical trend variable, increased significance is observed for one of the two cases, namely the bearish trend. In contrast, decreased significance is observed in the test data for the bullish trend case, leading to a mixed interpretation overall. The coefficient signs are again unchanged, however both the predicted effects as well as statistical significance is low in both cases. Overall, we can deduce that the trend variable is not effective in predicting intraday returns. This may be due to the structure of the market or the methodology used to create the trend variable. In either case, the trend variable can likely be excluded or reworked in future models without negatively affecting the predictive power of the model overall.

The trend variable is not necessarily the best predictor of intraday returns, and this appears to be the case for the 5-day cumulative prior return as well. Although the coefficient estimates are statistically significant at a 5% level in both cases, in both seen and unseen data the size of the effect is marginal. In addition, while most other coefficient signs have remained unchanged across the test and training data, this is not the case for this variable. A coefficient sign change across timeframes is not necessarily cause for concern, but in the context of other variables in this model measuring similar effects across these two periods, such as the gap and previous intraday return, there is reason to believe that the 5-day prior return is not a significant or accurate predictor variable of intraday returns in this case.

Lastly, the intraday range of the prior trading day is also predicted to be much larger and more significant in the unseen data, with an increase in size of more than 2x observed. In both cases, the coefficient estimates are also highly statistically significant. The positive coefficient sign aligns closely with the positive coefficient estimate of the prior intraday return. Despite the low correlations, the intraday range and intraday return are expected to move together in terms of their sign estimates. This is indeed observed, with a strong continuation effect for both variables visible, particularly on the unseen data. As a whole, the in and out-of-sample pooled regression has served as an excellent robustness check, as well as highlighted the changes in both the size and significance of the predictors across training and unseen data in a quantifiable manner. Based on these results, the three most influential variables are identified as: the overnight gap, prior intraday return and prior intraday range.

Overall, coefficients in the unseen data appear to be larger in size and statistical significance on average. This is an excellent sign that the model has performed well at capturing and explaining the underlying relationships between the intraday return and the explanatory variables. Quite often, models tend to perform worse on unseen data in comparison to the training data. This is because various explicit and implicit optimisations occur during model construction, where the data and variables are observable in their totality by the analyst. In this case, for both the model and many of the coefficients to have increased in significance using unseen data is a welcome surprise, and is likely indicative of two things – firstly, that the market characteristics may have changed over time such that gaps have played an increasingly significant role in predicting intraday returns, and secondly that the model built on the first 60% of the dataset has accurately identified the intrinsic underlying relationships that exist in the market across the entire dataset.

Table 5.6 on page 66 provides the regression output for Model 3, the full model built after M1 and M2 on the entire period of analysis, from 2000 – 2021.

5.1. Model 3

Table 5.6: Regression Coefficient Estimates and Model Summary Statistics for M3

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
(Intercept)	-0.00246	5.25E-04	-4.68	< .001
Overnight Gap (t-1)	-0.46304	0.00643	-72.04	**< .001**
Overnight Gap (t-2)	-0.0687	0.01672	-4.11	**< .001**
Intraday Return (t-1)	0.17541	0.01256	13.97	**< .001**
Trend: DOWN (t-1)	-0.00171	6.37E-04	-2.69	**0.007**
Trend: UP (t-1)	0.00116	6.22E-04	1.87	0.062
5-Day Total Return (t-1)	-0.00705	0.00325	-2.17	**0.03**
Intraday Range (t-1)	0.1046	0.01151	9.09	**< .001**

R-Squared	Adjusted R-squared	Akaike Information Criterion (AIC)	Model Test Statistic	Test Statistic p-value	Degrees of Freedom	Observations
39.5%	39.5%	-47392	1002	<0.001	7	10757

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Linear	
Regression Timeframe: 2012/03/01 – 2021/04/30	

While differences between in and out-of-sample data have been analysed from the data in Table 5.5 on page 62, Model 3 is also useful as an extension of this analysis, as M3 constitutes a combination of M1A and M2A. Overall, coefficients are largely in line with expectations based on M1A and M2A, as well as the analysis of the pooled regression. It is interesting to note that all coefficients, bar the trending variable, are statistically significant at a minimum significance level of 5% in this model. This is potentially partly due to the large sample size decreasing standard error values and thus increasing the significance of individual coefficient estimates, however there is also evidence to suggest that, overall, the model and the independent variables have captured and explained a significant portion of the variance in intraday returns. This can best be seen by the models adjusted R-squared value of 39.5%. While not as high as M2A's adjusted R-squared value of 46.8%, this is still a comparatively high value, and indicative of an effective and well specified model.

Table 5.7 below as well as Table 5.8 on page 68 provide the regression output for models M4 and M5, the regression models run specifically on the 2008 global financial crisis and 2020 Covid-19 recession sub-periods, respectively.

5.2. Model 4

Table 5.7: Regression Coefficient Estimates and Model Summary Statistics for M4

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
(Intercept)	-5.33e−4	0.00331	-0.161	0.872
Overnight Gap (t-1)	-0.56429	0.02977	-18.952	**< .001**
Overnight Gap (t-2)	-0.16227	0.09586	-1.693	0.091
Intraday Return (t-1)	0.0349	0.06385	0.547	0.585
Trend: DOWN (t-1)	-0.00126	0.00391	-0.323	0.747
Trend: UP (t-1)	0.00104	0.00371	0.279	0.78
5-Day Total Return (t-1)	-0.02508	0.03646	-0.688	0.492
Intraday Range (t-1)	0.10691	0.05743	1.861	0.063

R-Squared	Adjusted R-squared	Akaike Information Criterion (AIC)	Model Test Statistic	Test Statistic p-value	Degrees of Freedom	Observations
49.1%	48.2%	-1754	58.0	<0.001	7	430

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Linear	
Regression Timeframe: 2008/07/01 – 2009/12/31	

5.3. Model 5

Table 5.8: Regression Coefficient Estimates and Model Summary Statistics for M5

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
Intercept (Test)	-0.00831	0.00363	-2.291	0.022
Overnight Gap (t-1)	-0.39006	0.02978	-13.097	**< .001**
Overnight Gap (t-2)	-0.13354	0.07298	-1.83	0.068
Intraday Return (t-1)	0.25515	0.05893	4.33	**< .001**
Trend: DOWN (t-1)	-0.01113	0.00415	-2.685	**0.007**
Trend: UP (t-1)	0.0071	0.00453	1.565	0.118
5-Day Total Return (t-1)	0.02503	0.02938	0.852	0.395
Intraday Range (t-1)	0.22089	0.04814	4.589	**< .001**

R-Squared	Adjusted R-squared	Akaike Information Criterion (AIC)	Model Test Statistic	Test Statistic p-value	Degrees of Freedom	Observations
45.8%	45.0%	-1889	60.8	<0.001	7	512

Legend

Variable	Description
Overnight Gap (t-1)	The gap formed from the prior days close to today's open
Overnight Gap (t-2)	The gap formed from the 2-day prior close to open difference.
Intraday Return (t-1)	The prior days intraday return, from open to close.
Trend: Up & Down	A dummy variable with 3 levels, where the base case is non-trending. The trend is estimated based on a set of moving average conditions from the 9 & 20 SMA's.
5-Day Total Return (t-1)	The total return from 5 days ago to yesterdays close.
Intraday Range (t-1)	The trading range (in %) from the previous intraday session.
Regression Type: Multiple Linear	
Regression Timeframe: 2020/02/01 – 2020/09/30	

The regression timeframes for M4 and M5 correspond to the two largest periods of market turmoil in the dataset, characterised by increased trading activity, volatility and overall a large enough market decline to constitute a financial recession. It has been well documented that economic and financial crashes are subject to interesting variations in otherwise stable relationships, and in particular it has been theorised as well as observed that periods of market panic can be subject to increased levels of market inefficiency. As a result, following the results of the analysis of gaps so far, it is logical to expect that these relationships may be even stronger during market crashes.

Between the two models, results are mixed. While M4 has the highest adjusted R-squared of all the models created, the only statistically significant predictor in this model is the extreme overnight gap. This variable also has the largest estimated effect out of all the models at -0.56, indicating that gaps did appear to have a larger effect on intraday returns during the 2008 global financial crisis, when compared to ordinary market conditions. In contrast, the coefficient estimate for the main gap in M5 is substantially lower at -0.39, indicating that the increased predictive power of gaps observed for the 2008 financial crisis does not appear to repeat in the Covid-19 market crash, however it should be noted that there is less data for this period. Coefficient estimates from the other independent variables are largely in line with expectations.

While the two market crash periods have been separated into two models (M4 & M5) for the purpose of fidelity, the focus of this analysis will be on the pooled regression estimates, provided in Table 5.9 on page 70 that compare market crash conditions to ordinary trading conditions. For the purposes of this pooled regression, the two market crash sub-periods have been combined. In this case, the training data is the full dataset, excluding financial crash data, and the unseen data is data from financial crash periods only.

5.4. M3 & M4&5 Comparison

Table 5.9: Regression Coefficient Estimates for Combined Comparison Regression of M3 and M4&5

Term	Estimate	Standard Error	Test Statistic	P-value of Test Statistic
Intercept (Training)	0.00215	0.000272	7.923	< .001
Intercept (Training)	3.47E-04	0.000975	0.3555	0.722
Overnight Gap (t-1) (Training)	-0.46377	0.00675	-68.69	< .001
Overnight Gap (t-1) (Test)	0.00677	0.0174	0.3888	0.697
Overnight Gap (t-2) (Training)	-0.06269	0.01756	-3.57	< .001
Overnight Gap (t-2) (Test)	-0.0698	0.04669	-1.4949	0.135
Intraday Return (t-1) (Training)	0.16866	0.01343	12.56	< .001
Intraday Return (t-1) (Test)	-0.00274	0.03541	-0.0775	0.938
Trend: DOWN (t-1) (Training)	-9.51E-04	0.000641	-1.48	0.139
Trend: DOWN (t-1) (Test)	0.00756	0.00227	3.3354	**< .001**
Trend: UP (t-1) (Training)	0.00107	0.000623	1.72	0.085
Trend: UP (t-1) (Test)	0.00925	0.00259	3.5788	**< .001**
5-Day Total Return (t-1) (Training)	-0.0084	0.00319	-2.64	0.008
5-Day Total Return (t-1) (Test)	0.029	0.01729	1.677	0.094
Intraday Range (t-1) (Training)	0.1	0.01266	7.9	< .001
Intraday Range (t-1) (Test)	0.07264	0.03044	2.3864	0.017

In contrast to models M1 and M2, the differences shown in this pooled regression are much smaller. Notably, the coefficient estimates are similar and statistical significance levels are consistently lower in the unseen data. This is possibly due to the sample sizes, where the financial crash dataset is comparably much smaller than the dataset of ordinary trading conditions. All else equal, the larger the sample size, the more likely a coefficient is to be statistically significant, by nature of the calculation of significance levels and associated p-values. In the case of the main overnight gap, the predictor is expected to increase intraday returns by 0.0067% during financial recession periods. Considering that the estimated effect of this variable is a strong inverse relationship, an increase represents a weakening of the predictive power, albeit very slightly at 0.0067%. This predicted additional effect is also not statistically significant with a p-value of 69.7%.

The gap from 2 days ago does yield some interesting results, however. The predicted effect of this gap on intraday returns is estimated to be twice as large during periods of market turmoil. That is to say, while a 1% increase in the size of the gap from 2 days ago is predicted to decrease intraday returns by 0.063% during normal market conditions, the same 1% increase in gap size is predicted to decrease intraday returns by 0.13% during periods of market turmoil. Although this estimated additional effect is not very significant at a significance level of only 13.5%, it is interesting to note, and may hint at higher levels of return persistence during market crashes.

The previous days intraday return as well as the 5-day prior trading return also show some minor differences in coefficient estimates and statistical significance. Although the coefficient signs differ across the two periods, in relative terms the differences are negligible, and the statistical significance levels for these estimates are again low. The notable exception in terms of significance levels is the trend. Whereas this variable has largely been insignificant in other models, the additional influence of this variable during market crash periods is statistically significant in this case, however the size of the effect itself is very small. Overall, there is weak evidence to show that core relationships between the predictive variables and the intraday return changes significantly during periods of market instability.

5.5. Conclusion

The results of the analysis performed in this chapter have been very valuable in answering the core question of this research. Through the use of 9 regression models in total, a strong and persistent inverse relationship between the overnight gap and intraday returns has been observed. An in and out-of-sample analysis allowed for the unbiased creation and forward testing of these models, with results showing that the core relationships in the data have correctly been captured. This is evidenced by a greater level of predictive power and statistical significance observed, on average, when unseen data is used. Importantly, this was also observed to be the case for the primary variable of interest in this study, the overnight gap.

Out of the 6 independent variables that were originally specified, not all were found to be statistically significant predictors of intraday returns, with the overnight gap being by far the strongest predictor of intraday returns in all cases. An additional aim of this research was to identify if gap behaviour and predictive power changes during periods of market instability, such as global financial recessions. The results of the comparative models built do not support this theory, with the conclusion being that the predictive power of gaps does not change significantly during periods of market stability. Rather, it appears that changes in market structure that may influence the predictive power of the variables analysed here will likely take a considerable amount of time to develop, as shown by the increased predictive power of the overnight gap in the latter half of the dataset, a change which took over a decade to materialise fully.

Chapter 6: Trading Strategy Analysis

6.1. Introduction

Based on the results of the models built in Chapter 5, there is clear evidence of the mean-reverting behaviour between extreme overnight gaps and intraday returns. One of the additional objectives of this research is to analyse what the results of a possible trading strategy built around gaps would be. A simple trading strategy is proposed, one that aims to isolate the performance of trades based only on the extreme gaps that occur in the market. As opposed to the multiple regression models built with several independent variables, the trading strategy proposed here will only use extreme gaps as an input criterion and open a trade in the opposite direction of the gap. The results of this trading strategy will be analysed from a quantitative and qualitative perspective, as well as the implications of this approach. The results aim to replicate a real-world scenario, and as such transaction costs and other real-world limitations will also be considered extensively.

6.2. Transaction Costs

The JSE is Africa's largest stock exchange, however it is still considered to be part of a developing economy, and as such both the depth and breadth of the exchange are substantially lower than many international alternatives, particularly in the USA (Desjardins, 2016). One of the major drawbacks of trading on a smaller exchange, such as the JSE, is that the market structure is considerably less efficient than developed markets and exchanges. There is, quite simply, less competition in all aspects of the market, from market participants such as traders to market intermediaries, like brokers and clearing firms. An unfortunate consequence of this is that the average trading/transaction costs are much higher on the JSE than they are at many other international exchanges. While extreme competition has driven the average total fee a round-trip trade down to less than 0.05% of the market value of the trade from a reputable broker based in the USA, the same cannot be said for South African stockbrokers. Transaction costs for South Africa equities range from 0.5% to nearly 3% per round trip trade, on average. As such, the single biggest limiting factor affecting a potential trading strategy is its ability to generate an edge large enough to also overcome the steep transaction costs of trade on the JSE.

6.2.1. Transaction cost alternatives

Given the steep transaction costs associated with trading equities in South Africa, an alternative method of stock trading is also proposed, one that carries significantly lower fees overall: Contracts for Difference (CFD) trading.

As opposed to directly buying or selling the underlying asset/equity, as is the case with traditional stockbroking, CFD's are a derivative that allow a trader to gain the same exposure to the underlying asset through an artificial tradable instrument whose value is derived from the underlying asset in question. A CFD is "an agreement to exchange the difference in the price of an asset from the point at which the contract is opened to when it is closed." (IG, n.d.). Similar to other some other derivatives, trading this instrument allows a trader to avoid many of the other fees associated with trading the underlying asset, such as Securities Transfer Tax (STT).

The trading strategy formulated here will be subject to both the fees of CFD's as well as outright stock trading, with performance comparisons to the main benchmark of the AllShare Index also present.

