School of Economics and Finance
Faculty of Commerce, Law and Management

The Bias ratio: An effective fraud identification tool

Master of Commerce (50% Research) in Finance

Remon Haddad **[2336055]**

Supervisor: Prof. Gary van Vuuren

School of Economics and Finance

June 2025

Abstract

Financial fraud poses significant risks with far-reaching consequences, particularly in the context of growing assets under management and expanding equity markets. This thesis underscores the urgent need for robust measures to safeguard investors from fraudulent activities by exploring the consequences of notorious fraud cases such as Bernie Madoff's Ponzi scheme. Through analyses of hedge fund, index fund and stock price return data in the US and SA, over various periods starting in 1997 to 2024, it becomes evident that tools such as the Bias ratio, kurtosis, and skewness can serve as effective mechanisms for detecting fraudulent behaviour. The Bias ratio emerges as a dual-purpose tool. Beyond its fraud detection capabilities, it functions as a performance measurement metric akin to the Sharpe ratio, offering additional value during security analysis. By highlighting suspicious historical outperformance and signalling securities with unusual performance patterns, the Bias ratio enriches the evaluation process, enabling investors to make informed decisions and avoid fraudulent investments. This thesis demonstrates the efficacy of the Bias ratio by examining its application in the notorious Madoff case, where it successfully flagged fraudulent activity that was overlooked by traditional measures like the Sharpe ratio. The findings emphasize the critical role of the Bias ratio in validating the legitimacy of returns and enhancing investor protection.

Table of Contents

1. Introduction

1.1 Background information

Investors are always looking for new opportunities to earn abnormal profits, often turning to alternative investment strategies that can create exposure to highly risky investments in the quest for ever higher returns. These lucrative investments can lead to exposures to various forms of financial crime, highlighting the undeniable importance of addressing this issue. Financial crime, such as fraud in the financial sector is typically action based manipulation, which is the inflating of stock prices through the release of false information to signal fictitious positive growth in a company (Felixson & Pelli, 1999). These are often unpredictable and can result in substantial losses. Managers and investors thus require robust metrics to identify potential financial crimes, such as fraud, before making investment decisions.

The most common form of financial fraud is stock market manipulation. This type of fraud is also the most far reaching as barriers for entry to the stock market are easy to overcome. Despite stocks having a stricter regulatory framework, there are still notable cases of fraud embedded in the stock markets history. The Bias ratio is not exclusively used to measure hedge fund performance, it can also be used to analyse the performance of any security. It is important to measure the Bias ratios performance on stock market and in hedge funds. The Bias ratio – a metric which its Abdulali (2006) claim can signal fraudulent reporting of investment returns in the early stages of such fraud – has been affirmed as a method to detect potentially fraudulent activity in investment return reporting and can serve as a crucial tool to evaluate (but not confirm beyond doubt) the legitimacy of investment performance.

Hedge funds, like any other investment strategy, are inherently risky; therefore, investors seek to obtain as much information as possible about the hedge fund before making decisions. This due diligence is conducted to ensure that the fund is genuinely earning the returns it claims. Hedge funds are private investment pools structured as partnerships or limited liability partnerships (Fung & Hsieh, 1999). Due to the inherent characteristics of hedge funds, this investment strategy is highly exposed to financial crime. The substantial initial capital outlay by investors can potentially leave them vulnerable to significant losses. Unlike mutual funds, hedge funds operate with less stringent regulatory oversight, allowing for a wide array of investment styles employed by hedge fund managers (Zask, 2020). Given the capital intensity and lower regulatory scrutiny, then, hedge funds are more susceptible to engaging in fraudulent activities. Both investors and regulators thus require effective tools to analyse and assess hedge fund performance.

The Bias ratio provides investors with valuable information on how an asset performs relative to a benchmark. Initially, the average and standard deviation of the benchmark portfolio's performance are computed (Abdulali, 2006). The Bias ratio is then calculated and compared to a chosen benchmark. Once the Bias ratio is determined, it aids in predicting the expected returns of the asset under assessment. Ideally, the asset's performance should align with the anticipated variance of the benchmark's performance. This methodology is applicable across various asset classes and respective benchmarks. To apply the Bias ratio to hedge funds and at a firm level, a specific equity benchmark must be specified. As previously noted, hedge funds are vulnerable to fraudulent activities due to lax regulations. The Bias ratio acts as a tool to scrutinise hedge fund and firm level performance for potential fraud. By comparing the biased returns of the asset to the unbiased returns of the benchmark, discrepancies can be identified. Persistent significant outperformance of the asset compared to the benchmark over multiple periods may raise red flags and warrant further investigation into potentially fraudulent activities.

Fraud is defined as the misrepresentation or omission of amounts or disclosures aimed at deceiving investors and creditors (Blythe, 2020). It constitutes a serious crime and can result in substantial financial losses, as illustrated by cases like Bernie Madoff's Ponzi scheme, which cost investors millions (U.S. securities and exchange commission, 2008). The Madoff case underscores the utility of the Bias ratio. Before the adoption of the Bias ratio, the Sharpe ratio was commonly employed to detect fraudulent activity (Douady, Abdulali & Alderberg, 2009). However, when applied to known fraud cases, the Sharpe ratio often fails to provide compelling evidence of fraudulent activity, as it is primarily used to evaluate an asset's return relative to its risk (Kircher & Rösch, 2021). Upon the introduction of the Bias ratio, it became evident that the extraordinary performance in the Madoff case was attributable to fraud. According to the Douady *et al.* (2009), the Bias ratio has proven effective in identifying known fraud cases in the United States (US).

1.2 Motivation

Braun, Riutort and Roche (2024) report that the assets under management (AUM) of US hedge funds more than doubled over a 14-year period from 2008 to 2022. Starting at an estimated \$2.1tn in 2008, AUM increased to \$4.8tn by 2022. This approximately 130% rise in AUM raises concerns, particularly when considering the relatively weak regulatory oversight faced by hedge funds. This underscores the necessity of employing the Bias ratio, given its proven efficacy in ensuring that reported returns are not the result of fraudulent activities. Similarly, in the South African (SA) context, there has been a notable increase in AUM, with a 30% rise from 2021 to 2022 (Association for saving and investment South Africa (ASISA), 2023). While the SA hedge fund market of R113.0bn (≈ \$6.7bn) in 2022 is smaller compared to the US market of \$4.8tn, such significant growth in a single year raises questions about the factors driving this performance. Throughout this study, the Bias ratio will be used to assess whether this abnormally high growth reflects legitimate investment decisions or potential fraudulent practices.

When analysing the consequences of the Madoff case, it becomes clear that addressing fraudulent hedge funds proactively is crucial to preventing large financial losses. According to Homer, Jalain and Hoover (2023), the Madoff case affected approximately 15 000 victims, illustrating its widespread impact across thousands of individuals. The financial consequences were profound, with estimated monetary losses totalling around \$65bn (Quisenberry, 2017). The magnitude of these losses underscores the significant repercussions of hedge fund fraud. Detecting such fraud earlier could have potentially mitigated these losses. The Madoff case serves as just one example among a multitude of instances of fraud within hedge funds. Cumulatively, losses attributed to fraud highlight the importance of employing tools like the Bias ratio to protect investors. Despite the SA hedge fund and stock market being significantly smaller than that of the US, there is still a susceptibility to potential financial losses from fraud. Detecting and preventing fraud through mechanisms like the Bias ratio is essential for enabling informed investment decisions and safeguarding investor interests.

1.3 Research questions

1. Can the Bias ratio detect anomalous activity associated with known historical fraud cases?

2. To what extent do skewness and kurtosis validate the results produced by the Bias ratio?

1.4 Hypothesis

H1: There is a statistically significant difference in the Bias ratio between known fraudulent and conventional investment returns in the US and SA, at a significance level of $p < 0.01$

H2: Skewness and kurtosis contribute significantly to validating the robustness of the Bias ratio in identifying fraudulent versus conventional investment returns, with statistical significance at $p < 0.01$.

1.5 Research objectives

This study aims to contribute to the expanding body of literature surrounding the efficacy of the Bias ratio as a tool for detecting fraud and forecasting asset returns. It seeks to further enrich the literature on financial crime by offering practical methods for fraud detection. Currently, to the best of the authors knowledge, there is lack of robust measures available to investors for assessing the legitimacy of asset returns. The Bias ratio offers a promising solution by enabling the benchmarking of an asset's performance against that of a market benchmark portfolio. This comparative analysis helps identify whether an asset's reported performance may be attributed to illegitimate practices. Beyond fraud detection, the Bias ratio serves to predict expected returns aligned with both the benchmark and the asset's risk profile. It thus empowers investors to make more informed investment decisions based on comprehensive and reliable performance metrics.

1.6 Limitations and Benefits

In addition to its predictive power and capacity to identify fraudulent activity, the Bias ratio does not rely on any assumptions of the efficient market hypothesis. The Bias ratio is adaptable and can be applied across various asset classes. This ratio exclusively focuses on return distributions in relation to a benchmark. This study aims to provide investors with a robust framework for reliably identifying fraudulent activity within different asset classes. However, it is important to note that the Bias ratio may not outperform the Sharpe ratio as a performance metric. Since the Bias ratio is based on historical data, making future inferences could potentially lead to incorrect investment decisions. The results of the Bias ratio should not be extrapolated as this could lead to inaccurate results. The Bias ratio functions as a fraud detection tool rather than a fraud prevention measure. Additionally, due to data availability, there are differences of the periods of each dataset. This possibly reduces the comparability of results amongst each asset class. The remainder of this study is organised as follows: Literature Review, Data & Methodology, Results & Discussion, and Conclusion.

2. Literature Review

2.1 Fraud cases

There have been numerous fraud cases in both the US and SA at a hedge fund and a firm level that emphasise the need for robust fraud detection tools, such as the Bias ratio. Through analysis of both US and SA fraud cases at a hedge fund and a firm specific level it is evident that there is a need for a fraud detection tool. The most notorious fraud case to date is that of Bernie Madoff, who defrauded investors of billions of dollars. Madoff's case stands as one of the largest in history, with total losses amounting to $65bn (Quisenberry, 2017). In this instance, Madoff constructed a Ponzi scheme to sustain the unrealistic returns advertised by the hedge fund. Madoff overstated the hedge fund returns, requiring a continuous introduction of new investors to cover these fabricated returns. As Quisenberry (2017) explains, for a Ponzi scheme to function, a fraudster must continuously attract new investments from a fresh investors pool, which are then redistributed to previous investors as returns.

The collapse of a Ponzi scheme occurs when there is either a lack of new investment or significant redemption requests from investors. It was the latter that led to the downfall of Madoff's hedge fund (Quisenberry, 2017). A redemption request signifies that investors wish to withdraw funds from the hedge fund for various reasons. Redemption calls operate similarly to a bank run; a legitimate fund would possess the liquid capital necessary to fulfil these transactions. In Madoff's situation, investors demanded $7bn worth of redemptions (Quisenberry, 2017). Due to the fraudulent nature of his business, Madoff's hedge fund lacked the capital to execute these transactions, prompting investigations that ultimately uncovered the funds fraudulent activities.

