MCOM DISSERTATION

An investigation into reference-day risk-free metrics in the context of modern portfolio theory on the JSE

by

Samuel G Feinstein (FNSSAM001)

Department of Finance and Tax, University of Cape Town, South Africa, 7700

SUBMITTED TO THE UNIVERSITY OF CAPE TOWN

In partial fulfilment of the requirements for the degree Masters of Commerce Specialising in Finance in the field of Investment Management.

Supervisor: Associate Professor Kanshukan Rajaratnam

Department of Finance and Tax, University of Cape Town

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the partial fulfilment of the requirements for the degree Master of Commerce Specialising in Finance in the field of Investment Management. It has not been submitted before for any other degree or examination at any other university.

Signed by candidate

____

10th February 2019

Abstract

Modern portfolio theory (MPT), asset pricing models and broader financial modelling are dependent upon the accuracy of input parameters. For example, the accuracy of expected returns, standard deviations and correlations as an input into MPT will result in a more efficient selection of the optimal portfolio. These metrics are exposed to reference-day risk which is the variation in input estimation due to the selection of initial reference-day in calculations. This paper examines whether a change in reference-day, the day on which a metric is calculated, significantly affects estimates of risk-return metrics on the Johannesburg Stock Exchange (JSE). Thereafter, it applies these findings to the asset allocation problem of constructing a maximum Sharpe portfolio. The objective of this paper is to further prior research through the evaluation of an alternative simulation method and an extension of the range of tested metrics. The advancement of this prior research is achieved through the use of the Cholesky decomposition and a nonparametric bootstrapping procedure to generate reference-day risk-free estimates for average returns, standard deviations, correlations and betas. Furthermore, this paper applies the reference-day risk-free metrics to the construction of optimal multi-asset portfolios in the mean-variance framework. The findings suggest that through the use of a five-year period of monthly returns, the selection of a reference-day materially affects risk-return metrics and the subsequent portfolio characteristics that are based upon these metrics. The performance of portfolios, optimised on each reference-day, ranged between 10% during the out-of-sample period. Additionally, using traditional end of month data resulted in underperformance of out-of-sample, overstated average returns, understated standard deviations and lower correlations between asset classes. Based on these findings we propose an alternative bootstrapping method for calculating reference-day risk-free metrics which reduces the effect of reference-day risk. The purpose of this methodology is to use these estimates for portfolio construction, risk management and asset pricing. The results of this paper indicate that reference-day risk makes a material difference in portfolio construction.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

Throughout the history of financial markets, economic agents have sought to understand the nature and behaviour of asset returns. The adoption of quantitative methods from disciplines such as mathematics and statistics provides a formal framework from which to study asset returns and it forms the basis for many of the core theories of modern finance. Over the last century, these adoptions have revolutionised portfolio construction, risk management and asset pricing models. The statistical properties of asset returns are core components of these fields, providing a standardised framework for asset risk-return analysis and portfolio construction. The perceived benefits of this framework are dependent on an understanding of the underlying stochastic process of returns. Financial return distributions, their associated moments and the relationships between assets, provide a framework for this understanding and allows for the study of asset's behaviour under different environments over time. In tandem to this is the exposure of these returns to specific factors, for example, systematic risk. This exposure, measured by beta, is a key input for many theories of empirical finance such as the capital asset pricing model (CAPM) and multifactor models. Beta, in combination with the distribution moments and correlations, provides a toolbox of risk-return metrics for investment practitioners to allocate capital and manage risk under conditions of uncertainty efficiently.

To make accurate inferences regarding asset prices and their interrelationships it is important the best estimate of inputs are used. Examples of these inferences are; to determine if equities are correlated with bonds; if a particular stock is considered aggressive or defensive; for testing market efficiency; to perform event studies, or to construct portfolios. Thus, the reliability of any analysis is dependent on the accuracy of the model inputs which are exposed to the core concept of this paper – reference-day risk. Reference-day risk, as defined by Acker and Duck (2007), is the variation in input estimation due to the selection of initial reference-day in calculations and forms the centre of this paper's research problem. The concept of reference-day risk has been documented across global markets by Acker and Duck (2007); Feinstein, Polden, Richardson and Rajaratnam (2016); Dimitrov and Govindaraj (2007); Sahadev, Ward and Muller (2018); Gonzalez, Rodriguez and Stein (2014), as well as Baker, Rajaratnam and Flint (2016), with various methods being used to adjust for its influence on statistical variables. The predominant focus of the literature has been on estimates of systematic risk or beta, with no standard method to both measure and adjust for reference-day risk across other metrics, such as average returns, standard deviations and correlations. Additionally, the research in the South African context has been particularly limited, with mixed results in proving the existence of and adjusting for reference-day risk across metrics. Furthermore, to the authors own knowledge no research into the implications of reference-day risk on portfolio theory, which is sensitive to input estimation error is available (Michaud and Michaud, 2008). Consequently, investment practitioners and academics who use statistical inputs to understand or test financial phenomena should account for the extent that the selection of initial reference-day influences their results. This is accounted for primarily because these inputs are based on partial datasets (last closing price of each month); while excluding the other trading day returns during the same period which could provide material information (Acker and Duck, 2007).

As a result of these problems, the objectives of this paper are first, to establish the extent to which reference-day risk impacts commonly used risk-return metrics on the JSE. Secondly, to provide a nonparametric bootstrapping method to both; measure and adjust for reference-day risk in average returns, standard deviations, correlations and betas. Thirdly, to understand the nature and interrelationship of reference-day risk across these variables. Lastly, to apply these adjusted metrics to the construction of a reference-day risk-free efficient portfolio using the mean-variance framework as per Markowitz (1952).

This paper is structured as follows; Chapter 2 provides relevant empirical literature on reference-day risk and estimation errors in the international and South African context. Chapter 3 presents the theory and method for estimating reference-day risk-free, risk-return metrics through the use of simulated data and outlines the data and assumptions. Chapter 4 provides a discussion on the degree of reference-day risk on the JSE and estimates reference-day risk-free average returns, standard deviations, correlations and betas. Chapter 5 applies these findings through a case study, which focuses on the construction of a reference-day risk-free portfolio. These findings illustrate how reference-day risk translates from single stocks and metrics to asset classes and portfolios. The section concludes with a comparative analysis and a discussion regarding the performance of the reference-day risk-free portfolio. In Chapter 6, we conclude our findings, summarise the results and highlight the limitations present.

Chapter 2: Literature review

This literature review begins with an introduction to estimation error followed by an overview of the empirical literature on reference-day risk from both an international and South African context. A discussion around the estimation of reference-day risk-free metrics on the Johannesburg Stock Exchange (JSE) is then undertaken with a focus on the inputs to modern portfolio theory (average returns, standard deviations and correlations). This literature review concludes by discussing the simulation and nonparametric bootstrapping method that is employed in this paper. While the seminal work provides sufficient evidence of the existence of reference-day risk across metrics, such as average returns; median returns; standard deviations; betas and correlations, as well as useful methods of adjusting estimates for reference-day risk, there has been minimal focus on the application of these findings to modern portfolio theory. Specifically, in the South African context research has focused exclusively on systematic risk or beta (Baker et al., 2016; Sahadev et al., 2018).

The concept of mean-variance optimisation developed by Markowitz (1952), forms the foundation of modern portfolio theory (MPT) and serves as an instrument to systematically allocate resources to various investment alternatives. The optimisation process includes investors preferences, expectations of return, risk and correlations (either ex-post or ex-ante), under the assumption that through combining uncorrelated investments diversification can reduce the risk for a given amount of expected return. Furthermore, Markowitz (1952), outlined the assumption of a rational investor who aimed to maximise expected return per unit of risk, and through the mean-variance framework created the efficient frontier – the portfolio set that maximised return for a specified level of risk. On this efficient frontier sits the optimal portfolio commonly referred to as the tangency portfolio. This tangency portfolio represents the highest level of utility the investor could achieve with a set of constraints and a risk-free rate (Greig, 2016). Despite the empirical support and practical implications of MPT investment practitioners have shown reluctance regarding its use in practice which Michaud (1989) termed the “Markowitz optimisation enigma”. Fisher and Statman (1997), reiterated this finding, concluding that the optimal portfolio reflected the constraints imposed by the investor more so than the optimisation itself. When Markowitz (1952) introduced his seminal work, he made an assumption that investors knew exactly the input parameters and distributions of assets to be included in the portfolio. However, in practice, portfolio optimisation techniques face risks of incorrect input estimation – one example being reference-day risk, which has until recently been considered by academics and investment practitioners as negligible (Acker and Duck, 2007).

2.1. Estimation error

In the literature, estimation risk refers to uncertainty due to sampling variation in the estimation of variables and is thought to apply to independent samples drawn from the same population (Acker and Duck, 2007: 1). Estimation risk has been researched extensively by Stein (1955), showing that traditional sample statistics are not suitable for multivariate analyses. Barry (1974) and Michaud and Michaud (2008) provided methods to identify and correct for estimation risk, whereas Chopra and Ziemba (1993) performed empirical tests to show that errors in return estimates are more important than errors in risk estimates. The Monte Carlo simulation was used by Bey, Burgess and Cook (1990) who took sample data, randomised it using bootstrap resampling and then recalculated the input parameters. They ran this simulation several times and found the optimal portfolio for each scenario, eliminated outliers and chose the most appropriate portfolio subject to prespecified investor constraints. Michaud (1998), performed a similar approach using a parametric resampling technique. Correspondingly, reference-day risk is by definition a form of estimation risk which this paper aims to measure and adjust for, in effect contributing to the broader estimation risk literature.

2.2. Evidence of the existence of reference-day risk

Evidence of reference-day risk has been found across geographies, variables, time and asset classes. Acker and Duck (2007), presented the notion of reference-day risk through the example of considering an estimate of a stocks' sensitivity to the market (beta), its volatility (standard deviation), its return (average return) or its relationship with other assets (correlation). In practice, these variables are estimated using monthly historical price data, and the returns are calculated using the closing price from each month. However, given daily data, any initial reference-day could be chosen, for example, the 1st day of the month, the 3rd, the 4th, or the 10th, which cover the same period (five-years); consequently, it can be assumed to be taken from the same population. Through this process, a number of monthly return series can be generated using different initial reference-days, which can result in a number of different estimates for the same period. Acker and Duck (2007), termed this type of variation, 'reference-day' variation and the related estimation risk as 'reference-day risk'. This literature review continues by providing evidence of reference-day risk in both the international and South African markets.

2.2.1. Evidence from international markets

Acker and Duck (2007), investigated the presence of reference-day risk in a sample of stocks drawn from the S&P 500 (459 companies) throughout fifteen-years which ended on 31 December 2005. Parameters including betas, average returns, medians, standard deviations and correlations were estimated over three non-overlapping five-year periods, and the results were checked for data dependency using both CRSP and Datastream. The geographic dependency of the results was tested on the United Kingdom's FTSE All Share Index.

The estimates of means, medians and variances of the stock's monthly returns for each of the twenty-eight reference-days were calculated and the highest, lowest and range difference across days were recorded for each stock's corresponding metric. The results indicated that the selection of initial reference-day could have noticeable effects on the estimates of population variables. The selection of reference-day can double a stocks' average monthly return or change the estimate from positive to negative (Acker and Duck, 2007: 7). The effect on median returns and variances was more extreme, with larger ranges and more frequent sign changes in medians than means. In some cases, the variances more than doubled for individual companies and the broader market (Acker and Duck, 2007: 8). Within these outcomes, the highest and lowest estimates were frequently separated by a large number of reference-days, and the results indicate certain reference-days might have a higher association to low or high estimates within the range. As an example, in the five-year period, ended December 2005, a stock's average monthly return (unannualised) was -0.239% moving up to +0.934% using a different reference-day, the median return (unannualised) on another stock over the period, ended December 1995, was estimated to be +6.07% using one reference-day but -1.49% using a reference-day a week later.

Furthermore, to test the effect of reference-day risk on correlations Acker and Duck (2007), used monthly returns between nine groups of fifty-one stocks, calculating correlations between each of the stocks in a selected group. These findings confirmed the aforementioned parameters with the choice of initial reference-day causing a change in the sign of correlations as well as changes of up 0.8. The average change in correlation between consecutive reference-days was 0.04 to 0.05 with an absolute maximum change of 0.64 over the three periods analysed (Acker and Duck, 2007: 9).

Acker and Duck (2007), further investigated the sensitivity of both the capital asset pricing model (CAPM) and the Fama and French (1993) betas using the same sample of stocks, the S&P 500 index as a market proxy and the United States 3-month Treasury Bill as the risk-free rate. The results indicated CAPM betas had a large degree of sensitivity to the reference-day in which they were calculated changing by as much as 4 (from -2 to +2) (Acker and Ducker, 2007: 11). For the four companies with the least sensitivity, altering the reference-day could determine if the stock was categorised as aggressive (above 1) or defensive (below 1) (Acker and Duck, 2007: 11). Overall 53% of the companies had betas varying from below 1 to above 1 as the reference-day was altered and 8% moved from negative to positive. Furthermore, when the difference between the smallest and largest beta was high, the smallest beta was frequently not significantly different from zero, while the largest was statistically significant. The results from the Fama and French (1993) betas indicate a similar degree of sensitivity to reference-day risk, with the size and valuation factor betas as sensitive as the market betas, displaying similar patterns (Acker and Duck, 2007: 12). These results demonstrate that reference-day risk, in individual inputs, can influence asset pricing models or portfolios that are constructed using these inputs.

To demonstrate that this was not the result of idiosyncratic influences on single stocks and to provide long-term data that are not open to survivorship bias Acker and Duck (2007), investigated the correlation between international equity indices, which are important for international diversification. The correlation was analysed through the lens of three prisms, firstly their level, secondly their variability over time and thirdly their variability in extreme market conditions. To support their previous findings a different fifteen-year period was used (1975 – 1989) for the first two lenses and a thirty-year sample for the third. The results demonstrated that reference-day risk was as pertinent for the analysis of the behaviour of international equity markets as it was for the analysis of the behaviour of individual company's stock prices within a single market. The estimate of the correlation with the United States (US) market changed by as much as 0.46 when the initial reference-day changed. For Italy and Japan, the highest correlation was twice that of the lowest, and in the last five-year period, more than five times higher. Additionally, the change in correlations between consecutive periods was particularly dependent on the selection of initial reference-day. Where for the countries analysed the selection of reference-day determined if the correlation had strengthened or weakened over time. The range across countries was typically between 0.2 to 0.3 (Acker and Duck, 2007: 14).

Moreover, to judge the presence of reference-day risk in the tails of the distributions is to answer the question as to whether correlations increase or decrease when markets experience volatile behaviour (Longin and Solnik, 2001). Similarly, there was considerable variation across reference-days, where the variation tended to be more severe the fatter the tail and the fewer the observations present in these periods (Acker and Duck, 2007: 15). For example, in the correlations between the United States (US) and United Kingdom (UK), each reference-day displayed a pattern akin to the reference-days surrounding it. The correlation based on the 5th reference-day falls in bull markets and rises with increasing truncation in bear markets, while on the 12th they fall but only in bull markets. Additionally, on the 19th they fell in bear markets and rose in bull markets (Acker and Duck, 2007: 16). These results have implications for global investors who aim to build globally diversified portfolios as their estimates of correlations are influenced by the choice of reference-day rather than the underlying co-movement. Thus, it is imperative that a consistent method is used to estimate such inputs which account for this risk.

Dimitrov and Govindaraj (2007), investigated the data dependency of the Acker and Duck (2007) findings' through 439 stocks from the Centre for Research in Security Prices (CRSP) and expanded the research through the use of daily rather than monthly returns over the period of January 1995 to December 1999. They constructed several sixty-month return series covering nineteen reference-days, which used daily adjusted returns and confirmed the findings of Acker and Duck (2007), that the selection of reference-day significantly influences the estimates of beta. These results were confirmed for both individual companies as well as broad market indices in the United States market.

These findings were in turn extended by Gonzalez et al. (2014), by expanding on the amount of data and the methods used in the prior studies. Through a sample of 1563 companies that were traded on the NYSE, AMEX and NASDAQ exchanges from 2007 to 2011, they confirmed that the selection of a reference-day resulted in differences in betas that were statistically significant. Moreover, Gonzalez et al. (2016), extended these findings by utilising a dataset of daily equity mutual fund returns to examine the influence that reference-day risk had on regression alphas, a statistic that is broadly employed to calculate fund performance. The outcomes acted in accordance with the behaviour noticed in other financial asset statistics: for a substantial portion of funds that exhibit significant alphas; significance is contingent on the reference-day utilised. This conclusion generates uncertainty regarding the inferences formerly attained employing this statistic. Additional tests demonstrate that the fluctuation of the alpha and its significance does not have a recognisable order and is not associated with fund attributes. Some methodologies have been conducted which could reduce the issue; these incorporate factor models and regression. However, utilising Student-t distributed errors did not produce positive results (Gonzalez et al., 2016).

This evidence on reference-day risk across parameters, time and geographies illustrate that its' effects cannot be seen as negligible as previous research suggested. Additionally, the selection of a reference-day can have varying results on the estimates of metrics generated from monthly return data. The inferences pulled from these types of estimates should be considered as sensitive when they are founded on monthly returns that are calculated from a single reference-day in a period. Following this, these findings will be extended to the South African context.

2.2.2. Evidence from South Africa

Baker et al. (2016), investigated the presence of reference-day risk on the South African Stock Exchange (JSE) over a five-year period, 2010 – 2015, and demonstrated its persistence after altering estimates of systematic risk (beta) for common adjustments. These findings confirm the geographic scope of reference-day risk found by Acker and Duck (2007), into emerging markets and specifically in the South African context. Their findings suggest the effect of changing reference-days on beta estimates were more pronounced for some companies (AngloGold Ashanti, Anglo American Platinum and Brait), yet remained relatively constant for others (Mondi, Remgro and Standard Bank). They further tested for significant differences in betas across reference-days using the Pillai-Barlett trace from Fox, Friendly and Weisberg (2013) under two hypotheses:

H_1: That all the betas of a company are equal across reference-days

H_2: That the largest and smallest betas for the same company are significantly different

Their results conclude that H_1 could not be rejected for any companies; while H_2 could be rejected at the 10% level for sixteen of the forty stocks and at the 5% level for six of the companies analysed. This indicates that when the reference-day is altered; there are companies with statistically different betas. Baker et al. (2016), using a bootstrapping method established that reference-day risk was prevalent on the JSE for betas of the top forty companies. This prevalence generates further uncertainty for investors who plan to use estimates for portfolio construction, risk management, or asset pricing.

The findings of Baker et al. (2016) were extended by Feinstein et al. (2016) to include tests for the presence of reference-day risk in metrics other than beta for the components of the JSE Top 40. Their findings were that the three metrics; average returns, standard deviations and betas were influenced by the choice of initial reference-day. Betas and standard deviations had higher levels of sensitivity to reference-day risk than average returns based on the range of the highest and lowest estimates across days. They also showed the extent to which this sensitivity could translate into measures of performance such as Sharpe Ratios, with 18% of the forty company's Sharpe Ratios going from negative to positive when the reference-day was varied.

An additional study by Sahadev et al. (2018), investigated the extent that reference-day risk led to differences in betas on the JSE All Share Index. They confirmed that across twenty reference-days the unadjusted betas varied significantly. Their findings were that fifteen out of the one hundred and thirty-six companies exhibited betas, that could be categorised as positive or negative depending on the selection of day. In order to identify the level which altering the reference-day leads to a variation in betas, Sahadev et al. (2018), performed an ANOVA and found that at the 5% level of significance there was inadequate proof to validate a trading day impact on beta when the reference-day was changed (p-value of 0.501). The insufficient evidence could be caused by the covariance between returns cancelling each other out and is surprising given the ranges of beta estimates across reference-days.

Similarly, these results are consistent with Baker et al. (2016), who were not able to statistically validate an impact on beta across their twenty estimates. Although, when the difference between the minimum and maximum estimate was used it yielded different results, where a t-test comparing the betas between firms demonstrated that the min-max range significantly varied at the 5% level (p-value < 0.0001). Subsequently, evidence is provided regarding the effect of common adjustments made to account for reference-day risk.

2.3. Adjusting for reference-day risk

The presence of traditional sampling variation and the estimation risk that it results in is identified in academic literature, and several efforts have been made to deal with it, mainly in terms of beta estimation and reference-day risk. These beta methods include that of Blume (1971), Vasicek (1973) and Dimson (1979) in the international context and Baker et al. (2016) and Sahadev et al. (2018) in the South African context. Outlined below are the key findings that these methods have when they adjust for reference-day risk.

Acker and Duck (2007), found that the adjustments recommended by Blume (1971) and Vasicek (1973), reduced the variation of beta estimates for companies across reference-days. The adjusted betas were sensitive to the selection of a reference-day and reduced the average range between the highest and lowest beta for all companies. However, the selection of a reference-day could cause the betas to be double those attained from other reference-days and could result in them being re-categorised as aggressive or defensive depending on the choice of the day. The reference-day variation using the Dimson (1979) method made no improvement over the unadjusted alternative, and in some cases was worse. These results show that adjusting beta estimates for thin trading using the Dimson (1979) method does not make a difference to reference-day variation. While these adjustments have limited usefulness for beta estimation, they are less useful regarding the estimation of reference-day risk-free correlations, means and standard deviations – which are inputs in modern portfolio theory.

In addition, an alternative method was employed by Fama and French (1993), by structuring portfolios of stocks and building estimates on portfolios rather than single stocks. This technique is often utilised in tests of asset-pricing models as well as in the examinations of stock market anomalies. Acker and Duck (2007), found this method to considerably reduce the range between minimum and maximum values across reference-days but did not eliminate the impact of the reference-day choice on estimation. Although the results indicated using portfolio-level returns was the most successful method, this was primarily due to the averaging property embodied in forming portfolios, which emphasises that diversification reduces portfolio risk and reference-day risk. Despite this averaging property, there remained an amount of variability in estimation, and thus reference-day risk has important implications regarding the implementation of modern portfolio theory based on these inputs.

Furthermore, as a result of most institutional investors making adjustments to beta, Baker et al. (2016), investigated whether reference-day risk persisted even after making the adjustments by Blume (1971), Dimson (1979) and Vasicek (1973), on JSE listed companies. They found that the Blume (1971) betas increased the difference between the highest and lowest betas for nineteen of the thirty-one companies with a single stock's beta increasing by 81%, which is consistent with Acker and Duck (2007) – that the adjustment failed to improve estimation when reference-day risk was present. These findings implied that the Blume (1971) adjustment failed to reduce the high-low range among betas when reference-day risk was present. When the Dimson (1979) thin trading adjustment method is used, the average beta range was 0.608, which was more extensive than the average range of OLS betas, of 0.449. Where the beta ranges expanded for thirty of the forty companies. This expansion demonstrates that when the Dimson (1979) method is used, it increases reference-day risk. Under the Vasicek (1973) adjustments, Baker et al. (2016) found that it reduced the occurrences of extreme differences with thirty-five of the forty companies displaying lower average ranges than their unadjusted betas. These findings suggest that after making the adjustments from Blume (1971), Dimson (1979) and Vasicek (1973), reference-day risk in betas continued to persists when calculation was based on only a single reference-day.