6.3. Equity Trading: Transaction Costs

In total, the transaction costs associated with trading equities come from 6 components, namely the:

- Broker Fee
- Securities Transfer Tax (STT)
- Investor Protection Levy
- Strate Levy
- Value Added Tax (VAT)
- A liquidity/slippage factor

6.3.1. Broker Fee

The largest portion of the total trading costs, the broker fee constitutes a percentage of the market value of the trade that the stockbroker charges for each transaction that gets processed related to the purchase or sale of equities. This fee is the primary source of income for a stockbroking firm and covers their costs of accepting orders from clients and routing them to the market to get executed. Table 6.1 below provides the broker fees for some of South Africa's most popular stockbrokers (as ranked by JSE recommendation and internet search priority, listed in descending order by fee).

Table 6.1: Comparison of South African Broker Fees Associated with Equity Trading

BROKER	BROKER FEE
EASYEQUITIES	0.25%
SHARENET SECURITIES	0.30%
ABSA STOCKBROKERS	0.40%
SANLAM SECURITIES	0.42%
AFRIFOCUS	0.50%
MOMENTUM SECURITIES	0.50%
STANDARD BANK SECURITIES	0.50%
NEDGROUP PRIVATE WEALTH	0.50%
FNB SECURITIES	0.50%
BP BERNSTEIN	0.60%
THEBE STOCKBROKING	0.63%
PSG SECURITIES	0.75%

**Where sliding scale rates are applicable, the average of these rates was used*

(EasyEquities, n.d., Sharenet, n.d., Absa, n.d., Sanlam, n.d., Afrifocus, n.d., Momentum Securities, n.d., Standard Bank Securities, n.d., Nedbank Private Wealth (2022), FNB Securities, n.d., BP Bernstein, n.d., TSB Securities, n.d., PSG Securities, n.d.)

While the median fee of 0.5% is comparatively high, the lowest available broker fee at the time of writing is 0.25%. For the purposes of inclusion, the broker fee to be used in further analysis will be calculated as an average of the two lowest broker fees, which equates to an effective 0.275% fee.

6.3.2. Other Transaction Costs

In addition to the broker fee, the cumulative effect of other costs associated with trading equities in South Africa can also be substantial. Whenever a share is purchased on the JSE, a Securities Transfer Tax is levied on this transaction. Set at 0.25% for buyside transactions only, this fee is simply an additional form of revenue collection from the South African Revenue Service (SARS). In addition to the STT, SARS also collects VAT on the remaining transaction costs (excluding the STT), charged at 15% of the transaction cost value. These costs include the broker fee, Investor Protection Levy and Strate Levy. The Investor Protection Levy is collected directly by the JSE at a rate of 0.0002% of the market value of the trade and is collected in order to fund investigations into market manipulation and insider trading. Lastly, the Strate Levy, charged at 0.005787% is paid to the central securities depository and central collateral platform, Strate, which is a facilitator of trading on the JSE.

6.3.3. Liquidity & Slippage Factor

The last transaction cost to be considered here is implicit, rather than explicit as those considered above have been. A crucial factor that needs to be taken into account whenever placing a trade is the potential market impact that the trade may have. Particularly on developing exchanges like the JSE where volume and liquidity are lower, there is a limit to the size of the position that can be traded without obtaining a worse price. This would occur when an order is sent to the market that exceeds the current cumulative quantity of shares listed at a given bid/ask price, and the remaining quantity of shares gets bought/sold at the higher/lower bid/ask price. In addition to this liquidity constraint, slippage can also occur when executing an order. Slippage refers to a positive or negative impact on the price obtained in a trade resulting from the time delay of sending an order from a trader, through the broker and to the market. At market open, both liquidity limitations as well as slippage are typically higher than average levels, and thus certainly need to be taken into account as part of transaction costs. There is no prior literature in this context that sets a precedent or guideline for how such a factor may be determined. As such, based on the authors prior experience, a discretionary approach was taken in calculating this factor. For the purposes of this research, the illiquidity premium and slippage factor combined will be estimated at 5% of the average daily range of the stocks under analysis. This corresponds to a levy of 0.2% per trade.

The summary of transaction costs is provided in Table 6.2 below.

Table 6.2: Total Estimated Transaction Costs Associated with Trading Equities on the JSE

TRANSACTION COST	FEE	TRADE SIDE (BUY/SELL)	EFFECTIVE COST (ROUND TRIP)
BROKER FEE	0.275%	Both	0.55%
SHARE TRANSFER TAX	0.25%	Buy Only	0.25%
INVESTOR PROTECTION LEVY	0.0002%	Both	0.0004%
STRATE LEVY	0.0058%	Both	0.0116%
VAT (15%)	0.0421%	Both	0.0843%
LIQUIDITY LEVY	0.20%	Both	0.40%
TOTAL	0.77%		1.30%

6.4. CFD Trading: Transaction Costs

In contrast to equity trading, the transaction costs for CFD trading are both simpler and lower overall. Given that CFD's are derivatives, most of the equity costs do not apply, including the STT, Investor Protection Levy and the Strate Levy. Thus, only the broker fee for CFD transactions and associated VAT cost along with the liquidity levy will apply to CFD transactions. Table 6.3 on page 78 provides the list of brokers considered earlier for equities trading that also offer CFD trading, along with the fees charged for CFD trades.

Table 6.3: Comparison of South African Broker Fees Associated with CFD Trading

BROKER	BROKER FEE
SHARENET SECURITIES	0.20%
STANDARD BANK SECURITIES	0.20%
SANLAM SECURITIES	0.35%
FNB SECURITIES	0.35%
NEDGROUP PRIVATE WEALTH	0.4%
PSG SECURITIES	0.4%
AFRIFOCUS	0.50%
MOMENTUM SECURITIES	0.50%

(Sharenet Securities, n.d., Standard Bank Securities, n.d., Sanlam Securities, n.d., FNB Securities, n.d., Nedgroup Private Wealth (2022), PSG Securities, n.d., Afrifocus n.d., Momentum Securities, n.d.)

As with the equities broker fee, the average of the 2 lowest broker fees will be used as an indication of the typical CFD broker commission, which corresponds to a fee of 0.2%. Table 6.4 below summarises the total transaction costs associated with CFD trading.

Table 6.4: Table 16: Total Estimated Transaction Costs Associated with Trading CFD's on the JSE

TRANSACTION COST	FEE	TRADE SIDE (BUY/SELL)	EFFECTIVE COST (ROUND TRIP)
BROKER FEE	0.20%	Both	0.40%
VAT (15%)	0.03%	Both	0.06%
LIQUIDITY LEVY	0.20%	Both	0.40%
TOTAL	0.43%		0.86%

6.5. Gap Reversion Strategy

This hypothetical trading strategy is comparably simple and aims to isolate the effects of extreme overnight gaps and the real-world trading performance that can be achieved by trading with these gaps. Based on the strong inverse relationship between overnight gaps and intraday returns, this strategy involves taking a position in the opposite direction to the overnight gap at market open and closing this position at market close. Given that multiple gaps may occur on a daily basis, this strategy will equally weight the fund allocation to each trade based on the number of gaps that occur. For example, if there are 5 gaps on a given day, each trade that day will have a 20% weighting. By doing this, it is possible to compare the results of this trading strategy with various other commonly used benchmarks, such as the AllShare Index buy-and-hold return over the same period.

The gross trade return can be calculated as follows:

If $Gap_{i,t-1} > 0\%$ then

$$\mathrm{R_{i,t}(Short)} = -\frac{\mathrm{(Close\ Price - Open\ Price)}}{\mathrm{Open\ Price}} \times \frac{1}{\mathrm{\#Gaps_{t-1}}} \times 100 \qquad (6.1)$$

(Short position for positive gaps)

If $Gap_{i,t-1} < 0\%$ then

$$\mathrm{R_{i,t}(Long)} = \frac{\mathrm{(Close\ Price - Open\ Price)}}{\mathrm{Open\ Price}} \times \frac{1}{\mathrm{\#Gaps_{t-1}}} \times 100 \qquad (6.2)$$

(Long position for negative gaps)

The total daily gross return can then be calculated as:

$$\sum_{\mathrm{t}} \mathrm{R_{i,t}(Long)} + \mathrm{R_{i,t}(Short)} \qquad (6.3)$$

Having calculated the trade and daily gross returns, these returns can be summed in one of two ways: either cumulatively compounded or cumulatively non-compounded.

Although the closest and most accurate comparison to the buy-and-hold Index return calls for a compounded return calculation, given that this strategy falls firmly in the realm of active management strategies, it is quite likely that, at an institutional level, significant restrictions will be placed on the percentage of funds allocated to such a strategy as well as the individual exposure to any given stock. As such, cumulative non-compounded returns will be used as the primary method of illustrating and comparing returns.

6.6. Trading Strategy Results

Having identified the real-world limitations that will form an integral part of any trading strategy and its results, as well as outline the entry and exit criteria for the proposed strategy, we can now analyse the results of this trading strategy, both in isolation as well as in comparison to the AllShare Index buy-and-hold return, one of the most commonly used South African equity return benchmarks. Figures 6.1 and 6.2 on page 81 illustrate the annual cumulative net returns of this trading strategy, for the transaction costs associated with equities and CFD's, respectively. It is worth emphasising that returns are non-compounded and cumulative, hence the possibility of annual returns below -100%, as was the case for the early years of the trading strategy under the high equity transaction costs.

Figure 6.1: Gap Reversion Strategy Annual Non-Compounded Returns, Equity Transaction Costs

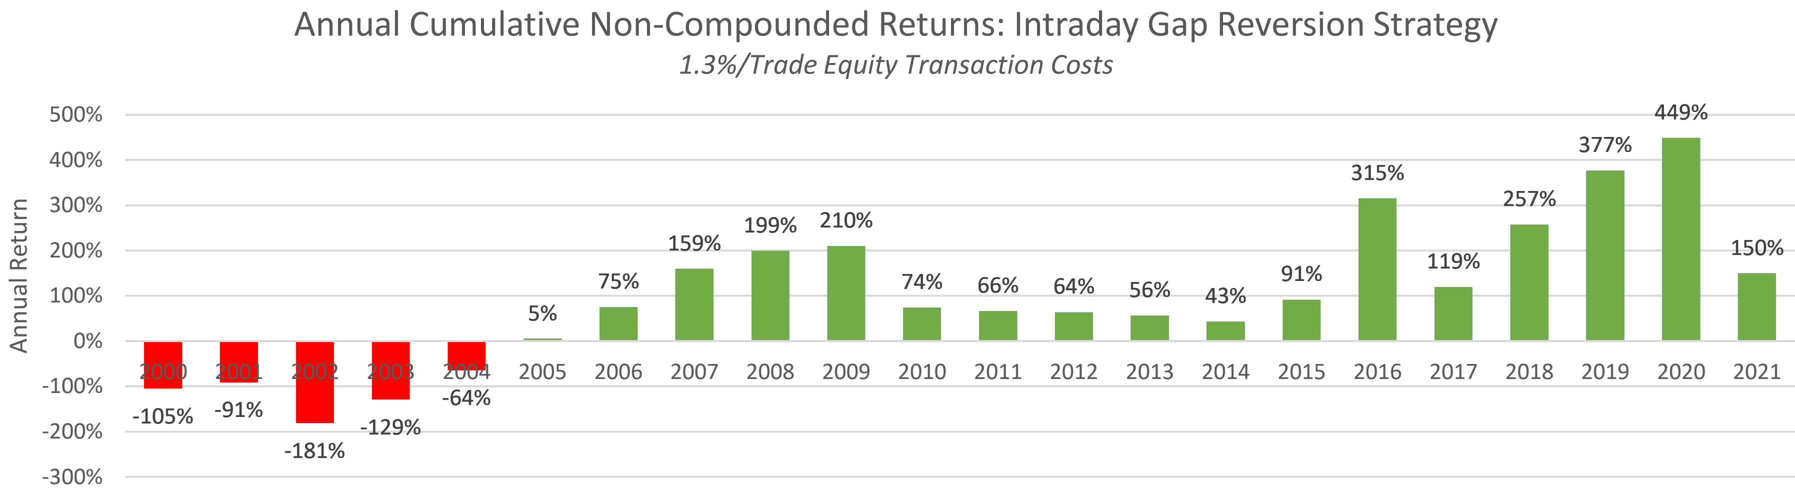

Figure 6.2: Gap Reversion Strategy Annual Non-Compounded Returns, CFD Transaction Costs

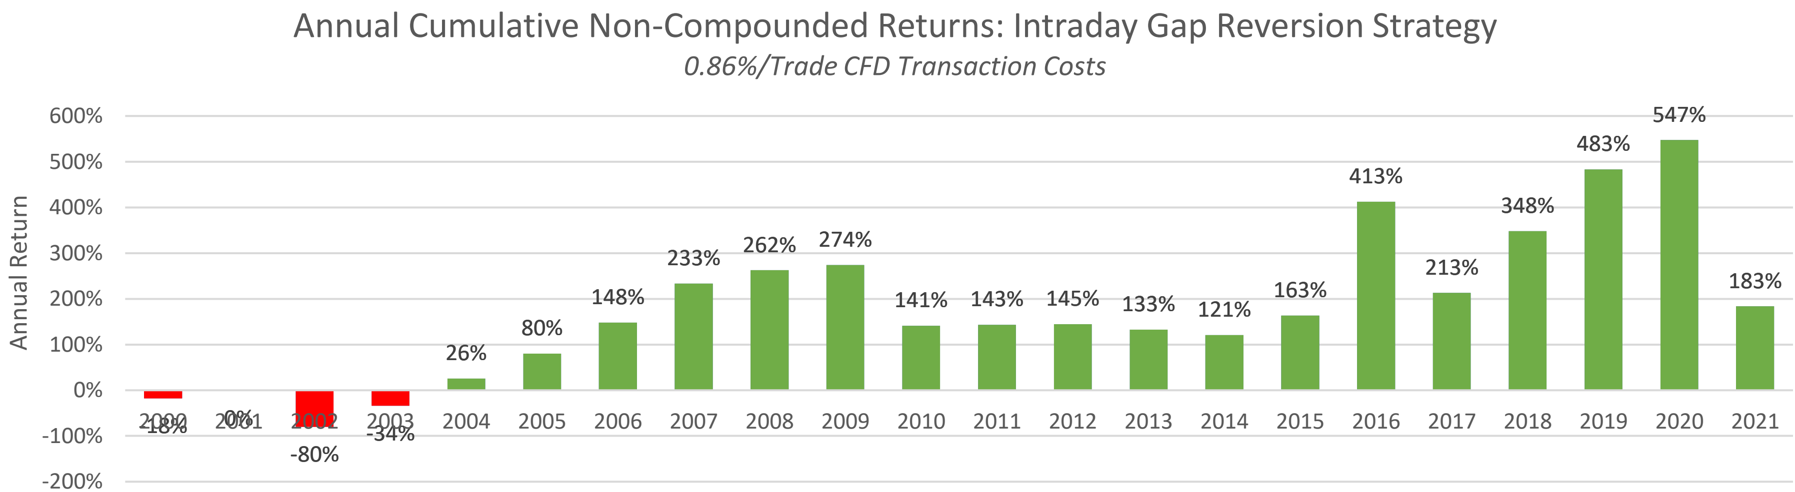

Looking at the overall trends between the two sets of trading returns, the extent to which returns are affected by transaction costs is clearly visible, with returns under CFD costs clearly superior to those under equity costs. There are also four distinct subperiods of returns visible, with the initial 4 years from 2000 to 2003 being characterised by poor returns across the board. From 2004 to 2009 there is also a clear uptrend, with annual returns increasing year on year for 6 consecutive years. This is followed by a period of positive but stagnated returns from 2010 to 2014, and another burst of average year on year growth from 2015 to 2021. Although 2021s returns may seem to revert to lower levels, this is due to the data for this analysis only extending to April of 2021. Extrapolating these results, it is quite likely that returns in 2021 would have matched or exceeded the average returns of the prior few years.