Madoff's case is one of the most severe instances of fraud, necessitating the critical need for effective fraud detection techniques. The need for enhanced detection methods is, however, not the sole lesson to be gleaned from the Madoff case. It also highlights the importance of investors making investment decisions with discernment. In retrospect, the returns that Madoff purportedly "earned" do not align with historical market behaviour. For several years, Madoff's hedge fund substantially outperformed the market, a feat that, in itself, should have raised suspicions. This abnormal performance is a significant red flag. Madoff fabricated auditing firms that published fictitious results, allowing him to evade scrutiny from the Securities and Exchange Commission (SEC) (Quisenberry, 2017).

The application of the Bias ratio would have flagged the abnormal performance of Madoff's returns within the first few months (Douady *et al.*, 2009). This would be done by comparing the Bias ratio of Madoff's returns to that of the broader market. This detection is possible because the Bias ratio evaluates an asset's performance in relation to a benchmark. By doing so, it would have highlighted the inconsistencies in Madoff's reported returns, potentially preventing the extensive fraud that ensued. Use of other statistical methods such as skewness and kurtosis will substantiate the findings of the Bias ratio.

Similarly to the Madoff case, Bayou Management saw Samuel Israel III and Daniel Marino defraud investors of ≈ $450mn (U.S. securities and exchange commission, 2005). Bayou Management exaggerated the fund's profits, despite the fund never actually posting a year-end profit (U.S. securities and exchange commission, 2005). Through the fictitious publishing of results, investors were misled into investing based on misrepresented performance disclosures by Bayou Management. In addition to overstating profits, Marino created a fictitious accounting firm to publish independent audits. Both Marino and Israel illicitly withdrew funds from investors' accounts, and Bayou Management continued to disseminate financial statements to investors that falsely described a profitable hedge fund (U.S. securities and exchange commission, 2005). As previously mentioned, hedge funds operate under less stringent regulatory frameworks, which can potentially make them targets for fraudulent activities.

Corporate scandals also necessitate the need for the Bias ratio, as evidenced by the cases of Enron and WorldCom, which illustrate how fraud can manifest in large corporate firms. Enron, a major energy company based in the US, created Special Purpose Entities (SPEs) to inflate financial results and exercise stock options (Di Miceli da Silveira, 2013). These SPEs were operated by Enron but, due to an accounting loophole, appeared not to be owned by Enron and, therefore, did not appear on the company's balance sheet. Enron then transferred bad assets to the SPEs to keep debt off the balance sheet and present a more profitable appearance (Di Miceli da Silveira, 2013). The SPEs would pay the investment bank for the bad assets, and Enron would subsequently "buy back" these bad assets from the SPEs using the company's stock. The debt thus remained with the

SPEs, which paid the investment bank for the assets while Enron retained the use of the assets (Di Miceli da Silveira, 2013). This method would only remain viable if the stock price continued to perform well. Albeit the Bias ratio would have flagged abnormal stock price performance. This condition is problematic as markets undergo cycles; during a downturn, the SPE method would fail, and Enron's losses would be exposed. Through Enron's creative accounting practices, the company managed to conceal its losses.

A significant lesson from the Enron scandal is the development of the Sarbanes-Oxley (SOX) Act (Di Miceli da Silveira, 2013). The SOX Act is characterised by stricter accounting standards to which firms and hedge funds must adhere (Di Miceli da Silveira, 2013). The SOX Act serves as a method of fraud prevention, distinct from fraud detection. When a firm or individual circumvents or exploits loopholes in fraud prevention laws, fraud detection becomes a crucial second layer of defence.

The WorldCom scandal bears similarities to the corporate scandal involving Enron. WorldCom, a telecommunications company, achieved significant success during the dot-com bubble of 2001 (Granschow, Kulemann, Ripp & Zastrow, 2020). However, due to mismanagement of funds, WorldCom ultimately faced collapse. The company was involved in institutional investments that produced false and misleading equity research. Furthermore, WorldCom's auditors also issued false and misleading audit reports (Granschow *et al*., 2020). Cynthia Cooper, the Vice President of Internal Audit, exposed the fraud, which led to a dramatic decline in the company's stock price, falling from \$65 to \$1 (Granschow *et al.*, 2020). The SEC imposed a penalty of \$2.25bn, including \$0.50bn in cash and a transfer of 10 million shares of common stock (Granschow *et al*., 2020). The WorldCom scandal played a significant role in the development of the SOX. It serves as a notable example of capital mismanagement that necessitated accounting manipulations.

When comparing the financial markets of the US and SA, it is evident that the US market is significantly larger than that of SA (ASISA, 2023) and (Braun *et al.*, 2024). Due to this disparity in size, the high-profile fraud cases in SA are likely to be less prominent compared to major scandals such as those involving Bernie Madoff and Enron. The US is also a developed economy, whereas SA is classified as an emerging economy. Despite its status as an emerging market, SA boasts a sophisticated financial system that can be compared to that of the US (Ojah, Muhanji & Kodongo, 2020).

One of the most notorious cases of hedge fund fraud in SA involves Barry Tannenbaum who defrauded investors of nearly R12bn (Kasipo, 2016). The Tannenbaum Ponzi scheme was marketed as a profitable hedge fund and promised significant returns to early investors, which were subsequently paid out using funds from later investors. Tannenbaum, the grandson of the founder of Adcock Ingram, a prominent pharmaceutical company (Kasipo, 2016), exploited his connection with Adcock Ingram to gain the trust of investors. Tannenbaum promised returns of up to 200% for investments in anti-retrovirals (Kasipo, 2016). The scheme ultimately collapsed when Tannenbaum was unable to meet the obligations promised to investors, mirroring the Madoff case.

Sharemax also orchestrated a Ponzi scheme. Sharemax was an alleged real estate management company, involved in the rental, sale, and management of properties (Woker, 2013). Sharemax defrauded approximately 40 000 investors, amassing a total of R4.5bn (Woker, 2013). The funds were used to acquire empty companies, potentially to perpetuate the Ponzi scheme. Much like the Enron scandal, the Sharemax Ponzi scheme led to a revision of the Companies Act and updates to various regulatory acts governing companies in SA (Woker, 2013).

Steinhoff represents another instance of fraud in SA. The company engaged in accounting manipulation, including the overstating of profits and asset values. Concerns regarding Steinhoff’s financial health emerged in Germany in 2017 (Rossouw & Styan, 2019), prompting an investigation by the company’s auditors who requested additional evidence. The CEO at that time, Markus Jooste, was expected to provide the evidence to the Steinhoff offices in the Netherlands (Rossouw & Styan, 2019). Jooste failed to deliver the evidence and subsequently resigned as CEO, raising suspicions of potentially fraudulent activity. This led to a market reaction resulting in an 80% loss in the stock price (Van Der Linde, 2022). This sharp decrease in the stock price is attributed to stock price inflation caused by the falsification of accounting results. Steinhoff's existence was terminated due to accounting manipulation and fraudulent activities.

Fraud can manifest in various forms, with Ponzi schemes and accounting manipulation being among the most damaging. Ponzi schemes operate in a pyramid-like structure, where funds are required at each round of investment to pay returns to earlier investors (Jory & Perry, 2011). These schemes are destined to fail as the capital requirements increase geometrically, eventually leading to the collapse of the fund (Jory & Perry, 2011).

For example, if a fund manager promises an investor a 50% return on an investment of $100 000, the investor will expect to receive $150 000 at the end of the investment period. To meet this obligation, the manager must attract another investor and promise the same 50% return, this time requiring $150 000 upfront to pay the initial investor. At the end of the subsequent period, the second investor will expect a return of $225 000, and this cycle continues. The scheme ultimately collapses when the fund manager is unable to provide the promised returns, as the required capital becomes unsustainable.

Due to the unsustainable nature of these returns over time, it is crucial to have methods in place to identify fraudulent investments such as Ponzi schemes before they cause significant damage. One such method is the Bias ratio, which can act as an indicator of returns that deviate significantly from a benchmark. While the Bias ratio alone cannot detect fraud definitively, it can suggest potentially fraudulent activity by highlighting outliers in investment returns.

Investors and institutions are not the only victims of Ponzi schemes; feeder funds also play significant roles in these schemes. A feeder fund is a vehicle through which a group of investors pool money and invests in a mutual fund or hedge fund (Spalding, 2012). The impact of a Ponzi scheme collapse extends beyond the immediate financial losses experienced by investors. Employees involved in these schemes are affected not only by job and income loss but also by the reputational damage associated with being linked to a known fraud case, often without the employee's prior knowledge. Auditing and accounting firms face reputational risks, which can tarnish credibility and affect future business prospects. The broader financial markets are also impacted, with the potential to affect entire benchmark portfolios, especially in high-profile fraud cases.

There are also significant tax considerations that arise from these fraud cases. The financial losses sustained by investors and institutions may lead to complex tax implications, including the potential for deductions or write-offs. The repercussions can also impact the tax positions of the entities involved, including any potential liabilities for tax evasion or fraudulent activities. It is essential to address these tax considerations comprehensively to understand the full impact of such fraud cases. Taxpayers may be able to claim a theft loss if the loss can be proven to be a result of a loss from theft rather than being lost or misplaced (Shechtman, Wilensky, & Fusfeld, 2009). Losses incurred from a Ponzi scheme may qualify as a theft loss if it can be proven that the investment manager intended to and did deprive the investor of money through criminal acts (Shechtman *et al.*, 2009). If such evidence is provided, affected individuals may be eligible for reimbursement. This reimbursement is facilitated by the government and can potentially lead to significant financial burdens on government resources.

The key point here is that the consequences of a Ponzi scheme extend beyond the immediate financial losses experienced by investors. The impact of financial crimes is profound and far-reaching, affecting various stakeholders, including employees, auditing and accounting firms, and the broader financial markets. The severity of financial crimes should not be underestimated. It is imperative for investors to have effective tools and mechanisms to identify and protect themselves against fraud.

2.2 Performance metrics

The Bias ratio serves as a performance metric before it can be applied to investigations of financial crimes. Before the invention of the Bias ratio, other metrics were used as performance metrics and fraud detection methods. As a standalone metric, the Bias ratio has a moderate effectiveness in detecting potential fraud. Its utility increases significantly when used in conjunction with other indicators, such as kurtosis and skewness, which can enhance the detection of potential financial crimes. To appreciate the efficacy of alternative methods in identifying potential fraud, it is essential to understand the origins and development of the Bias ratio as a performance metric.

$$Sharpe\ ratio = \frac{R_P - R_f}{\sigma_p} \tag{1}$$

The Sharpe ratio is a widely utilised performance metric in financial markets. R_P is the return of a specific asset or portfolio, R_f is the risk free rate of return. It is calculated by determining the excess return of an asset over the risk-free rate ($R_P - R_f$) and then dividing this by the standard deviation of the excess return (σ_p) (Sharpe, 1966). The Sharpe ratio measures an asset's performance relative to the risk taken. A higher Sharpe ratio indicates that the asset provides a favourable return for each unit of risk. For instance, a Sharpe ratio of 2 means that for each unit of risk, the asset is expected to generate twice as much return. Sharpe ratios < 1 are considered less favourable, as the asset's risk is not adequately compensated by its return.