Moreover, Gonzalez et al. (2014), compared a technique based on the t-distribution for adjusting betas; established by Cademartori, Romo, Campos and Galea (2003), to the standard OLS regression. They proposed the student t-distribution was more appropriate for beta estimation in the presence of reference-day risk as it compensated for the residual in the regression method. Additionally, they proved that the t-distribution technique was better able to integrate the effect of outliers when calculating beta. Gonzalez et al. (2014), also re-tested the Blume (1971) adjustment and concluded as per the research by Acker and Duck (2007), that it was the technique that accounted for the impact of reference-day risk the best. However, they also confirmed that betas exhibited significant variation across reference-days even after adjustment.

Therefore, after making adjustments in estimating statistical variables from returns data, the effects of reference-day risk persist, in both the international and the South African context. Following this, the literature review introduces and discusses the proposed methods for estimating reference-day risk-free metrics for average returns, standard deviations, correlations and betas.

2.4. Estimating reference-day risk-free metrics

There are two primary methods proposed in the literature not only to adjust a given metric for reference-day risk; but to estimate a reference-day risk-free version for a given metric. The first was a nonparametric bootstrapping technique proposed by Baker et al. (2016), and the second was a volume-weighted-average-price (VWAP) proposed by Sahadev et al. (2018). Both of these techniques applied to estimate a reference-day risk-free version of beta.

Baker et al. (2016), proposed a bootstrapping method to estimate a beta independent of reference-day risk. By calculating a beta estimate for each reference-day for instance; day one, day two until day twenty, the average of these days could serve as an approximation of the reference-day risk-free estimate. However, this estimate would create errors as a result of the small sample size (Ader and Ader, 2008: 373). Baker et al. (2016), proposed a bootstrapped distribution method for estimating a reference-day risk-free estimate. The basis of this method was that the underlying distribution of any metric is unobservable, and the expected value of this distribution served as the most accurate metric value. Hence their method focused on the estimation of this point estimate rather than the distribution itself.

To test the robustness of their method, they simulated returns for a hypothetical stock and the market index with a predefined beta relationship and ran their bootstrapping method to test if they were able to estimate the true value of beta for the dataset. Both of the series were based on Brownian motion to control for the level of correlation between the two series. The data is sorted into a table of sixty by twenty (sixty months, twenty days) and a random day is chosen from each month, for example, day one in month one, day five in month two until day n of month sixty, along with the corresponding day in the market sequence. An estimate of beta is then calculated, and the process is repeated with replacement 100 000 times to generate a distribution of betas, which is graphed. This process was tested assuming both normal and uniform distributions. If normality was assumed, then the simulated bootstrapped average was predominately within half of a standard deviation from the predefined value. When the distribution was presumed to be uniform, the bootstrapped distribution remained relatively normal and became more condensed around the mean. The results remained consistent regardless of the assumed beta estimate before the simulation was run. Baker et al. (2016), noted that the point estimate of beta from the distribution was generally equal to the average of the twenty betas calculated across the twenty reference-days, despite the risk of introducing errors because of a small sample size. This provides a quick method for practical applications or for those without access to relevant software for estimating a reference-day risk-free estimate using more advanced methods.

Furthermore, this method was tested by Feinstein et al. (2016) who found betas to behave in a manner consistent with Baker et al. (2016) across days. However, average returns and standard deviations bootstrapped estimates were found to be slightly lower than the average across the twenty reference-days. This phenomenon translated into Sharpe Ratios based on these bootstrapped values which were also understated relative to their twenty-day average.

Alternatively, Sahadev et al. (2018), investigated a volume-weighted-average-price (VWAP) beta. They used the constituents of the JSE All Share Index (136 shares) from DataStream for the period from 31 December 1992 to 30 June 2017. This investigation involved testing the statistical relevance of utilising the VWAP estimate relative to the standard beta method. Sahadev et al. (2018), used a sixty-day (reference-day inclusive) ex-ante VWAP, which was estimated using method from Ting (2006), to calculate beta estimates rather than closing prices alone. In order to comprehend whether the VWAP beta was an improved estimate of systematic risk relative to the unadjusted beta in the presence of reference-day risk, they were both estimated utilising the same sample period and tested for statistical differences using a paired t-test. The two results were confirmed as statistically different.

Firstly, Sahadev et al. (2018), aimed to estimate the distributions of VWAP betas for the 136-company sample, and if these distributions exhibited less dispersion around the mean relative to estimates of traditional betas. Secondly, they sought to identify if a specific beta in the twenty-day range tended to be a more precise estimate of systematic risk. Both graphical and statistical analysis was undertaken akin to Baker et al. (2016), with a focus on the skew and dispersion of the VWAP adjusted and unadjusted betas distributions. The normality of these distributions was tested using statistical tests (Kolmogorov-Smirnov, Cramer-Von Mises and Anderson-Darling) and the variances of each distribution were analysed as per Gonzalez et al. (2014), to test the fit around the mean. Lastly, to validate these findings a Levene's test was conducted, confirming the graphical outcome. The results of the Levene's test indicated that the 20th reference-day exhibited the most suitable betas. In comparison, The VWAP calculation resulted in a less statistically suitable set of betas. In some of the more extreme cases, the VWAP method resulted in slightly more robust estimates, but this could not be directly attributed to the method as much as the inherent volatility in the share prices.

Consequently, this paper aims to build on the bootstrapping method proposed by Baker et al. (2016), due to its statistical reliability and consistency in results. This literature review continues by addressing the choice for a risk-free rate to be used in performance evaluation and for portfolio optimisation.

2.5. Choice of risk-free rate

The risk-free rate is a significant input into the empirical testing of financial theory and forms the foundation for key aspects of finance, for instance portfolio theory, asset pricing models and corporate financial modelling as per Damodaran (2008), Vaihekoski (2014) and Ernst and Young (2009). Risk-free rates predominantly matter due to their use in computing discount rates to discount future cash flows. The cost of equity adds a risk-adjusted equity risk premium to the risk-free rate, the cost of debt adds a default spread, and thus ceteris paribus higher risk-free rates result in higher discount rates, higher risk premiums and lower present values when using discounted cash flow valuation (Damodaran, 2008: 4). This section first defines a risk-free rate; then it addresses common issues in determining an appropriate rate and concludes with a review of relevant literature on a risk-free rate in the South African context.

A risk-free asset is one in which the realised return is always equal to the expected return and meets the conditions outlined in Damodaran (2008), of having cash flows specified at initiation, no reinvestment risk and no default risk in the entity from which the cashflows originate. Further, Damodaran (2008), provides a framework for deciding an appropriate risk-free rate which addresses the idea of real versus nominal rates, finding a risk-free proxy and the nonstationary of the risk-free rate.

Inflation volatility or consistently high inflation can cause nominal variables to become too unstable to use in valuation and modelling. To account for this instability Damodaran (2008) suggests using real variables for inputs such as betas and expected returns. South Africa’s headline inflation rate as calculated by the Consumer Price Index has ranged from three to seven percent over the last eight years, which this paper considers as stable enough to use as inputs into our calculations.

The South African government issues both foreign and domestic currency denominated debt and thus avoids the issue of having no long-term risk-free proxy issued by the government in local currency. South African government bonds, as of July 2018, are rated at Baa3 by Moody’s and BB by S&P, which implies that the returns are not entirely risk-free. This paper accepts rating agencies assessment of country risk. Furthermore, South Africa has historically defaulted numerous times on sovereign debt, and it would thus be incorrect to use unadjusted government bond yields as representative of risk-free rates. This is because there is an element of default risk embedded in the rate, causing the default risk in risk premiums to be double counted, resulting in higher discount rates and lower than actual estimates of value (Damodaran, 2008: 24). To account for this effect Damodaran (2008) suggests two methods of adjusting for risk: using the Credit Default Swap (CDS) spread method or using credit ratings to imply a default spread. In both cases, the spread is subtracted from the ten-year government bond yield to arrive at a better reflection of a risk-free rate. This paper adopts the CDS spread method to adjust for default risk.

In practice, most firms in South Africa use long-term government bonds that have a smoother volatility profile and result in a better fit with the terms of investments of companies (Correia and Cramer, 2008: 12). The survey data provided below aims to capture the most common proxy for a risk-free rate, for the purposes of valuation and financial modelling in the context of South Africa given the conditions outlined by Damodaran (2008). Correia and Cramer (2008), indicated that most firms used a long-term government bond yield as a proxy for the risk-free rate. They found a large disparity in the maturity of the instrument used but found greater use of the R153 bond.

Correia and Cramer (2008), also found that 80% of firms do not adjust for tax when determining the risk-free rate for CAPM. This finding corresponds with the standard form of CAPM from the academic literature. The latest valuation method survey by PwC (2017), shows that these trends have persisted over time with the preference shifting to the 10-year government bond yield. This yield was used by 33% of respondents followed by the R208 by 13% and the R213 by 10%, which is consistent with both the 2016 and 2015 survey results as well as surveys from before 2010. Thus, over the last two decades in South Africa, the risk-free rate has consistently been calculated using longer-dated government bonds – with the 10-year being one of the most liquid and referenced. This paper uses the South African government 10-year bond yield net of the appropriate CDS spread as per the Damodaran (2008) method, as a representation of the risk-free rate for the application to portfolio construction.

2.6. Conclusions drawn from the literature

The preceding literature substantiated the existence of reference-day risk across time, geographies and metrics with reference-day risk remaining prevalent even after making common adjustments. While the literature has focused predominantly on beta, no research has been done into the application of these findings to other metrics such as average returns, standard deviations or correlations in South Africa. A further gap in the literature is with regards to the implementation of a consistent method to both measure reference-day risk and to calculate a reference-day risk-free estimate. Furthermore, no research has been done into the implications of these results for modern portfolio theory, which uses these metrics as inputs. The remainder of this paper aims to extend the above literature in these areas by providing a robust method to both measure and account for reference-day risk in estimation and apply the findings to the construction of an optimal reference-day risk-free portfolio.

Chapter 3: Data and Assumptions

This research utilises the end of day closing prices of the largest fifty-six qualifying companies by market capitalisation on the JSE and the closing level of the JSE All Share Index (JSE), the South African Composite Bond Index (ALBI), the MSCI World Index (MSCI), the FTSE/JSE Property index (SAPY) and the Bloomberg All Commodity Index between 1st January 2013 to 31st December 2018. A qualifying company is one which has six-years of historical returns data during these defined dates and is in the top eighty companies ranked by market capitalisation on the JSE. This price data was sourced from Bloomberg. The choice of fifty-six companies aims to extend previous reference-day risk research on the JSE which has focused on the largest forty companies (Baker et al., 2016); while maintaining a liquid enough sample to minimise the impact of illiquidity as noted by Bradfield (2003). A full list of the fifty-six qualifying companies used in this paper is provided in Appendix A.

In accordance with Baker et al. (2016), Sahadev et al. (2018), Gonzalez et al. (2014), prices were adjusted for corporate actions such as unbundling's, share splits or mergers as well as for dividend pay-outs. All the prices were denominated in South African Rand converted at the appropriate daily closing exchange rate. This paper does not account for survivorship bias in its method; however, it acknowledges that this bias could be introduced through the inclusion of the chosen companies as listed in the All Share Index throughout this time. According to Baker et al. (2016) and Dimson (1979), this paper performs no thin trading adjustment on beta because the sample includes only large capitalisation liquid companies. As a consequence of the variability in the price data available, we standardised our calculations to twenty trading days in a given month. Companies that did not have the full six-years of required returns data were removed from our analysis. For instances where there were greater than twenty trading days in a given month, we used the first twenty days and removed all trading days thereafter. Whereas, for the companies which had fewer than twenty trading days; the last day's price was kept constant in accordance with Acker and Duck (2007). The risk-free rate used in this paper is the average ten-year yield on the South African government bond during the sample period of 8.37%. This rate is adjusted using the average ten-year credit default swap spread of 2.84% as described in section 2.5, resulting in a risk-free rate of 5.53%. In Chapters 4 and 5, we assume that when one calculates average returns, standard deviations, correlations and betas that we are drawing a random variable from a distribution that estimates the real underlying variable. For example, the true underlying systematic risk measured by beta and the total risk by standard deviation.

This paper adopts the assumption that stock and index returns are lognormally distributed when performing simulation, inferential analysis and bootstrapping (Dimson, 1979; Baker et al., 2016). This paper's research is centred around the presumption that companies price returns have verifying degrees of dependence on the market (index) based on the market model:

$$R_t = \alpha + \beta M_t + \varepsilon_t \,,$$

where R_t is stock return at time t, α is the component of return not explained by the market, M_t is the market return at time t, ε_t is the error of the return estimate, and β is the exposure to market movements explained by:

$$\beta_i = \frac{cov(R_i; R_m)}{var(R_m)} \,,$$

where β_i is the beta for company i, $cov(R_i; R_m)$ is the covariance between the returns of company i and the returns of the market, and $var(R_m)$ is the variance of the market returns. This metric, which this paper refers to as beta, reflects an approximation of the systematic risk embodied in a stock's return. Furthermore, this paper assumes that volatility is represented by standard deviation, that ex-post average returns represent the best estimate of ex-ante returns and that all stocks have a relationship with all other assets based on the co-movement of returns. Furthermore, it is assumed that these variables are nonstationary, and thus in the process of measurement are exposed to reference-day risk.

In addition, excess kurtosis is defined as any value in excess of 3, and a correlation greater than 0.5 or less than -0.5 is assumed to represent a strong positive or negative linear relationship between variables respectively. Unless otherwise stated, this paper makes use of annualised returns and standard deviations. When testing for reference-day risk in correlations and betas, it is assumed that the reference index is the ALSI, unless otherwise stated. This assumption allows for comparative analysis against the same underlying benchmark for each of the companies.

Chapter 4: Methodology, Results and Discussion

For the purpose of performing multi-metric analyses, this paper organises five-years of returns into a sixty by twenty table, with the sixty rows representing the sixty months in a five-year period and the twenty columns representing the twenty trading days in each of the sixty months. The first column represents the first day of the month, the second column the second day, and the first two rows represent the first and second month respectively, counting up to the sixtieth month. As discussed in Chapter 3, each month has approximately twenty trading days, where weekends and local public holidays; when the market is closed, are excluded. Each metric is calculated based on the respective reference-day, for example, the first average return is calculated using the returns from the first day from each of the sixty months, the second average return using the second day until all twenty averages are calculated. Correlations and betas are calculated using the same method, where for instance the correlation using the first reference-day is based on the first day of each of the sixty months for both the stock and market returns. Inspecting the average returns, standard deviations, correlations and betas calculated across reference-days, all of these metrics are affected by varying degrees of reference-day risk. Below, we give an overview of the variation in these metrics across reference-days and provide the results of statistical tests to illustrate the significance of this variation.

4.1. Evaluating the existence of reference-day risk in South Africa

This section provides evidence of the existence of reference-day risk in average returns, standard deviations, correlations and betas. Furthermore, the interrelationship of reference-day risk across metrics and the practical implications of reference-day risk for each metric is discussed.

4.1.1. Evidence found in average returns and standard deviations

Average returns over the five-year period for the fifty-six companies had an average range from the lowest to the highest day of 3.9%, a median range of 3.5% and standard deviation of 2% across reference-days. Kumba Iron Ore, Anglo American platinum and Capitec had the highest ranges of

12.5%, 7.4% and 7% respectively. Companies who had the lowest sensitivity to reference-day risk in their average returns were Richemont, Fortress A and MTN who had ranges of 1.5%, 1.6% and 1.7% respectively. Overall there was wide dispersion across both companies and sectors depending on the choice of initial reference-day, with the high-low ranges exhibiting a positive skew of 1.81 and excess kurtosis of 3.2 in the distribution of the fifty-six ranges for average returns.

To investigate if there were significant differences in average returns across the twenty reference-days, we performed an analysis of variance (ANOVA) with the hypotheses:

$H_0: \mu_1 = \mu_2 = \cdots = \mu_{20}$

H_1: Two or more average returns are significantly different accross days

While the underlying sample distributions for each average return may not conform to the ANOVA normality assumption, according to Kuzma and Bohnenblust (2005) as long as the sample sizes are large and equal or nearly equal, moderate departures from normality are not a problem. The samples in this paper conform to this requirement, as they are of equal size for this test. In addition, because the ANOVA test is robust, it is acceptable if the largest standard deviation is less than double the smallest standard deviation in the sample, which was the case for the fifty-six companies analysed in this paper (Sullivan, 2011).

Despite the varying high-low ranges of average returns, the ANOVA results indicated that none of the fifty-six companies had average returns that were significantly different across reference-days, even at the 10% significance level. To reaffirm these results, we investigated if the highest average return was significantly different from the lowest average return across reference-days using a two-sample t-test with the hypotheses:

$H_0:\ \mu_{high} = \mu_{low}$

$H_1:\ \mu_{high} \neq \mu_{low}$

As with the ANOVA results, no companies exhibited significant differences in their highest and lowest average returns across the twenty reference-days, with an average p-value of 0.4606 across the fifty-six companies. The companies with the lowest p-values were Rand Merchant Insurance, INTU properties and Barclays; while the highest were Exxaro, Richemont and MTN. The companies with the lowest p-values, which had the most significant difference between the highest and lowest average returns, were not the companies with the highest absolute high-low range differences. In contrast; however, both MTN and Richemont had two of the lowest absolute ranges but the highest p-values. While the p-values are high relative to significance levels of 5% and 10%, their variation across stocks and sectors; along with the fact that on average they were less than 0.5, imply that the impact of reference-day risk should be considered in empirical tests. This impact should be accounted for during any application that uses average returns as an input.

The initial choice of reference-day in calculating average returns has implications for both portfolio construction and asset pricing. Modern portfolio theory is based on using average returns as an input in estimating efficient portfolios (Elton and Gruber, 1997: 1745). Therefore, if the estimate of return varies considerably across reference-days, the conclusions drawn from using one day in the reference period creates a bias in the end portfolio. This bias results in conclusions that are either not consistent ex-post or are a function of the choice of reference-day and not the true return profile of the underlying asset over the reference period. For example, if the 20^{th} reference-day was used for INTU properties, there was an average return of 0.37%, but if the 1^{st} day was used, there was an average return of -4.98%. The same applies to AngloGold Ashanti whose average return on the 15^{th} reference-day was -2.4%; while it was 2.81% on the 4^{th} day. This phenomenon occurred for three of the fifty-six companies (5%). From a different perspective, reference-day risk applies not only for changes from negative to positive returns but for large deviations in positive returns. For example, Kumba Iron Ore and Rand Merchant Investments had ranges of 12.5% and 6.3% across the twenty reference-days respectively. These results have an impact on the return assumptions an investor makes about a particular company, and on the weight applied to a stock in a portfolio optimised using the mean-variance framework.

The practical implications for asset pricing can be illustrated through the equity risk premium, which is a core component of discount rates used in equity valuation models (Damodaran, 2015). The historical equity risk premium is equal to the excess return of a stock above a risk-free asset over a given period that is represented by the difference between the average return and the risk-free rate. The choice of reference-day influences the average return, and thus the spread above or below the risk-free rate the stock exhibits. Should the choice of reference-day yield high excess returns, then the equity risk premium will be higher, the discount rate higher, and future cash flows that are discounted back would result in a lower present valuation ceteris paribus. The inverse applies if the choice of reference-day understates the true average return for a period.

There are also implications for performance management where for example, if a performance analyst is looking for stocks which had an average return that was higher than the JSE All Share Index over the last five years. For illustrative purposes, the 20th reference-day is used as the average return of the index of 9.32%. For seven of the fifty-six companies (12.5%) the choice of reference-day impacted the conclusion that a stock on average under or outperformed the benchmark. For example, when Shoprite was analysed using the 16th reference-day, there was underperformance of -3.53%; while when using the 1st reference-day there was outperformance of 1.48%.

Therefore, it is important for purposes of portfolio construction, asset pricing and performance measurement that a reliable estimate of the average return for a stock over a specified period be estimated. A full breakdown of the above analysis of average returns can be found in Appendix B.

Following the analysis of average returns, this paper continues by focusing on standard deviations. Standard deviations over the five-year period for the fifty-six companies had an average range from the lowest to the highest day of 6.14%, a median range of 5.3% and a standard deviation of 3%. Anglo American PLC, Glencore and AngloGold Ashanti had the highest ranges of 12.22%, 11.59% and 10.26% respectively. Companies who had the lowest sensitivity to reference-day risk in their standard deviations were AVI, Spar and Richemont who had ranges of 2.86%, 3.33% and 3.42% respectively. Relative to average returns there was wider dispersion in standard deviations across companies over this period, with excess kurtosis of 10.6 and skewness of 3.06 for the distribution of standard deviation ranges of the fifty-six companies.

To investigate whether the standard deviations were significantly different across reference-days, we tested to see if the highest standard deviation was significantly different to the lowest, using an F-test with the hypotheses:

$H_0\!:\ \sigma_{high} = \sigma_{low}$

$H_1\!:\ \sigma_{high} \neq \sigma_{low}$

A requirement for the use of the F-test is that there is an assumption of normality for the sample being analysed. Considering that we are comparing two standard deviations from two reference-days, we have two samples of sixty-days for each company. To test the normality of the sample, a Shapiro-Wilk test was used with the hypotheses:

H_0: the data is normally distributed

H_1: the data is not normally distributed

When the fifty-six companies were analysed on the reference-day which had the highest standard deviation, thirty-eight of the companies (68%) were normally distributed at the 5% level. By comparison, forty-eight of the fifty-six companies (86%) were normal at the 5% level on the reference-day with the lowest standard deviation. These results show that the normality assumption of returns is not consistent across reference-days, with nine of the fifty-six companies (16%) exhibiting normality in one reference-day and not the other. The F-test results point to an average p-value of 0.12 across the fifty-six companies, with fifteen companies (27%) having standard deviations significantly different at the 5% level, thirty-three (60%) at the 10% level and forty-six (82%) companies at the 20% level. These results indicate a strong presence of reference-day risk in standard deviation across reference-days and were significantly higher than average returns.

Given these results and since not all the samples were normal, we did not rely solely on the F-test result. In addition, this paper tested for homogeneity of standard deviations across the twenty reference-days using a Levene's (1960) and Brown-Forsythe (1973) test with the hypotheses:

H_0: $\sigma_1 = \sigma_2 = \cdots = \sigma_{20}$

H_1: At least one standard deviation is significantly different.

Despite the wide high-low ranges in standard deviations across reference-days and the findings of the F-test above, both the Levene and Brown-Forsythe test failed to reject the null hypothesis that standard deviations were equal across reference-days at both the 5% and 10% level. These findings are analogous to those of Baker et al. (2016) and Sahadev et al. (2018), who used a similar test to analyse the dispersion of betas across trading days. They could not statistically verify a reference-day impact across their twenty betas for the largest forty companies on the JSE; however, they found significant differences when comparing the highest and lowest betas. These results correspond with the above results.