Considering the extreme levels of annualised returns observed, as well as the distinct periods of growth or stagnation, the question arises over what the driving factor behind these return variations may be. Broadly speaking, we can identify that a change in the average rate of returns can be driven by one or both of two factors: the quantity and/or quality of the trades.

Figure 6.3 on page 83 illustrates the total number of extreme, tradeable gaps per year, while Figure 6.4 shows the average profit per trade over the same timeframe. These two graphs in combination allow us to much more easily break down the variations in returns. Evidently, the initial surge in returns in the first half of the dataset is attributable largely to the increase in average profitability per trade. A large drop in profitability is observed in 2010, as evidenced by the stagnation of returns. The last period of rising returns is arguably driven to a greater degree by the increased number of trades, as opposed to the average profitability of each trade.

Figure 6.3: Total Frequency of Tradeable Gaps Per Year

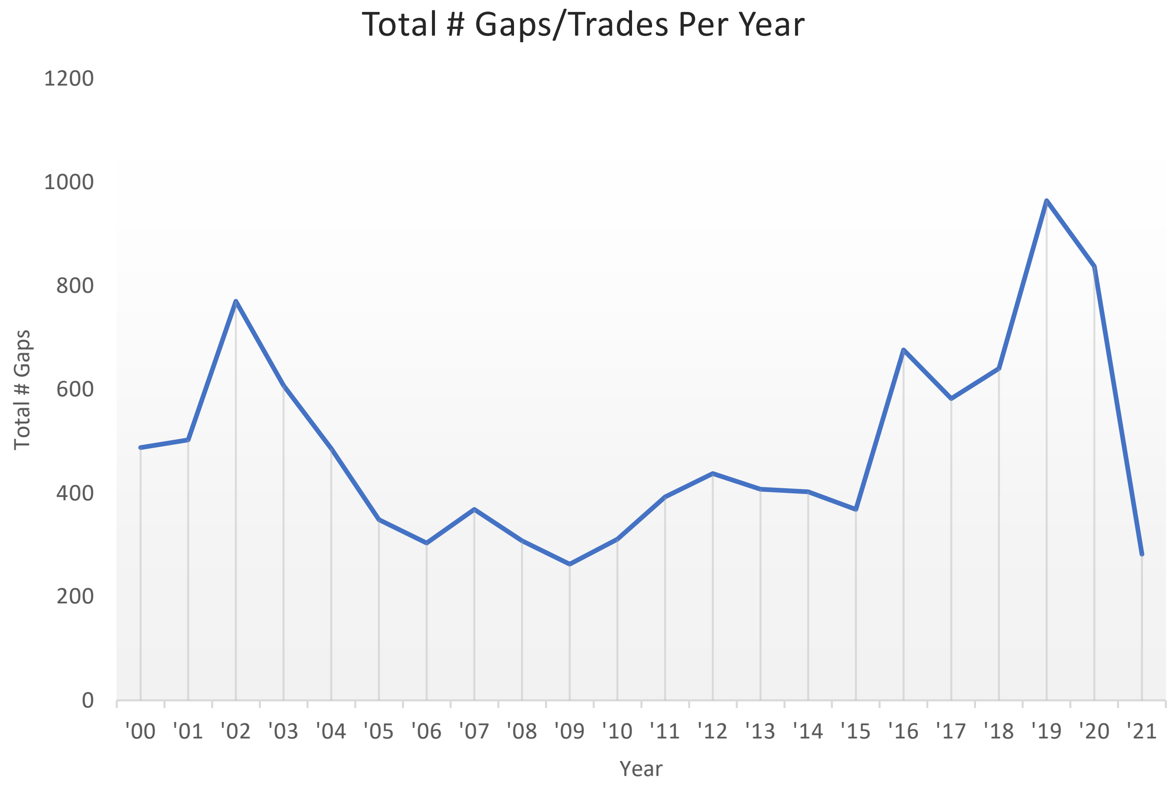

Figure 6.4: Average Profit per Gap Reversion Trade per Year

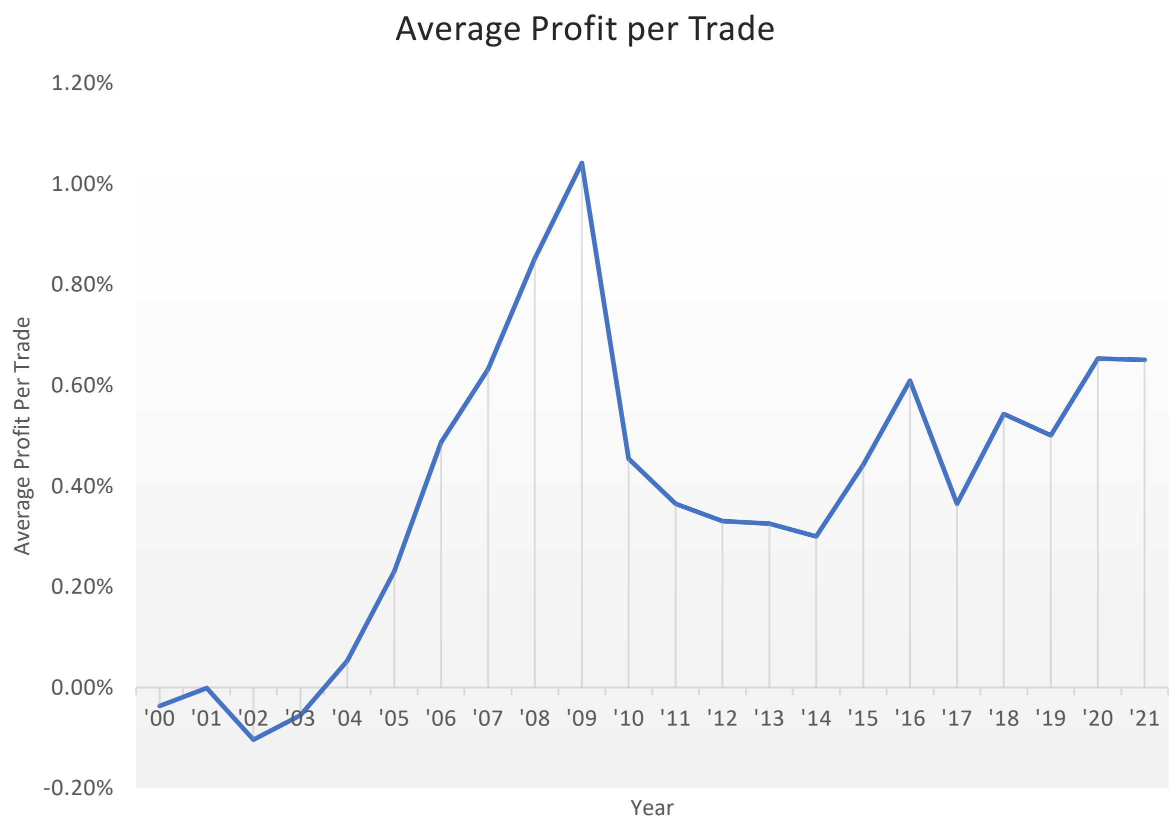

Subsequently, it is also worth looking at the annualised returns for the trading strategy and compare them with South Africa's buy and hold benchmark, the JSE AllShare Index. When considering the AllShare Index, the Total Return Index was used (ALSI TRI) to better illustrate the total performance that would be expected from holding this asset. The Total Return Index is calculated as the net total of capital appreciation as well as dividend reinvestment, and thus captures the total returns better than the AllShare Index price changes alone. Data for the ALSI TRI is only available from the primary data source of this research, Bloomberg, from 2002. As such, returns for the gap reversion trading strategy are only displayed for the years of 2000 and 2001, with subsequent years including both the ALSI TRI return as well as the gap reversion strategy returns. Table 6.5 below provides a comparison between the trading returns and the AllShare Index total return.

Table 6.5: Annual Net Returns Comparison: Benchmark vs. Trading Strategy Returns

		ANNUAL NET RETURNS	
YEAR	ALSI TRI	Trading (1.3% Equity Transaction Costs)	Trading (0.86% CFD Transaction Costs)
2000	-	-105%	-18%
2001	-	-91%	0%
2002	-17%	-181%	-80%
2003	16%	-129%	-34%
2004	24%	-64%	26%
2005	40%	5%	80%
2006	37%	75%	148%
2007	19%	159%	233%
2008	-20%	199%	262%
2009	31%	210%	274%
2010	19%	74%	141%
2011	4%	66%	143%
2012	24%	64%	145%
2013	20%	56%	133%
2014	11%	43%	121%
2015	6%	91%	163%
2016	4%	315%	413%
2017	20%	119%	213%
2018	-7%	257%	348%
2019	12%	377%	483%
2020	11%	449%	547%
2021	14%	150%	183%

**Red: Lowest Return   Orange: Intermediate Return Green: Highest Return*

6.7. Validity & Statistical Analysis

Before considering the results and returns obtained further, it is essential that the accuracy of these results be verified in a quantitative method through a set of statistical tests and analyses. The method of quantifying the relative differences in portfolio performance is a paired Student's t-test. This test will be performed with a single and two-tailed distribution. The paired t-test is used because of the commonality of the timeframe between the two series being compared, the AllShare Index total return and the Gap Reversion strategy return. In this case, a two-tailed paired t-test will determine whether the mean difference between these two sets of returns is equal to 0. If the difference in means is not equal to 0, we can conclude that the two series are not equal. In the case of the single-tailed test, this statistic will indicate whether the returns associated with the Gap Reversion strategy are statistically significantly higher than the buy-and-hold return of the AllShare Index.

As with any statistical analysis, it is imperative that the assumptions of the statistical test be met before any reliable inference can be drawn from the results. In the case of the paired t-test, the assumptions are that:

- The dependent variables are measured on a continuous scale.
- Observations are independent of each other.
- The dependent variable/s are normally distributed.
- The dependent variable/s do not contain outliers.

Table 6.6 below provides some summary statistics of the two data series under analysis, as well as the p-values associated with the Shapiro-Wilks test of normality.

Table 6.6: Descriptive Statistics and Normality Test for Benchmark vs Trading Strategy Returns

Descriptive Statistics

	Sample Size (Days)	Mean	Median	Standard Deviation	Standard Error	Shapiro-Wilks p-value
ALSI TRI Return	4934	5.45e-4	2.98e-4	0.0118	1.68e-4	**< .001**
Trading Return	4934	0.00796	3.32e-4	0.0208	2.97e-4	**< .001**

Although assumptions 1 & 2 of the t-test are met, the assumption of normality and no outliers is not met. As a result, the t-test may no longer be the most powerful test, and a non-parametric test is preferred. In both cases, we can test for differences between the two series under the following hypothesis:

Two-tailed

$H_0: \mu_{Trading} = \mu_{ALSI\ TR}$

$H_1: \mu_{Trading} \neq \mu_{ALSI\ TR}$

Where:

$\mu_{Trading}$ is the return of the Gap Reversion trading strategy and

$\mu_{ALSI\ TR}$ is the total buy-and-hold return of the AllShare Index.

One-tailed

$H_0: \mu_{Trading} \leq \mu_{ALSI\ TR}$

$H_1: \mu_{Trading} > \mu_{ALSI\ TR}$

In the parametric case, the t-statistic can be calculated with the following formula:

$$t = \frac{\sum d}{\sqrt{\frac{n(\sum d^2)-(\sum d)^2}{n-1}}} \tag{6.4}$$

Where:

d is the difference per paired value (between the AllShare Index total return and the Gap Reversion trading return), and

n is the number of samples/trading days.

Table 6.7 on page 87 provides the parametric t-test statistic as well as the non-parametric alternative, the Wilcoxon Signed-Rank Test, along with p-values.

Table 6.7: Two-tailed Parametric & Non-parametric Tests for Significance of Trading Strategy Returns

Two-tailed Tests			
	Test Statistic	**Degrees of Freedom**	**p-value**
Student's t	-21.7	4933	**< .001**
Wilcoxon W	3.56e+6		**< .001**

Despite the t-test no longer being the most powerful, the non-parametric equivalent test and associated p-value also indicate with compelling evidence that the true difference between the AllShare Index buy-and-hold return and the Gap Reversion trading return is statistically significant. Table 6.8 below provides the one-sided version of the same tests.

Table 6.8: One-tailed Parametric & Non-parametric Tests for Significance of Trading Strategy Returns

One-tailed Tests			
	Test Statistic	**Degrees of Freedom**	**p-value**
Student's t	-21.7	4933	**< .001**
Wilcoxon W	3.56e+6		**< .001**

Overall, the two-sided tests indicate that there is clear evidence of a significant difference between the buy-and-hold return of the AllShare Index and the Gap Reversion trading return. The one-sided version of these test allows for a finer analysis, and the highly significant test statistics associated with these tests indicate that the difference in returns can be attributed to the Gap Reversion trading return being higher than that of the AllShare Index's total return. With these results in mind, it is worth analysing the relative performance of each of these returns in more depth. Figures 6.5 and 6.6 on page 88 illustrate the yearly net returns under the different transaction cost levels, while Figure 6.7 on page 89 and Figure 6.8 on page 90 show the total cumulative net returns over time, visually illustrating the structure of returns.

Figure 6.5: Benchmark vs. Trading Strategy Returns Comparison with Equity Transaction Costs

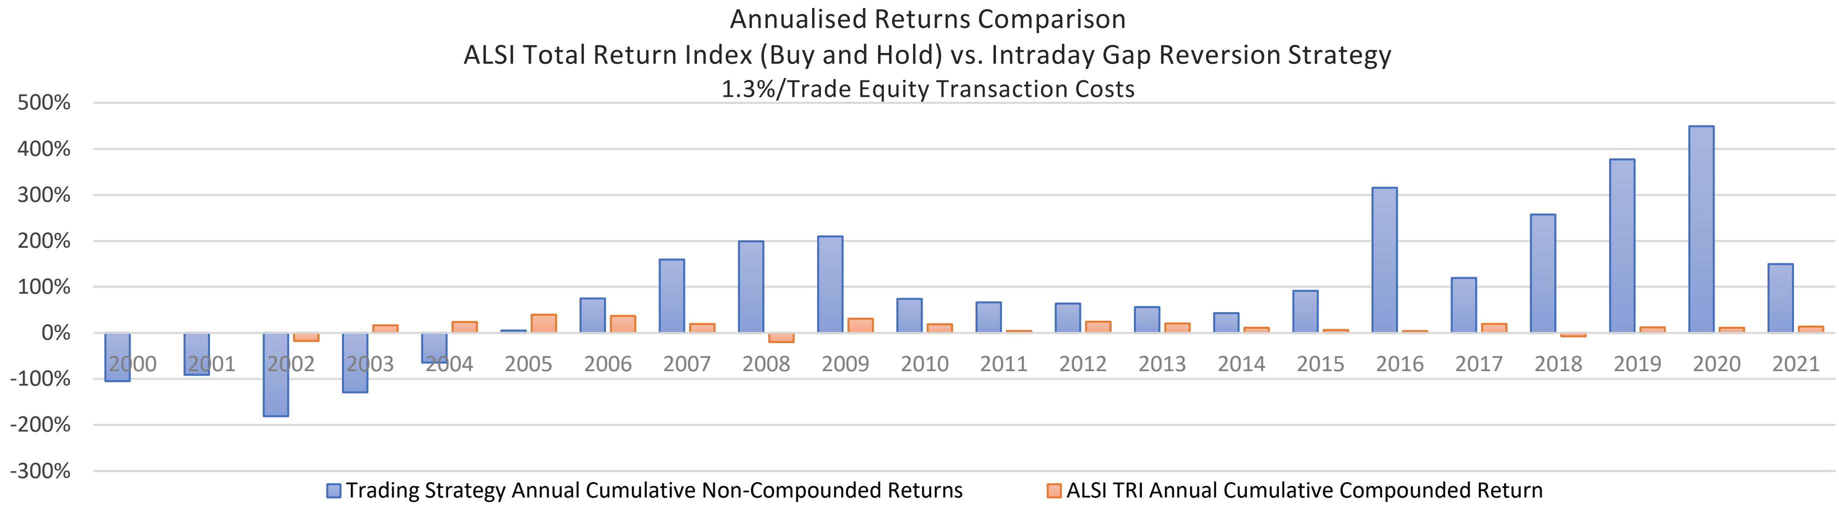

Figure 6.6: Benchmark vs. Trading Strategy Returns Comparison with CFD Transaction Costs