Wang, Chen, Lian, and Chen (2022) aim to achieve abnormal returns that exceed market performance by constructing portfolios based on high Sharpe ratios. Once these portfolios are constructed, Wang *et al.* (2022) employs high-frequency trading strategies to exploit market opportunities. The study reveals that portfolios with higher Sharpe ratios tend to outperform the benchmark. This finding supports the use of the Sharpe ratio as both a performance metric and an investment strategy, further highlighting its relevance.

Lassance (2022) seeks to maximise the out-of-sample Sharpe ratio to generate returns above the benchmark. The out-of-sample approach is chosen for three reasons: first, mean-variance utility provides a practically useful interpretation as the certainty-equivalent return; second, when including a risk-free asset in the portfolio, all portfolios on the efficient frontier exhibit an equal out-of-sample Sharpe ratio; and third, out-of-sample utility varies between two random variables (Lassance, 2022). Lassance (2022) finds that mean-variance portfolios designed to maximise the out-of-sample Sharpe ratio outperform the respective benchmarks. The Sharpe ratio increased from 0.99 to 1.11 before transaction costs; after considering transaction costs, the Sharpe ratio adjusted to 1.02 (Lassance, 2022).

The Sharpe ratio can be a valuable tool in portfolio construction, providing insights into performance relative to risk. As a performance metric, the Sharpe ratio is a viable option; however, it is not without its limitations, as is the case with any model. There exists a trade-off between simplicity and accuracy, which pertains to the risk measure utilised. The standard deviation of returns may not provide a sufficiently accurate risk measure and might not be the optimal choice compared to a benchmark (Sharpe, 1966). Van Heerden (2020) supports the view that the standard deviation is an inadequate measure of risk. Van Heerden (2020) sought to identify an alternative denominator for the Sharpe ratio suitable for use in a momentum strategy. Through examination of 24 different variations of the Sharpe ratio, Van Heerden (2020) discovered that the Value at Risk

(VaR) Sharpe ratio emerged as the most effective of these variations. The VaR Sharpe ratio is adapted to use the value at risk as the denominator rather than the standard deviation (Van Heerden, 2020). This adaptation not only outperforms the benchmark portfolio but also underscores the significance of active management. Van Heerden (2020) also argues that the choice of Sharpe ratio variant should align with the specific investment strategy employed. Thus as a performance metric the Sharpe ratio is a good proxy for the Bias ratio. However when looking at the fraud detection capabilities of the Sharpe ratio, the Bias ratio is far superior.

The Sharpe ratio is founded on the assumption that the risk-return relationship is positive (Sharpe, 1966), a premise derived from the concept of the efficient frontier. According to this assumption, for every unit of increased risk, investors would anticipate an increased unit of return. The Sharpe ratio serves as a method to quantify this relationship. However, research by Van Rensburg and Robertson (2003) indicates that on the Johannesburg Stock Exchange (JSE), lower levels of risk were associated with higher returns compared to high-risk assets. This finding challenges the traditional notion of a positive risk-return relationship. Auret and Vivian (2014) subsequently retested Van Rensburg and Robertson's (2003) method and found that the value factor indeed persists on the JSE. These contrasting results prompted a further inquiry into whether the risk-return relationship is truly linear and positive.

Van Heerden (2020) delves into this investigation and reveals that the risk-return relationship is, in fact, time-varying. This means that market conditions significantly influence how the risk-return relationship behaves. During periods of low volatility, the risk-return relationship tends to be positive. During high volatility periods, an inverse relationship between risk and return is observed (Van Heerden, 2020). This raises a critical question for investors: How should investors adjust investment strategies during periods of high volatility?

Volatility increases uncertainty, which will affect stock prices. Volatility is not compensated for in the market (Xiong, Idzorek, & Ibbotson, 2014) therefore volatility cannot be built into the price as added risk. This is possibly what investors are expecting during these high volatility periods. Hence the inverted shape relationship. Volatility risk is still included in the riskiness of the investment but not compensated for. Hence the risk of the asset increases but the return remains at the same level.

It is plausible that alternative methods may provide a more accurate explanation of the risk-return relationship and enhance portfolio performance assessment. One such method is the Sortino ratio, which offers a nuanced perspective compared to the Sharpe ratio. While both ratios are designed to measure performance relative to risk, the Sortino ratio addresses a key limitation of the Sharpe ratio, particularly concerning its risk measurement. The Sortino ratio, like the Sharpe ratio, uses a similar numerator, which is the excess return of the asset over the risk-free rate $(R_p - R_f)$ (Sortino & Van Der Meer, 1991). However, the denominator (σ_d) in the Sortino ratio differs as it exclusively incorporates the downside deviation of returns, rather than the total standard deviation (Sortino & Van Der Meer, 1991).

$$Sortino\ ratio = \frac{R_p - R_f}{\sigma_d} \tag{2}$$

This modification reflects a more focused approach to risk assessment by considering only the negative volatility associated with an investment. The Sortino ratio acknowledges the risks inherent in investing by providing a risk-adjusted return metric. A higher Sortino ratio indicates that an investment delivers favourable returns for a lower level of downside risk. Although the interpretation of the Sortino ratio is akin to that of the Sharpe ratio, the primary distinction lies in its exclusive emphasis on downside risk, thus offering a potentially more refined measure of performance.

To evaluate the efficacy of the Sortino ratio as a performance metric, Chen, Yang and Peng (2014) investigated trading strategies for mutual funds based on the Sortino ratio. Like the approach taken by Lassance (2022), Chen *et al.* (2014) then constructed portfolios featuring the highest Sortino ratios and simulated the metrics performance. The study found that the model developed not only mitigated portfolio risk but also enhanced returns. According to the results of Chen *et al.* (2014), investors can achieve excess returns by employing the Sortino ratio, particularly when it is used alongside the Sharpe ratio. This synergy likely arises from the Sortino ratio’s ability to isolate downside risk, differentiating it from the total standard deviation used in the Sharpe ratio.

Despite the Sharpe ratio's widespread use and popularity, the Sortino ratio provides additional insights into an asset's performance by focusing specifically on downside risk. the Sortino ratio can thus offer a more detailed assessment of performance compared to the Sharpe ratio alone.

Srivastava and Mazhar (2018) aimed to compare the effectiveness of the Sharpe ratio and the Sortino ratio in assessing the performance of mutual funds. The findings presented by Sortino and Van Der Meer (1991) clearly indicate that the Sortino ratio outperforms the Sharpe ratio in providing a more nuanced understanding of fund performance. The Sortino ratio is particularly useful for assessing high-volatility portfolios. This does become problematic as the Sortino is more effective during periods of high volatility, but the market does not stay in a constant state of high volatility. Despite this fact, the Sortino ratio is unable to act as a proxy for the Bias ratio in fraud identification.

Alternatives to the Sharpe and Sortino ratios include the Information ratio and the Calmar ratio. The Information ratio, like the Sharpe ratio, uses excess returns as its numerator ($R_p - R_b$), but it differs in its risk measure. Instead of standard deviation, the Information ratio employs tracking error (TE) as the risk measure (Sharpe, 1994).

$$Information\ ratio = \frac{R_p - R_b}{TE} \tag{3}$$

Tracking error is the difference between the portfolio return and the benchmark return. The Information ratio evaluates the return per unit of risk, but focuses on *excess* volatility rather than *total* volatility (Sharpe, 1994).

The Information ratio is particularly valuable for assessing performance in scenarios of high volatility because it measures excess risk relative to excess volatility, making it a suitable substitute for the Sharpe ratio in such conditions. The ability of the Information ratio to account for excess volatility provides a more tailored measure of performance during periods of increased market turbulence (Sharpe, 1994). Like the Sortino ratio the Information ratio is more effective during periods of high volatility. When comparing the Information ratio to the Sharpe ratio, it is evident that the Sharpe ratio is still the best metric for performance evaluation and fraud detection.

The Calmar ratio is another performance analysis metric that offers a different perspective compared to the previously discussed ratios. Unlike the Sharpe and Sortino ratios, the Calmar ratio is based on the annual rate of return of an asset relative to its maximum drawdown (MD) (Sen, 2022). Maximum drawdown is defined as the greatest loss from the peak to the trough of an asset's performance over a specified period. Once again like the previous ratios, the Calmar ratio has excess return in the numerator ($R_p - R_f$) (Sen, 2022).

$$Calmar\ ratio = \frac{R_p - R_f}{MD} \tag{4}$$

Sen (2022) explored the effectiveness of maximising portfolio returns using the Sharpe, Sortino, and Calmar ratios. In this study, portfolios were constructed with the goal of optimising each of these ratios. The findings highlighted that while the Calmar ratio effectively captures periodic and cyclical fluctuations in returns, the Sharpe ratio was still observed to outperform the other two ratios. This indicates that despite its ability to reflect fluctuations, the Sharpe ratio remains a more robust metric for performance assessment in the context of the study (Sen, 2022). When evaluating the Calmar ratio in comparison to the Sortino, Information and Sharpe ratios, it is evident that each ratio has its strengths and weaknesses. However do not act as proxies for the Bias ratio as none of the ratios effectively identify fraudulent activity. For this study the generalisability of the results is critical. This is because fraud detection over a multiperiod sample requires a consistent measure to fairly and accurately evaluate the data. Prior to the use of the Bias ratio, ratio analyses were conducted to identify fraudulent activity and compare results to a benchmark, much like the method of the Bias ratio.

Despite its limitations, the Sharpe ratio has consistently proven to be the most effective measure for portfolio construction in comparison to the other performance metrics. It generally performs best under conditions of low volatility. During periods of higher volatility, alternative ratios such as the Sortino ratio or the Information ratio may be more appropriate. Both the Sharpe ratio and the Sortino ratio have been shown to outperform portfolios evaluated with the Information ratio, which explains why the Sharpe and Sortino ratios are more commonly utilised than the Calmar ratio. When compared to the Bias ratio, the Sharpe ratio underperforms in fraud detection. The Sortino ratio may be preferred over the Information ratio due to its simplicity and ease of use, making it a practical choice for many investors. However, none of these performance ratios serve a dual purpose of functioning both as performance metrics and fraud detection tools, as the Bias ratio does. While the Bias ratio cannot directly identify fraud, it is designed to flag potentially fraudulent returns by highlighting those that are outliers compared to a benchmark. To effectively identify fraud, the Bias ratio should be used in conjunction with additional metrics such as kurtosis and skewness. These supplementary metrics help to provide a clearer picture of potentially fraudulent activity by analysing the distribution and shape of returns.