These findings have practical implications for performance measurement metrics that use standard deviation as an input, such as the Sharpe ratio, overall risk management processes and portfolio management processes that target specific volatility levels. The value of the Sharpe ratio moves inversely with standard deviation ceteris paribus. This means that when a portfolio's performance is measured based on a reference-day; where the portfolio's standard deviation was high, the Sharpe ratio would be understated relative to the ratio that would be independent of reference-day risk. In contrast, if the reference-day that was chosen resulted in a lower estimate of the standard deviation relative to other reference-days the inverse would apply. The same logic can pertain to portfolios constructed on the basis of targeting a specific volatility level, for example, that of the ALSI. The ALSI volatility ranged from 8.89% to 10.2% across reference-days. Four (7%) companies had their lowest standard deviations below 8.89% but their highest standard deviation above 10.2% across the twenty reference-days. Therefore, targeting a specific volatility level using one reference-day leads to portfolio volatility that deviates significantly from expectations. Thus, it is important for risk management, performance measurement and portfolio construction techniques, that rely on standard deviations as an input, to be based on an estimate that reflects the true volatility profile across reference-days. A breakdown of standard deviation ranges can be found in Appendix C.

4.1.2. Evidence found in betas and correlations

In line with the findings of Gonzalez et al. (2014), Baker et al. (2016) and Sahadev et al. (2018), betas changed depending on the choice of initial reference-day for the fifty-six companies. Betas over the five-year period had an average range from their highest to lowest day of 0.58, a median range of 0.49 and a standard deviation of 0.34. Glencore, Kuma Iron Ore and AngloGold Ashanti had the highest ranges of 1.82, 1.68 and 1.44 respectively. Companies with the lowest exposure to reference-day risk in their betas were Redefine, Sanlam and Investec with ranges of 0.18, 0.19 and 0.25 respectively. While there was a dispersion in beta ranges across days, the excess kurtosis of the distribution of ranges was lower than the average returns and the standard deviations at 1.4. The skewness corresponded with the average returns of 1.96. Out of the fifty-six companies, twenty-five (45%) had betas that ranged from below 1 (defensive) to above 1 (aggressive), and three (5%) had betas moving from negative to positive when the reference-day was changed. Out of the ten companies with the largest high-low beta ranges, there were no companies from the financial sector, with the ten being a mix of resource, property and retail companies. This outcome is counter to the findings of Sahadev et al. (2018), who found nine out of the top ten companies were resources companies and could be attributed to the wider sample size used in their study.

Since the statistical significance of reference-day risk in OLS betas on the JSE has been investigated extensively by Baker et al. (2016) and Sahadev et al. (2018), this paper aimed to confirm these findings for a more recent period under the hypotheses and test statistic of:

H_0: $\beta_{high} = \beta_{low}$

H_1: The highest beta across reference-days is significantly different from the lowest beta

$$\frac{\beta_{high} - \beta_{low}}{\sqrt{SE_{high}^2 + SE_{low}^2}},$$

where β_{high} is the highest beta across reference-days, β_{low} is the lowest beta across reference-days, SE_{high} is the standard error of the highest beta, and SE_{low} is the standard error of the lowest beta.

Our findings confirmed the presence of reference-day risk in OLS betas using the ALSI as the reference index, with fourteen (25%) of the fifty-six-companies exhibiting betas that were not statistically different across reference-days. Six (11%) of the betas were significantly different at the 5% level, fourteen betas (25%) at the 10% level and twenty-two (39%) at the 20% level. These results confirmed the extent of the differences in betas across companies, sectors and reference-days. The differences between betas have implications for the application of beta-style portfolio construction, estimating appropriate exposure to risk premiums and the CAPM (which applies beta to the equity risk premium).

To illustrate this, consider an investor who responds adversely to risk, seeking to construct a defensive portfolio that exhibits less sensitivity to market movements with a beta below 1. Assume that the investor estimated the beta of a stock using a reference-day that exhibited a lower beta, for example, AngloGold Ashanti had a beta of 0.76 on the 18th reference-day, MMI had a beta of 0.82 on the 1st reference-day, and Standard Bank had a beta of 0.83 on the 17th reference-day. The investor could be both underestimating beta and constructing a portfolio that is in effect the opposite of that intended at the offset. For instance, using the same three companies, and using the 12th, 14th and 2nd reference-days, yielded betas that moved from below to above 1 of 2.19, 1.34 and 1.12 respectively. This type of portfolio exhibits stronger fluctuations in its value when the market moves relative to choosing a reference-day exhibiting lower systematic risk leading to the portfolio deviating from the risk-return objectives and constraints outlined at initiation when the portfolio was constructed.

The same example could be extended to an investor wanting to construct a portfolio that moves inversely to the market achieved, by having a beta less than 0. When Intu properties was analysed on eight out of the twenty reference-days, it yielded negative betas with a range of -0.22 to 0.27. This example demonstrates how the choice of reference-day undermined the results of a portfolio that was constructed using metrics that calculated using one reference-day. These findings have implications for institutional strategies, such as market neutral funds, as well as academic studies that rely on a single reference-day to make inferences about relationships. Furthermore, it illustrates the need for a robust estimate of beta that is adjusted for the presence of reference-day risk. A full breakdown of beta ranges can be found in Appendix D.

Subsequently, this paper continues by focusing on reference-day risk in correlations. Correlations over the five-year period had an average range from their highest to lowest day of 0.24, a median range of 0.22 and a standard deviation of 0.09. Glencore, MTN and Growthpoint had the largest ranges of 0.46, 0.44 and 0.42 respectively. Companies with the lowest ranges were Naspers, Standard Bank and Old Mutual of 0.09, 0.12 and 0.13 respectively. The tendency for Naspers to have significantly lower reference-day risk in its' correlation than other companies is explained by the higher weight Naspers holds in the All Share Index (18%) as of September 2018. Of the fifty-six companies three (3%) correlations; Glencore, Intu and Truworths, changed signs from negative to positive from -0.02 to 0.44, -0.13 to 0.18 and -0.02 to 0.35 respectively. Furthermore, assuming a meaningful positive relationship between two variables is evidenced with a correlation coefficient greater than 0.5 as per Chapter 3. Eighteen of the fifty-six companies (32%) had correlations moving from below to above 0.5 when the reference-day was varied. To understand the relationship of correlations across reference-days we undertake two statistical tests. First, we test the statistical significance of the correlation on the twenty reference-days for each of fifty-six companies, using a t-test with the hypotheses and test statistic:

H_0: that the true correlation in the population is 0 (ρ = 0)

H_1: the correlation in the population is different from 0 (ρ ≠ 0)

$$t = \frac{r_{xy}\sqrt{n-2}}{\sqrt{1-r^2}},$$

where t is the t-test statistic, r_{xy} is the correlation between sample x and sample y, and n is the sample size. Secondly, we use a Fisher r-to-z transformation and a two-tailed z-test to investigate whether the highest correlation coefficient for the twenty reference-days is significantly different from the lowest correlation.

Across the fifty-six companies, on average seventeen of the twenty reference-days had correlations that were significant at the 5% level, and eighteen at the 10% significance level. The median number of reference-days with significant correlations was nineteen at the 5% level and twenty at the 10% level. Out of the fifty-six companies twenty-six (47%) had all their correlations significant at both the 5% and 10% level across reference-days; while one company – Intu Properties, had no significant correlations at either the 5% or 10% level. The property and resource sectors had the fewest significant correlations at the 5% level, with an average of thirteen and fifteen significant correlations across reference-days respectively. Financials and telecommunications had the highest number of significant correlations, with averages of twenty and nineteen days respectively. There was no consistency of statistical significance across reference-days. Where, for example, Pick 'n Pay had fifteen correlations significantly different from zero across reference-days; while Shoprite had twelve significant correlations across reference-days with the 1st, 3rd, 4th, 5th, 6th, 7th, 8th and 20th days not significantly different from zero. Another example was Telkom, which across the twenty reference-days had correlations on the 17th and 19th day that were not significantly different from zero at the 5% level. These results show that the significance of individual correlations varies across reference-days, where the significance is not consistent across companies or sectors.

We used a Fisher r-to-z transformation and a two-tailed z-test to investigate if the highest correlation coefficient for the twenty reference-days were significantly different to the lowest correlation, using the hypotheses and formula:

H_0: Max r_{xy} = Min r_{xy}

H_1: Max $r_{xy} \neq$ Min r_{xy}

$$Z_{xy} = \frac{0.5\ln\left(\frac{1+r_x}{1-r_x}\right) - 0.5ln\left(\frac{1+r_y}{1-r_y}\right)}{\sqrt{\left(\frac{1}{N_x-3}\right)+\left(\frac{1}{N_y-3}\right)}} ,$$

where Z_{xy} is the z test statistic, r_x is the sample x correlation, r_y is the sample y correlation, N_x is the size of sample x and N_y is the size of sample y.

Despite the high ranges observed in the returns data for correlations, no stock's highest correlation was significantly different from the lowest correlation across reference-days at the 10% level. However, the companies with the lowest p-values were MTN, Glencore and Growthpoint, with p-values of 0.1211, 0.1499 and 0.1738 respectively. These three companies happened to have the largest high-low ranges across reference-days for correlations.

These findings have implications for understanding the interrelationship between assets and the statistical significance of such relationships. Since these relationships are a fundamental input into the construction of efficient portfolios, they are of particular importance to this paper. For example, an investor looking to add diversifying exposure to a long index portfolio and using the 20th reference-day of the month would find Glencore to exhibit a negative correlation of -0.02. The result of including Glencore in the portfolio would be counterintuitive, as across the other nineteen reference-days the correlation between the index and Glencore was positive and could be as high as 0.44 if the 14th reference-day was used. Thus, on average, the investor's portfolio would move more in-line with the market movements than initially thought. The same logic applies to Truworths, where the correlation ranged from -0.02 to 0.35. Even if the correlations did not vary from positive to negative, they had implications for both risk management and portfolio construction. Suppose a portfolio manager has a mandate which allows for stocks with a certain correlation threshold to be added to the portfolio, for example, stocks with a correlation of 0.3 or less to the market. When the 6th reference-day was used, the Pioneer Food Group had a correlation of 0.42 against the market and would be excluded from this process. However, on the 1st reference-day, the company had a correlation of 0.25 and thus would be eligible for inclusion. These examples highlight the need for a correlation that is not impacted by the choice of reference-day to base sound investment, risk management and portfolio construction decisions on. A full breakdown of correlation ranges is found in Appendix E.

4.1.3. Interrelationships of average returns, standard deviations, correlations and betas in the presence of reference-day risk

This paper proceeds by addressing the interrelationships between the different metrics exposure to reference-day risk. Below are regressions comparing the maximum-minimum ranges for each metric relative to each other. They are compared in order to assess the dependency of reference-day risk across metrics.

Average returns and standard deviations exhibited the closest relationship in high-low ranges, with 41% of the variation in standard deviation ranges being explained by the variation in average return ranges. Average returns and correlations exhibited the weakest relationship with an R^2 of 0%. This implied that companies which displayed a high range in average returns across reference-days, would also exhibit a high range in standard deviations, but not with correlations. The relationship between average returns and betas were between the two, with an R^2 of 28%.

Overall these regressions (charted below) suggested that the exposure of a company's statistical properties to reference-day risk is metric specific. Furthermore, this showed that having high reference-day risk in one metric did not result in having high reference-day risk in others. For example, while average returns and standard deviation ranges across references days had a relatively strong relationship, the relationship between standard deviation and correlation ranges only had an R^2 of 2%. Furthermore, the relationship between beta and correlation ranges was strong relative to others, with an R^2 of 25%. This result is explained because of the inclusion of correlations in the calculation of beta. In conclusion, there were varying degrees of explanatory power between the ranges of different metrics across reference-days. Overall there was not a strong explanatory relationship in the degree of reference-day risk for one metric relative to another. Therefore, a method for estimating reference-day risk-free versions of each metric needs to be developed to ensure a robust toolbox of statistical measures for practical application. The method, as proposed in this paper, must be able to be consistently applied across metrics.

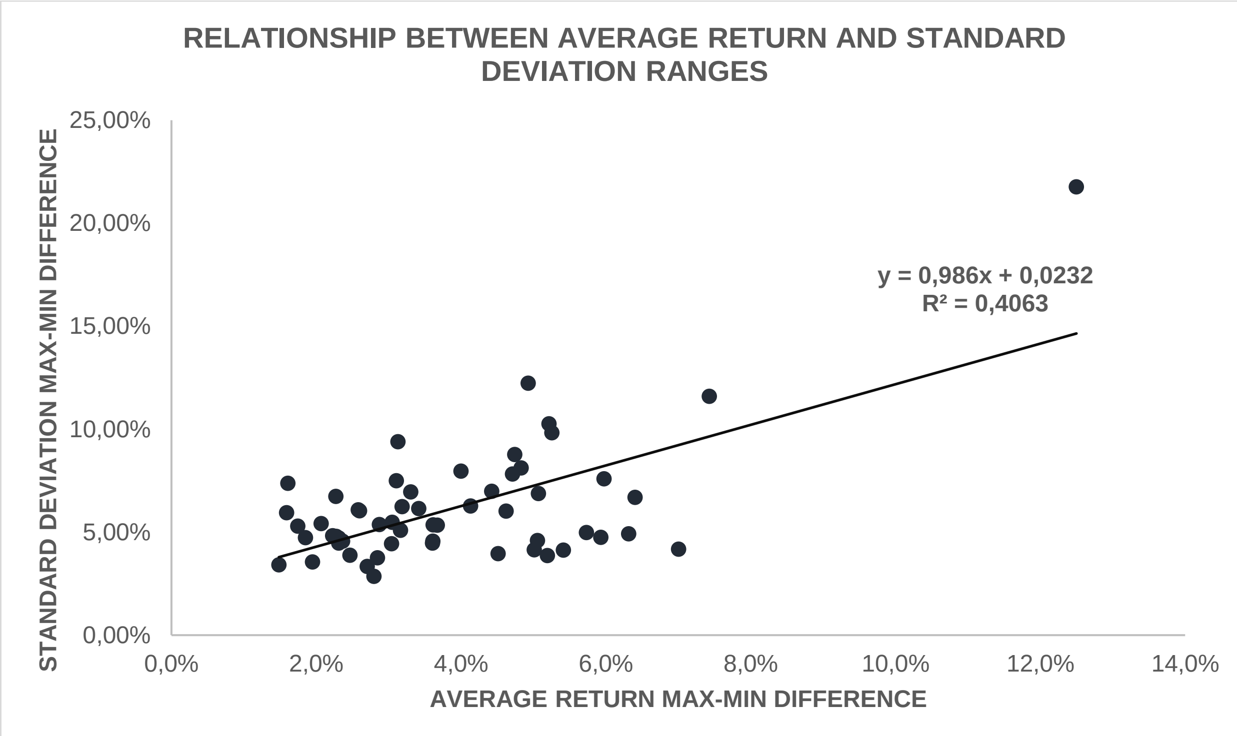

Figure 1: The OLS regression of average return and standard deviation ranges across reference-days.

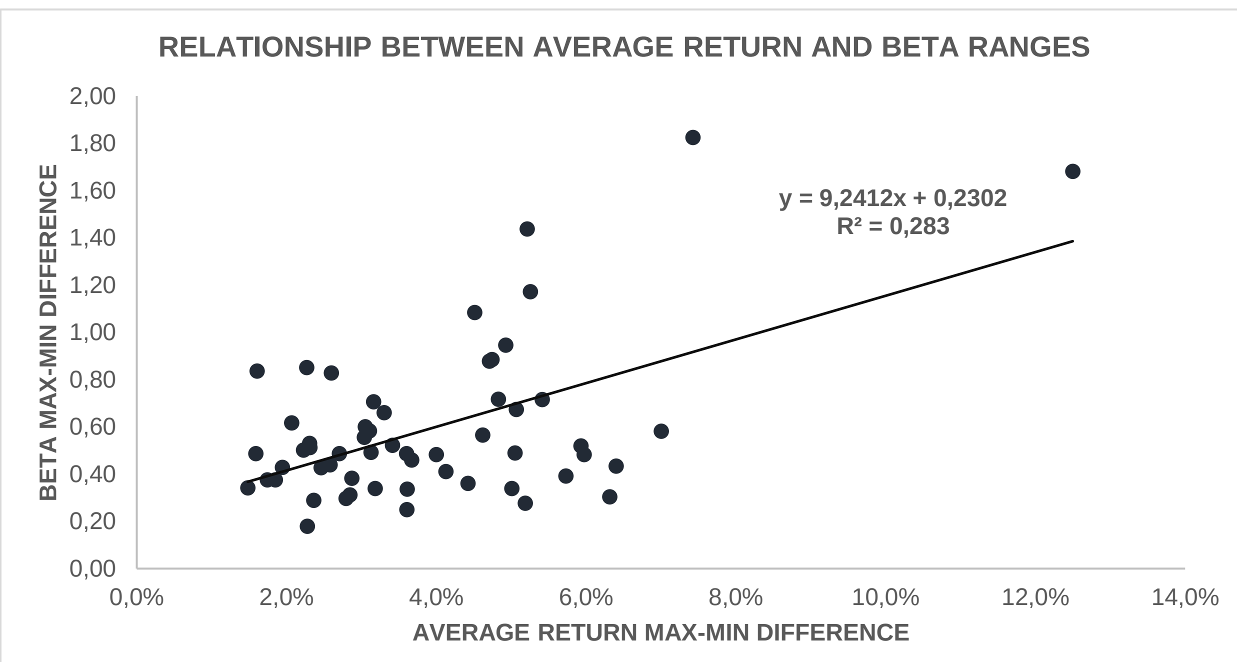

Figure 2: The OLS regression of average return and beta ranges across reference-days.

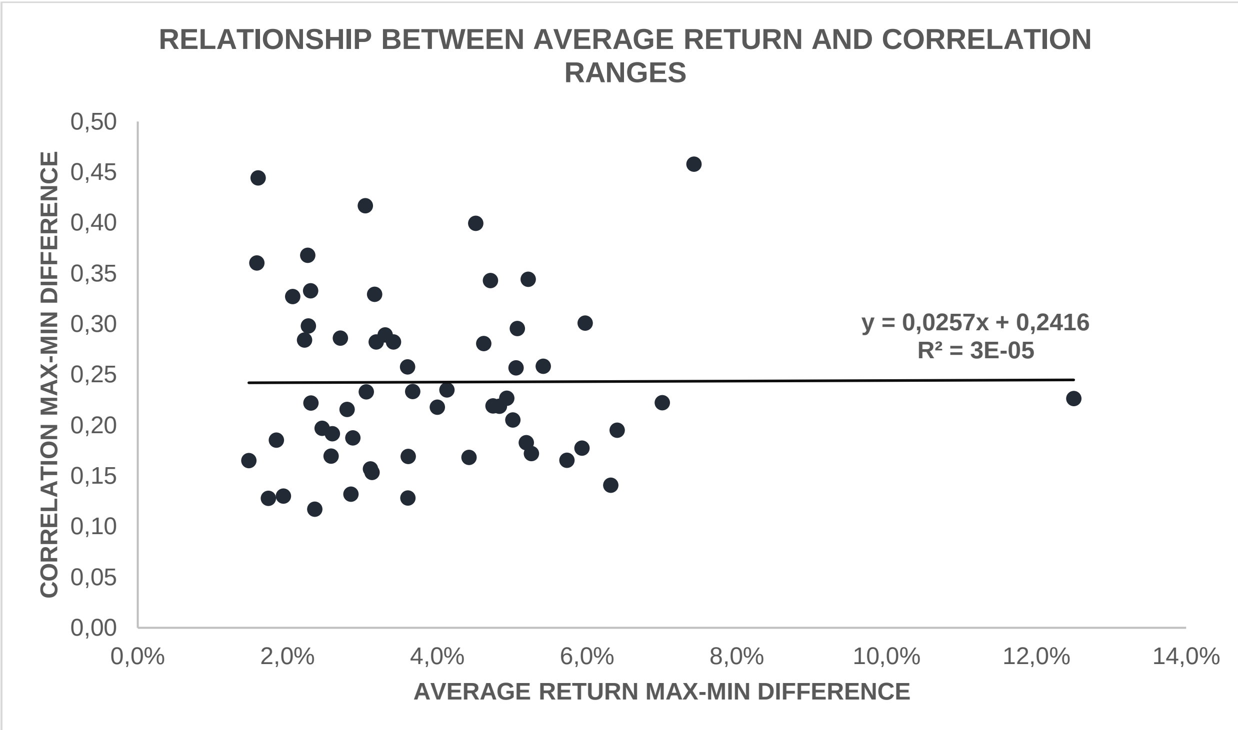

Figure 3: The OLS regression of average return and correlation ranges across reference-days.

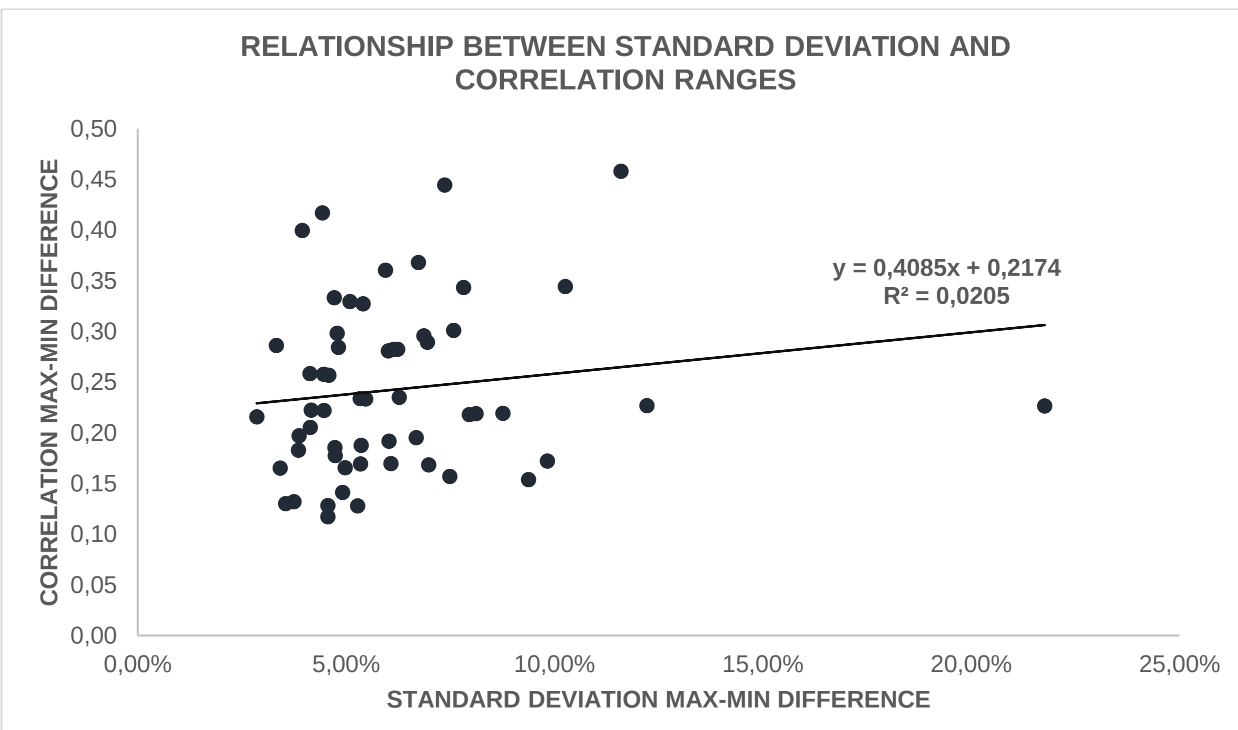

Figure 4: The OLS regression of standard deviation and correlation ranges across reference-days.