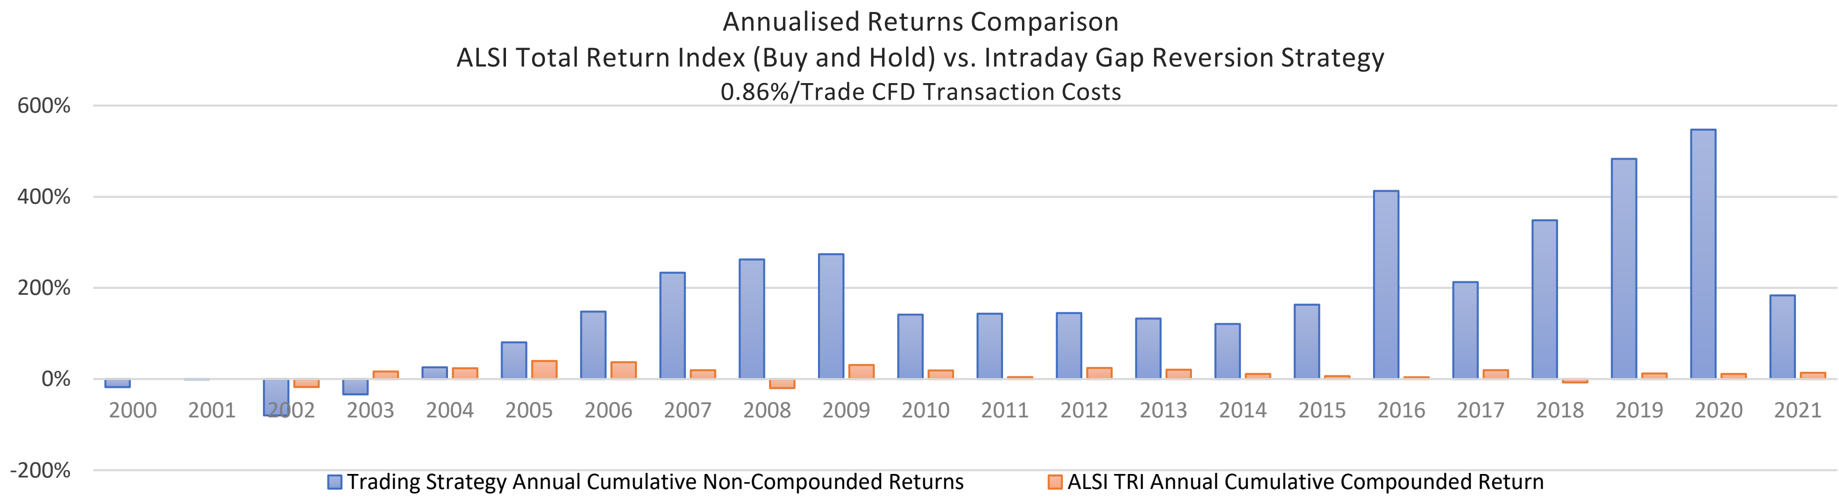

Figure 6.7: Benchmark vs. Trading Strategy Cumulative Returns Comparison with Equity Transaction Costs

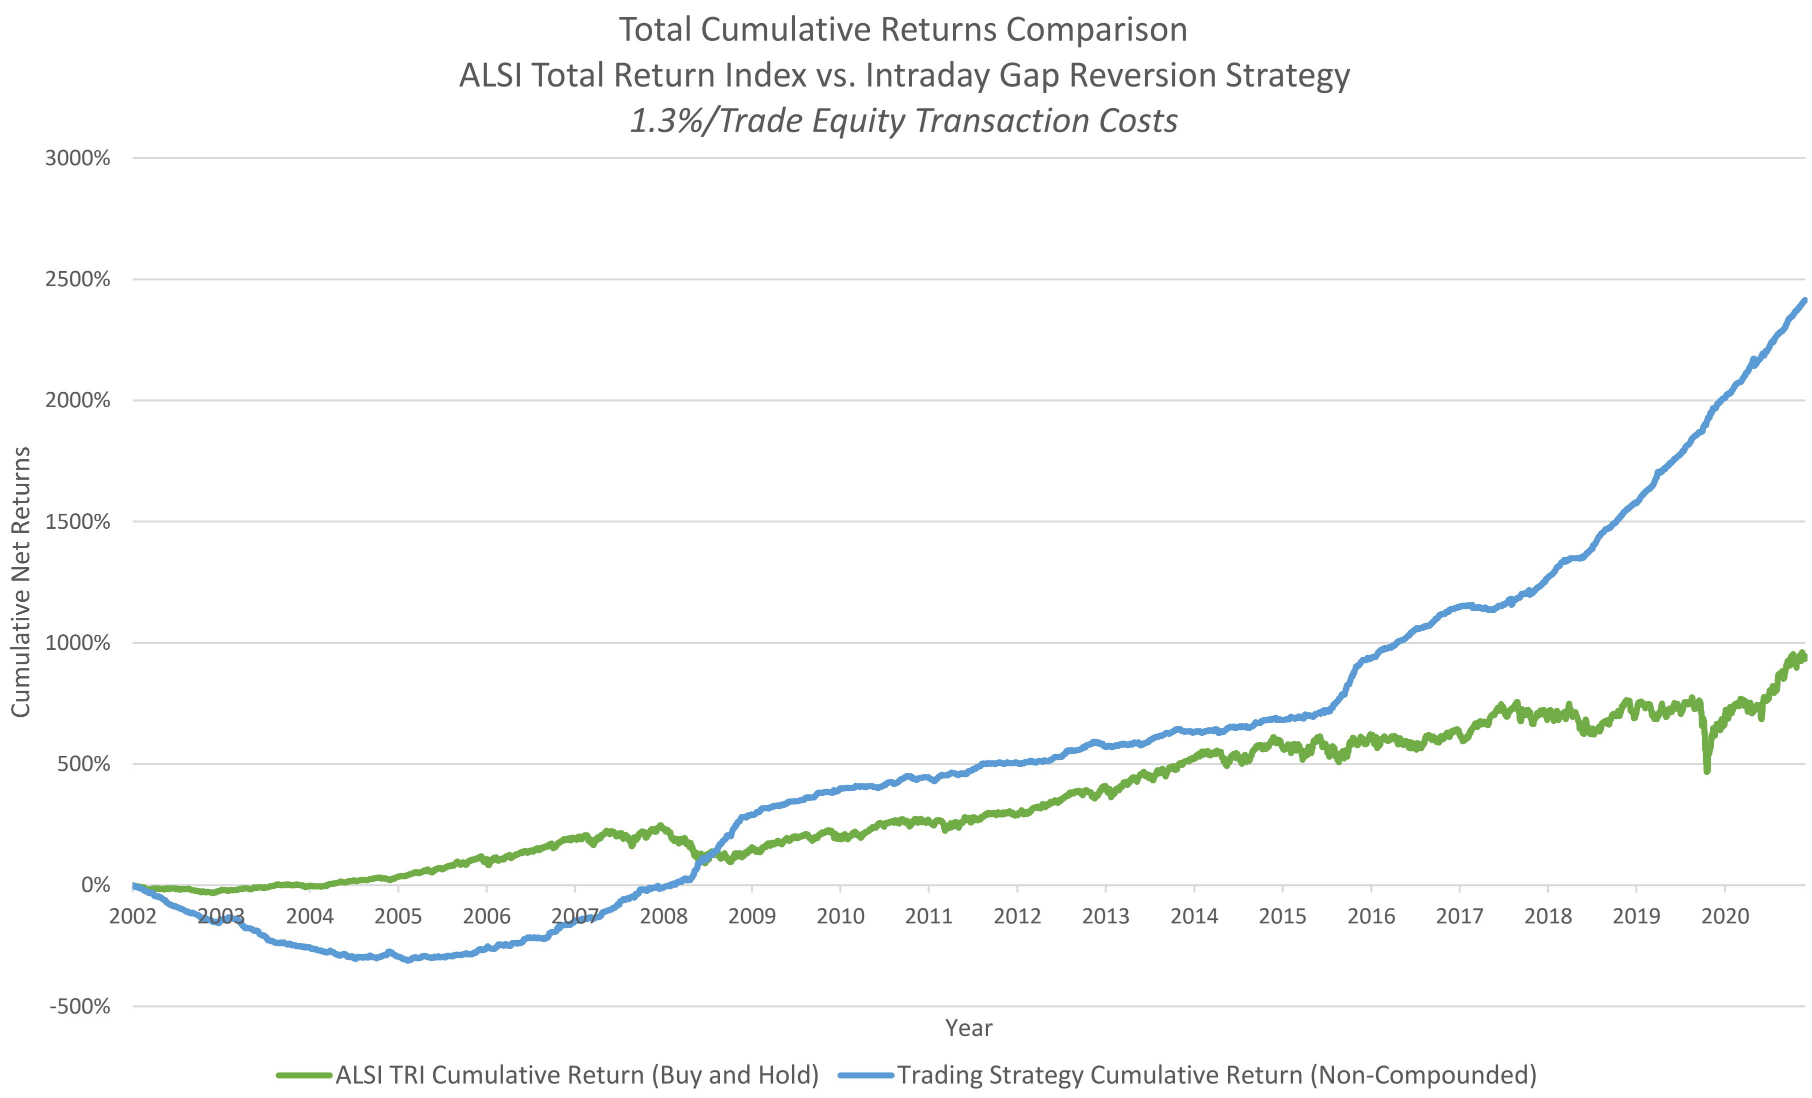

Figure 6.8: Benchmark vs. Trading Strategy Cumulative Returns Comparison with CFD Transaction Costs

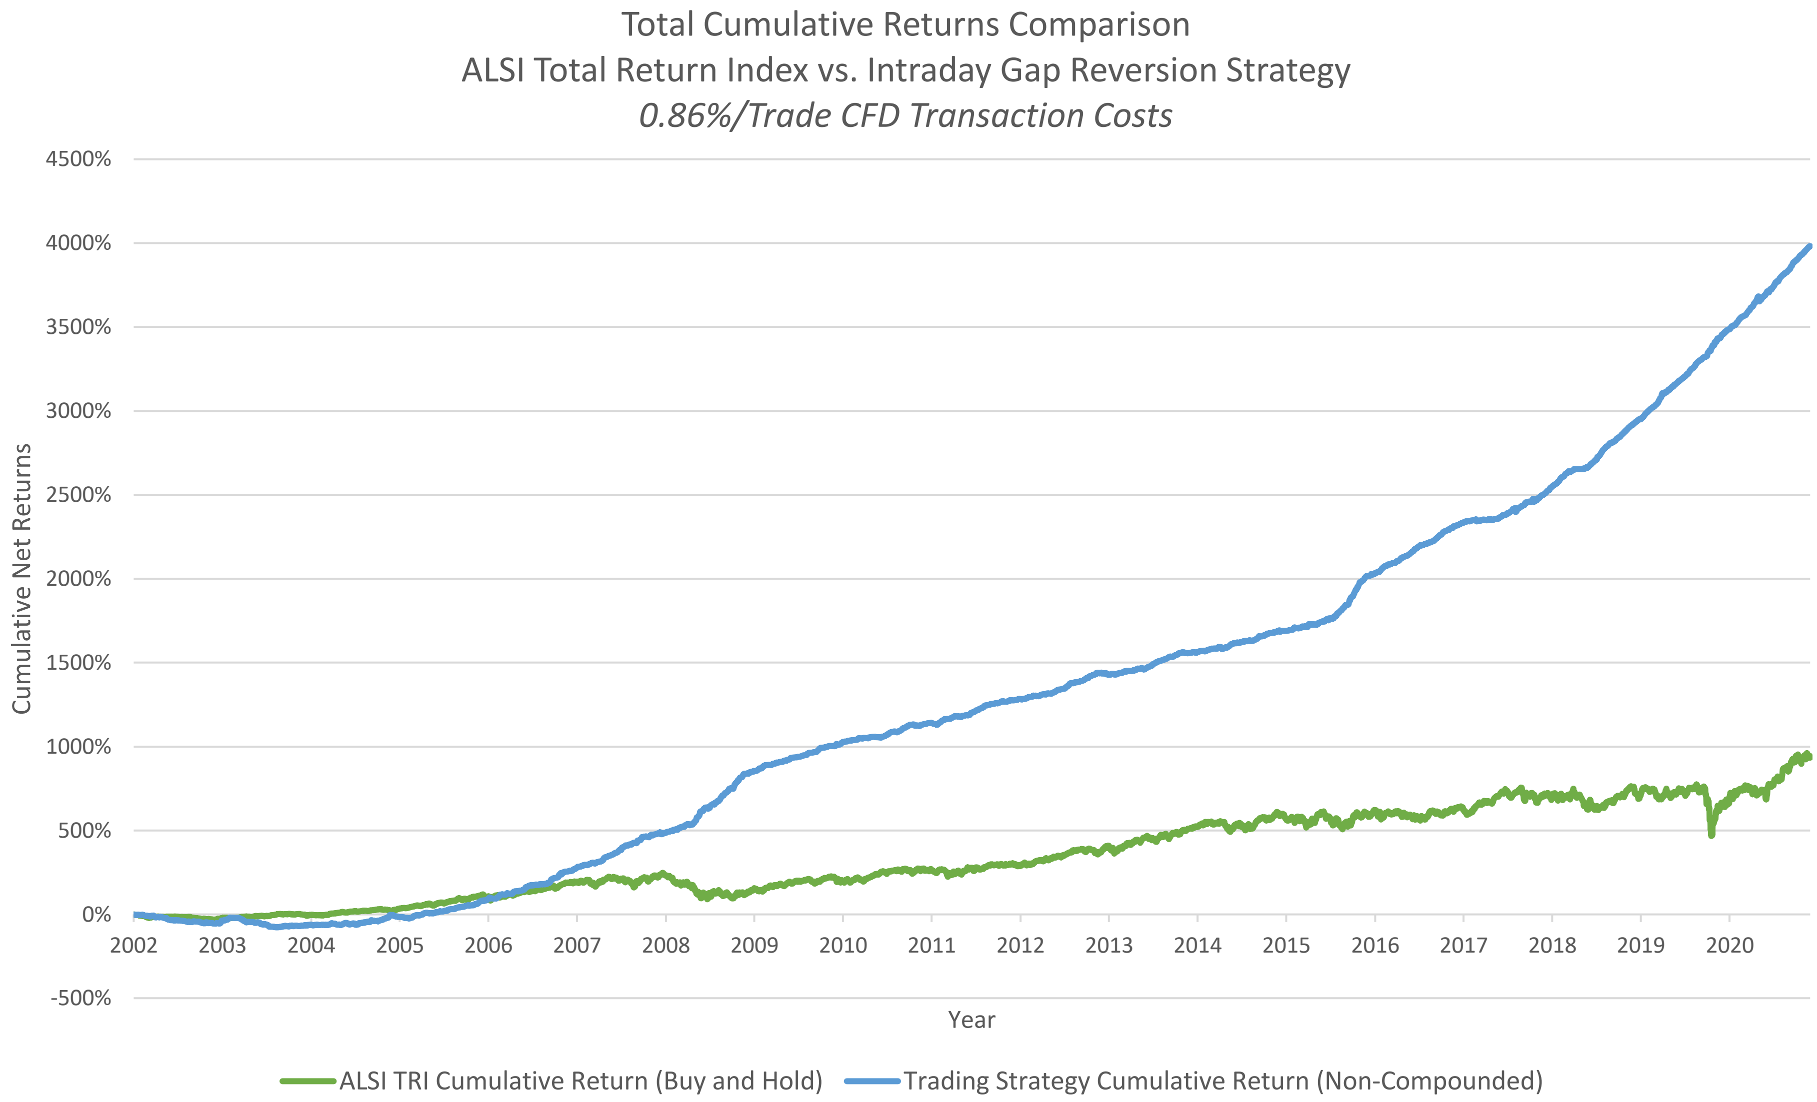

6.8. Performance Analysis

Figures 6.5 and 6.6 on page 88 are useful graphical representations of the annual net performance of the three sets of returns under analysis, however Figures 6.7 and 6.8 are particularly important in illustrating finer details between these returns. These graphs clearly also show the aforementioned four different return periods, particularly in the case of equity transaction costs. An initial period of weak performance, then accelerating returns followed by stagnation, and finally a period of increasing returns again. These graphs also illustrate the relatively consistent and persistent performance of this trading strategy. That is to say, in comparison to the AllShare index, there is a far greater tendency for trading strategy negative returns to be followed by more negative returns, and conversely for positive returns to be followed by more positive returns. Overall, there is also a far greater consistency of positive returns in the trading strategy. Whereas a large amount of near random variation occurs at the Index level on a day to day and even month to month basis, there is less variation present in the trading returns, with an easily observable positive trend shown.