2.3 The Bias ratio

The Bias ratio can be a valuable tool for flagging potential fraud, as illustrated in studies such as Douady *et al.* (2009). In the context of the Madoff case, the Bias ratio effectively identifies the fraudulent nature of the scheme. Douady *et al.* (2009) showed that while the Sharpe ratio does not raise any red flags for the Bayou and Madoff cases, presenting them as top performers, the Bias ratio distinctly highlights these cases as fraudulent (Douady *et al.*, 2009). To strengthen the results of the Bias ratios detection of fraudulent activity, skewness and kurtosis will be used to accurately identify abnormal activity. Van Vuuren and Van Vuuren (2022), also substantiate the performance of the Bias ratio as a fraud detection tool and performance metric. Although Van Vuuren and Van Vuuren (2022), look at fraud in the context of the US, the study suggests to look at other countries for future work. Govan (2024), looks at the Bias ratio in a SA context and finds that the Bias ratio is able to detect fraudulent activity in markets other than the US.

The Bias ratio theoretically averages ≈ 1, suggesting that any result near this value is considered normal, and will not raise any concerns of fraudulent activity. Any returns that are in the top 1st percentile can be flagged as 'suspicious.' For known fraudulent cases such as Bayou and Madoff, with Bias ratios of 6.0 and 6.5 respectively, which were in the top 0,8% of a sample of 2281 observations (Douady *et al.*, 2009). According to the results from Douady *et al.* (2009), a Bias ratio ≥ 6.0 should be regarded with concern. To further reduce subjectivity in identifying potential fraud, it is advisable to incorporate additional metrics such as skewness and kurtosis. These metrics help to accurately pinpoint outliers by analysing the distribution and shape of returns, thus providing a more comprehensive assessment of potentially fraudulent activity. Skewness and kurtosis can help validate the results obtained from the Bias ratio. This claim is supported by the findings in Van Vuuren and Van Vuuren (2022), in that statistical measures, such as skewness and kurtosis, can help verify the results of the Bias ratio by providing context to the potential fraud cases.

2.4 Alternative flagging method

Through statistical analyses the challenges of identifying fraudulent activity in this study are addressed. To this end, in addition to using the Bias ratio, kurtosis and skewness will be incorporated into the analysis. The rationale behind including these metrics is to conduct a thorough examination of the data and accurately identify outliers. Once the distributions are accurately measured the top percentile of Bias ratio results can accurately be measured and identified as outliers.

Kurtosis measures the tails of a distribution relative to its overall shape, providing insight into the extremity of outlier data (Ketenci *et al.*, 2021). Ketenci *et al.* (2021) explored the application of machine learning to detect money laundering, utilising both kurtosis and skewness as preliminary indicators or flags for identifying financial crimes. Skewness quantifies the asymmetry of the distribution (Ketenci *et al.*, 2021). Huang, Liu, Li, Guo and Chen (2024) emphasise the use of Pearson's coefficient of skewness, which the authors assert is the most significant measure of skewness.

Integrating both kurtosis and skewness into the analysis, will enhance the interpretation of the data and derive precise conclusions from the Bias ratio calculations. Kurtosis, which measures the tails of the distribution, and skewness, which assesses the asymmetry, offer valuable insights into the distributional characteristics of the data. This combined approach will help identify deviations and outliers more accurately.

Both Ketenci *et al.* (2021) and Huang *et al.* (2024) emphasise the significance of outlier detection in the investigation of financial crime, underscoring the necessity of a data analyst's perspective. By incorporating these metrics, the analysis will benefit from a more comprehensive understanding of data anomalies. This systematic approach will facilitate a more meaningful and accurate interpretation of results, ultimately improving the effectiveness of fraud detection efforts.

3. Data & Methodology

3.1 Hypotheses

Hypothesis 1: The Bias ratio of *known* fraud cases will be in the 99^{th} percentile of values. This is because abnormal returns (generated by fraudulent cases) will, by definition, earn more than the market and will be outliers of the broader market.

Hypothesis 2: The skewness and kurtosis of the return series of known fraud cases will be statistically significant from the values expected from a normal distribution at the $p > 0.01$ level.

3.2 Problem statement

Financial fraud has become an increasingly prevalent issue, leading to significant losses for investors, financial institutions, and the economy. Despite advances in regulatory frameworks and auditing practices, many fraudulent activities remain undetected until they have caused irreparable damage. Traditional methods of fraud detection often fail to adequately identify early warning signs. The need for more effective quantitative tools to detect anomalies and potential fraud is critical.

The Bias ratio offers a promising approach by evaluating the skewness and kurtosis of a fund's returns to identify deviations from normal market behaviour. However, its application in real-world scenarios remain underexplored, and its effectiveness in consistently identifying fraudulent activities is not well-established. This thesis seeks to address this gap by investigating how the Bias ratio can be applied to detect financial fraud, thereby offering a more robust method for early identification of suspicious financial practices.

3.3 Data

The study utilises US monthly hedge fund return data from January 1997 to February 2020, which have been sourced from Barclays hedge fund database. Included in this dataset are the data on the Madoff and Steinhoff cases. These cases in particular are selected as they represent the biggest fraud scandals in both the US and SA markets respectively. This sample period is used as hedge fund data are difficult to obtain and these are the only freely available data for a sample period reaching 2024. Because hedge fund data are expensive to acquire, other sources of return data are used to investigate the efficacy of the Bias ratio. The Steinhoff data has been obtained from Bloomberg. Index fund data are then used which are sourced from EDHEC business school, with a sample period of 1997 to 2021. In addition to the hedge and index fund data, monthly stock data for both the US and SA are used. Monthly stock return data are sourced from Bloomberg for the period beginning in January 2000 to September 2024. The sample period for these data stretch to 2024 to make use of the most recently available data; and to ensure that the major fraud cases are included in the dataset. To provide a basis for comparison, the K value, benchmark data need to be introduced. For the US, historical return data for the S&P 500, Dow Jones and Nasdaq were sourced from Iress. The average of the 3 benchmarks were used as this gives the greatest insight in to the average of the market data, as the S&P 500, Dow Jones and Nasdaq are the most successful indices in the US. Similarly for the SA data, the JSE all share index data are sourced from Iress to serve as a benchmark for the study. The reason for choosing the JSE all share index is similar that of the US benchmark. The all share index represents the broader market and is a good proxy for the market benchmark.

Both US and SA data are used as the US can be classified as one of the biggest economies in the world and SA has a highly sophisticated financial market, and one of the largest in Africa. The inclusion of both US and SA fund and company data provides an insight into how fraud operates in both developed and emerging economies. Both active and delisted securities will be included to avoid survivorship bias. Missing data will be replaced with 0 values to avoid breaks in the data. Histograms will be used to check the normality of the data.

3.4 Testing the Bias ratio

To test the Bias ratio, it is essential to identify known fraud cases within the period of the data. The Madoff case will be considered, as this case falls within the period from 2000 to 2009. To evaluate the efficacy of the Bias ratio in identifying fraud, the Bias ratios of the Madoff and Steinhoff cases will be compared to that of the average benchmark Bias ratio. A similar methodology will be applied to the rest of the data, where the Bias ratio of fund and stock market data will be computed and compared to the Bias ratios of the average benchmarks. Thereafter, the relationship between the Bias ratio against both kurtosis and skewness will be evaluated. Skewness and kurtosis only serve as measures to provide context to the distribution of the data, to inform if a higher Bias ratio is justified by the shape of the return data. Lastly, using all the available information, specific funds or stocks will be flagged as potentially fraudulent.

Once the Bias ratio has been computed for the respective data, p-values will be computed to validate the results. In order to compute p-values, the mean and standard deviation of each dataset is extracted, this will then be used to calculate random normal values with the same mean and standard deviation as the sample data. Thereafter, portfolios will be created to match the same number of observations as the sample data. The Bias ratio of the theoretical portfolios will be computed and compared to observed Bias ratio's. The theoretical portfolios serve as return data without fraudulent returns and act as a basis for comparison.

3.5 The Bias ratio

The Bias ratio is defined as:

$$BR = \frac{Count(r_i): r_i \varepsilon [0, +1.0\sigma]}{K + \ Count(r_i): r_i \varepsilon [0, -1.0\sigma]} \tag{5}$$

or

$$BR = \frac{\int_0^{1.0\sigma} rdr}{K + \ \int_{-1.0\sigma}^{0} rdr} \tag{6}$$

where r is the asset return, rdr is the return distribution of that specific asset and sigma is the standard deviation (risk) of rdr between 0 and 1 in the numerator, and -1 and 0 in the denominator (Abdulali, 2006). K is a constant such as a benchmark, in the case of US stock and fund data, the benchmark will be an average of the S&P 500, Dow Jones and Nasdaq. For the SA case, the benchmark will be the JSE all share index. The numerator is the area underneath a histogram in the first interval (Abdulali, 2006). In the numerator, the return distributions of an asset are a function of the asset's standard deviation (risk) in the first period. Similarly, the denominator is the area underneath a histogram in the second interval (Abdulali, 2006). Like the numerator, the return distributions in the denominator are a function of the asset's standard deviation at period 2. Return variations as a function of risk from T_1 are a function of the return variations as a function of risk in T_2 plus a constant.

When applying this model to the above datasets, the constant *K* will need to be defined. Time intervals are defined as 1 month, hence the use of monthly performance data. Each asset's Bias ratio will then be computed using the monthly return data, and scores will be assigned to each respective asset in relation to the market benchmark. It is expected that a Bias ratio of 1 will be observed. This is due to the theoretical premise that the market is a zero sum game. A Bias ratio of 1 signals that there are the same amount of gains as there are losses. Thereafter, kurtosis and skewness scores will be given to each respective dataset. Assets with high Bias ratios in relation to the kurtosis and skewness of the data can therefore be flagged as potentially fraudulent.

4. Results & Discussion

4.1 Hedge fund analysis

Theoretically, the Bias ratio is expected to average around 1.00. Such a Bias ratio corresponds to skewness and kurtosis values of 0. This scenario assumes a perfectly normal return distribution, which inherently precludes the possibility of outperformance. However, in practical contexts, this idealised condition is rarely observed, as return data often exhibit negative skewness. Due to these characteristics of return data the Bias ratio will differ from 1, as seen in Figure 1.

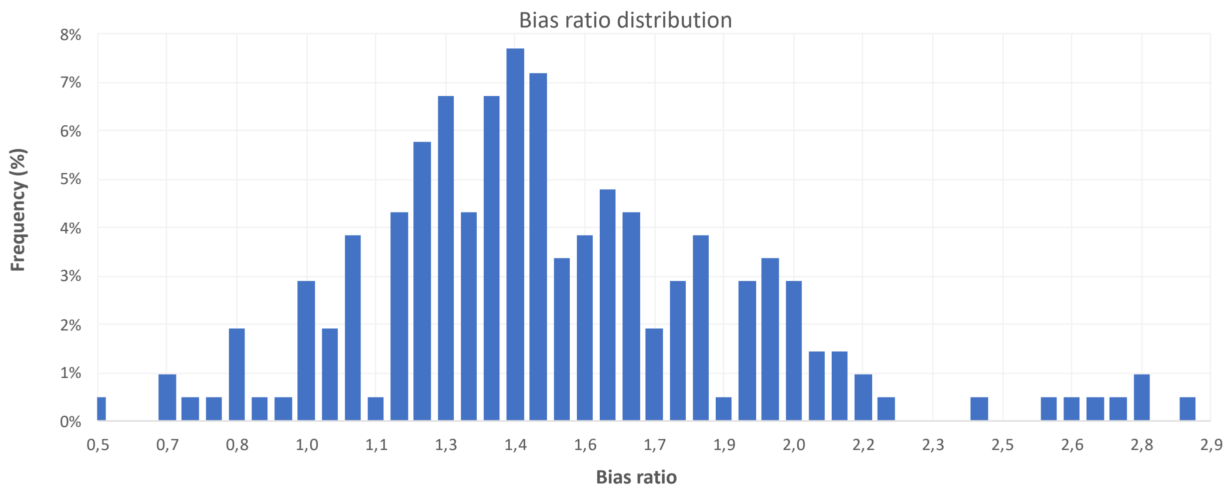

Figure 1: Frequency distribution of the Bias ratio of the US hedge fund data.