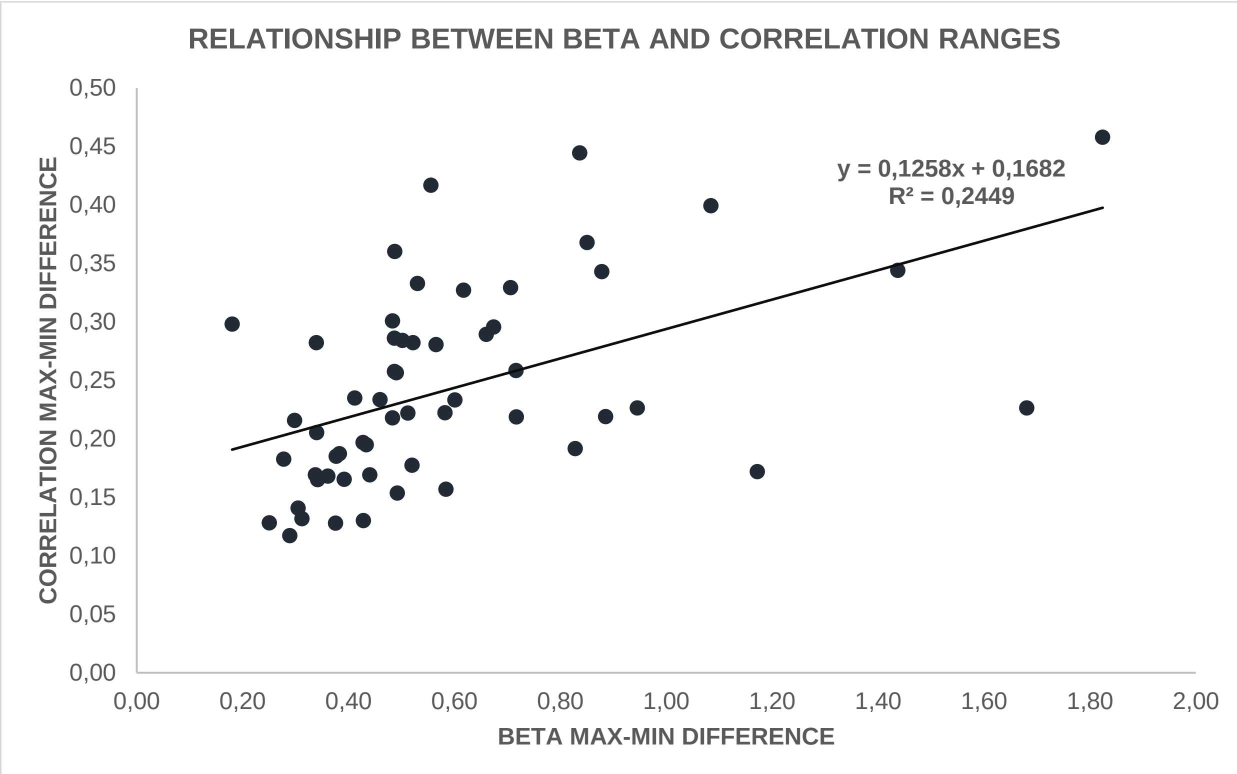

Figure 5: The OLS regression of beta and correlation ranges across reference-day.

4.2. Investigation into using a bootstrapping method to estimate average returns, standard deviations, correlations and betas in the presence of reference-day risk

While the primary goal of this paper is to apply reference-day risk-free metrics to the construction of efficient portfolios, this application begins with the accurate estimation of model inputs across reference-days. A technique to accomplish this, recommended by Baker et al. (2016), was to average the values across reference-days. In this section, we investigate the use of a bootstrapped distribution method to estimate average returns, standard deviations, correlations and betas that are independent of reference-day risk. First, we use simulated data to test the robustness and accuracy of the method. We then apply the process to the fifty-six-company sample in order to estimate average returns, standard deviations, correlations and betas independent of reference-day risk. As per Baker et al. (2016), we use the expected value of this distribution as estimate of the real underlying metric, and the dispersion of the distribution as a measure of the degree of reference-day risk. The standard deviation of the bootstrapped distribution along with the high-low range of estimates across reference-days, serve as two measures of reference-day risk for a given metric. We assume each metric is pulled from an underlying unobservable distribution. This paper aims to estimate the reference-day risk-free metric as the expected value of this underlying distribution for a period of time, as per Chapter 3.

4.2.1. Data simulation and evidence in simulated data

Through using Geometric Brownian motion, we simulated stock returns data for sixty months (1200 trading days). The simulation was started from a predefined base price, and subsequent price movements were determined from the assumption that the price the following day is a function of the previous days using the formula:

$$S_t = S_{t-1} e^{\left(\mu - \frac{\sigma}{2}\right) \times dt + \sqrt{dt} \times W(t)},$$

where S_t is today's price, e is the exponential variable, μ is the assets annualised average return, σ is the annualised volatility, dt is the change in time, and W(t) is the stochastic random component.

To imbed the relationship between the prices of the stock and the market index, we adjust the simulation by imposing a Cholesky Decomposition to the random error terms. Furthermore, we predefine the volatility levels for both the stock and the index returns for simulation.

We repeat this simulation process for average returns, standard deviations and correlations where lognormal returns for both the index and hypothetical stock were calculated. The same Cholesky Decomposition matrix is then used to modify the returns of the index into a sequence correlated with the hypothetical stock returns. We then convert these returns into a sixty by twenty matrix for easier computational analysis and in accordance with the format used for empirical data.

A random reference-day, from one to twenty, is then chosen from each of the sixty months, for example, the 3rd of the first month, the 18th of the 2nd month, up to the 60th month. The average return, standard deviation, beta and correlation to the simulated index is calculated using these sixty returns. This process is repeated 100 000 times with replacement to generate a bootstrapped distribution of each metric (Baker et al., 2016). The point estimate of this distribution is assumed to be the true underlying value that is independent of reference-day risk for the period under analysis. This simulated data allowed us to test the underlying method by assuming values for average returns, standard deviations, correlations and betas over various intervals. Following this, the nonparametric bootstrapping method is applied to test if the point estimate of the bootstrapped distribution equates to the predefined value for each metric. However, we first tested if the simulated data exhibits reference-day risk across the four metrics.

After simulating the data in accordance with the above method, we calculated the average return, standard deviation, correlation and beta across the twenty reference-days for the hypothetical company and index return series. Consistent with the empirical observations in section 4.1, we found varying degrees of reference-day risk across the four metrics using repeated samples of simulated data.

Through simulation, we created daily price data whose returns were normally distributed. A breakdown of predefined values and the resulting distributions are found in Appendix F, G and H. For example, we varied betas from negative five to positive five, and correlations from negative one to positive one. We found that for all the metrics the distribution centred around the predefined value with varying degrees of dispersion depending on the sample. In the context of reference-day risk, the process resulted in the best estimate being found.

The results indicated that as the predefined values tended towards zero for betas, correlations and standard deviations, the bootstrapped distributions became more normal and exhibited a higher standard deviation. When the values were closer to extremes on either side, for example, 5 and -5 for betas, and -0.75 and 0.75 for correlations, the distributions were more peaked around the mean and displayed higher levels of skewness and kurtosis (Appendix F, G and H). Although the higher order moments varied as the predefined values changed, the mean – which is the focus of the method – remained almost exactly centred on the predefined value for all the metrics. However, average returns did not exhibit the same behaviour as the other three metrics, with fairly consistent measures of dispersion and normality as the predefined value was varied. These results were evidence of the robustness of the method for both the estimation of reference-day risk through the dispersion of the bootstrapped distribution, as well as in estimating reference-day risk-free metrics from the mean of the distribution. This method is applied to empirical price data from the JSE, in order to estimate reference-day risk-free metrics, measure the degree of reference-day risk, and perform a comparative analysis, to better understand the nature and behaviour of reference-day risk in the South African context.

4.3. Application of bootstrapped average returns, standard deviations, correlations and betas

A bootstrapped distribution was generated for each of the fifty-six company's underlying metrics – average returns, standard deviations, correlations and betas. For the four metrics, the bootstrapped distributions were approximately normally distributed, and the estimate was equal to the mean value across the twenty reference-days. To confirm these findings, we regressed the average value across reference-days against the bootstrapped values for the fifty-six companies. Across all of the metrics, there was an almost perfect linear relationship with an R^2 of 99% for each instance. This relationship implied that the average value across reference-days was a good proxy for the reference-day risk-free estimates but lacked the statistical robustness and testing of the full bootstrapping method.

From the data, it can be inferred that wider bootstrapped distributions - as measured by the standard deviation of the bootstrapped distribution - exhibited a higher degree of reference-day risk for a given metric. For example, Kumba Iron ore had the highest standard deviation of bootstrapped average returns of 1.42%, while Fortress A had the lowest at 0.36%. These results implied that Kumba had a higher degree of reference-day risk in average returns than Fortress A. The company which had the largest standard deviation of its bootstrap distribution, did not always have the highest range across the twenty reference-days. For instance, Capitec had a range of 7% in average returns across days, but a standard deviation of 0.51% in the bootstrapped distribution. In contrast, AngloGold Ashanti had a smaller range of 5.2% across days but had a dispersion over two times higher, of 1.14%. To further this point, we regressed the standard deviation of the bootstrapped distribution for each company against the corresponding high-low range for each metric across reference-days. This was to evaluate the extent to which the high-low range could be used as a proxy to measure reference-day risk. Average returns resulted in a $R^2$of 34%, standard deviations an R^2 of 79%, betas an R^2 of 0% and correlations an R^2 of 46%. Moreover, the results demonstrated that the relationship was not consistent across metrics and that each metric needed to be evaluated individually, when testing and adjusting for reference-day risk as per section 4.1. These results highlighted that although the range provided a good proxy of the degree of reference-day risk, a more robust estimate was the standard deviation of the bootstrapped distribution due to a larger sample size and more empirical support of the method.

The expected value of the distribution was equal to or within half a standard deviation of the average across reference-days. These findings were consistent across all the metrics and is evident from the data provided in appendices B, C, D and E. In addition, this paper provides the bootstrapped distributions of the highest and lowest ranges for each metric. It highlights the distribution mean and standard deviation as measures of the reference-day risk-free estimate and the degree of reference-day risk respectively. Figure 6 shows the bootstrapped average return distribution of Kumba Iron Ore and Richemont, who had the largest and smallest ranges in average returns respectively.

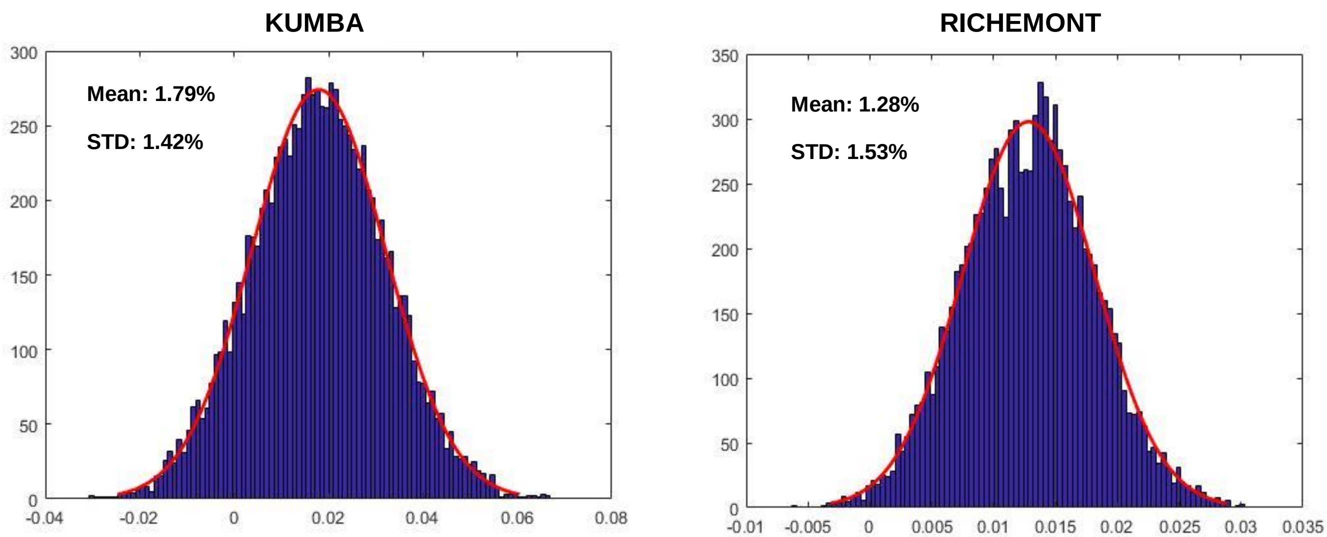

Figure 6: The bootstrapped average return distributions for the companies with the highest and lowest ranges of unadjusted average returns. Where STD represents the standard deviation of the distribution.

Figure 7 shows the bootstrapped standard deviations of Kumba Iron Ore and AVI, who had the highest and lowest ranges in standard deviations respectively.

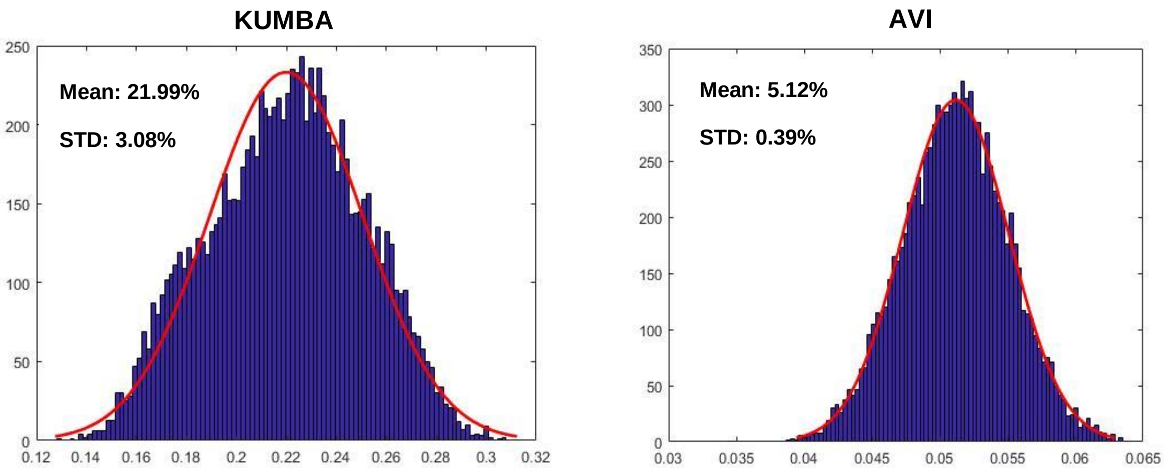

Figure 7: The bootstrapped standard deviation distributions for the companies with the highest and lowest ranges of unadjusted standard deviations. Where STD represents the standard deviation of the distribution.

Figure 8 shows the bootstrapped correlations of Glencore and Standard Bank, who had the highest and lowest ranges in correlations respectively.

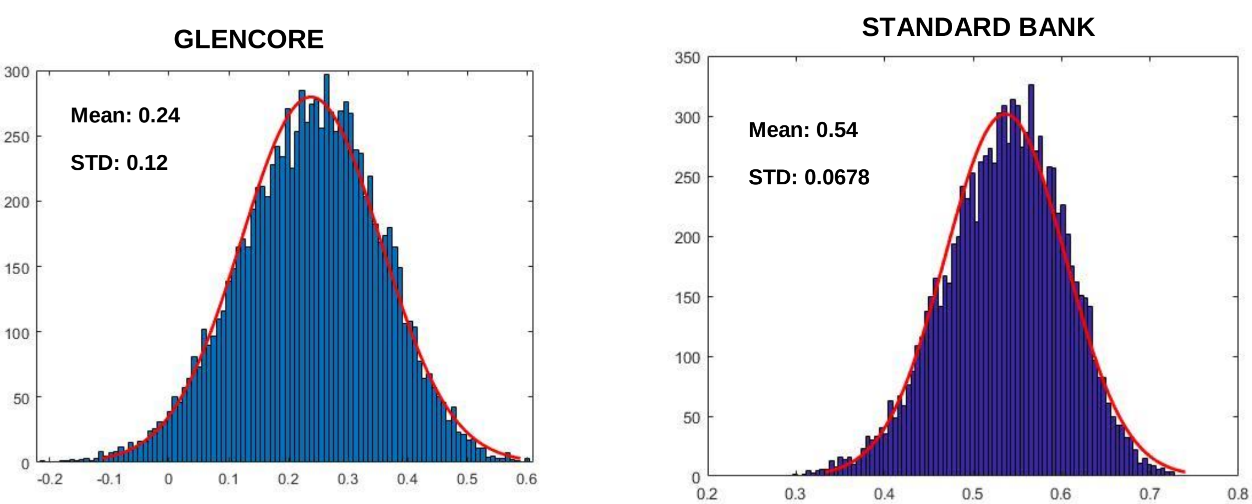

Figure 8: The bootstrapped correlation distributions for the companies with the highest and lowest ranges of unadjusted correlations. Where STD represents the standard deviation of the distribution.

Figure 9 shows the bootstrapped betas of Glencore and Redefine, who had the highest and lowest ranges in betas respectively.

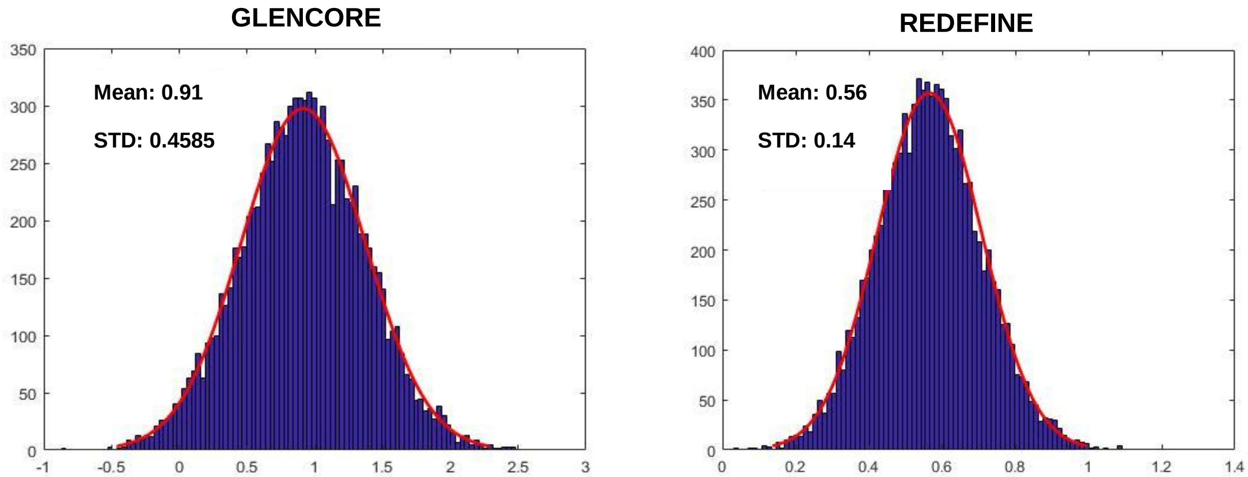

Figure 9: The bootstrapped distributions for the companies with the highest and lowest ranges. Where STD represents the standard deviation of the distribution.

To analyse the extent to which the reference-day risk-free bootstrapped betas were different from the unadjusted end of month betas from Bloomberg, we compared the differences across the fifty-six companies. There were two companies, Woolworths and Netcare that had bootstrapped betas that were the same as their Bloomberg comparable. Furthermore, sixteen companies (27%) had betas that were larger than the Bloomberg betas, and forty (71%) with betas smaller. Glencore and Intu had the largest differences with bootstrapped betas, of 0.91 and 0.04 compared to the Bloomberg betas of 1.83 and 0.77 respectively. The average difference between the bootstrapped and Bloomberg betas was 0.1, the median difference 0.12, and a standard deviation of 0.27. Lastly, we regressed the bootstrapped betas against the Bloomberg comparable. The relationship was weak considering the wide use of these Bloomberg betas in practice, with an R^2 of 61% (Figure 10). These results showed that there is room for improvement in the calculation of betas that are pulled from such platforms, to account for phenomena such as reference-day risk. The result of not making such adjustments could cause misestimation of systematic risk and inaccurate inferences to be drawn.

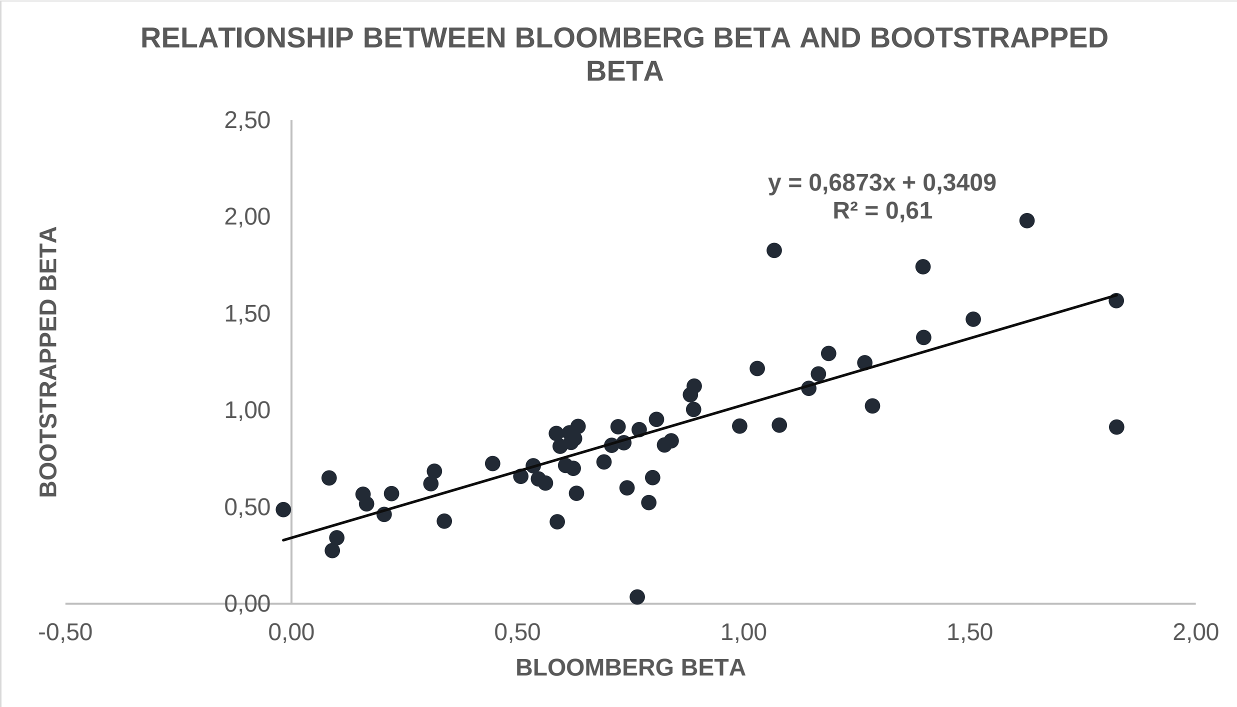

Figure 10: A regression of Bloomberg OLS betas against bootstrapped betas.