While the graphical representations of returns clearly illustrate the superior risk-adjusted returns of the Gap Trading strategy, particularly in the case of lower CFD transaction costs, this performance can be analysed in more depth by comparing commonly used trading performance metrics across the three sets of returns. For this comparison, the win rate, average winning and losing trade size, risk/reward ratio, expectancy, profit factor and maximal drawdown will be calculated and analysed. These performance metrics are calculated as follows:

Win Rate (%): $\frac{Total\ \#\ Winning\ Trades}{Total\ \#\ Trades} \times 100$ (6.5)

Average Winner: $\frac{\sum Winning\ Trade\ Returns}{Total\ \#\ Winning\ Trades} \times 100$ (6.6)

Average Loser: $\frac{\sum Losing\ Trade\ Returns}{Total\ \#\ Losing\ Trades} \times 100$ (6.7)

Risk/Reward Ratio: $\frac{Average\ Winner}{|Average\ Loser|}$ (6.8)

Expectancy: $[(Win\ Rate \times Average\ Winner) - ((1 - Win\ Rate) \times |Average\ Loser|)] \times 100$ (6.9)

Profit Factor $\frac{Gross\ Profits}{Gross\ Losses}$ (6.10)

Max Drawdown: $\frac{Capital\ Peak\ High - Capital\ Trough\ Low}{Capital\ Peak\ High}$ (6.11)

Table 6.9 below presents the trading metrics described above for further comparisons. It should be noted that these calculations were performed on a daily aggregate level, in order to facilitate an equal comparison. That is to say, where there may be none, one or many trades per day, the returns of these trades are summed and the performance metrics calculated based on these daily totals, in order to match the daily returns associated with the ALSI TR Index for an accurate comparison to be drawn.

Table 6.9: Summary Table of Trading Performance Metrics for Intraday Gap Trading Strategy

METRIC	ALSI TRI	TRADING RETURN (1.3% EQUITY TX COSTS)	TRADING RETURN (0.86% CFD TX COSTS)
Win Rate	51%	43%	51%
Average Win	0.87%	1.97%	2.12%
Average Loss	-0.80%	-0.69%	-0.56%
Risk/Reward Ratio	1.09	2.83	3.82
Expectancy	0.05%	0.48%	0.80%
Profit Factor (Net)	1.14	2.14	3.90
Profit Factor (Gross)	1.14	6.19	6.19
Max Drawdown	-45%	-99%	-63%

The results of both the graphical and metric analysis show the overall superiority of the returns associated with the Gap Trading strategy. Considering metrics such as the win rate in isolation can be misleading, as the conclusion may falsely be drawn that the simulated trading strategy has only performed as well as, or underperformed, the AllShare Index. For this reason, the use of the profit factor and expectancy are preferred. While the profit factor indicates the degree to which profits have exceeded losses historical, the expectancy is forward looking, and uses the historical characteristics of the returns to calculate en expected value associated with the strategy. In the case of the trading strategy, the relatively low win rate is more than offset by the winning trades/days being far larger than the losing trades/days. In this case, an expected value of 0.05% for the AllShare Index indicates that, assuming the historical characteristics continue going forward, your average daily return in the long run is expected to be 0.05% per day. In contrast, the expected daily return for the trading strategy is 0.48% or 0.8%, depending on the level of transaction costs. This corresponds to expected daily returns that are roughly 10 and 16 times greater than those of an AllShare Index, a substantial difference.

One major drawback to these exceptional returns does exist, however. The far superior trading returns are not without risk. One of the most commonly used measures of risk evaluation for historical performance is the maximal equity drawdown associated with that strategy. The drawdown measures the percentage difference between the highest equity peak and the following lowest equity trough, and can be considered to illustrate the largest percentage reduction in capital suffered over the trading history of the strategy. As has been discussed, the returns of this trading strategy have 4 distinct periods. Unsurprisingly, the initial period of underperformance and negative returns between 2000 and 2004/2005 is the sole contributor to the large drawdown values observed for this trading strategy. A drawdown of 99% for equity transaction costs indicates that, at the lowest point, 99% of the initial trading capital would have been lost between 2000 and 2005. The subsequent recovery to still far exceed the return of the AllShare index is remarkable, but not the focus of risk analysis. While the trading drawdown under CFD transaction costs is still large at 63%, this is more reasonable, particularly considering the 45% drawdown associated with the AllShare Index, an investment considered to be much safer than active trading.

Naturally, it is essential that the full timeframe be considered when calculating both risk and returns, however it has been illustrated that returns under the trading strategy are highly persistent. A case can be made that such a strategy only be implemented once returns have turned and stayed positive for a certain amount of time. If this were the case, the story of drawdowns also changes drastically. If only the period of 2006 to 2021 is considered, the maximal equity drawdowns associated with equity and CFD trading are only 30% and 29%, respectively. Not only are these comparatively low drawdown values, but the other metrics focussing on returns only will also be greatly improved under this timeframe of analysis.

6.9. Conclusion

Based on the results of the statistical analysis performed in Chapter 5, a strong and persistent relationship was identified between the overnight gap and following intraday returns. This chapter has explored the possibility of creating a trading strategy based on this relationship, with a simplified strategy created that places a trade in the opposite direction of an extreme overnight gap once one occurs. Considering the focus on the real-world application of such a strategy, extensive consideration was given to the transaction and other costs associated with implementing such a strategy, with 2 levels of transaction costs considered. Overall, the results and performance metrics indicate that this trading strategy has offered superior returns in comparison to the AllShare Index from 2000 to 2021, however a higher level of risk is also attached. Particularly in the first 5 years under analysis, an overall underperformance was observed, leading to extremely large drawdowns, particularly under the highest transaction costs. Following this initial period of underperformance, an extreme outperformance of the AllShare Index was observed instead, with risk metrics such as maximal drawdown now falling below that of the AllShare Index as well, leading to superior risk-adjusted returns. In terms of the objectives of this research, the results of this analysis make it clear that a profitable trading strategy based on gaps can indeed be created, even with excessive transaction costs.

Chapter 7: Limitations & Areas of Future Research

7.1. Limited Analysis of Behavioural Aspects

One of the core objectives of this research is to establish whether or not an exploitable market anomaly exists with regards to the phenomenon of overnight gaps and subsequent intraday returns. While the analysis and research of this thesis has given compelling evidence that this is the case, there is far less evidence and analysis of the underlying reasons behind what drives this market anomaly to exist. While Chapter 1 & 2 briefly outline similar market anomalies and their proposed behavioural origins, such literature is particularly scarce as it relates to gaps. An analysis of behavioural finance on the underlying reason for the creation and exploitable nature of gaps could give crucial insights into the core mechanism that drives market anomalies; however, this is beyond the scope of this research. It should also be noted that an increasing volume of transactions are executed algorithmically, and thus an analysis of behavioural finance also relates to the learnt behaviours of these algorithms (Hendershott, Jones & Menkveld, 2011).

7.2. Non-linear Models

This research has presented a wide range of statistical tests and analyses that have been used to investigate the predictive power of overnight gaps. Specifically, the use of multiple linear regression models is the primary method of analysis. Although this set of tests and analyses is comprehensive, this analysis does not extend to non-linear models. Based on the high degree of significance discovered using linear models, the possibility of non-linear models being able to capture even more of this variance does exist. Similar research conducted in the past, such as that by De Gooijer, Diks & Gatarek (2009) shows that non-linear models can perform better than linear models in certain cases. Although the core aim of this research was to identify the potential relationship between gaps and overnight returns, the use of non-linear models as a potential method of analysis to enhance the understanding of this relationship is identified as an area of potential future research.

7.3. Applicability & Scalability of Models

The results of this research have given compelling proof of the existence of exploitable gap behaviour as well as the profitable trading strategy that could be built around these gaps. Although this is an exciting conclusion to this research, these findings cannot necessarily be translated to markets outside of South Africa, which operates at a much lower level of liquidity and trading activity than many international markets, such as the New York Stock Exchange (NYSE). Several key limitations also exist to the implementation of such a strategy on the JSE as well, most notably the liquidity constraints of the stock during the first minute of the day during which trades must be entered. Thus, while the results are promising, the successful implementation of such results in both the domestic and foreign markets is uncertain and untested. While the real-world trading of such a strategy falls outside the scope of this research, the findings of such an endeavour would add substantial credibility to the findings of this research, particularly with regards to the trading strategy.

7.4. Limited Timeframe of Analysis for Price Impact

This study has only considered the intraday return following the creation of a gap. This essentially means that only the 1-day impact of gaps can be assessed. Prior research as well as a preliminary analysis of the data both present a compelling case for the largest effects to realised immediately following the gap, as opposed to serval days later. An extended analysis of periods longer than one day also adds complexity to the analysis. For example, if two gaps occurred on the same stock during a 5-day period of analysis, the effect of each gap would need to be isolated by estimation. While it may be interesting to analyse the effect of a gap on returns longer than one day, there are valid reasons for this research not to do so within the presented constraints. It is unlikely that more significant results would be observed, however any result obtained via the scientific method is valuable in expanding the body of knowledge on gaps, or more broadly market anomalies, and as such, this extended analysis can certainly be identified as an area of future research.

7.5. Consideration of News & Other Fundamental Data

Throughout the analysis of gaps, one key element of their creation has not been considered, namely a fundamental catalyst. As described briefly in Chapter 1, this could be anything from an earnings report to breaking news, or any catalyst of this nature. Given the extremely large amount of data that has been considered here, it is practically impossible to obtain all relevant information catalysts around the date of the gap, as well as filter out irrelevant information to only obtain the causal catalyst. In addition, such data is largely unavailable, as headlines and company press releases on the JSE only go back a few years. As such, it was not possible to incorporate news and other fundamental catalysts into the analysis of gaps. It can be argued that if this information were available, the nature of the study would shift such that it becomes predominantly an analysis of new information released to the market, with a secondary focus on gaps. Substantial research has been conducted on the topic of what impact new information has on subsequent market prices, and thus such a study would add comparatively less to the body of research. In addition, when we look at gaps in isolation, we are able to use gaps as proxies for “all relevant information” that has been released overnight, including news releases, pricing information from the prior day and similar relevant factors. By looking at gaps as a whole, without attempting to isolate the one or many influences that lead to the creation of the gap, the focus shifts to what the impact of the gap itself is, rather than what the impact of the catalyst is. Nonetheless, a future study on a narrower and smaller dataset with the inclusion of possibly relevant catalysts could meaningfully enrich our understanding of stock market gaps by focussing on their creation rather than their consequences.

Chapter 8: Summary & Conclusion

8.1. Introduction

Research on the predictive power of overnight gaps on subsequent returns is underdeveloped internationally, and nearly non-existent in the South African context. Findings from research on international exchanges has shown promising results in this field, and it is based on these findings and the identified gap in the body of knowledge that the primary objective of this research was formulated. Specifically, the purpose of this research is to conclusively answer the question of whether or not overnight gaps exhibit any influence or predictive power on subsequent returns on the JSE. Further research objectives include looking at whether additional technical and/or fundamental variables can influence subsequent intraday returns, as well as identifying if gap behaviour and predictive power changes during periods of market instability, such as a financial recession and lastly what the performance of a hypothetical trading strategy centred around gaps would look like. Results from a wide range of statistical tests and analyses reveals a strong and persistent inverse relationship between large gaps and subsequent intraday returns. Intraday returns are predicted to decrease by up to 0.5% for a 1% increase in the size of the preceding gap, a very strong inverse relationship. The empirical results of this research will now be summarised in more depth.

8.2. Summary of Results

The data used in this research has been collected from Bloomberg, with various pricing, technical and fundamental data included in the dataset. The period of analysis is January 2000 to April 2021, and in total more than 720,000 unique observations were analysed. A combination of visual analysis and the Tukey 2.2. rule were used to identify and exclude extreme outlier values and corrupted data, while missing data fields were also removed from the dataset. Additional variables created for each observation include the overnight gap, intraday returns at different lagged periods, cumulative preceding returns, the total index return of the AllShare Index, the values of the 9 & 20 simple moving averages, the total daily value traded and the 1 day and

1 week average intraday ranges. Summary statistics revealed a large number small gaps in the dataset that are not meaningful in the context of this research. Based on this, an additional filter was applied in order to only include sufficiently large gaps in the analysis dataset. The criterion for a sufficiently large gap is based on the 5-day average intraday range of the stock in question. If a gap is larger than this 5-day average range, it constitutes a meaningful and extreme gap, and is thus included in the subsequent analyses. Lastly, a liquidity filter was applied in order to remove illiquid shares with irregular trading patterns. This was done by calculating the average daily value traded for each stock for the first 2 weeks of the year. These stocks were then ranked from highest to lowest daily value traded, with the bottom 50% excluded from the portfolio of applicable stocks. This process was repeated yearly to maintain a dynamically adjusting portfolio of only the most liquid shares, while avoiding survivorship bias.

The primary method of statistical analysis in this research is a series of multivariate regression models. Specifically, a combination of multiple linear and multiple logistic models were created and analysed, along with a pooled regression. Variables of interest identified from the preliminary analysis of Chapter 3 & 4 include the overnight gap the prior day as well as 2 days ago, the prior intraday return, categorical trend variable and 5-day cumulative prior return and prior intraday range. These variables have been included in all subsequent regression models.

In order to remove look-ahead bias, regression models are initially split into 2 periods. The first 60% of the dataset is used as training data for Model 1 to build and evaluate the initial performance of the regression models, while the remaining 40% of the data is then used for Model 2 in order to forward test these models and observe the performance on fresh, unseen data. In addition, both Models 1 and 2 also included a secondary multiple logistic regression. The primary purpose of this logistic regression is to serve as a further robustness check against the results of the multiple linear models. Model 3 was then run on the full 20+ year dataset, and lastly Models 4 and 5 were run on market crisis subperiods, namely the 2008 global financial recession and 2020 Covid-19 market crash, respectively. In order to accurately evaluate differences between the various subperiods, the use of pooled regressions was also employed. The pooled regressions add an interaction term of each independent variable multiplied by a dummy variable that indicates what period an observation is in (0 for training and 1 for unseen). This pooled regression allows for the additive influence of each variable to be analysed in the unseen vs. seen data, and also includes an associated test statistic and p-value that can be used to quantifiable determine whether these differences are statistically significant.

Before the results of any regression model were analysed, the assumptions of an OLS regression were identified and explored in depth. These assumptions are: the linear relationship and the additive property of covariates, the assumption of no multicollinearity, homoscedasticity of the error term, normally distributed residuals and no serial autocorrelation. These assumptions were tested by means of a combination of statistical tests and visual analysis, and the conclusions for all models were that these assumptions have not been violated.