The skewness observed in the data results in higher Bias ratios as well as elevated skewness and kurtosis values. As illustrated in Figure 1, most Bias ratios exceed the theoretical value of 1.00. This distribution is anticipated, given that return data are not perfectly normally distributed. The mode of the data is 1.40, as opposed to the theoretical value of 1.00. Over 80% of the data fall within the range of Bias ratios between 1 and 2. The mean Bias ratio of the dataset is 1.46, which is significantly higher than the theoretical value of 1.

To contextualise this average, the data must be compared against an appropriate benchmark. In this study, the benchmark comprises the average Bias ratios of the S&P 500, Dow Jones, and Nasdaq indices, which collectively yield an average value of 1.68. Compared to this benchmark, the average hedge fund Bias ratio of 1.46 indicates an underperformance relative to the broader market. This disparity suggests that the hedge fund sector exhibits a bearish outlook on market conditions.

The maximum Bias ratio in this dataset is 2.80. To determine whether 2.80 constitutes an outlier, it is necessary to compute and evaluate the upper percentiles of the data distribution. Specifically, the upper fifth and tenth percentiles of the hedge fund data are calculated as 2.12 and 1.97, respectively. For meaningful interpretation, these values are compared to the corresponding upper percentiles of the benchmark data, which are both 1.80 for the top five and ten percentiles.

The upper limits of the hedge fund data exceed those of the benchmark data, indicating the presence of outliers. This observation raises an initial concern regarding the abnormally high Bias ratios observed in the hedge fund dataset. To substantiate or dismiss this concern, further analysis is required by comparing the Bias ratio to the skewness of the data as seen in Figure 2.

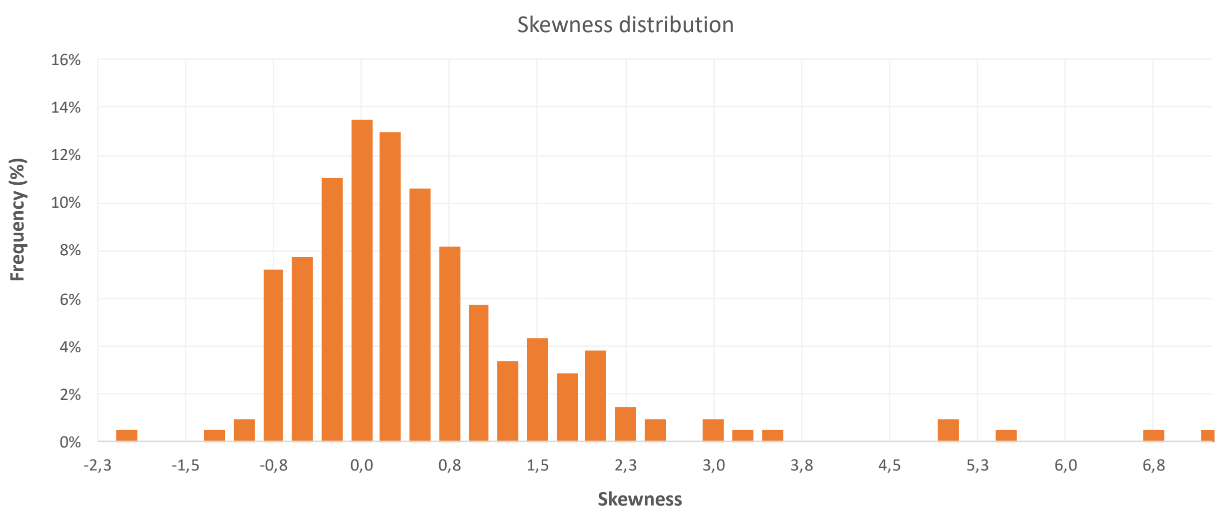

Figure 2: Frequency distribution of the skewness of US hedge fund return data.

It is essential to analyse and understand the skewness and kurtosis of the data. For a distribution to achieve perfect skewness, its skewness value would need to be zero. While the mode of this data is zero, this accounts for only 13% of the data. Roughly 28% of the data are negatively skewed, while the remaining 59% exhibit positive skewness.

The presence of positively skewed data indicates less clustering of values on the positive side of the distribution, resulting in more extreme positive outliers. These outliers are consistent with the observed higher Bias ratios. In contrast, negative skewness reflects clustering of data around small negative values, with most of the negative skewness falling within the range of -2 to 0. Positive skewness spans a wider range, from 0 to 7.50.

The combination of higher positive outliers and clustered small negative values influences the calculation of the Bias ratio. Specifically, it leads to larger positive numerators and smaller denominators, thereby resulting in higher Bias ratios. Greater skewness values, particularly positive skewness, are associated with elevated Bias ratios. Figure 3 shows the high level of kurtosis in the hedge fund data, which is associated with higher Bias ratios.

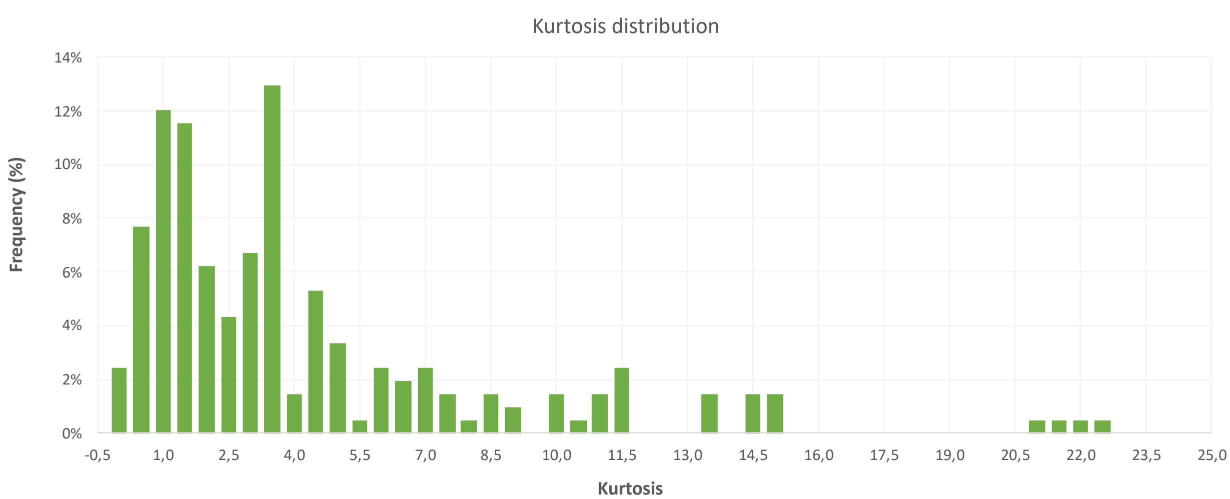

Figure 3: Frequency distribution of the kurtosis of US hedge fund data.

Like skewness, theoretical kurtosis should also be centred around zero to achieve the theoretical Bias ratio of one. However, the hedge fund data exhibit only slightly over 2% of the values at zero kurtosis. The distribution is positively skewed, as evidenced by a mode of 3.50. There are outliers observed between the values of 21 and 23, which reflect a leptokurtic distribution. This type of distribution is characterised by sharper peaks and fatter tails, suggesting higher probabilities of extreme values.

Given this leptokurtic nature, the presence of outliers in the upper percentiles of the Bias ratio is concerning but not unexpected. As with skewness, leptokurtic data increase the likelihood of extreme values, supporting the plausibility of higher Bias ratios. The distribution's characteristics point to larger positive numerators, which in turn result in elevated Bias ratios. Both the skewness and kurtosis analyses converge to support the finding of higher Bias ratios, as the probability of extreme events is amplified with increasing skewness and kurtosis. This relationship is illustrated in Figure 4.

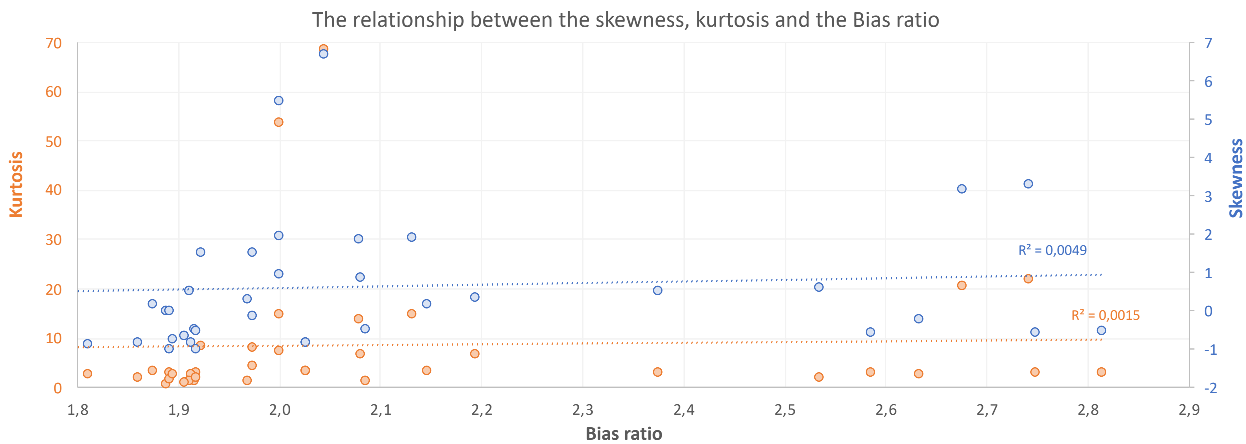

Figure 4: Scatterplot showing the relationship between both skewness and kurtosis against the Bias ratio with a threshold of 1,68.

When the benchmark is set to the average of the US indices (1.68), the model yields low R^2 values, with both the R^2 values for skewness and kurtosis tending towards zero. Specifically, the Bias ratio explains only 0.50% of the variability in skewness, suggesting a weak or negligible relationship between the variables. The Bias ratio accounts for just 0.20% of the variability in kurtosis, further indicating that the relationship between the Bias ratio and kurtosis is minimal. These results imply that both skewness and kurtosis are largely independent of the Bias ratio, thereby allowing for independent conclusions to be drawn from each measure. Had the Bias ratio exhibited a stronger R^2 relationship with skewness and kurtosis, it could potentially influence the result derived from these two variables. For the R^2 relationship to be stronger the benchmark would need to lower as shown in Figure 5.