This process was repeated with correlations, which contrasted our bootstrapped values with those pulled from Bloomberg for the comparable period. None of the Bloomberg correlations were the same as the bootstrapped values. Forty-eight companies (86%) had Bloomberg correlations that were lower than the bootstrapped alternatives, and eight companies (14%) had a higher Bloomberg correlation. Intu had the largest difference between the bootstrapped correlation (0.02) and the Bloomberg value (0.33), of 0.31 for the companies where the Bloomberg correlation was greater than the bootstrapped comparable. Old Mutual had the largest difference out of the companies, where the bootstrapped correlation was higher than Bloomberg's of 0.76. The average difference between the two estimates across the fifty-six companies was -0.1, the median -0.09 and the standard deviation 0.14. Lastly, we regressed the bootstrapped correlations against the Bloomberg comparable. At 40% the R^2 was slightly lower than the equivalent beta regression. As discussed in section 4.1.2, this has important implications for the objective of this paper, and it continues to highlight the need to find a correlation that is adjusted for reference-day risk. Furthermore, researchers should not primarily trust correlations that are published by data providers, as they could lead to diversification effects that are dependent more on the choice of reference-day than the underlying relationship.

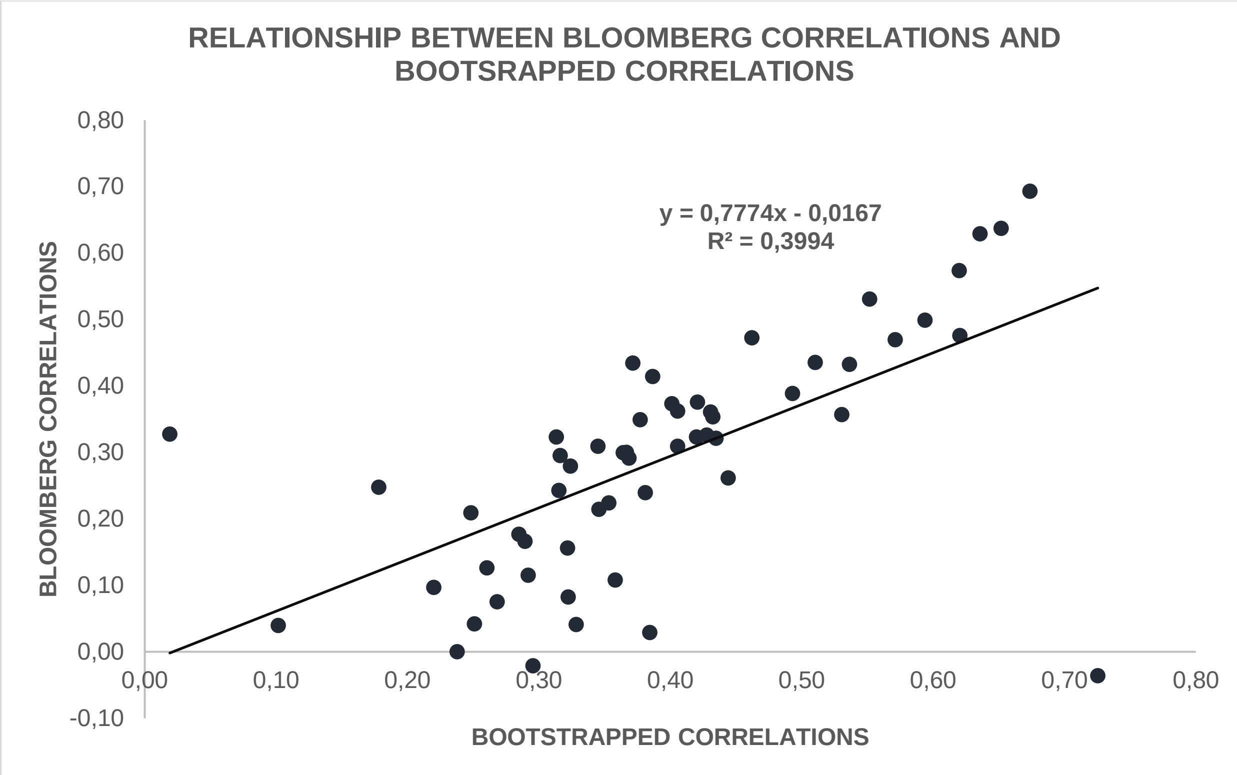

Figure 11: A regression of Bloomberg correlations against bootstrapped correlations.

These findings emphasised the importance of adjusting for reference-day risk when calculating inputs for the modelling of optimised portfolios. Without these adjustments, metrics are exposed to being biased to the choice of initial reference-day and could cause the incorrect relationship between risk and return to be assumed. Furthermore, it would have a notable effect on the measurement of inter-asset relationships based on correlations and betas. A summary table of the simulated point estimates for average returns, standard deviations, correlations and betas for the fifty-six companies as well as relevant statistics is found in Appendix B, C, D and E respectively.

Chapter 5: A case study on the reference-day risk-free portfolio

This Chapter extends the findings of section 4.3 to broad asset classes. Initially, it is shown that reference-day risk exists in broadly diversified indices, which represent these respective asset classes. Following this, we will demonstrate that the choice of reference-day can significantly influence the optimal maximum Sharpe portfolio. Lastly, the optimal maximum Sharpe portfolio is constructed, which is independent of reference-day risk. This is achieved by calculating the reference-day risk-free model inputs as per section 4.2.1. We then compare this portfolio to the unadjusted alternative using the end of month data that was sourced from Bloomberg, which represents the portfolio exposed to reference-day risk. Equities are represented by the JSE All Share Index; international equities by the MSCI World Index; Commodities by the Bloomberg All Commodity Index; Property by the JSE property index; and fixed income by the JSE All Bond Index, over the period of the 1st January 2013 to 31st December 2017 as per Chapter 3. This period is referred to as the in-sample period. Due to data availability, the out-of-sample period runs from the 1st January 2018 to the 31st December 2018.

5.1. The application of reference-day risk-free metrics to constructing efficient portfolios

As with simulated data and individual companies, there is evidence of reference-day risk across all broad asset classes for average returns, standard deviations and correlations. A summary of high-low ranges for average returns and standard deviations are provided in the table below. It is evident that the degree of reference-day risk differs across asset classes, with locally focused asset classes, such as local equities and property, having the highest sensitivity to the choice of reference-day.

Asset Class	Average Return High-Low Range	Standard Deviation High-Low Range
Local Equities	1.33%	3.02%*

International Equities	0.54%	1.98%***
Fixed Income	0.82%	1.69%**
Commodities	1.32%	1.38%
Property	1.9%	4.64%*

Table 1: The difference between the maximum and minimum average return and standard deviation across reference-days for the major asset classes. * is significant at the 5% level, ** at the 10% level and *** at the 20% level.

These ranges were relatively lower than the single company ranges we found in section 4.1. This could be due to the diversifying effect of using an index as addressed by Fama and French (1993), who found a smaller degree of reference-day risk in portfolios relative to single stocks. Additionally, this links to the empirical literature on risk, which states that through combining uncorrelated assets the volatility of the portfolio can be reduced. These results indicated that this principle applies to the reduction of reference-day risk and total risk, as measured by standard deviation. However, constructing portfolios did not mitigate or remove the problems caused by reference-day risk. Therefore, a method to both measure and adjust for its effects is useful for investment practitioners looking to use the mean-variance framework in portfolio construction.

As with individual companies, we tested for statistical significance in the difference between the highest and the lowest average return and standard deviation for each asset class. We did these tests across reference-days with the use of a two-sided t-test and an F-test respectively.

The results indicated no asset class had their highest average return being significantly different from their lowest average return at the 20% level. Property had the lowest p-value of 0.49 whereas international equities had the highest at 0.76. This was based on the null hypothesis of equality of the highest and lowest average returns across reference-days. Moreover, property had the most significant difference in its standard deviations (p-value of 0.015) that were significant at the 5% level, followed by local equities (p-value of 0.04). Fixed income’s difference in standard deviations across reference-days was significant at the 10% level (p-value of 0.09); while the international equities difference was significantly different at the 20% level (p-value of 0.125). Commodities standard deviations were not significantly different (p-value of 0.38).

Correlations between the asset classes exhibited large degrees of reference-day risk; with a full breakdown of cross-correlations between asset classes provided in Appendix I and illustrated in Figure 12 and 13 below. These figures highlight the large ranges of correlations between different asset classes depending on the choice of reference-day.

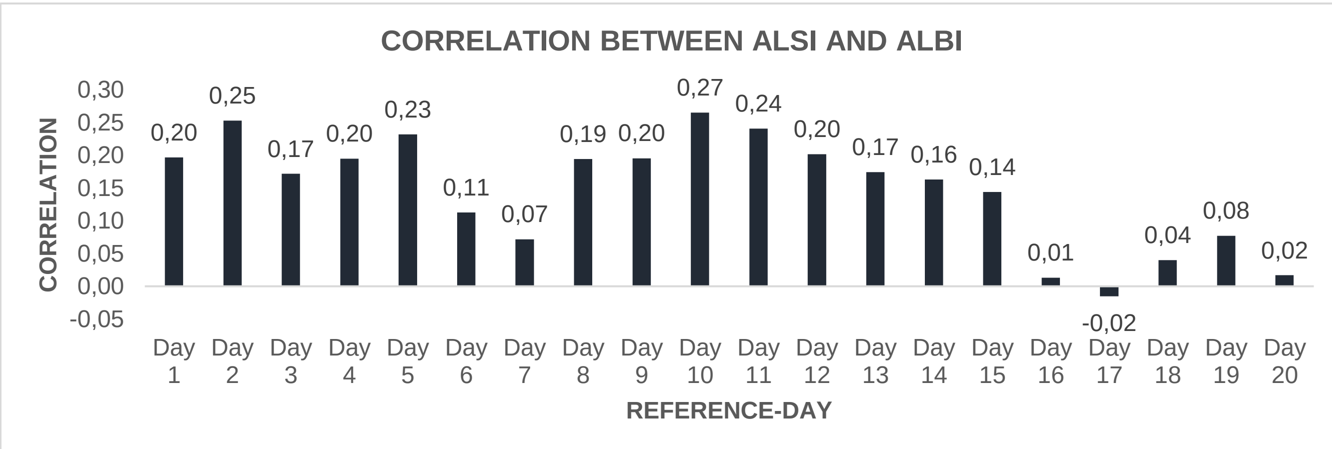

Figure 12: The correlations across twenty reference-days between equities (ALSI) and bonds (ALBI).

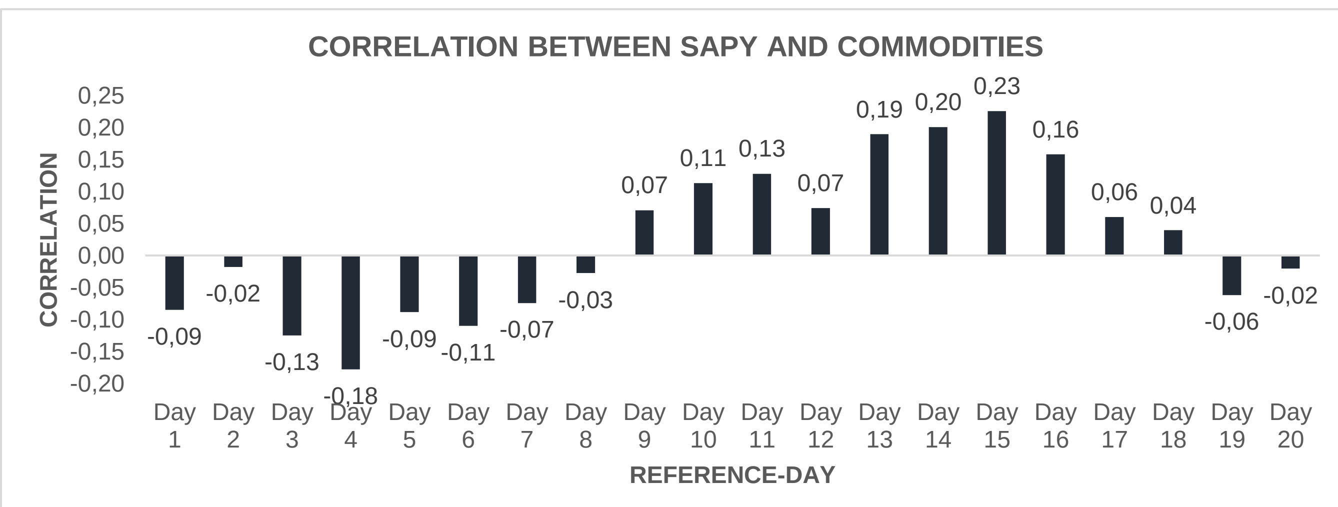

Figure 13: The correlations across twenty reference-days between property (SAPY) and commodities (Bloomberg commodity index).

The charts show that correlations are dynamic and time-varying, depending on when the two related series are analysed. For example, the correlation between property and commodities in Figure 13 was negative for the start of the month; became positive during the middle; and then became negative again towards the end. Anyone wanting to build a portfolio of negatively correlated assets, using any of the days from the 9th to the 18th day would have assumed a positive relationship between the two asset classes. While in fact, half of the reference-days resulted in a negative correlation. Even if the correlations did not vary from positive to negative, there were large variations in the degree of positive correlation. For example, international equities and commodities reached a correlation of 0.55 on the 15th and 16th reference-day and 0.19 on the 4th day. Therefore, it is important to use an estimate of correlation that is not dependent on the choice of reference-day to ensure that the most accurate inferences are drawn regarding the relationship between two investments and the diversification implications thereof.

These findings provide evidence of the existence of reference-day risk across asset classes. We continue by addressing the implications that this variation in estimates has on the implementation of the mean-variance framework of Markowitz (1952). In accordance with the empirical literature, the optimal portfolio is defined as the portfolio which achieves the highest excess return above a risk-free alternative per unit of risk, referred to as the Sharpe ratio.

We begin by calculating the optimal portfolio for each of the twenty reference-days using the average returns, standard deviations and correlations of each asset class for that specific reference-day. Below, Figure 14 highlights the wide range of portfolios which could be considered optimal for the five-year period based on the initial choice of reference-day and the Markowitz (1952) mean-variance framework.

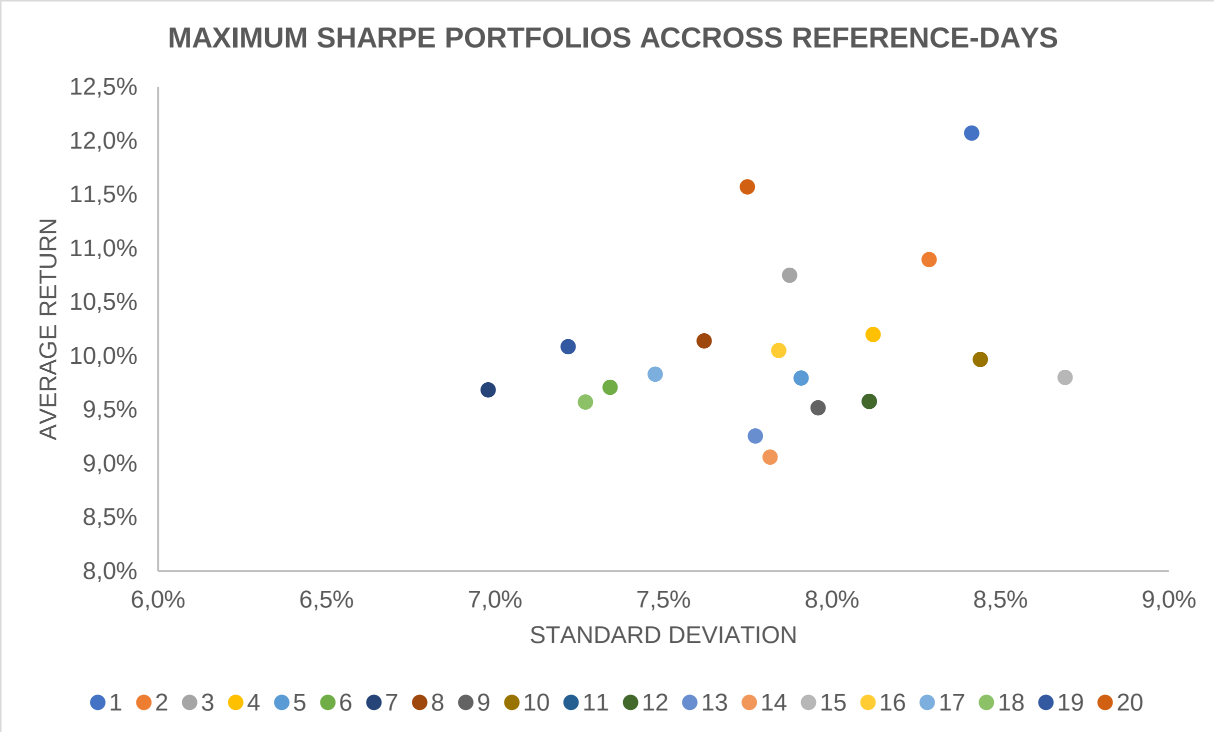

Figure 14: The maximum Sharpe portfolios using the mean, standard deviation and correlation of each reference-day.

Not only did the average return across portfolios vary from 9.1% on the 14th reference-day to 12.1% on the 1st reference-day, but the volatility of these returns ranged from 6.98% on the 7th reference-day to 8.7% on the 15th day. These results have implications for both optimising and selecting the most efficient portfolio based on historical data. For example, the portfolio from the 20th reference-day achieved an average return of 12% with a volatility of 7.7%; while the 13th reference-day achieved an average return of 9.3% for the same level of volatility as the 20th day. A rational investor would prefer the portfolio from the 20th day as it yields a higher average return for the same level of risk. Although by using only one reference-day they would be unaware of more efficient alternatives, or know which portfolio was truly optimal. From an alternative perspective, the portfolio optimised on the 15th day yielded an average return of 10% with a volatility of 8.7%; while the 7th day yielded the same average return but with a 6.78% volatility. Therefore, a rational investor would prefer the portfolio calculated on the 7th day to the 15th day. However, if they used the end of month data, the investor would be unaware of the large divergences in model inputs and subsequent portfolios due to reference-day risk.

In conclusion, the portfolio optimised using the inputs from the 20th, and 7th reference-days are preferred as they yielded a higher or the same average return for the same or lower level of volatility. Both these portfolios were optimised over the same period. Thus, there is a need to identify the model inputs that are independent of reference-day risk. to ensure that the most efficient portfolio is constructed based on the inputs that represent the most accurate underlying risk-return dynamics for each asset.

To emphasise the divergence across reference-days we track the performance of each reference-day's optimal portfolio over the in-sample and out-of-sample periods in Figure 15 and 16 below. Graphical inspection showed that the portfolios moved together over time in-sample, with an average correlation between portfolios of 0.97, a median of 0.97 and a standard deviation of 0.025. The performance ranged from a 53% gain to a 75% gain depending on the choice of initial reference-day and assuming monthly rebalancing. The average return for portfolios across reference-days for the full five-year period was 60%, the median 59% and the standard deviation was 5.37%. However, this has limited practical implications for investors as the optimised weights are only known ex-post. Subsequently, we assess the out-of-sample returns based on the in-sample optimisation.

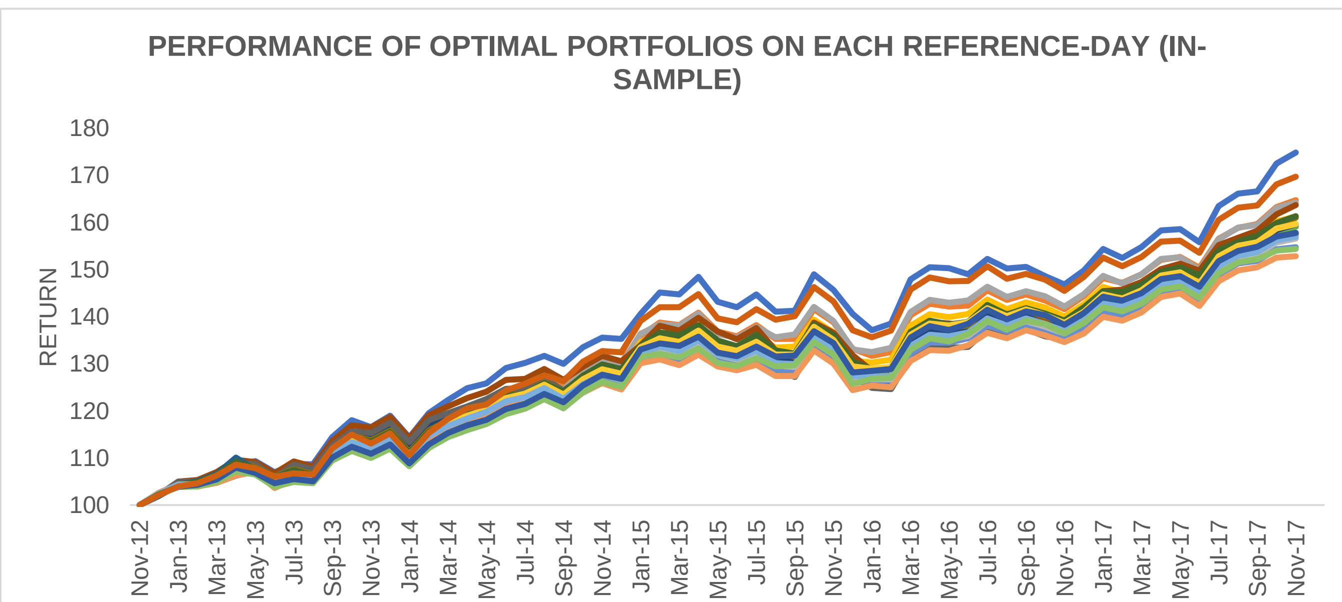

Figure 15: The performance of portfolio's optimised on each of the twenty reference-days during the in-sample period.