Chapter 5 presents the output and discussion of the specified models. The identification of the overnight gap as a highly significant predictor of intraday returns persisted across all models and timeframes, and interestingly, the strength of this relationship was found to increase substantially in the more recent, unseen data, as opposed to the first 12 years of the dataset, with coefficient estimates of -0.4 and -0.5 observed, respectively. Conclusions from the effectiveness of the additional variables are mixed, with only the intraday return and intraday range of the previous trading day found to be effective predictors over time. It is worth noting that both of these variables have an estimated positive effect on intraday returns, in contrast to that of the gap, where the effect is negative. The total intraday return following a large overnight gap is thus the product of two conflicting factors: continuation from the prior days intraday return and a reversal of the overnight return. The analysis of variations in predictive power over different sub-periods did not yield compelling evidence for a significantly different relationship being present during periods of market instability. As such, in conjunction with the earlier analysis of in and out-of-sample periods, the conclusion with regards to the variables analysed here is that changes in market characteristics that may affect the predictive power of these variables appears to occur slowly, over a period of years or decades.

Lastly, Chapter 6 deals exclusively with the additional research objective of identifying the performance of a trading strategy built around gaps. A simple strategy is proposed, whereby a position is taken in the opposite direction of a gap at market open, and this position is closed at market close. Where multiple gaps occur on a given day, each gap is equally weighted and the respective returns are summed to obtain the daily total.

The aim is to highlight the real-world performance of such a strategy, and as such extensive consideration is given to limitations that may affect the performance of such a strategy, such as transaction costs and slippage. Two levels of transaction costs are identified, namely for direct stock trading and for CFD trading. The estimates of the total costs associated with these methods of trading are 1.3% and 0.86% per round-trip trade, respectively. Despite the high levels of transaction costs, the results of the trading strategy are very positive. A substantial average annual outperformance of the buy-and-hold return of the AllShare Index is observed with excess returns of 85% and 166% recorded, depending on the transaction costs. The major caveat to these returns is the drawdown observed from these strategies, however. The first 5 years, from 2000 to 2004 yielded consistently poor returns and year-on-year losses, leading to extreme drawdown values of 99% and 63%, respectively, in comparison to the AllShare Index's 45%. Nonetheless, despite this initial drawdown period, the trading strategy outperformed every year thereafter, as well as over the full 20 year period of analysis. The conclusion here is that not only can a trading strategy based on gaps be built, but such a strategy is also very lucrative in terms of outperformance and profit potential.

8.3. Conclusion

The research of this thesis has presented compelling evidence of the non-random behaviour of gaps for JSE-listed companies. Building on the body of knowledge of market anomalies, the overnight gap is identified as an additional market anomaly that poses a serious challenge to the efficient market hypothesis. In addition to the theoretical findings, the practical applications have also been considered and shown to be very lucrative, illustrating the potential real-world performance advantage that can be gained by correctly exploiting the predictive power of overnight gaps.

Reference list

Absa, n.d. *Trade shares and ETFs online with a Smart Trading Account.* Available: https://www.absa.co.za/personal/save-invest/stockbroker-and-portfolio-management/smart-account/. [January 23, 2023]

Aboody, D., Even-Tov, O., Lehavy, R. & Trueman, B. 2018. Overnight Returns and Firm-Specific Investor Sentiment. *Journal of Financial and Quantitative Analysis; J.Financ.Quant.Anal.* 53(2):485-505. DOI: 10.1017/S0022109017000989.

Afrifocus Securities, n.d. *Trading Account (ZAR).* Available: https://afrifocus.tradedesk.co/markets/TradingAccountZAR/#SouthAfricanMarkets. [January 23, 2023]

Alexandroi, S. 2019. *Cash Flow as a Predictor of Share Returns: Evidence from the Johannesburg Stock Exchange.* M.Fin. dissertation. University of Cape Town. Available: http://hdl.handle.net/11427/31430.

Alur, R. 2016. *Short-term share price overreaction: evidence from the Johannesburg Stock Exchange*. M.Fin. dissertation. University of Cape Town. Available: http://hdl.handle.net/11427/21748.

Atsin, J. & Ocran, M. 2015. Calendar effects and market anomalies on the Johannesburg Stock Exchange. *IDEAS Working Paper Series from RePEc.* Available: https://mpra.ub.uni-muenchen.de/87448/.

Baker, M. & Wurgler, J. 2006. Investor Sentiment and the Cross-Section of Stock Returns. *The Journal of Finance (New York).* 61(4):1645-1680. DOI: 10.1111/j.1540-6261.2006.00885.x.

Bala, D.A. & Takimoto, T. 2017. Stock markets volatility spillovers during financial crises: A DCC-MGARCH with skewed-t density approach. *Borsa Istanbul Review.* 17(1):25-48. DOI:10.1016/j.bir.2017.02.002.

Barber, B.M. & Odean, T. 2008. All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *The Review of Financial Studies.* 21(2):785-818. DOI: 10.1093/rfs/hhm079.

Barclay, M.J. & Hendershott, T. 2003. Price Discovery and Trading after Hours. *The Review of Financial Studies.* 16(4):1041-1073. DOI: 10.1093/rfs/hhg030.

Berkman, H., Dimitrov, V., Jain, P.C., Koch, P.D. & Tice, S. 2009. Sell on the news: Differences of opinion, short-sales constraints, and returns around earnings announcements. *Journal of Financial Economics.* 92(3):376-399. DOI: 10.1016/j.jfineco.2008.04.009.

Berkman, H., Koch, P.D., Tuttle, L. & Zhang, Y.J. 2012. Paying Attention: Overnight Returns and the Hidden Cost of Buying at the Open. *Journal of Financial and Quantitative Analysis; J.Financ.Quant.Anal.* 47(4):715-741. DOI: 10.1017/S0022109012000270.

BMJ Publishing Group. 2021. *The t tests.* Available: https://www.bmj.com/about-bmj/resources-readers/publications/statistics-square-one/7-t-tests. [January 23, 2023]

Bogousslavsky, V. 2021. The cross-section of intraday and overnight returns. *Journal of Financial Economics.* 141(1):172-194. DOI: 10.1016/j.jfineco.2020.07.020.

BP Bernstein, n.d. *Open an Account – Fee Schedule*. Available: https://www.bpbernstein.co.za/application/. [January 23, 2023]

Charles Schwab. 2021. *Extended Hours Overview*. Available: https://www.schwab.com/public/schwab/nn/qq/about_extended_hours_trading.html. [January 23, 2023]

Choi, S. 2021. Analysis of stock market efficiency during crisis periods in the US stock market: Differences between the global financial crisis and COVID-19 pandemic. *Physica A: Statistical Mechanics and its Applications.* 574:125988. DOI: https://doi.org/10.1016/j.physa.2021.125988.

Christian, L.D., Laws, J. & Rudy, J. 2011. Profitable mean reversion after large price drops: A story of day and night in the S&P 500, 400 MidCap and 600 SmallCap Indices. *Journal of Asset Management.* 12(3):185. DOI: 10.1057/jam.2011.15.

Cooper, M.J., Cliff, M.T. & Gulen, H. 2008. *Return Differences between Trading and Non-Trading Hours: Like Night and Day*. DOI: 10.2139/ssrn.1004081.

De Gooijer, J.,G., Diks, C.G.H. & Gatarek, L.T. *Information Flows around the Globe: Predicting Opening Gaps from Overnight Foreign Stock Price Patterns.* DOI: 10.2139/ssrn.1510069.

Delong, J.B., Shleifer, A., Summers, L.H. & Waldmann, R.J. 1990. Noise Trader Risk in Financial Markets. *The Journal of Political Economy.* 98(4):703-738. DOI:10.1086/261703.

Desjardins, J. 2016. *All of the World's Stock Exchanges by Size*. Available: http://money.visualcapitalist.com/all-of-the-worlds-stock-exchanges-by-size/. [January 23, 2023]

Dicle, M.F. & Levendis, J. 2013. The impact of technological improvements on developing financial markets: The case of the Johannesburg Stock Exchange. *Review of Development Finance.* 3(4):204. DOI:10.1016/j.rdf.2013.09.001.

Dorador, A.C. 2017. *Do Stop-Loss rules stop losses? An analytical framework based on modeling overnight gaps and the Stationary Bootstrap.* M.CompSc dissertation. Universitat Politècnica de Catalunya. Available: https://upcommons.upc.edu/bitstream/handle/2117/106091/memoria.pdf.

Dunis, C., Laws, J. & Rudy, J. 2011. Profitable mean reversion after large price drops: A story of day and night in the S&P 500, 400 MidCap and 600 SmallCap Indices. *Journal of Asset Management.* 12(3):185-202. DOI:10.1057/jam.2011.15.

EasyEquities, n.d. *Pricing Breakdown*. Available: https://www.easyequities.co.za/pricing. [January 23, 2023]

Ellington, M. 2018. Financial market illiquidity shocks and macroeconomic dynamics: Evidence from the UK. *Journal of Banking and Finance.* 89:225-236. DOI:10.1016/j.jbankfin.2018.02.013.

Eloff, F.N. 2014. *Momentum trading strategy on the Johannesburg Stock Exchange.* M.Fin dissertation. University of Cape Town. Available: http://hdl.handle.net/11427/8557.

Enowbi, M.B., Guidi, F. & Mlambo, K. 2010. Testing the weak-form market efficiency and the day of the week effects of some African countries. *African Finance Journal.* 2010(si-1):1-26. DOI: 10.10520/EJC33729.

Fama, E.F. 1998. Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics.* 49(3):283-306. DOI: 10.1016/S0304-405X(98)00026-9.

Fama, E.F. 1970. EFFICIENT CAPITAL MARKETS: A REVIEW OF THEORY AND EMPIRICAL WORK. *The Journal of Finance (New York).* 25(2):383-417. DOI:10.1111/j.1540-6261.1970.tb00518.x.

FNB Securities, n.d. *Share Investor.* Available: https://www.fnb.co.za/share-investing/postLogin/shareInvestor.html. [January 23, 2023]

Frisch, T., Kolaric, S. & Schiereck, D. 2014. Returns on large stock price declines and increases in the South African stock market: A note on market efficiency. *International Business & Economics Research Journal (IBER).* 13(3):581-592. DOI: 10.19030/iber.v13i3.8595.

González-Hermosillo, B. & Hesse, H. 2009. Global market conditions and systemic risk. *IMF Working Papers*. 2009(230):1-22. DOI: 10.5089/9781451873771.001.

Guidolin, M., Hyde, S., Mcmillan, D. & Ono, S. 2009. Non-linear predictability in stock and bond returns: When and where is it exploitable? *International Journal of Forecasting.* 25(2):373-399. DOI:10.1016/j.ijforecast.2009.01.002.

Harrison, M.J & McCabe, B.P.M. 1979. A Test for Heteroscedasticity Based on Ordinary Least Squares Residuals. *Journal of the American Statistical Association.* 74(366):494-499. DOI: 10.2307/2286361.

Hendershott, T., Jones, C.M. & Menkveld, A.J. 2011. Does Algorithmic Trading Improve Liquidity? *The Journal of Finance.* 66(1):1-33. DOI: 10.1111/j.1540-6261.2010.01624.x

Heymans, A. & Santana, L. 2018. How efficient is the Johannesburg Stock Exchange really? *South African Journal of Economic and Management Sciences.* 21(1) DOI:10.4102/sajems.v21i1.1968.

Hribar, P. & McInnis, J. 2012. Investor Sentiment and Analysts' Earnings Forecast Errors. *Management Science.* 58(2):293-307. DOI:10.1287/mnsc.1110.1356.

Hoaglin, D., Iflewicz, B. 1987. Fine-Tuning Some Resistant Rules for Outlier Labeling. *Journal of the American Statistical Association.* 82(1):1147-1149. DOI: doi.org/10.2307/2289392.

Huang, A.Y., Hu, M. & Truong, Q.T. 2021. Asymmetrical impacts from overnight returns on stock returns. *Review of Quantitative Finance and Accounting.* 56(3):849-889. DOI:10.1007/s11156-020-00911-y.

IG. n.d. What is CFD trading? Available: https://www.ig.com/za/cfd-trading/what-is-cfd-trading-and-how-does-it-work. [January 23, 2023]

Kelly, M.A. & Clark, S.P. 2011. Returns in trading versus non-trading hours: The difference is day and night. *Journal of Asset Management.* 12(2):132-145. DOI:10.1057/jam.2011.2.

Laerd Statistics. n.d. *One-way ANOVA.* Available: https://statistics.laerd.com/statistical-guides/one-way-anova-statistical-guide-2.php. [January 23, 2023]

Laerd Statistics. n.d. *Mann-Whitney U Test using SPSS Statistics.* Available: https://statistics.laerd.com/spss-tutorials/mann-whitney-u-test-using-spss-statistics.php. [January 23, 2023]

Laerd Statistics. n.d. *Independent t-test for two samples.* Available: https://statistics.laerd.com/statistical-guides/independent-t-test-statistical-guide.php. [January 23, 2023]

Laerd Statistics. n.d. *Multiple Regression Analysis using SPSS Statistics.* Available: https://statistics.laerd.com/spss-tutorials/multiple-regression-using-spss-statistics.php. [January 23, 2023]

Leung, W.K. 2009. Price Reversal and Firm Size in the US Stock Markets, New Evidence. *Proceedings of the World Congress on Engineering.* Volume II. London, U.K. Available: http://www.iaeng.org/publication/WCE2009/WCE2009_pp1396-1399.pdf.

Liu, X., Wang, Y., Jiang, Z., Chen, Q., Li, H., Braha, D., de Aguiar, M.,A.M., Gershenson, C. et al. 2021. Does Exist Gap-Filling Phenomenon in Stock Markets? *Unifying Themes in Complex Systems X: Proceedings of the Tenth International Conference on Complex Systems.* 2020. Springer Proceedings in Complexity. Springer, Cham. DOI: 10.1007/978-3-030-67318-5_13

Lo, A.W. 2004. The Adaptive Markets Hypothesis. *Journal of Portfolio Management.* 30(5):15-29. DOI:10.3905/jpm.2004.442611.

Lockhart-Ross, S. 2016. *Cross-sectional and time-series momentum on the JSE.* M.Fin. dissertation. University of Cape Town. Available: http://hdl.handle.net/11427/20816.

Metghalchi, M., Hajilee, M. & Hayes, L.A. 2019. Return Predictability and Market Efficiency: Evidence from the Bulgarian Stock Market. *Eastern European Economics.* 57(3):251-268. DOI:10.1080/00128775.2018.1542601.

Momentum Securities, n.d*. Trading Account (ZAR).* Available: https://momentumsecuritiescfd.tradedesk.co/markets/TradingAccountZAR/#SouthAfricanMarkets. [January 23, 2023]

Nasdaq. n.d. *Extended Trading*. Available: https://www.nasdaq.com/market-activity/quotes/extended-trading. [January 23, 2023]

National Bureau of Economic Research. 2021. *Determination of the April 2020 Trough in US Economic Activity.* Available: https://www.nber.org/news/business-cycle-dating-committee-announcement-july-19-2021. [January 23, 2023]

Nedgroup Private Wealth. 2022. *Fees as at August 2022*. Available: https://www.nedbankprivatewealth.co.za/content/private-wealth-sa/south-africa/en/products-and-services/stockbroking/Ourstockbrokingfees.html. [January 23, 2023]

Page, D., McClelland, D. & Auret, C. 2020. Idiosyncratic momentum on the JSE. *The Investment Analysts Journal.* 49(3):180-198. DOI:10.1080/10293523.2020.1783864.

Page, M.J. & Way, C.V. 1992. Stock market over-reaction: The South African evidence. *Investment Analysts Journal.* 21(36):35-49. DOI: 10.1080/10293523.1992.11082314.

Plastun, A., Makarenko, I., Khomutenko, L., Shcherbak, S. & Tryfonova, O. 2019. Exploring price gap anomaly in the Ukrainian stock market. *Investment Management & Financial Innovations.* 16(2):150-158. DOI:10.21511/imfi.16(2).2019.13.

Plastun, A., Sibande, X., Gupta, R. & Wohar, M.E. 2020. Price gap anomaly in the US stock market: The whole story. *The North American Journal of Economics and Finance.* 52:101177. DOI:10.1016/j.najef.2020.101177.