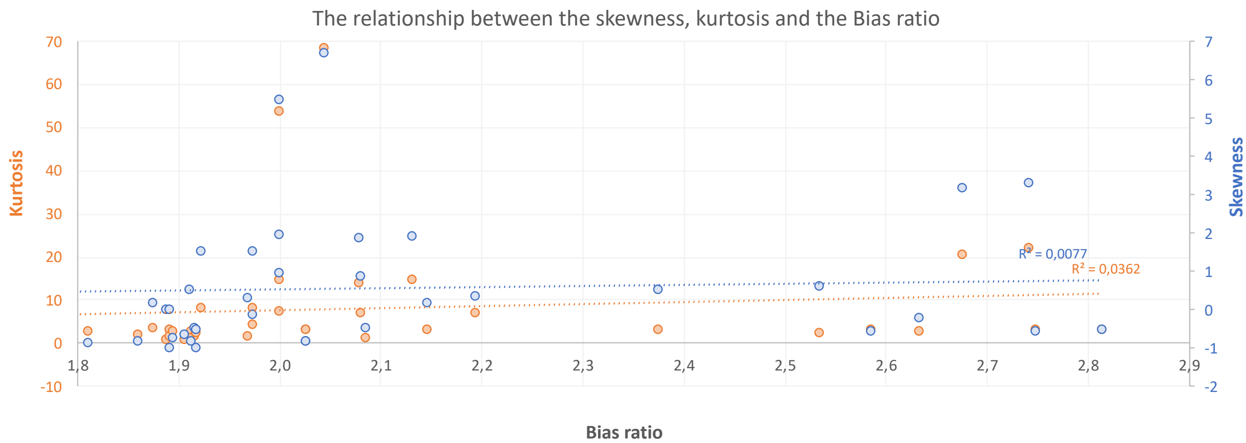

Figure 5: Scatterplot showing the relationship between both skewness and kurtosis against the Bias ratio with a threshold of 1.

The benchmark serves as the threshold for assessing the relationship between skewness, kurtosis, and the Bias ratio. Figure 5 presents the results of adjusting the benchmark to 1. Under this condition, the Bias ratio explains ≈ 1% of the variation in skewness and 4% of the variation in kurtosis. These findings reinforce the results presented in Figure 4, which indicate a minimal to non-existent relationship between skewness, kurtosis, and the Bias ratio. Due to skewness and kurtosis being independent of the Bias ratio, both metrics can be used to help create context of the high Bias ratio. Therefore the results presented above do not flag any fraudulent activity

4.2 Madoff analysis

Figure 6 illustrates an example of a theoretically perfect distribution. However, as previously discussed, such a distribution is idealised and not reflective of real-world return data. In practice, factors such as market sentiment and irrational investment decisions influence the skewness and kurtosis of the return distribution. These factors lead to variations in returns, with an increased likelihood of observing higher Bias ratios. This introduces the challenge of distinguishing legitimate market behaviour from potential fraudulent activity, which may be signalled by unusually elevated Bias ratios.

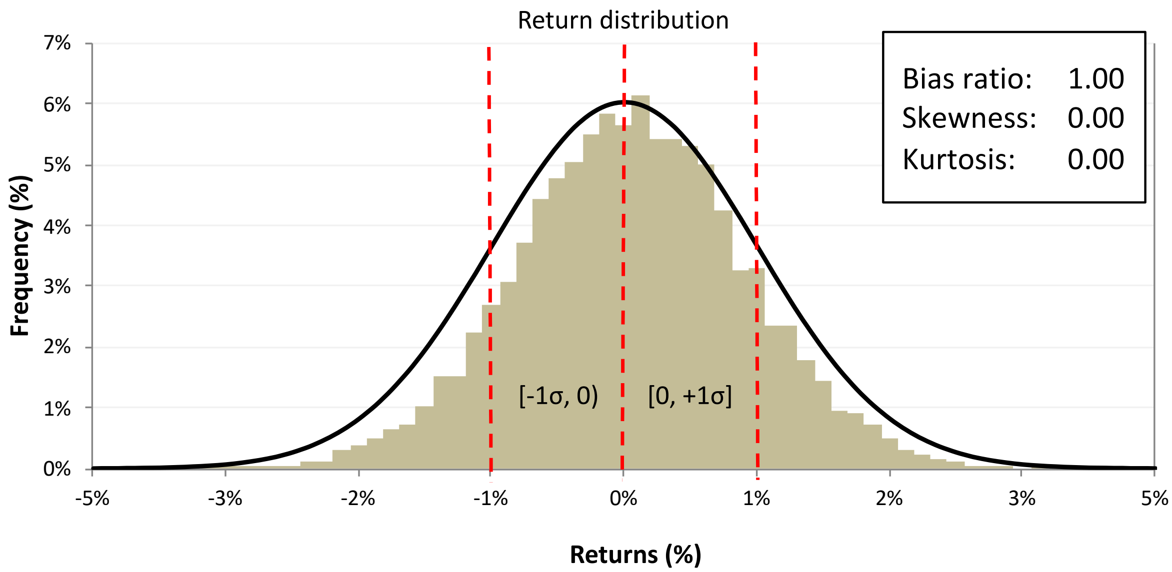

Figure 6: Frequency distribution of a theoretical Bias ratio, skewness and kurtosis.

A more accurate representation of the Bias ratio is presented in Figure 7. The return data exhibit characteristics of positive skewness and leptokurtosis, both of which are consistent with the properties of the hedge fund data. Despite a Bias ratio of ≈ 4, the accompanying data on skewness and kurtosis justify the elevated Bias ratio. The value of 4.08 in the Bias ratio does raise concerns, as it lies within the upper percentile of Bias ratios. This value can be explained by the observed skewness and kurtosis in the data. In contrast, the Madoff case presents a scenario where the supporting data do not sufficiently explain the extremely high Bias ratio, suggesting potential irregularities or fraudulent activity.

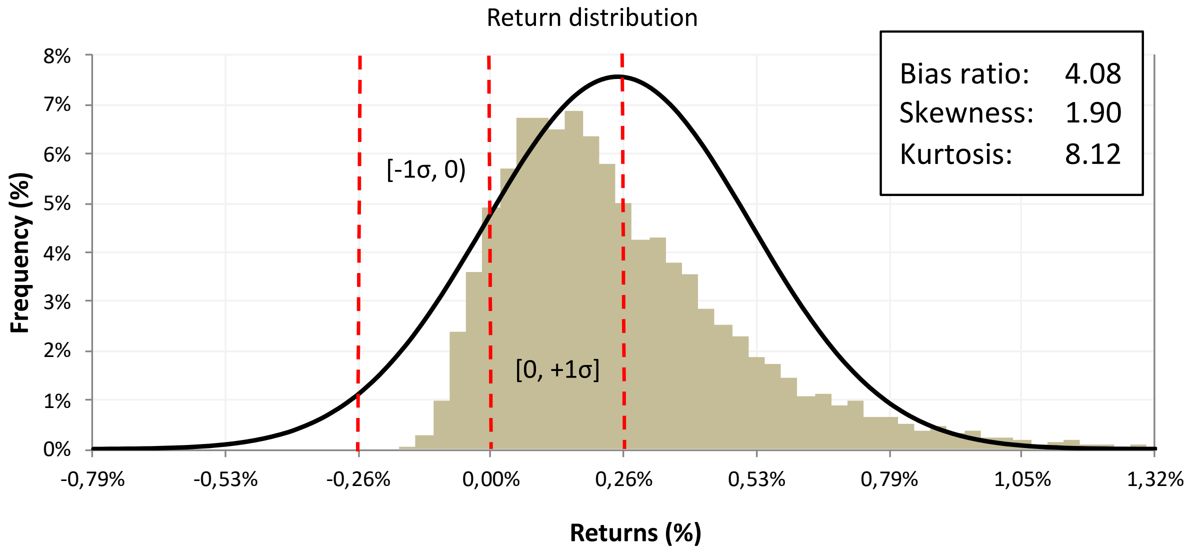

Figure 7: Frequency distribution of a Bias ratio = 4.08, skewness = 1.9 and kurtosis = 8.12.

Figure 8 provides a graphical representation of the Madoff case which is a known fraud case. Madoff's monthly return distribution, which clearly shows that most of the return data are positive. Comparing the skewness and kurtosis values to those in Figure 7, the values in Figure 8 are closer to zero, indicating that the distribution tends towards a normal distribution. This leads to the hypothesis that the Bias ratio should be ≈ 1, or possibly slightly higher. The observed Bias ratio of 8.09 is concerning, especially given that the skewness and kurtosis values are lower than what is typically expected in return data. The Bias ratio of 8.09 is inconsistent with the skewness and kurtosis measures, providing a compelling indication that the returns may be fraudulent. In hindsight, it was revealed that the returns were indeed fraudulent, and the elevated Bias ratio strongly suggested this outcome prior to its discovery. Had the Bias ratio been used in this instance; 15 000 victims could have avoided harm and the $65bn loss could have potentially been avoided.

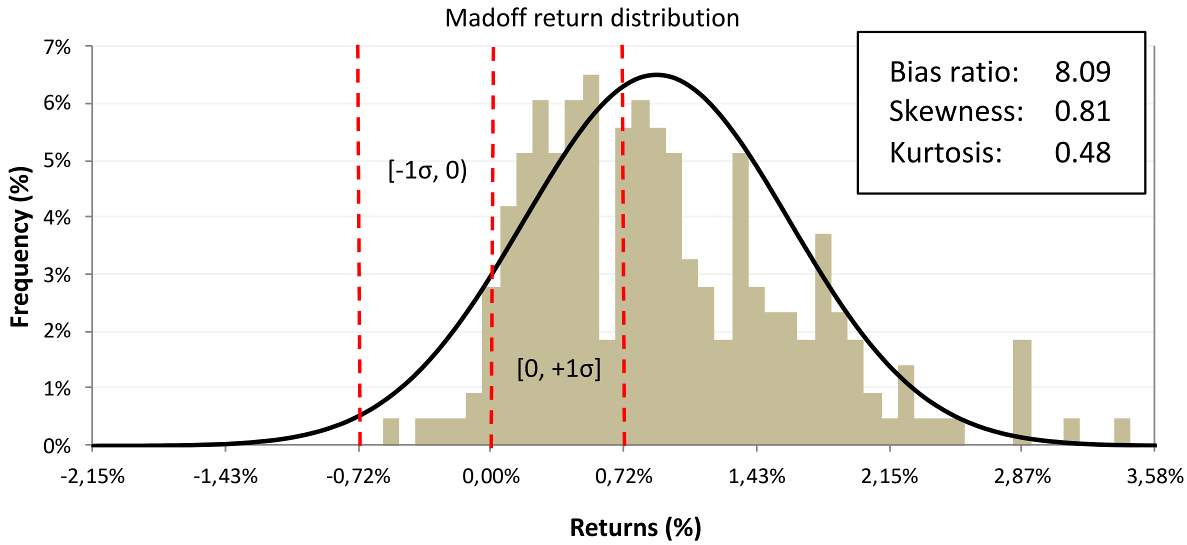

Figure 8: Frequency distribution of the Bias ratio, skewness and kurtosis for a known fraud case (the Bernie Madoff case).