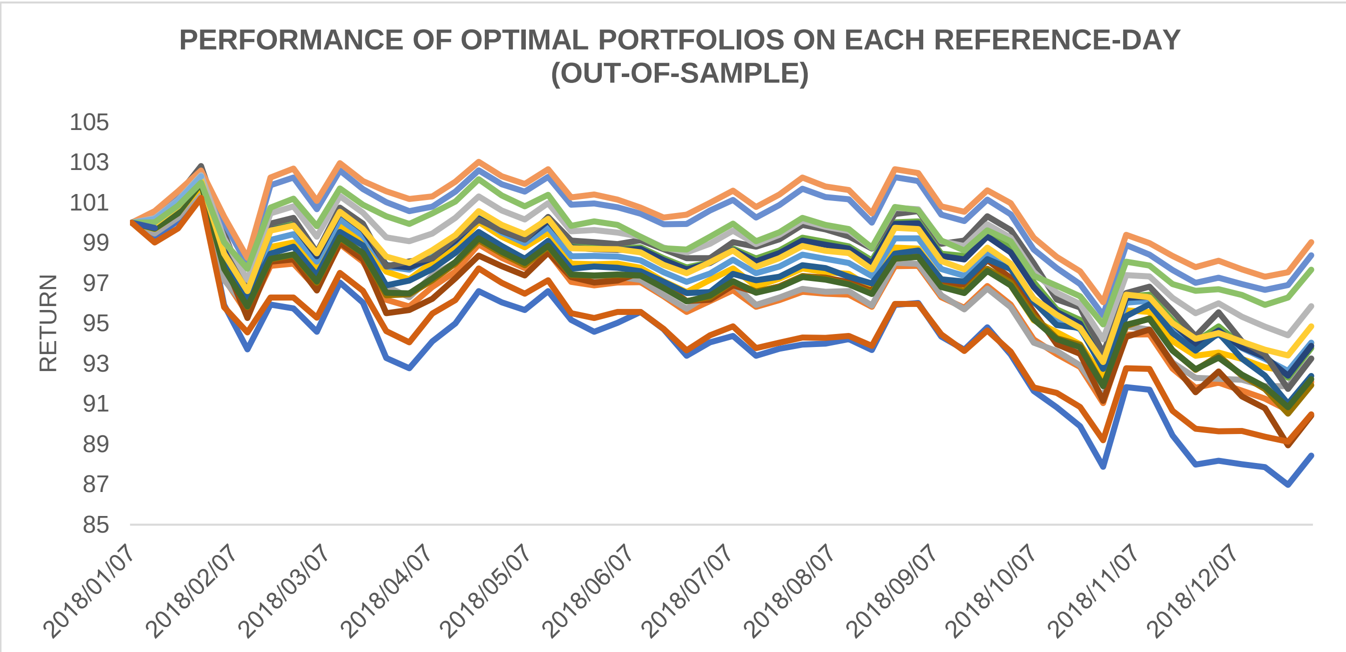

Figure 16: The performance of the portfolio's optimised on each of the twenty reference-days during the out-of-sample period.

The graphical inspection of Figure 16 showed that the portfolios tended to move together over time with an average correlation of 0.96, a median correlation of 0.97 and a standard deviation of correlations of 0.04 across reference-days. The performances ranged from -11.6% on the 1st reference-day to -1% on the 14th day. The average return for portfolios across the out-of-sample period was -6.1%, with a standard deviation of 2.71% across days. The results showed that the best performing reference-day portfolio in-sample was the worst performing in the out-of-sample period. Therefore, investment practitioners who optimise their portfolio using a single reference-day without adjusting model inputs to account for reference-day risk, are susceptible to large deviations in actual performance relative to expectations. We address these issues by estimating the optimal reference-day risk-free portfolio. This is the portfolio that represents the best estimate of return, risk and correlation after accounting for the influence of reference-day risk.

5.2. Estimating the reference-day risk-free optimal portfolio

This paper has provided evidence of the existence of reference-day risk across key risk-return metrics for individual companies and asset classes. We continue by combining both sections in order to estimate the optimal portfolio independent of reference-day risk. To begin this process, we repeat the method that was used for individual companies in section 4.3 and estimate the reference-day risk-free estimates of average returns, standard deviations and correlations across the previously discussed asset classes. Provided below in Figure 17 are four charts showing the distributional properties of these metrics for selected asset classes.

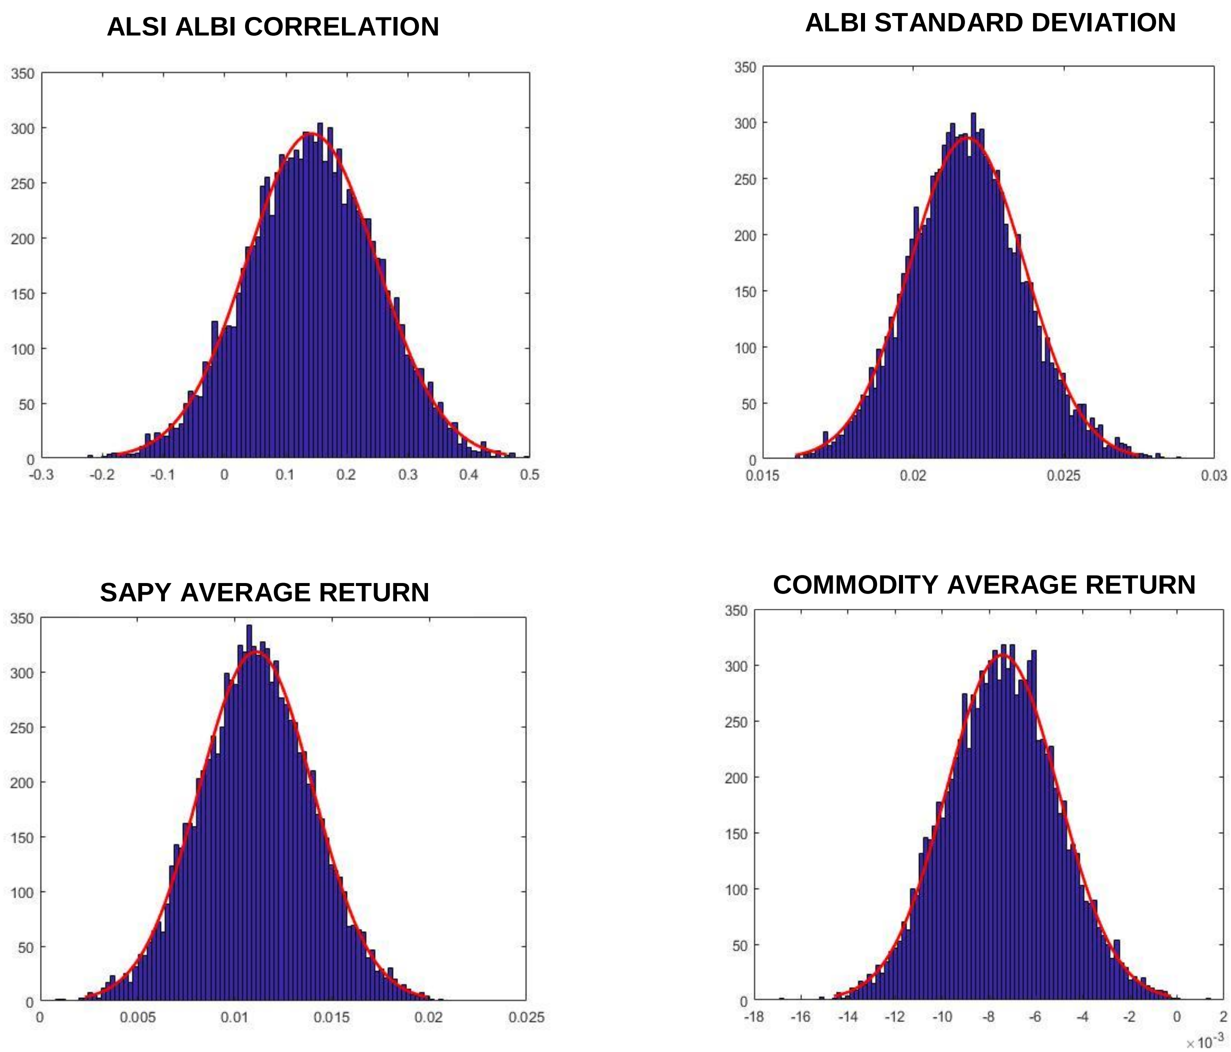

Figure 17: The selected distributions of bootstrapped average returns, correlations and standard deviations.

We calculated the optimal portfolio using the end of month data sourced from Bloomberg and our estimates of bootstrapped reference-day risk-free average returns, standard deviations and correlations for each asset class. The results indicated that the Bloomberg data overstated average returns and understated standard deviations. The average returns from Bloomberg were higher for the five asset classes, and the standard deviations were lower for the ALSI, SAPY and MSIC World, relative to the reference-day risk-free values. Correlations from Bloomberg were lower than their reference-day risk-free counterparts, for example, the Bloomberg correlation between the ALSI and SAPY was 0.12; while the reference-day risk-free correlation was 0.43. This combination of higher average returns, lower standard deviations and lower correlations resulted in an optimised portfolio that had higher local equity exposure, higher property exposure and lower offshore and fixed income exposure than the reference-day risk-free portfolio. When Bloomberg data is used an investor would inadvertently be allocating a higher exposure of the portfolio to the two asset classes with the largest levels of volatility. The table below summarises the optimised weights from each portfolio and provides a chart depicting the related performance over the in-sample period along with individual asset class performance.

Asset Class	Bloomberg Weight	Reference-day Risk-Free Weight	Difference
Local Equities	46%	32%	13%
International Equities	7%	18%	-23%
Fixed Income	14%	37%	20%
Commodities	0%	0%	0%
Property	33%	13%	11%

Table 2: The optimised weights based on Bloomberg and bootstrapped inputs.

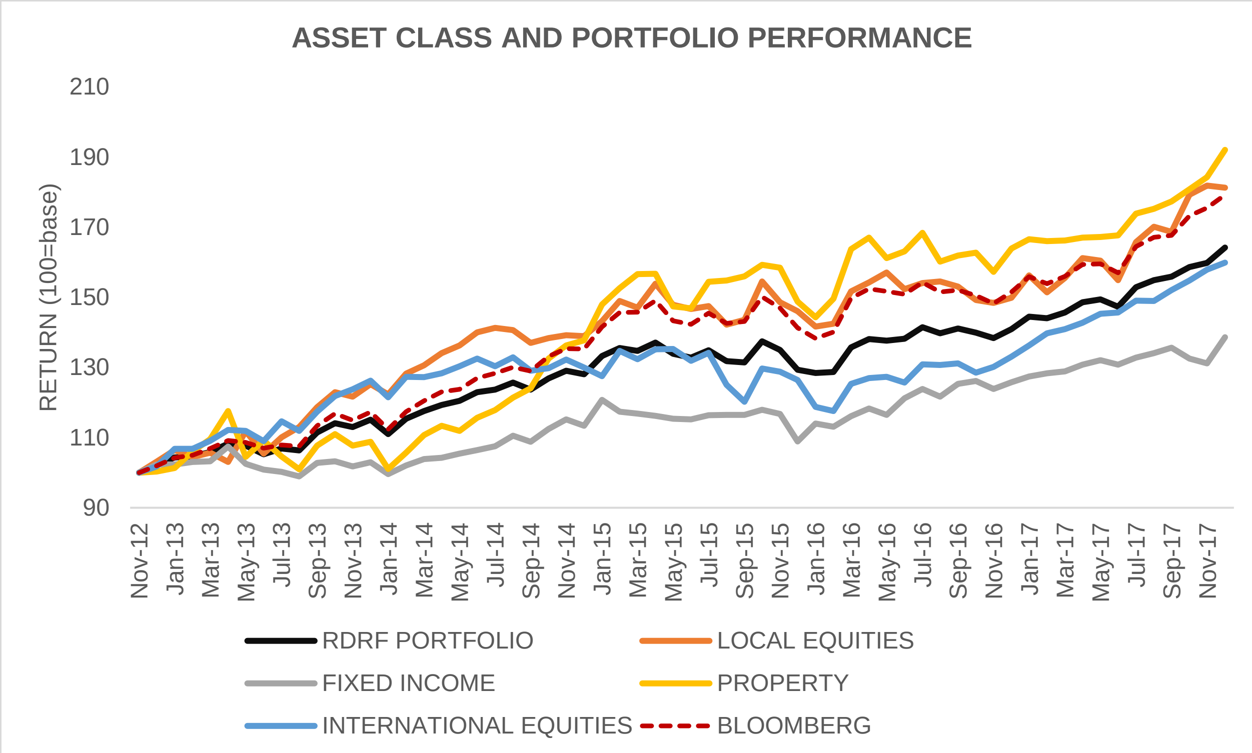

Figure 18: The relative performance of asset classes, Bloomberg end of month optimised portfolio and reference-day risk-free optimised portfolio. RDRF Portfolio represents the reference-day risk-free portfolio.

Overall Figure 18 shows that the end of month optimal portfolio returned 79%; while the reference-day risk-free portfolio returned 64% over the five-year period, with Sharpe ratios of 0.84 and 0.56 respectively. The higher risk-adjusted performance for the end of month portfolio is explained by the higher average return, lower standard deviation and lower correlation inputs relative to the reference-day risk-free portfolio. To better understand the relationship between the two portfolios we used regression analysis to investigate the extent to which the returns of the Bloomberg portfolio explained the returns of the reference-day risk-free portfolio (Figure 19). They were highly correlated with an R squared of 93%, and the reference-day risk-free portfolio moved 0.83% for every 1% increase in the Bloomberg portfolio.

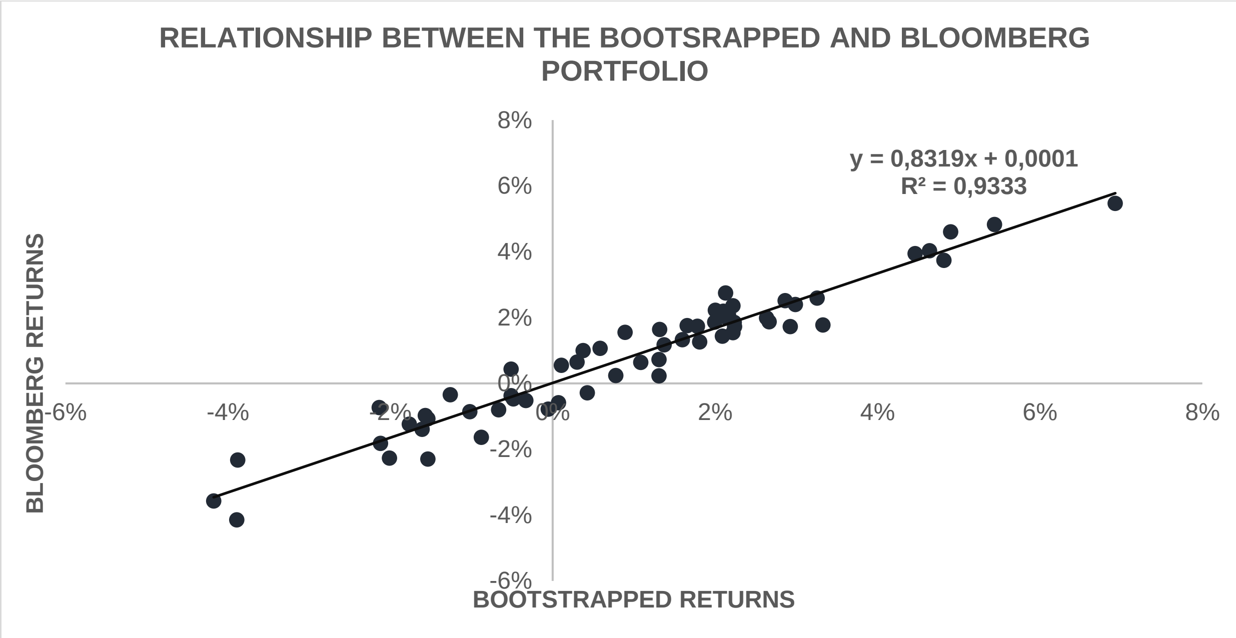

Figure 19: An OLS regression of reference-day risk-free portfolio returns against end of month Bloomberg portfolio returns.

Following this process, to understand the primary drivers of the returns for both the Bloomberg portfolio and the reference-day risk-free portfolio we regressed each asset class against both portfolios. The results are summarised in Table 3 below.

Asset Class	Beta to Bloomberg Portfolio	Beta to Adjusted Portfolio
Local Equities	0.56	0.45
International Equities	0.54	0.50
Fixed Income	0.53	0.56
Property	0.54	0.35

Table 3: The betas of each portfolio against individual asset classes.

Table 3 shows that the reference-day risk-free portfolio had a lower systematic risk, lower sensitivity to the most volatile property sector, less exposure to offshore market risk and a similar beta in relation to fixed income. This highlighted that the risk level of the Bloomberg portfolio was overstated based on using the end of month data relative to the reference-day risk-free underlying metrics for the period.

Similar to the portfolio's that were optimised on each of the twenty reference-days in Figure 16, we test the out-of-sample performance of both the end of month Bloomberg portfolio and our reference-day risk-free portfolio below in Figure 20. The weights of each portfolio were based on the optimised weights from the in-sample period. The reference-day risk-free portfolio outperformed by 6.2% during 2018 with a -6% return; while the end of month Bloomberg portfolio was down -12.20%.

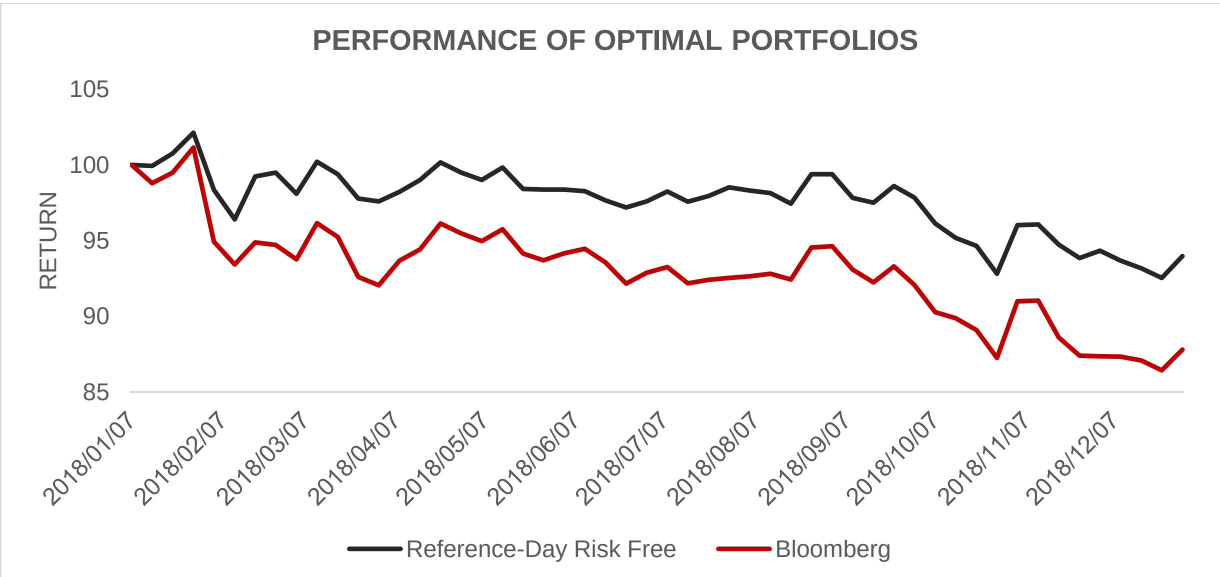

Figure 20: The out-of-sample performance of the end of month optimal Bloomberg portfolio relative to the reference-day risk-free optimal portfolio.

In conclusion, there were large difference in inputs, allocation weights and performance when the reference-day risk-free portfolio was compared to an unadjusted end of month alternative. While the unadjusted portfolio performed better over the in-sample period, the inputs that were used were not an accurate reflection of the risk-return dynamics of the underlying assets. The bootstrapping method outlined in section 4.2.1 provided the most accurate estimate of metrics in the presence of reference-day risk. By using these adjusted metrics as inputs for portfolio optimisation, the portfolio would more accurately represent the real underlying risk-return profile.

Chapter 6: Conclusion

The objectives of this paper were to establish the extent to which reference-day risk impacted commonly used risk-return metrics on the JSE as well as to establish a multi-metric method to both measure and adjust for the presence of reference-day risk. Lastly, it aimed to apply these findings to modern portfolio theory in order to construct the optimal reference-day risk-free portfolio. In accordance with the results of Baker et al. (2016) and Sahadev et al. (2018), we concluded that reference-day risk exists across average returns, standard deviations, correlations and betas at both a single stock and index level on the JSE as well as in simulated data. This has implications for valuation, portfolio theory, risk management, and it emphasises the need for a practical, robust and cross-metric method for estimating reference-day risk-free metrics. After proving the existence of reference-day risk across metrics, this paper used a nonparametric bootstrapping method to determine the values of average returns, standard deviations, correlations and betas adjusted for reference-day risk. First, we provided evidence of the method's accuracy by predefining values for each metric and generating a bootstrapped distribution composed of 100 000 iterations. The mean and standard deviation of the bootstrapped distribution was shown to be appropriate estimates of the predefined reference-day risk-free value and the associated level of reference-day risk respectively. We applied this method to the largest fifty-six companies on the JSE and found that the bootstrapped values were approximately equal to the average across reference-days with an R squared of 99% for all the metrics. These findings were extended from single companies to portfolios that showed the extent to which both general and optimised portfolios were highly dependent on the choice of initial reference-day. Lastly, we used these reference-day risk-free values to find the optimal reference-day risk-free portfolio and performed a comparative analysis to a traditional end of month portfolio using Bloomberg data. We found that the end of month portfolio underperformed out-of-sample, overstated average returns, understated standard deviations and had lower correlations than those adjusted for reference-day risk. Therefore, it is important that reference-day risk is measured and accounted for in the process of applying the Markovitz (1952) mean-variance optimisation, in addition to any analysis that is based on statistical inputs measured on a single reference-day.

For further research, the findings of this paper can be applied to the practice of valuation and asset pricing, or they could be tested against additional metrics such as duration and convexity for fixed income instruments. Additionally, the extent to which the findings apply to global markets can be investigated in order to examine the geographic dependency of the results. This will allow investors to understand the global impact of reference-day risk better. Lastly, the findings can be extended to other portfolio optimisation methods, such as multi-factor models.

This paper focused solely on the JSE and a resulting limitation is that the findings may have limited application to international markets. Furthermore, although this research included a large range of companies from different sectors, there were no smaller, less liquid companies presented. This means that these findings may not be generalisable to proportions of the population outside of the sample presented, as these samples may respond differently to conditions of reference-day risk.

Bibliography

Acker, D. & Duck, N.W. 2007. Reference-day risk and the use of monthly returns data. *Journal of Accounting, Auditing & Finance.* 22(4): 527-557.

Ader, H.J. & Ader, M. 2008. *Advising on research methods: a consultant's companion*. Huizen, Netherlands: Johannes Van Kessel.

Baker, C., Rajaratnam, K. & Flint, E.J. 2016. Beta estimates of shares on the JSE Top 40 in the context of reference-day risk. *Environment Systems and Decisions*. 36(2): 126-141.

Barry, C.B. 1974. Portfolio analysis under uncertain means, variances, and covariances. *Journal of Finance.* 29(2): 515-22.

Bey, R., Burgess, R. & Cook, P. 1990. Measurement of estimation risk in Markowitz portfolios. University of Tusla. Unpublished.

Blume, M.E. 1971. On the assessment of risk. *Journal of Finance.* 26(1): 1-10.

Bradfield, D. 2003. Investment basics XLVI. On estimation the beta coefficient. *Investment Analysts Journal.* 57: 47-52.

Cademartori, D., Romo, C., Campos, R. & Galea, M. 2003. Robust estimation of systematic risk using the t-distribution in the Chilean stock markets. *Applied Economics Letters*. 10(7): 447.

Chopra, V.K. & Ziemba, W.T. 1993. The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*. 19(2): 6-11.