PSG Securities, n.d. *Ready to take stock of your wealth?* Available: https://www.psg.co.za/files/wealth/forms/KID-Local-Share-Trading-Account-Direct-Clients.pdf. [January 23, 2023]

Raubenheimer, H. 2010. Concentration in the South African equity market and its implied restrictions on the long only equity fund manager's opportunity set. *South African Journal of Business Management*. 41(4):1-10. Available: http://hdl.handle.net/11427/27543.

Sanlam Securities, n.d. *Reliable, Simple, Innovative. Trade on your own or with our help using Sanlam iTrade.* Available: https://www.sanlamitrade.co.za/w/Costs.aspx. [January 23, 2023]

Serletis, A. & Rosenberg, A.A. 2009. Mean reversion in the US stock market. *Chaos, Solitons and Fractals.* 40(4):2007. DOI:10.1016/j.chaos.2007.09.085.

Seybert, N. & Yang, H.I. 2012. The Party's Over: The Role of Earnings Guidance in Resolving Sentiment-Driven Overvaluation. *Management Science.* 58(2):308-319. DOI:10.1287/mnsc.1110.1386.

Sharenet Securities, n.d. *Local Fees.* Available: https://www.sharenet.co.za/feeschedule/#:~:text=and%20experienced%20traders.-,An%20annual%20account%20fee%20of%200.15%25%20is%20payable%2C%20calculated%20pro,markets%20on%2060%20exchanges%20worldwide. [January 23, 2023]

Shleifer, A. & Vishny, R.W. 1997. The Limits of Arbitrage. *The Journal of Finance (New York).* 52(1):35-55. DOI:10.1111/j.1540-6261.1997.tb03807.x.

Standard Bank Securities, n.d. *Fees and costs*. Available: https://onlinesharetrading.standardbank.co.za/standimg/OST/fees-and-costs.html. [January 23, 2023]

Stübinger, J. & Schneider, L. 2019. Statistical Arbitrage with Mean-Reverting Overnight Price Gaps on High-Frequency Data of the S&P 500. *Journal of Risk and Financial Management.* 12(2):51. DOI:10.3390/jrfm12020051.

Tradeciety.com. n.d. How To Use Moving Averages – Moving Average Trading 101. Available: https://tradeciety.com/how-to-use-moving-averages/. [January 23, 2023]

TradingHours.com. n.d. JSE Trading Hours. Available: https://www.tradinghours.com/markets/jse/hours. [January 23, 2023]

TSB Securities, n.d*. Open an Account – Trade for Yourself*. Available: http://www.thebesec.co.za/downloads/Thebe_Stockbroking_Fee_and_Brokerage_Schedule.pdf. [January 23, 2023]

United States Securities and Exchange Commission. 2011. *Investor Bulletin: After-Hours Trading.* Available: https://www.investor.gov/introduction-investing/general-resources/news-alerts/alerts-bulletins/investor-bulletins-42. [January 23, 2023]

Uyanto, S.S. 2019. Monte Carlo power comparison of seven most commonly used heteroscedasticity tests. *Communications in Statistics - Simulation and Computation.* 51(4):2065-2082. DOI:10.1080/03610918.2019.1692031.

Van Heerden, J.D. & Van Rensburg, P. 2015. The cross-section of Johannesburg Securities Exchange listed equity returns (1994-2011). *Studies in Economics and Finance (Charlotte, N.C.).* 32(4):422-444. DOI:10.1108/SEF-09-2014-0181.

Wakandigara, V. 2019. *Testing adaptive market efficiency in the presence of non-Gaussian uncertainties.* M.Fin. dissertation. University of Cape Town. Available: http://hdl.handle.net/11427/31299.

Zach. 2020. *How to Perform a One-Way ANOVA by Hand*. Available: https://www.statology.org/one-way-anova-by-hand/. [January 23, 2023]

Zhong, Z. 2008. *The predictability of overnight information*. M.Fin. dissertation. Singapore Management University. Available: http://ink.library.smu.edu.sg/cgi/viewcontent.cgi?article=1045&context=etd_coll.

Appendices

Appendix 1 – Regression Assumption Checks for Model M2A

Model 2A (M2A)

Linear Relationship Analysis: Fitted/Residuals Plot

Figure A1: Plot of Residual Values Against Fitted Values for M2A

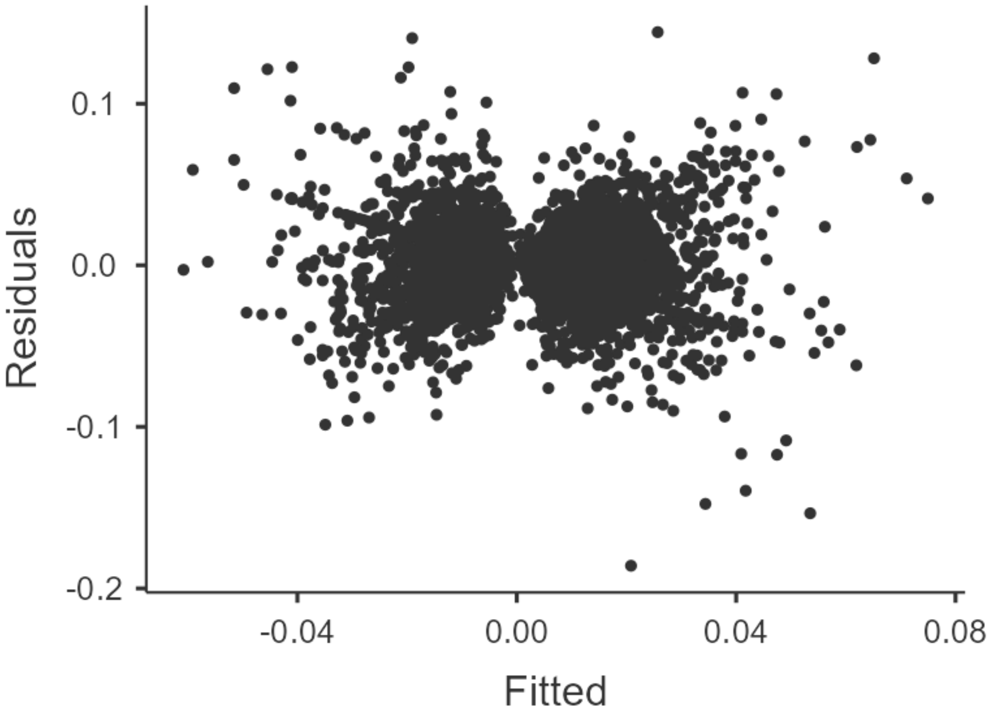

Multicollinearity Analysis: Variance Inflation Factors

Table A1: Collinearity Statistics for Predictor Variables in M2A

Collinearity Statistics

	VIF	Tolerance
Gap (t-1)	1.09	0.916
Gap (t-2)	1.17	0.852
Intraday Change (t-1)	1.27	0.787
Trend (t-1)	1.02	0.977
Cumulative Return (t-5)	1.05	0.955
Intraday Range (t-1)	1.01	0.986

Autocorrelation & Homoscedasticity Analysis Test Statistics

Table A2: Autocorrelation & Homoscedasticity Statistics for Predictor Variables in M2A

	Test Statistic	P-value
Homoscedasticity: Harrison-McCabe	0.419	<0.001
Autocorrelation: Durbin-Watson	1.8	<0.001

Normally Distributed Residuals Analysis: Q-Q Plot

Figure A2: Q-Q Plot of Standardized Residuals for M2A

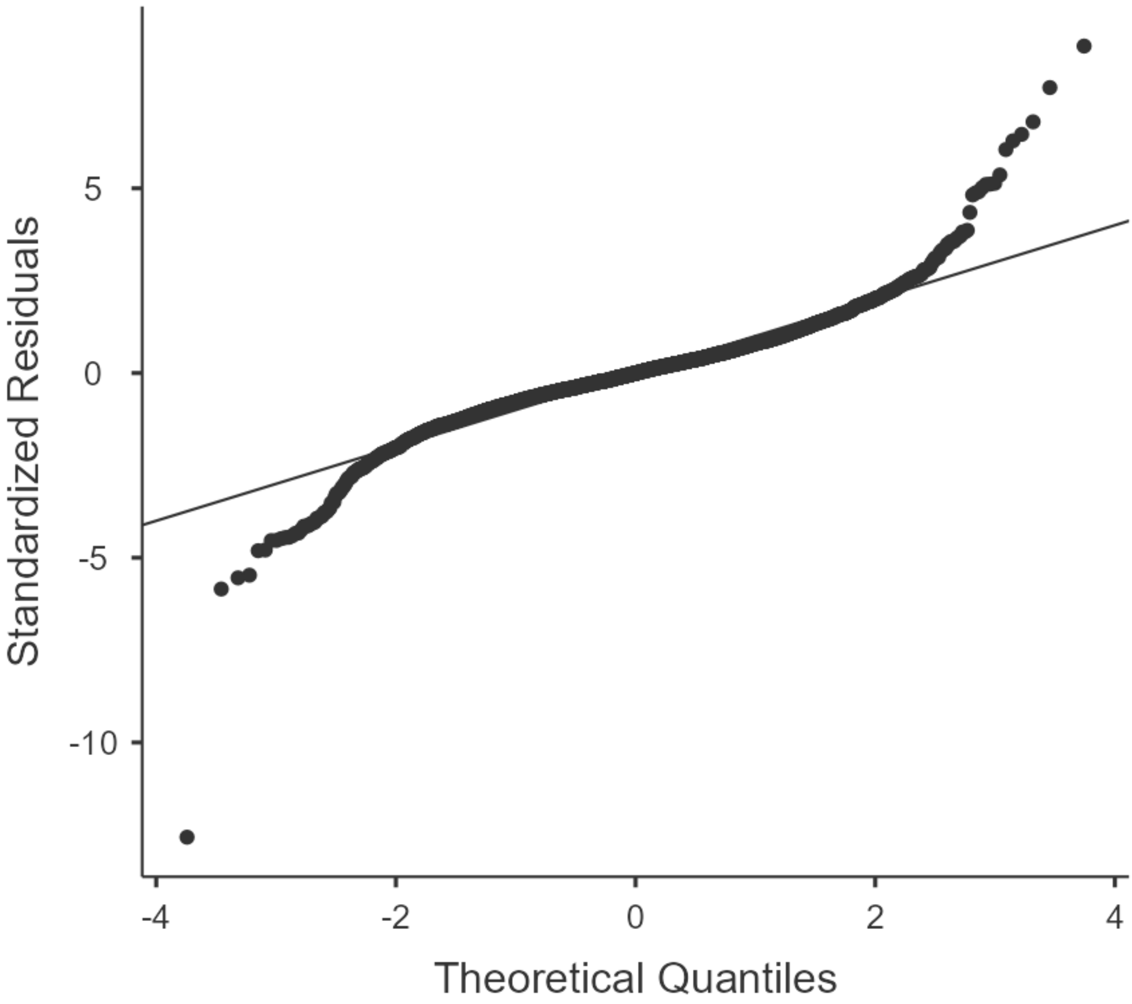

Model 3

Linear Relationship Analysis: Fitted/Residuals Plot

Figure A3: Plot of Residual Values Against Fitted Values for M3

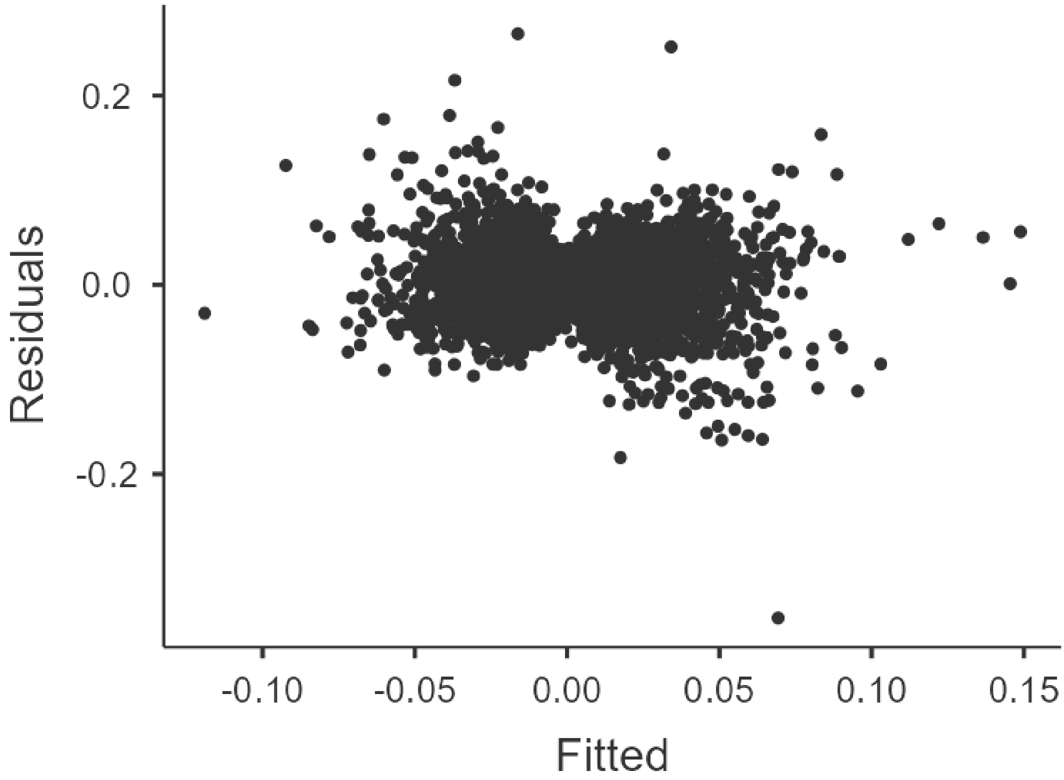

Multicollinearity Analysis: Variance Inflation Factors

Table A3: Collinearity Statistics for Predictor Variables in M3

Collinearity Statistics

	VIF	Tolerance
Gap (t-1)	1.05	0.954
Gap (t-2)	1.1	0.911
Intraday Change (t-1)	1.18	0.849
Trend (t-1)	1.03	0.975
Cumulative Return (t-5)	1.07	0.931
Intraday Range (t-1)	1.01	0.992

Autocorrelation & Homoscedasticity Analysis Test Statistics

Table A4: Autocorrelation & Homoscedasticity Statistics for Predictor Variables in M3

	Test Statistic	P-value
Homoscedasticity: Harrison-McCabe	0.410	<0.001
Autocorrelation: Durbin-Watson	1.79	<0.001