4.3 The Steinhoff case

In order to verify the efficacy of the Bias ratio, the Steinhoff case will be analysed. In contrast to the US fraud case, Steinhoff was one of the largest fraud cases in SA. Figure 9 shows the return distribution of the Steinhoff shares. Surprisingly, there is a low Bias ratio of 0.92. This result does not particularly raise any concerns. In line with what was observed in the Madoff case, the Steinhoff case has a skewness value of 1.48. This means that the data is fairly symmetrical. Interestingly, the kurtosis value of the Steinhoff case is 12.64 suggests that the data is peaked with narrow tails. When considering the distribution of returns it is unlikely that there will be an observed extreme value, hence the low Bias ratio. In contrast to what the observed Bias ratio suggests, more than 65% of the returns are positive.

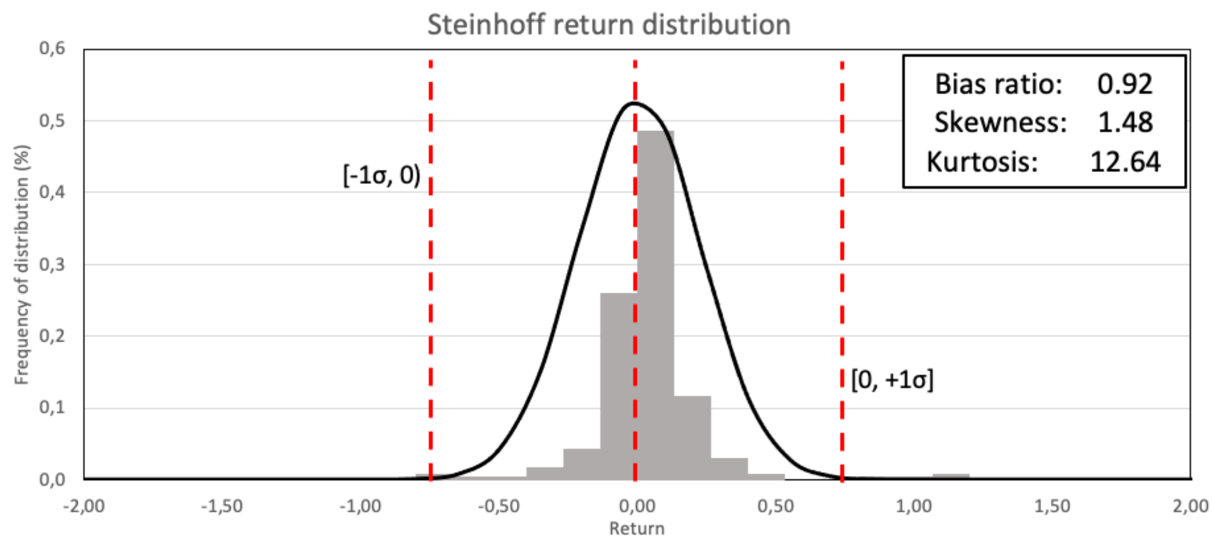

Figure 9: Frequency distribution of the Bias ratio, skewness and kurtosis for a known fraud case (the Steinhoff case).

Seemingly the Bias ratio did not detect the fraudulent activity of the Steinhoff case, however this is not the case. Steinhoff was active and listed for 10 years before the fraud occurred, these returns are taken into consideration when calculating the Bias ratio over the life of the stock. A further argument can ensue that the Bias ratio detected fraud in the Madoff case. This is likely due to the fact that the fraud had been fraudulent for many years. Another. Interesting caveat is that due to Steinhoff being listed, the stock price would adjust as information got released. The difference in Bias ratios between the Madoff and Steinhoff cases is likely due to the type of asset. Steinhoff being a listed stock, is subject to the information distribution of the market.

4.4 Index fund analysis

The inclusion of index fund data introduces an additional dimension to the study, owing to the distinctive characteristics of index funds. Unlike hedge funds, index funds typically require a smaller initial investment, which increases the accessibility of the instrument. This accessibility is likely to result in a higher trading volume compared to the hedge fund counterparts. However, the number of available index funds in the US is significantly smaller than the number of hedge funds.

Figure 10 shows the index fund data which comprise of only 13 indices, which is meagre compared to the hedge fund dataset, comprising 208 funds.

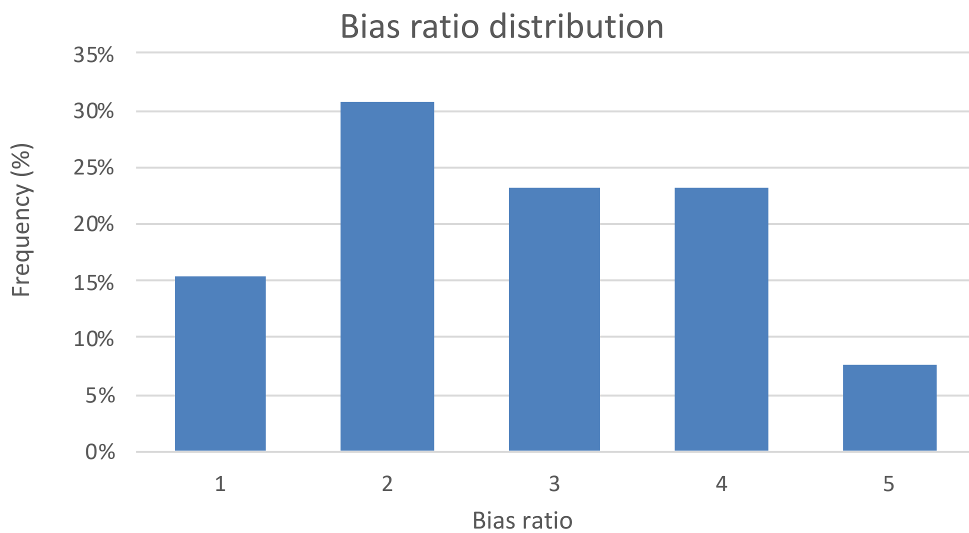

Figure 10: Frequency distribution of the Bias ratio in the index fund data.

The index fund data exhibit relatively high Bias ratios, with the majority falling within the range of 1 to 4. In isolation, such data could be cause for concern. However, as observed with the hedge fund data, higher Bias ratios are often substantiated by the underlying data characteristics. The maximum Bias ratio in this dataset is 4.23, with the upper fifth and tenth percentiles reporting values of 4.08 and 3.83, respectively. When compared to the benchmark, it is evident that the upper percentiles of the index fund data are significantly more extreme than the market percentiles of 1.80. The average Bias ratio of 2.37 exceeds that of the benchmark, indicating a bullish outlook among index fund indices regarding the future state of the market. The context of the index fund data support the elevated Bias ratios observed. To further validate these findings, the skewness of the data, as shown in Figure 11, must be analysed to corroborate the high Bias ratios.

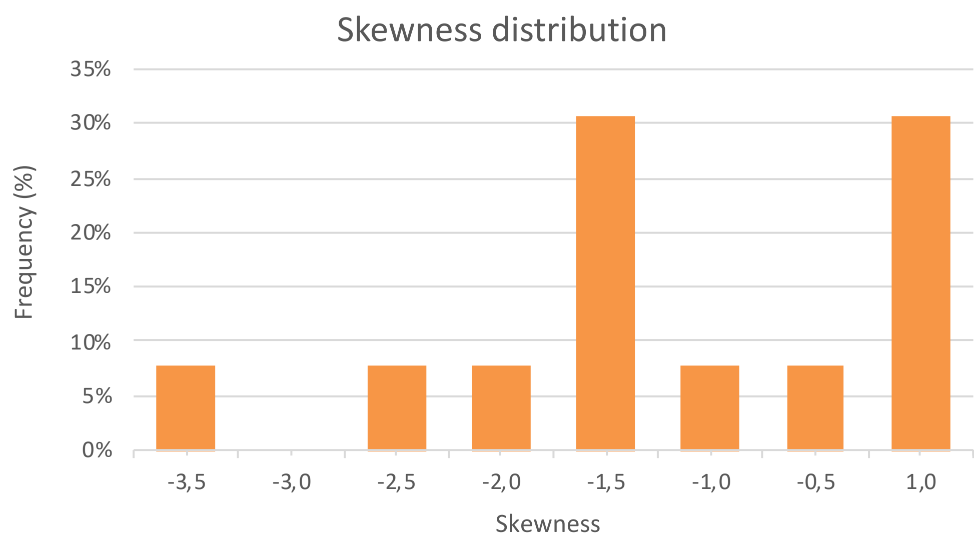

Figure 11: Frequency distribution of the skewness of the return data in the index fund data.

An intriguing finding is that the index fund data exhibit greater negative skewness. In contrast to the hedge fund data, which demonstrated positive skewness. The index fund data are characterised by a negative skew. This negative skewness results in more dispersed negative returns and a clustering of smaller positive returns. The Bias ratio then tends to be lower, as smaller, clustered positive returns lead to reduced numerators, while larger, more dispersed negative returns result in larger denominators. As a result, the skewness component of the data do not align with the observed high Bias ratios, which raises concerns regarding the underlying data characteristics. The kurtosis distribution, as shown in Figure 12, which may offer additional insights to help explain the elevated Bias ratios within the data.

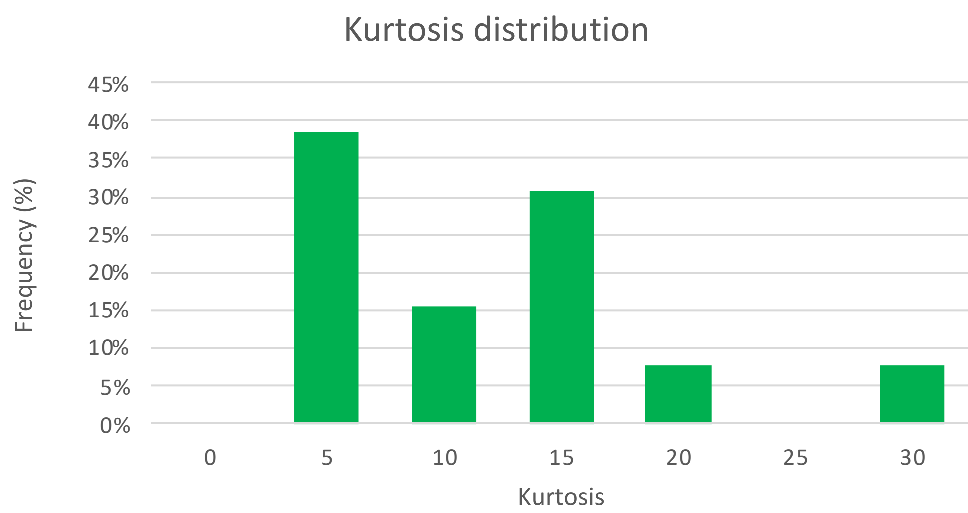

Figure 12: Frequency distribution of the kurtosis of returns in the index fund data.