Correia, C. & Cramer, P. 2008. An analysis of cost of capital, capital structure and capital budgeting practices: a survey of South African listed companies. *Meditari Accountancy Research.* 16(2): 31-52.

Damodaran, A. 2008. What is risk-free rate? A search for the basic building block. Available: http://ssrn.com/abstract=1317436. Unpublished.

Damodaran, A. 2015. Equity risk premiums: determinants, estimation and implications. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2581517. Unpublished.

Dimitrov, V. & Govindaraj, S. 2007. Reference-day risk: observations and extensions. *Journal of Accounting, Auditing & Finance*. 22(4): 559-572.

Dimson, E. 1979. Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics.* 7(2): 197-226.

Elton, E.J. & Gruber, M.J. 1997. Modern portfolio theory, 1950 to date. *Journal of Banking & Finance.* 21(1): 1743-1759.

Ernst and Young. 2009. *Estimating risk-free rates for valuations*. Available: https://www.ey.com/Publication/vwLUAssets/EY-estimating-risk-free-rates-for-valuations/$FILE/EY-estimating-risk-free-rates-for-valuations.pdf.

Fama, E.F. & French, K.R. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics.* 33(1): 3-56.

Feinstein, S., Polden, S., Richardson, L. & Rajaratnam, K. 2016. An investigation into a methodology for estimating reference-day risk free metrics on the JSE top 40. Unpublished.

Fisher, K.L. & Statman, M. 1997. Investment advice from mutual fund companies. *Journal of Portfolio Management.* 24(1): 9-25.

Fox, J., Friendly, M. & Weisberg, S. 2013. Hypothesis testing for multivariate linear models using the car package. *R Journal*. 5(1): 39-52.

Gonzalez, M., Rodriguez, A. & Stein, R. 2014. Adjusted betas under reference-day risk. *Engineering Economist.* 59(1): 79-88.

Gonzalez, M., Rodriguez, A. & Stein, R. 2016. Portfolio performance under reference-day risk. *Investment Analyst Journal.* 46(1): 32-43.

Greig, N. 2016. Risk parity as an asset allocation technique: evidence from the South African capital markets. Thesis. University of Cape Town. Unpublished.

Kuzma, J.W. & Bohnenblust, S.E. 2005. *Basic statistics for the health sciences.* Boston: McGraw-Hill.

Longin, F. & Solnik, B. 2001. Extreme correlation of international equity markets. *Journal of Finance.* 56(2): 649-676.

Markowitz, H. 1952. Portfolio selection. *Journal of Finance*. 7(1): 77-91.

Michaud, R.O. 1989. The Markowitz optimization enigma: is 'optimized' optimal?. *Financial Analysts Journal*. 45(1): 31-42.

Michaud, R. 1998. *Efficient asset management*. New York, United States: Harvard Business School Press.

Michaud, R.O. & Michaud, R. 2008. Estimation error and portfolio optimisation: a resampling solution. *Journal of Investment Management.* 6(1): 8-28.

PwC. 2017. *PwC valuation methodology survey 8th edition*. Available: https://www.pwc.co.za/en/publications/valuation-methodology-survey.html.

Sahadev, K., Ward, M. & Muller, C. 2018. The impact of reference-day risk on beta estimation and a proposed solution. *Investment Analysts Journal*. 47(4): 327 – 342.

Stein, C. 1955. *Proceedings of the 3rd Berkeley Symposium on Mathematical Statistics and Probability*. July-August 1955. University of California Press, Berkeley: Royal Statistical Society.

Sullivan, G.M. 2011. A primer on the validity of assessment instruments. *Journal of Graduate Medical Education.* 3(2): 119-120.

Ting, C. 2006. Which daily price is less noisy?. *Financial Management*. 35(3): 81-95.

Vaihekoski, M. 2014. A note on the calculation of the risk-free rate for tests of asset pricing models and event studies. *Contributions to Mathematics, Statistics, Econometrics and Finance*. 296(7): 291-306.

Vasicek, O.A. 1973. A note on using cross-sectional information in Bayesian estimation of security bias. *Journal of Finance*. 28(5): 1233-1239

Appendices

Appendix A: List of the fifty-six qualifying companies used for this paper

Companies	
Anglo American Platinum	Mr Price
Anglo American PLC	MTN
AngloGold Ashanti	Naspers
Aspen Pharma	Nedbank
AVI	Netcare
Barclays	Old Mutual
Barloworld	Pick n Pay
BHP Billiton	Pioneer Food Group
Bidvest	PSG Group
British American Tabaco	Rand Merchant Holdings
Capitec	Rand Merchant Investments
Clicks	Redefine
Coronations	Remgro
Discovery	Resilient
Exarro Resources	Richemont
FirstRand	Sanlam
Fortress A	Sappi
Glencore	Sasol
Goldfields	Shoprite
Growth Point	SPAR
Hyprop	Standard Bank
Imperial	Telkom
Intu	The Foshini Group
Investec	Tiger Brands
Kumba Iron Ore	Truworths
Life Healthcare	Vodacom
MMI	Vukile
Mondi	Woolworths

Appendix B: Average return high, low, ranges and reference-day risk-free estimates

Average Returns						
Company	**Min**	**Max**	**Range**	**Bootstrapped**	**20 Day Average**	**STD of Bootstrapped**
All Share Index	8.9%	10.2%	1.3%	9.4%	9.3%	0.3%
Anglo American Platinum	4.6%	9.9%	5.3%	7.2%	7.3%	0.9%
Anglo American PLC	12.7%	17.7%	4.9%	14.7%	14.8%	1.0%
Anglogold Ashanti	-2.4%	2.8%	5.2%	0.4%	0.4%	1.1%
Aspen Pharma	14.6%	19.3%	4.7%	16.9%	17.0%	0.6%
AVI	18.3%	21.1%	2.8%	19.6%	19.6%	0.4%
Barclays	8.0%	13.7%	5.7%	10.8%	11.0%	0.5%
Barloworld	20.0%	25.1%	5.1%	22.6%	22.5%	0.6%
BHP Billiton	3.8%	6.6%	2.8%	5.3%	5.3%	0.7%
Bidvest	10.7%	15.2%	4.5%	11.9%	12.0%	0.6%
British American Tabaco	19.5%	21.4%	1.9%	20.6%	20.6%	0.4%
Capitec	45.5%	52.6%	7.0%	48.5%	48.3%	0.5%
Clicks	27.2%	30.9%	3.7%	28.3%	28.3%	0.4%
Coronations	21.8%	26.4%	4.6%	23.9%	24.0%	0.5%
Discovery	27.5%	31.1%	3.6%	29.2%	29.3%	0.5%
Exarro	10.6%	13.2%	2.6%	12.1%	12.2%	0.9%
FirstRand	20.3%	25.3%	5.0%	23.0%	23.0%	0.5%
Fortress A	13.1%	14.7%	1.6%	13.9%	14.0%	0.4%
Glencore	10.4%	17.8%	7.4%	13.4%	13.2%	0.9%
Goldfields	-1.6%	3.2%	4.7%	0.8%	0.7%	1.1%
Growth Point	8.3%	11.3%	3.0%	9.6%	9.7%	0.4%
Hyprop	16.1%	18.3%	2.2%	16.9%	17.1%	0.4%
Imperial	12.5%	17.3%	4.8%	14.3%	14.4%	0.6%
Intu	-5.6%	0.4%	6.0%	-1.5%	-1.4%	0.5%
Investec	15.3%	18.9%	3.6%	16.5%	16.5%	0.5%
Kumba Iron Ore	16.3%	28.8%	12.5%	23.7%	23.8%	1.4%
Life Healthcare	1.4%	3.9%	2.5%	2.5%	2.5%	0.5%
MMI	6.5%	9.9%	3.4%	8.5%	8.5%	0.5%
Mondi	34.0%	36.0%	1.9%	35.3%	35.3%	0.5%
Mr Price	18.8%	24.2%	5.4%	20.8%	21.0%	0.7%
MTN	3.1%	4.7%	1.6%	3.8%	3.9%	0.6%
Naspers	49.4%	52.5%	3.1%	50.9%	51.1%	0.6%
Nedbank	11.6%	14.5%	2.9%	13.0%	13.0%	0.4%
Netcare	8.9%	11.0%	2.1%	10.0%	10.1%	0.5%
Old Mutual	14.5%	16.3%	1.7%	15.4%	15.3%	0.4%
Pick n Pay	13.5%	16.7%	3.2%	15.1%	15.1%	0.5%
Pioneer Food Group	17.9%	22.3%	4.4%	20.8%	20.7%	0.5%
PSG Group	38.3%	44.2%	5.9%	41.3%	41.3%	0.6%
Rand Merchant Holdings	17.0%	22.2%	5.2%	19.4%	19.6%	0.5%
Rand Merchant Investments	17.6%	23.9%	6.3%	21.1%	21.0%	0.4%
Redefine	10.5%	12.8%	2.3%	11.7%	11.7%	0.4%
Remgro	10.2%	12.7%	2.6%	11.5%	11.5%	0.4%
Resilient	29.4%	33.0%	3.6%	31.5%	31.6%	0.4%
Richemont	16.0%	17.5%	1.5%	16.5%	16.6%	0.5%
Sanlam	19.6%	23.7%	4.1%	21.3%	21.3%	0.5%
Sappi	28.9%	32.2%	3.3%	30.5%	30.3%	0.6%
Sasol	9.7%	12.8%	3.1%	11.0%	10.9%	0.6%
Shoprite	5.7%	10.8%	5.1%	7.8%	8.0%	0.6%
SPAR	14.5%	17.2%	2.7%	15.9%	16.0%	0.5%
Standard Bank	17.0%	19.3%	2.4%	18.3%	18.2%	0.5%
Telkom	34.7%	37.9%	3.1%	36.2%	36.1%	0.7%
The Foshini Group	11.8%	18.2%	6.4%	14.3%	14.4%	0.7%
Tiger Brands	11.2%	13.5%	2.3%	12.3%	12.3%	0.5%
Truworths	4.7%	6.9%	2.3%	5.7%	5.7%	0.6%
Vodacom	10.7%	13.0%	2.3%	11.7%	11.9%	0.5%
Vukile	11.5%	14.7%	3.2%	13.2%	13.2%	0.4%
Woolworths	1.8%	5.8%	4.0%	3.7%	3.7%	0.6%

Appendix C: Standard deviation high, low, ranges and reference-day risk-free estimates

Standard Deviations						
Company	**Min**	**Max**	**Range**	**Bootstrapped**	**20 Day Average**	**STD of Bootstrapped**
All Share Index	10.7%	14.1%	3.3%	12.3%	12.4%	0.2%
Anglo American Platinum	45.4%	55.2%	9.8%	50.1%	49.9%	1.2%
Anglo American PLC	44.5%	56.8%	12.2%	49.5%	50.0%	1.9%
Anglogold Ashanti	50.5%	60.7%	10.3%	55.7%	55.8%	1.0%
Aspen Pharma	24.8%	32.7%	7.8%	28.0%	28.1%	0.5%
AVI	16.1%	18.9%	2.9%	17.7%	17.8%	0.4%
Barclays	20.0%	24.9%	5.0%	22.2%	22.4%	0.5%
Barloworld	25.4%	32.2%	6.9%	28.6%	28.7%	0.7%
BHP Billiton	26.6%	30.4%	3.8%	28.6%	28.7%	0.5%
Bidvest	33.4%	37.3%	3.9%	35.6%	35.8%	1.2%
British American Tabaco	14.5%	19.2%	4.7%	16.6%	16.6%	0.3%
Capitec	25.8%	30.0%	4.2%	27.8%	27.8%	0.5%
Clicks	19.0%	24.4%	5.3%	21.5%	21.6%	0.4%
Coronations	24.2%	30.2%	6.0%	27.1%	27.1%	0.5%
Discovery	20.4%	25.7%	5.3%	23.5%	23.5%	0.4%
Exarro	41.1%	47.2%	6.0%	44.4%	44.8%	1.2%
FirstRand	21.3%	25.5%	4.1%	23.4%	23.5%	0.5%
Fortress A	11.1%	17.1%	5.9%	13.3%	13.3%	0.4%
Glencore	43.0%	54.6%	11.6%	47.6%	47.7%	1.2%
Goldfields	46.6%	55.4%	8.8%	50.9%	51.1%	0.9%
Growth Point	16.6%	21.0%	4.4%	18.3%	18.4%	0.4%
Hyprop	18.5%	23.4%	4.8%	20.8%	20.8%	0.4%
Imperial	26.3%	34.4%	8.1%	30.4%	30.5%	0.6%
Intu	17.4%	25.0%	7.6%	21.4%	21.4%	0.5%
Investec	18.8%	23.4%	4.6%	21.0%	21.0%	0.4%
Kumba Iron Ore	65.2%	87.0%	21.8%	76.2%	76.9%	3.1%
Life Healthcare	20.8%	24.7%	3.9%	22.2%	22.3%	0.5%
MMI	21.5%	27.6%	6.1%	24.1%	24.2%	0.5%
Mondi	22.7%	26.2%	3.5%	24.5%	24.5%	0.4%
Mr Price	31.1%	35.2%	4.1%	33.0%	33.2%	0.6%
MTN	22.6%	30.0%	7.4%	25.4%	25.5%	0.6%
Naspers	26.6%	32.1%	5.5%	29.2%	29.3%	0.5%
Nedbank	18.5%	23.8%	5.4%	21.0%	21.0%	0.4%
Netcare	20.5%	25.9%	5.4%	23.2%	23.3%	0.4%
Old Mutual	17.7%	23.0%	5.3%	20.2%	20.2%	0.4%
Pick n Pay	23.3%	28.3%	5.1%	25.5%	25.6%	0.5%
Pioneer Food Group	23.1%	30.1%	7.0%	26.5%	26.5%	0.5%
PSG Group	26.8%	31.6%	4.7%	29.2%	29.3%	0.5%
Rand Merchant Holdings	19.9%	23.8%	3.9%	21.7%	21.8%	0.4%
Rand Merchant Investments	17.9%	22.8%	4.9%	20.4%	20.4%	0.4%
Redefine	17.6%	22.4%	4.8%	19.3%	19.4%	0.4%
Remgro	17.0%	23.1%	6.1%	19.7%	19.8%	0.4%
Resilient	19.7%	24.1%	4.5%	21.4%	21.5%	0.4%
Richemont	21.0%	24.4%	3.4%	22.7%	22.8%	0.6%
Sanlam	20.9%	27.2%	6.3%	24.1%	24.2%	0.5%
Sappi	25.2%	32.2%	7.0%	28.5%	28.6%	0.6%
Sasol	23.5%	31.0%	7.5%	26.6%	26.7%	0.6%
Shoprite	24.1%	28.7%	4.6%	26.8%	26.9%	0.6%
SPAR	20.2%	23.5%	3.3%	22.2%	22.3%	0.5%
Standard Bank	20.8%	25.4%	4.6%	23.0%	23.1%	0.5%
Telkom	31.0%	40.3%	9.4%	35.8%	35.8%	0.6%
The Foshini Group	28.0%	34.6%	6.7%	30.6%	30.7%	0.7%
Tiger Brands	19.6%	24.3%	4.7%	21.7%	21.7%	0.4%
Truworths	26.1%	32.9%	6.7%	29.2%	29.3%	0.5%
Vodacom	17.9%	22.4%	4.5%	19.8%	19.9%	0.4%
Vukile	13.1%	19.4%	6.2%	15.6%	15.6%	0.4%
Woolworths	21.0%	29.0%	8.0%	25.8%	25.8%	0.5%

Appendix D: Beta high, low, ranges and reference-day risk-free estimates

			Betas			
Company	**Min**	**Max**	**Range**	**20 Day Average**	**Bloomberg**	**Bootstrappped**
Anglo American Platinum	1.11	2.29	1.17	1.74	1.40	1.74
Anglo American PLC	1.65	2.59	0.95	1.98	1.63	1.98
Anglogold Ashanti	0.76	2.19	1.44	1.28	1.19	1.29
Aspen Pharma	0.44	1.32	0.88	0.82	0.62	0.83
AVI	0.30	0.60	0.30	0.46	0.21	0.46
Barclays	0.56	0.95	0.39	0.74	0.69	0.73
Barloworld	0.26	0.93	0.67	0.67	0.32	0.68
BHP Billiton	1.40	1.71	0.31	1.57	1.82	1.57
Bidvest	0.18	1.27	1.08	0.60	0.31	0.62
British American Tabaco	0.35	0.72	0.38	0.53	0.79	0.52
Capitec	0.49	1.08	0.58	0.72	0.45	0.73
Clicks	0.43	0.89	0.46	0.65	0.55	0.65
Coronations	0.55	1.11	0.57	0.83	0.74	0.83
Discovery	0.65	0.98	0.34	0.82	0.59	0.81
Exarro	0.85	1.68	0.83	1.25	1.27	1.25
FirstRand	0.61	0.95	0.34	0.82	0.71	0.82
Fortress A	0.05	0.53	0.49	0.26	0.09	0.27
Glencore	-0.09	1.74	1.82	0.87	1.83	0.91
Goldfields	0.07	0.95	0.89	0.42	0.34	0.43
Growth Point	0.11	0.67	0.56	0.48	-0.02	0.49
Hyprop	0.39	0.89	0.50	0.64	0.08	0.65
Imperial	0.77	1.49	0.72	1.06	0.88	1.08
Intu	-0.22	0.27	0.48	0.04	0.77	0.04
Investec	0.98	1.23	0.25	1.11	1.14	1.11
Kumba Iron Ore	1.12	2.80	1.68	1.79	1.07	1.83
Life Healthcare	0.41	0.84	0.43	0.63	0.56	0.62
MMI	0.82	1.34	0.52	1.11	0.89	1.13
Mondi	0.72	1.15	0.43	0.93	1.08	0.92
Mr Price	0.55	1.27	0.72	0.86	0.63	0.85
MTN	0.35	1.19	0.84	0.90	0.63	0.92
Naspers	1.18	1.78	0.60	1.46	1.51	1.47
Nedbank	0.53	0.91	0.38	0.72	0.61	0.72
Netcare	0.45	1.07	0.62	0.82	0.83	0.82
Old Mutual	0.98	1.35	0.38	1.19	1.17	1.19
Pick n Pay	0.31	1.02	0.71	0.65	0.80	0.65
Pioneer Food Group	0.55	0.91	0.36	0.70	0.62	0.70
PSG Group	0.70	1.22	0.52	0.92	0.72	0.91
Rand Merchant Holdings	0.75	1.03	0.28	0.90	0.77	0.90
Rand Merchant Investments	0.73	1.03	0.30	0.88	0.62	0.88
Redefine	0.32	0.50	0.18	0.56	0.16	0.57
Remgro	0.76	1.20	0.44	0.96	0.81	0.95
Resilient	0.20	0.68	0.49	0.50	0.17	0.52
Richemont	0.79	1.13	0.34	1.02	1.29	1.02
Sanlam	0.96	1.37	0.41	1.21	1.03	1.22
Sappi	0.18	0.84	0.66	0.57	0.63	0.57
Sasol	1.12	1.70	0.58	1.38	1.40	1.38
Shoprite	0.35	0.84	0.49	0.56	0.22	0.57
SPAR	0.40	0.89	0.49	0.65	0.51	0.66
Standard Bank	0.83	1.12	0.29	1.01	0.89	1.00
Telkom	0.66	1.15	0.49	0.92	0.99	0.92
The Foshini Group	0.63	1.07	0.43	0.88	0.59	0.88
Tiger Brands	0.41	0.94	0.53	0.70	0.54	0.71
Truworths	-0.06	0.79	0.85	0.44	0.59	0.42
Vodacom	0.39	0.90	0.51	0.61	0.74	0.60
Vukile	0.17	0.51	0.34	0.34	0.10	0.34
Woolworths	0.57	1.05	0.48	0.84	0.84	0.84

Appendix E: Correlation high, low, ranges and reference-day risk-free estimates

Correlations							
Company	**Min**	**Max**	**Range**	**20 Day Average**	**Bloomberg**	**Bootstrapped**	**STD of Bootstrapped**
Anglo American Platinum	0.33	0.50	0.17	0.43	0.32	0.42	0.08
Anglo American PLC	0.37	0.59	0.23	0.49	0.39	0.49	0.06
Anglogold Ashanti	0.15	0.50	0.34	0.28	0.18	0.28	0.10
Aspen Pharma	0.19	0.53	0.34	0.36	0.29	0.37	0.09
AVI	0.21	0.43	0.22	0.32	0.16	0.32	0.09
Barclays	0.33	0.49	0.17	0.41	0.36	0.41	0.08
Barloworld	0.11	0.40	0.30	0.29	0.12	0.29	0.09
BHP Billiton	0.61	0.74	0.13	0.68	0.69	0.67	0.05
Bidvest	0.06	0.46	0.40	0.21	0.10	0.22	0.11
British American Tabaco	0.30	0.48	0.19	0.39	0.41	0.39	0.07
Capitec	0.24	0.46	0.22	0.32	0.08	0.32	0.09
Clicks	0.26	0.50	0.23	0.37	0.30	0.37	0.08
Coronations	0.25	0.53	0.28	0.38	0.35	0.38	0.08
Discovery	0.34	0.51	0.17	0.43	0.33	0.43	0.07
Exarro	0.25	0.44	0.19	0.35	0.31	0.34	0.09
FirstRand	0.30	0.51	0.21	0.43	0.36	0.43	0.08
Fortress A	0.04	0.40	0.36	0.24	0.04	0.25	0.11
Glencore	-0.02	0.44	0.46	0.23	-	0.24	0.12
Goldfields	0.01	0.23	0.22	0.10	0.04	0.10	0.10
Growth Point	0.07	0.49	0.42	0.32	0.04	0.33	0.10
Hyprop	0.22	0.50	0.28	0.38	0.03	0.38	0.08
Imperial	0.34	0.56	0.22	0.43	0.32	0.43	0.07
Intu	-0.13	0.18	0.30	0.02	0.33	0.02	0.09
Investec	0.59	0.71	0.13	0.65	0.64	0.65	0.06
Kumba Iron Ore	0.17	0.40	0.23	0.29	0.17	0.29	0.08
Life Healthcare	0.25	0.45	0.20	0.35	0.21	0.35	0.08
MMI	0.39	0.68	0.28	0.57	0.47	0.57	0.06
Mondi	0.40	0.53	0.13	0.46	0.47	0.46	0.07
Mr Price	0.20	0.45	0.26	0.32	0.24	0.32	0.09
MTN	0.16	0.60	0.44	0.44	0.26	0.44	0.09
Naspers	0.49	0.73	0.23	0.62	0.57	0.62	0.06
Nedbank	0.34	0.53	0.19	0.42	0.38	0.42	0.08
Netcare	0.23	0.56	0.33	0.43	0.35	0.43	0.08
Old Mutual	0.66	0.79	0.13	0.73	-0.04	0.73	0.05
Pick n Pay	0.15	0.48	0.33	0.31	0.32	0.31	0.08
Pioneer Food Group	0.25	0.42	0.17	0.33	0.28	0.32	0.08
PSG Group	0.29	0.47	0.18	0.38	0.24	0.38	0.08
Rand Merchant Holdings	0.40	0.58	0.18	0.51	0.44	0.51	0.07
Rand Merchant Investments	0.45	0.59	0.14	0.53	0.36	0.53	0.07
Redefine	0.20	0.50	0.30	0.36	0.11	0.36	0.08
Remgro	0.51	0.68	0.17	0.59	0.50	0.59	0.06
Resilient	0.11	0.37	0.26	0.29	-0.02	0.30	0.09
Richemont	0.47	0.64	0.17	0.55	0.53	0.55	0.07
Sanlam	0.49	0.73	0.23	0.62	0.48	0.62	0.06
Sappi	0.08	0.37	0.29	0.25	0.21	0.25	0.10
Sasol	0.55	0.71	0.16	0.64	0.63	0.64	0.05
Shoprite	0.14	0.39	0.26	0.26	0.13	0.26	0.08
SPAR	0.21	0.49	0.29	0.36	0.30	0.36	0.08
Standard Bank	0.47	0.58	0.12	0.54	0.43	0.54	0.07
Telkom	0.23	0.39	0.15	0.32	0.30	0.32	0.08
The Foshini Group	0.26	0.45	0.20	0.35	0.22	0.35	0.08
Tiger Brands	0.21	0.54	0.33	0.40	0.31	0.41	0.08
Truworths	-0.02	0.35	0.37	0.19	0.25	0.18	0.08
Vodacom	0.28	0.50	0.22	0.38	0.43	0.37	0.08
Vukile	0.12	0.41	0.28	0.27	0.08	0.27	0.10
Woolworths	0.26	0.48	0.22	0.40	0.37	0.40	0.08