Normally Distributed Residuals Analysis: Q-Q Plot

Figure A4: Q-Q Plot of Standardized Residuals for M3

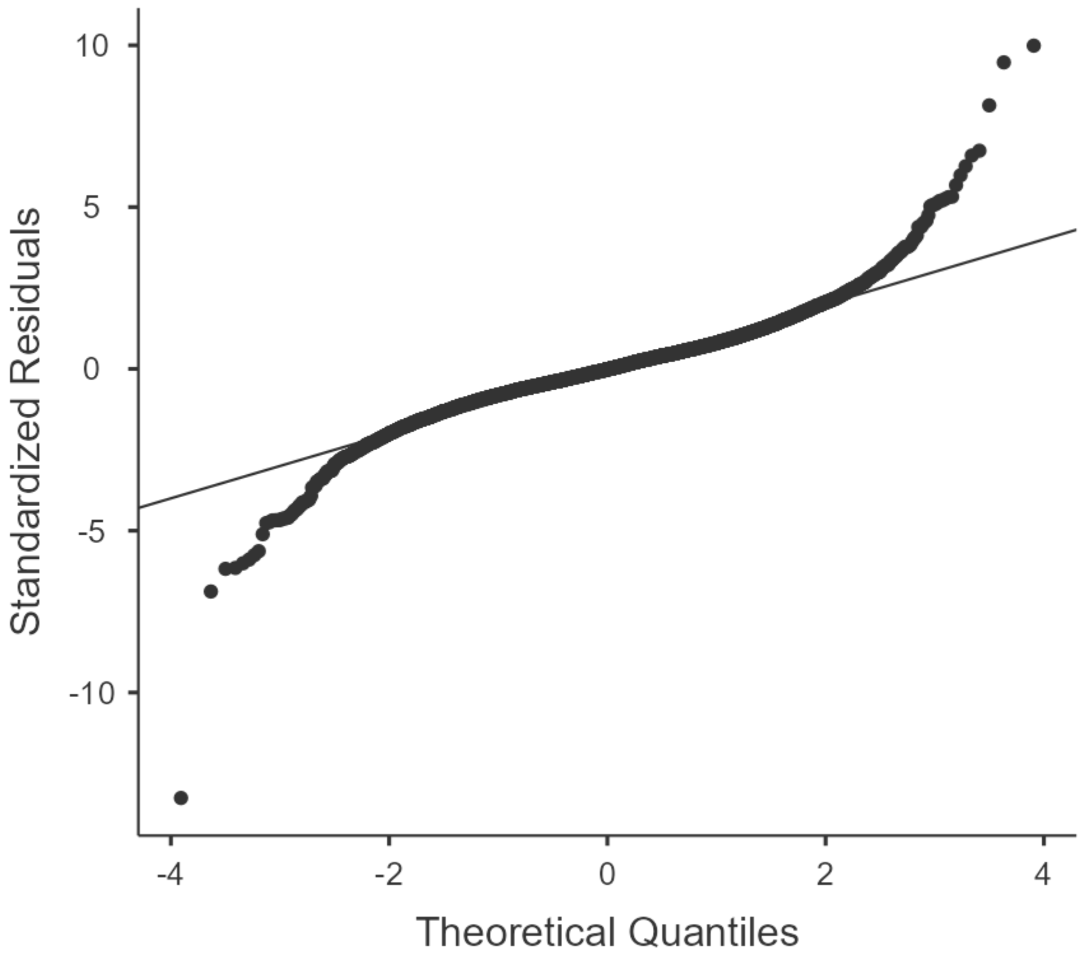

Model 4

Linear Relationship Analysis: Fitted/Residuals Plot

Figure A5: Plot of Residual Values Against Fitted Values for M4

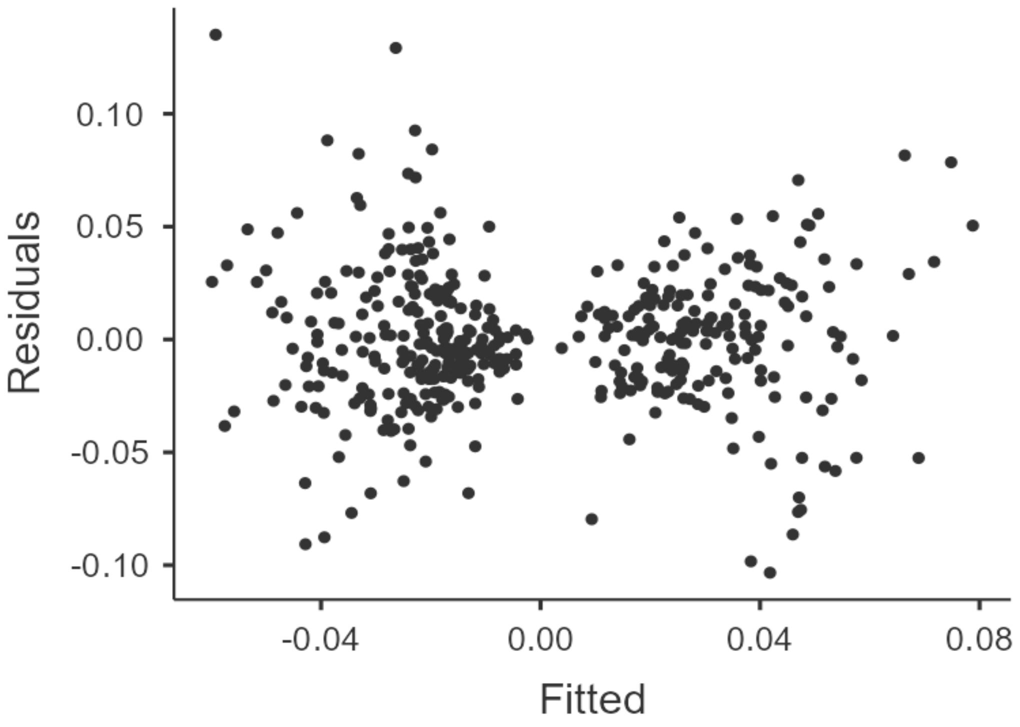

Multicollinearity Analysis: Variance Inflation Factors

Table A5: Collinearity Statistics for Predictor Variables in M4

Collinearity Statistics

	VIF	Tolerance
Gap (t-1)	1.04	0.957
Gap (t-2)	1.14	0.88
Intraday Change (t-1)	1.27	0.79
Trend (t-1)	1.08	0.929
Cumulative Return (t-5)	1.32	0.76
Intraday Range (t-1)	1.03	0.974

Autocorrelation & Homoscedasticity Analysis Test Statistics

Table A6: Autocorrelation & Homoscedasticity Statistics for Predictor Variables in M4

	Test Statistic	P-value
Homoscedasticity: Harrison-McCabe	0.682	1.000
Autocorrelation: Durbin-Watson	1.75	0.002

Normally Distributed Residuals Analysis: Q-Q Plot

Figure A6: Q-Q Plot of Standardized Residuals for M4

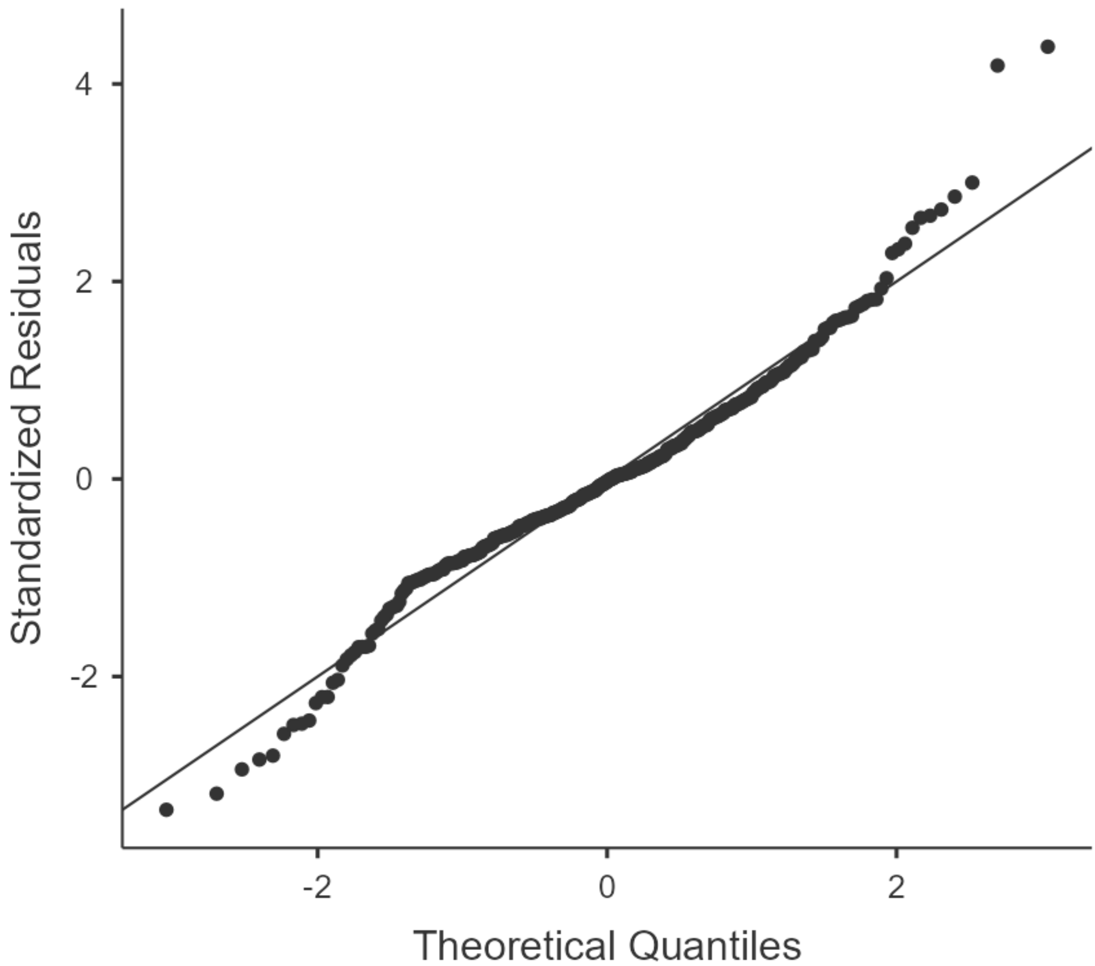

Model 5

Linear Relationship Analysis: Fitted/Residuals Plot

Figure A7: Plot of Residual Values Against Fitted Values for M5

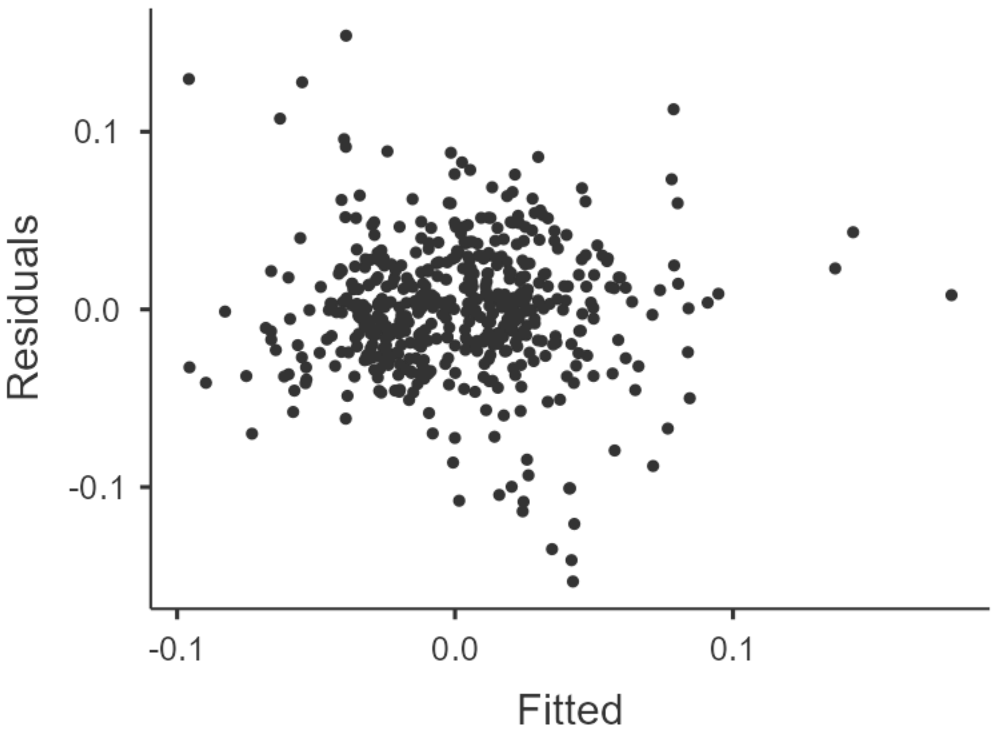

Multicollinearity Analysis: Variance Inflation Factors

Table A7: Collinearity Statistics for Predictor Variables in M5

Collinearity Statistics

	VIF	Tolerance
Gap (t-1)	1.11	0.899
Gap (t-2)	1.14	0.881
Intraday Change (t-1)	1.46	0.683
Trend (t-1)	1.07	0.935
Cumulative Return (t-5)	1.44	0.695
Intraday Range (t-1)	1.05	0.955

Autocorrelation & Homoscedasticity Analysis Test Statistics

Table A8: Autocorrelation & Homoscedasticity Statistics for Predictor Variables in M5

	Test Statistic	P-value
Homoscedasticity: Harrison-McCabe	0.7	1.000
Autocorrelation: Durbin-Watson	1.85	0.064

Normally Distributed Residuals Analysis: Q-Q Plot

Figure A8: Q-Q Plot of Standardized Residuals for M5

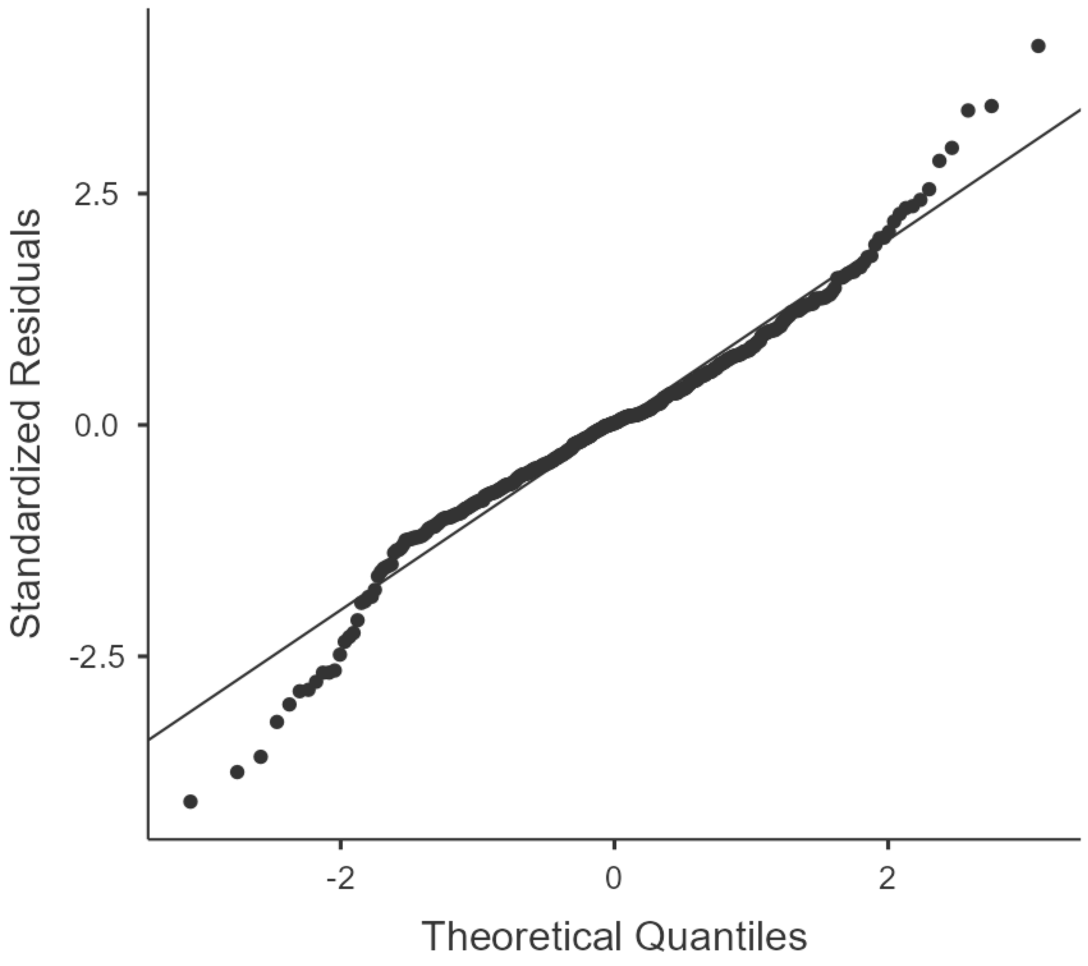

Turn-it-in Report

djgjoh005:Ian_de_Jager_Masters_Thesis.docx

ORIGINALITY REPORT

0%	0%	0%	0%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

Exclude quotes	On	Exclude matches	< 5%
Exclude bibliography	On		