Appendix F: Simulated correlation distributions

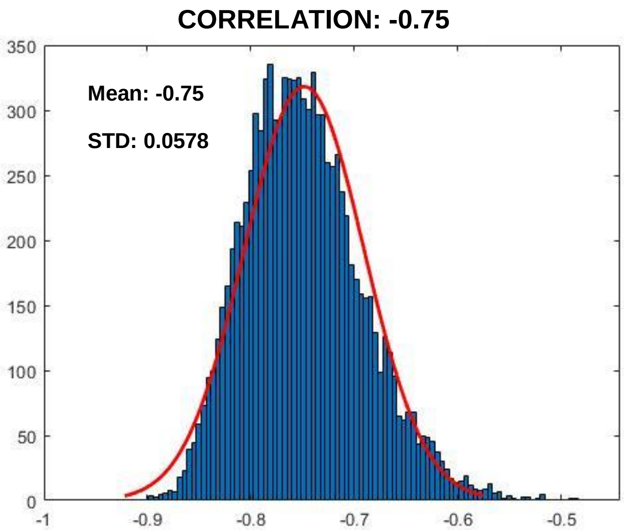

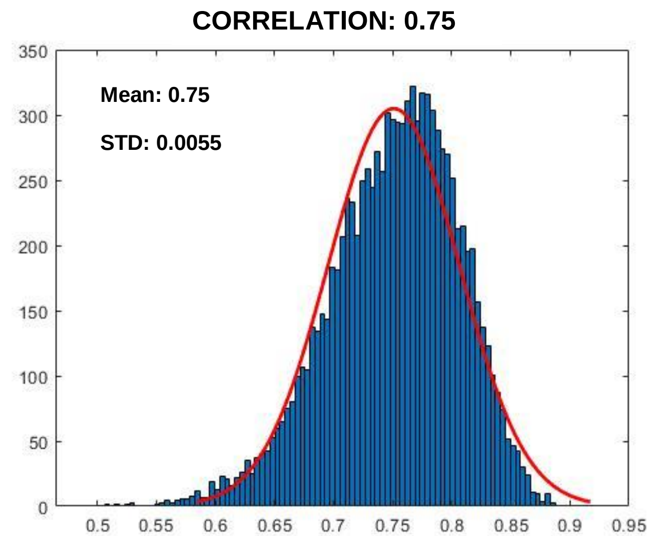

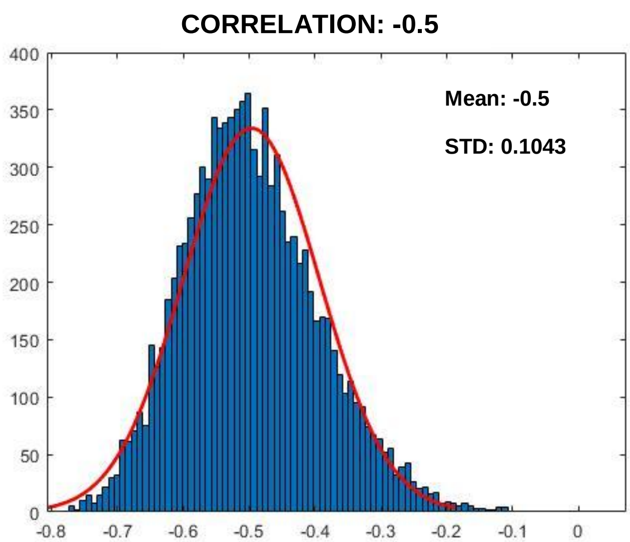

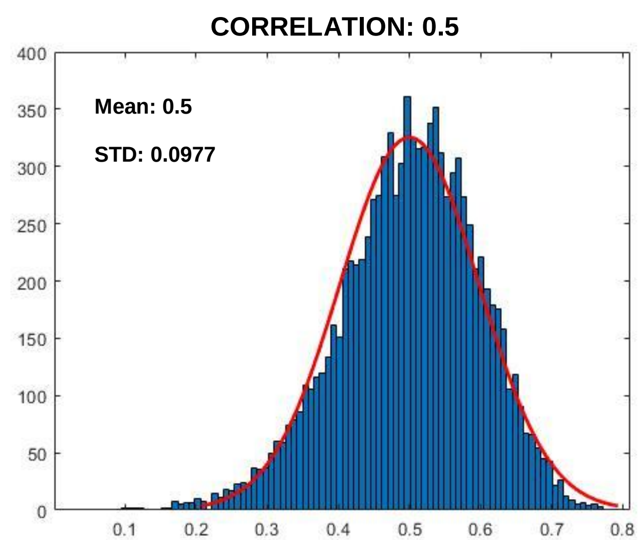

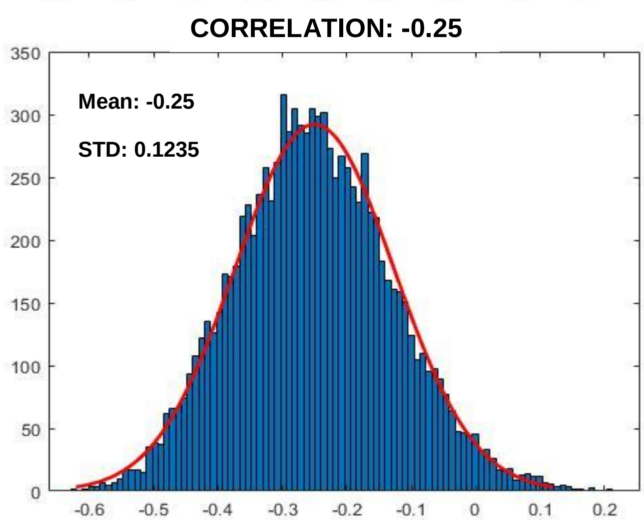

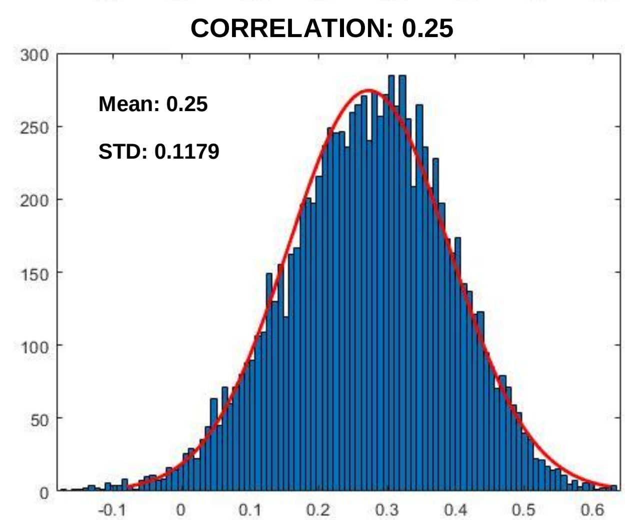

Appendix G: Simulated beta distributions

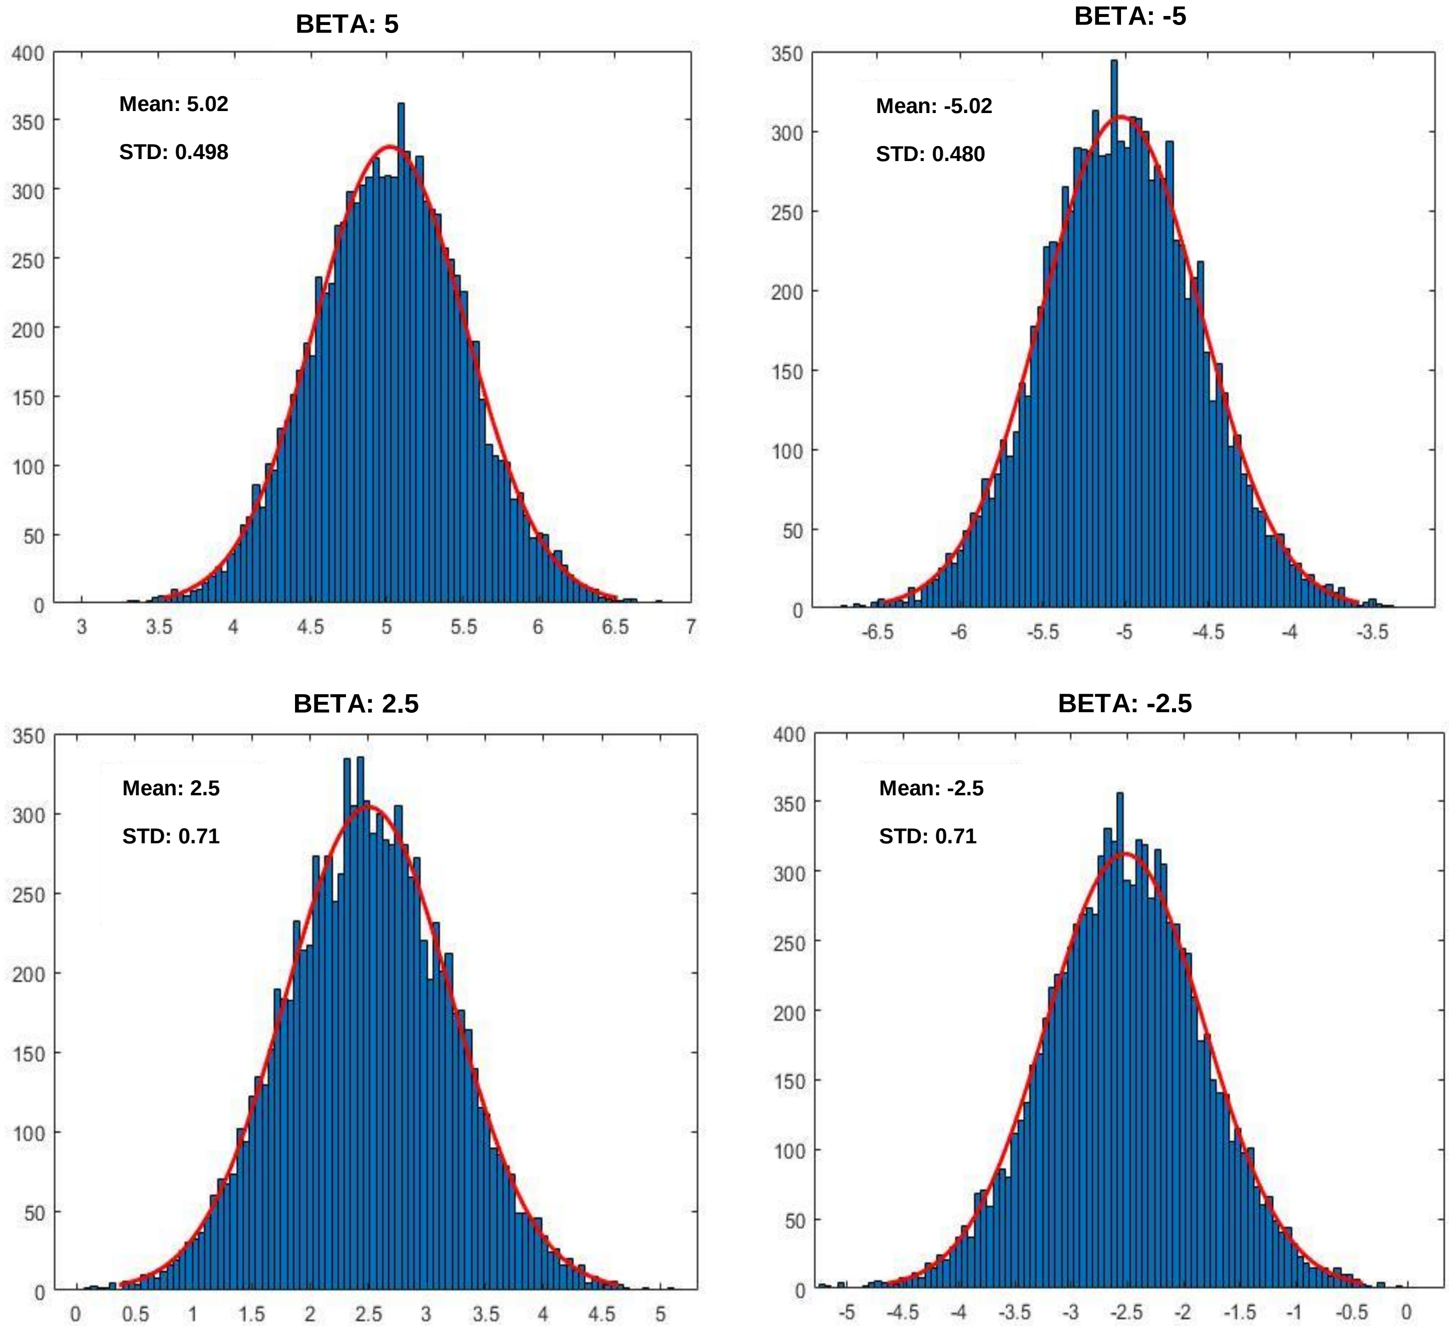

Appendix H: Simulated standard deviation (annual) and mean distributions (daily)

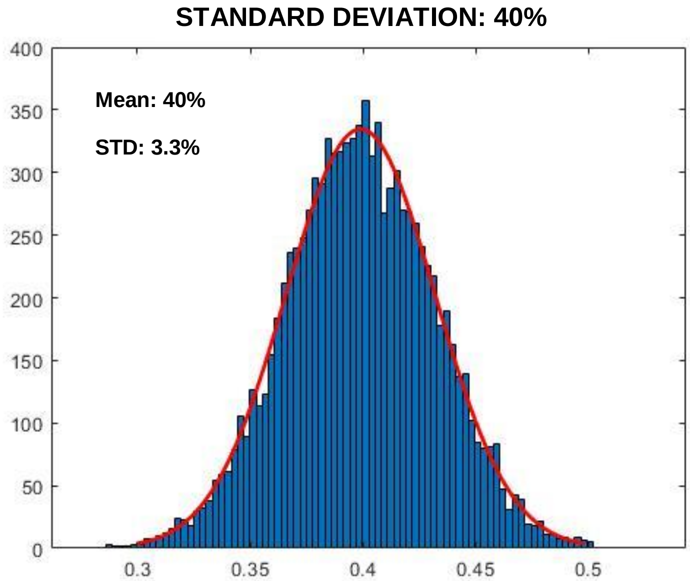

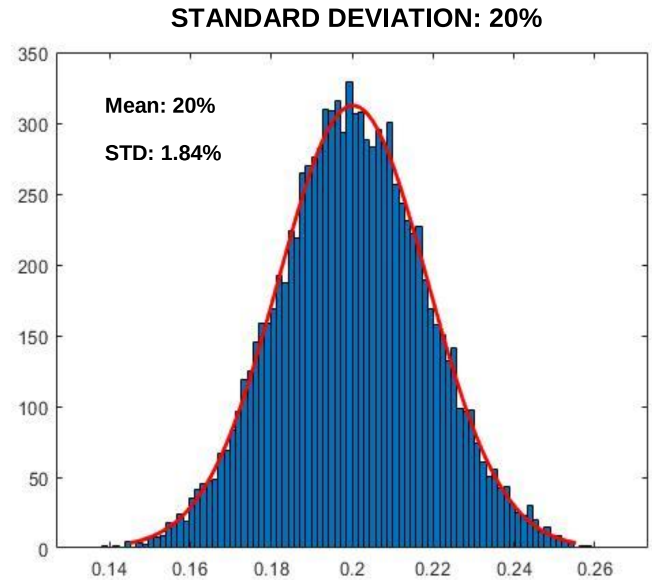

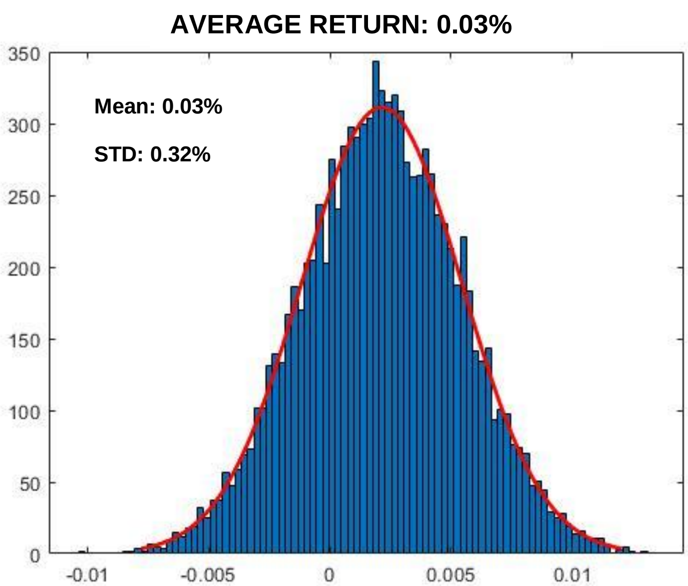

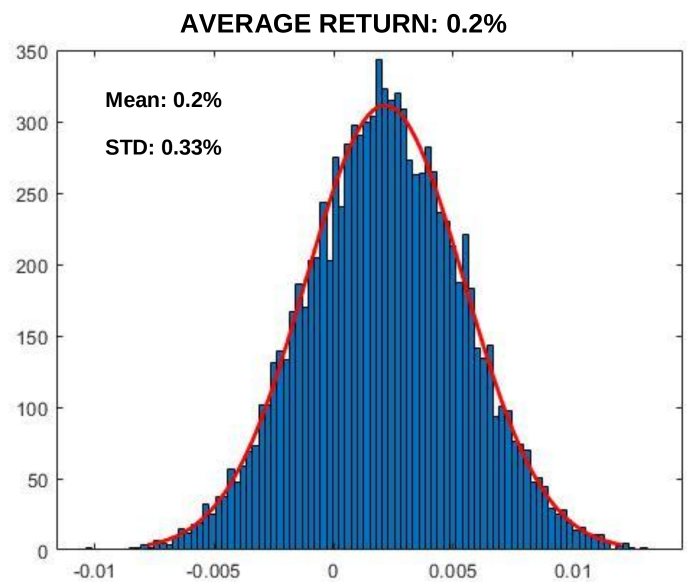

Appendix I: Asset class correlations across reference-days

Correlations	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	Day 10	Day 11	Day 12	Day 13	Day 14	Day 15	Day 16	Day 17	Day 18	Day 19	Day 20	Trend
ALSI | ALBI	0.20	0.25	0.17	0.20	0.23	0.11	0.07	0.19	0.20	0.27	0.24	0.20	0.17	0.16	0.14	0.01	-0.02	0.04	0.08	0.02	
ALSI | SAPY	0.26	0.43	0.39	0.47	0.50	0.41	0.34	0.39	0.48	0.51	0.55	0.47	0.50	0.53	0.46	0.38	0.37	0.42	0.37	0.27	
ALSI | STEFI	-0.05	-0.06	-0.09	-0.07	-0.08	-0.02	-0.03	0.03	0.02	-0.02	-0.02	-0.11	-0.14	-0.14	-0.14	-0.09	-0.16	-0.08	-0.08	-0.05	
ALSI | MSCI	0.61	0.69	0.70	0.67	0.67	0.63	0.57	0.58	0.61	0.68	0.67	0.66	0.65	0.72	0.70	0.65	0.68	0.71	0.66	0.62	
ALSI | BLGCOM	0.33	0.38	0.27	0.26	0.37	0.35	0.43	0.37	0.42	0.49	0.48	0.46	0.57	0.55	0.66	0.61	0.58	0.49	0.42	0.39	
ALBI | SAPY	0.66	0.66	0.67	0.66	0.67	0.71	0.70	0.73	0.78	0.72	0.66	0.66	0.70	0.68	0.67	0.62	0.63	0.60	0.57	0.64	
ALBI | STEFI	0.01	0.01	0.08	0.05	0.12	0.03	0.05	0.20	0.19	0.13	0.08	0.09	0.11	0.04	-0.05	-0.02	0.07	0.05	0.07	0.07	
ALBI | MSCI	0.29	0.31	0.32	0.40	0.38	0.31	0.31	0.21	0.15	0.29	0.28	0.36	0.45	0.39	0.38	0.27	0.22	0.24	0.14	0.17	
ALBI | BLGCOM	0.33	0.38	0.27	0.26	0.37	0.35	0.43	0.37	0.42	0.49	0.48	0.46	0.57	0.55	0.66	0.61	0.58	0.49	0.42	0.39	
SAPY | STEFI	-0.04	0.00	0.00	-0.01	0.02	0.02	0.01	0.08	0.11	0.04	0.00	-0.05	-0.07	-0.09	-0.17	-0.13	-0.07	-0.03	-0.02	-0.03	
SAPY | MSCI	0.30	0.46	0.47	0.51	0.55	0.46	0.46	0.36	0.36	0.48	0.48	0.52	0.62	0.64	0.60	0.52	0.46	0.55	0.43	0.41	
SAPY | BLGCOM	-0.09	-0.02	-0.13	-0.18	-0.09	-0.11	-0.07	-0.03	0.07	0.11	0.13	0.07	0.19	0.20	0.23	0.16	0.06	0.04	-0.06	-0.02	
STEFI | MSCI	-0.04	-0.08	-0.07	-0.08	-0.07	-0.02	0.00	0.05	0.02	-0.01	-0.03	0.00	-0.05	-0.07	-0.17	-0.15	-0.09	-0.05	-0.07	-0.06	
STEFI | BLGCOM	0.10	0.14	0.06	0.08	0.06	0.08	0.02	0.03	0.10	0.08	0.14	0.09	0.04	0.08	0.05	0.07	0.07	0.12	0.14	0.14	
MSCI | BLGCOM	0.26	0.22	0.19	0.21	0.27	0.25	0.35	0.30	0.30	0.33	0.31	0.33	0.46	0.53	0.55	0.55	0.48	0.37	0.31	0.30