The kurtosis distribution supports the findings of larger Bias ratios. The kurtosis values in the data are significantly elevated and deviate substantially from zero, indicating a distribution with pronounced fatter tails. These fatter tails correspond to a higher probability of extreme returns, which in turn increases the likelihood of a larger numerator in the Bias ratio calculation. The kurtosis measure substantiates the observed higher Bias ratios in the data. To further validate these findings, it is necessary to analyse the R^2 values.

Interestingly, the R^2 analysis of the index fund data reveal no significant relationship between the Bias ratio and skewness. The analysis, conducted using a benchmark threshold value of 1.68, indicates that the Bias ratio is highly unlikely to influence the skewness of the return data. In contrast, the R^2 value for kurtosis is 77%, signifying a strong relationship between the Bias ratio and kurtosis. This strong correlation likely accounts for the higher Bias ratios observed in the data. This relationship does not raise concerns, as it provides a clear explanation for the variability in the results. Based on this analysis, there is insufficient evidence to substantiate any meaningful concerns of fraudulent activity regarding the elevated Bias ratios in this data. Similarly to the hedge fund data, there was no fraudulent activity flagged in this data.

4.5 US stock analysis

Delving into the analysis of the US stock market adds an additional layer to the study. While both hedge funds and index funds are primarily focused on indices, analysis of individual securities provides valuable insight into how the Bias ratio behaves at the granular level. The stock market is characterised by greater liquidity, accessibility, and collectively higher trading volumes compared to hedge and index funds. Due to these characteristics, most Bias ratios for individual securities are observed to be greater than 1, as illustrated in Figure 13.

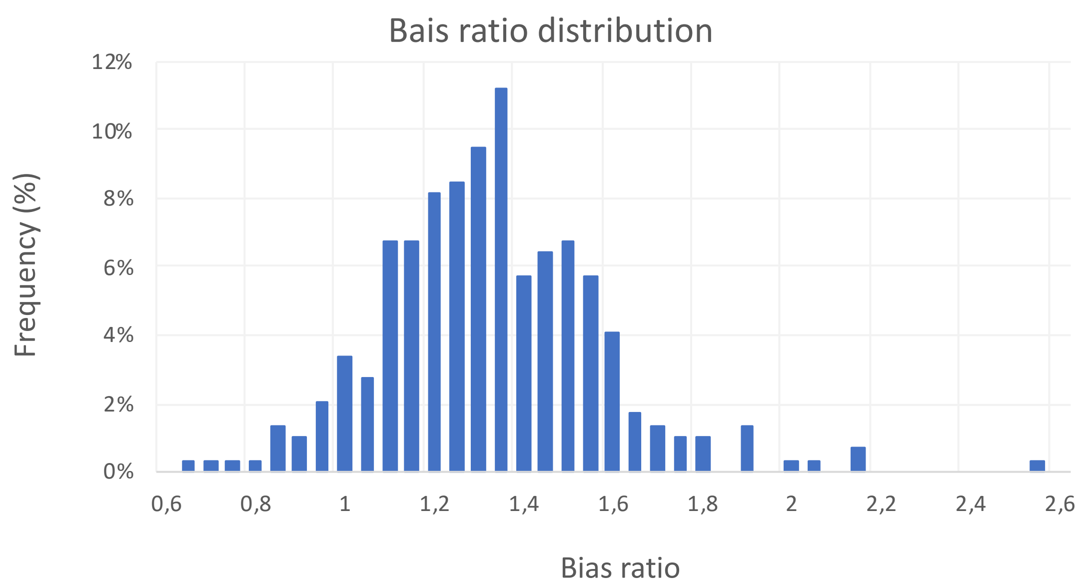

Figure 13: Frequency distribution of the Bias ratio in the US stock market data.

The Bias ratio distribution of the US stock market data closely resembles that of hedge fund returns, with a mode of 1.35. In contrast to the index fund data, the US stock market data exhibit a maximum Bias ratio of 2.55. When comparing the upper fifth and tenth percentiles of the US stock market data to the benchmark, the results are consistent. The upper fifth and tenth percentiles of the US stock market data are 1.70 and 1.57, respectively, both of which are below the benchmark value of 1.80. Furthermore, the average Bias ratio of the US stock market data is 1.30, compared to the benchmark average of 1.68. This comparison suggests that the US stock market reflects a bearish outlook on the future state of the market. The observed maximum Bias ratio of 2.55 does not raise significant concerns. Figure 14 illustrates the skewness of the US stock market return data, to substantiate the level of Bias ratio observed.

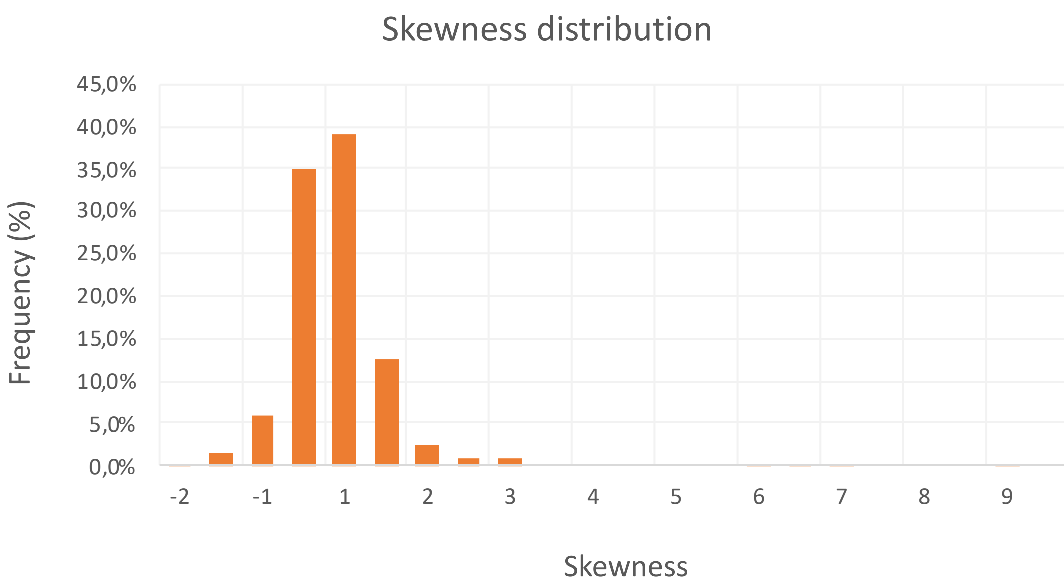

Figure 14: Frequency distribution of the skewness of the return data in the US stock market data.

As anticipated, the skewness measures are ≈ 1, indicating minimal positive skewness in the data. This slight positive skewness contributes to a modest increase in the Bias ratio. Unlike the hedge fund and index fund data, the US stock market data exhibit only mild positive skewness, which align with the observed average Bias ratio values. The lower skewness values suggest that the return distribution of the US stock market data are nearly symmetrical. This symmetry results in a relatively balanced relationship between the numerator and denominator of the Bias ratio, leading to lower Bias ratio values. The kurtosis distribution, as shown in Figure 15, can provide further substantiation for these findings, offering additional insight into why the Bias ratio for the US stock market data are lower compared to that of hedge and index fund data.

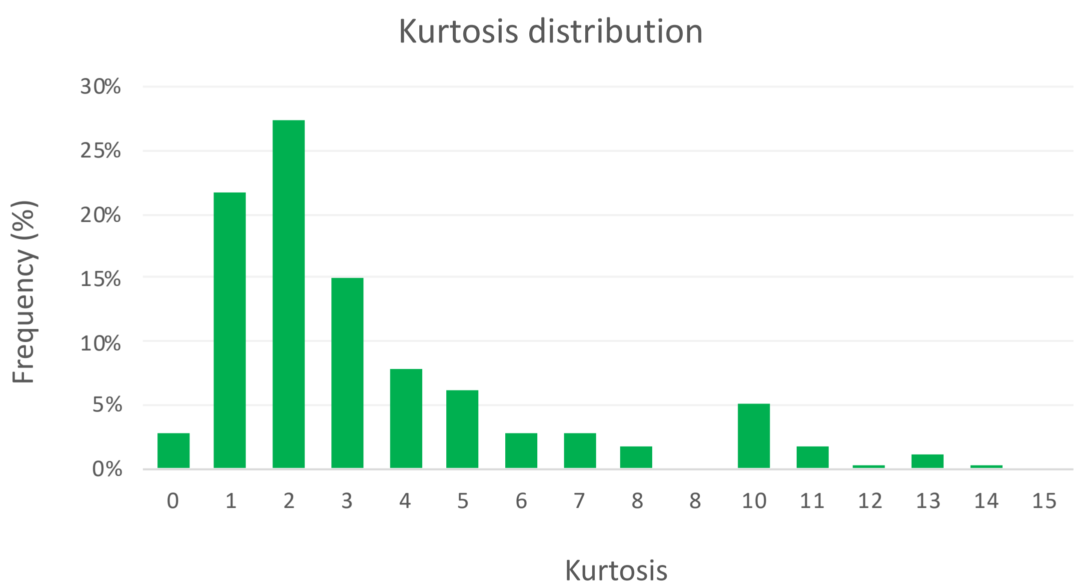

Figure 15: Frequency distribution of the kurtosis of returns in the US stock market data.

Most of the kurtosis values in the US stock market data are $1 \leq \kappa \leq 4$, indicating relatively low kurtosis. Compared to the kurtosis values observed in the hedge and index fund data, the US stock market kurtosis values are significantly smaller. This reflects narrower tails in the return distribution, which reduce the probability of extreme returns. This lower probability of extreme values in the numerator leads to a smaller average Bias ratio. To evaluate whether skewness or kurtosis influences the Bias ratio, it is essential to analyse the R^2 values.

Like the index fund results, skewness shows no significant relationship with the Bias ratio. The R^2 value is influenced by the selected benchmark threshold. A smaller threshold would likely increase the correlation between skewness and the Bias ratio. Kurtosis demonstrates a strong relationship with the Bias ratio, evidenced by an R^2 value of 80%. These results closely resemble those observed in the index fund dataset, where the variability in the Bias ratio is largely explained by kurtosis. Hence, when considering all statistical measures observed, it is apparent that there is no fraudulent activity in this data. This suggests that kurtosis plays a significant role in predicting the Bias ratio. Due to the lower kurtosis values observed in the US stock market data, the Bias ratio is correspondingly lower when compared to other datasets, further substantiating the findings.

4.6 SA stock analysis

The Bias ratio has been analysed within various contexts of the US market. To comprehensively understand its functionality, it is essential to evaluate the efficacy of the Bias ratio in an emerging market such as SA. One of the strengths of the Bias ratio lies in its universal applicability, meaning that, regardless of the data or market in question, the Bias ratio can be easily interpreted and compared. As illustrated in Figure 16, the distribution of the Bias ratio for the SA stock market closely mirrors that of the US instruments, suggesting similar underlying dynamics.

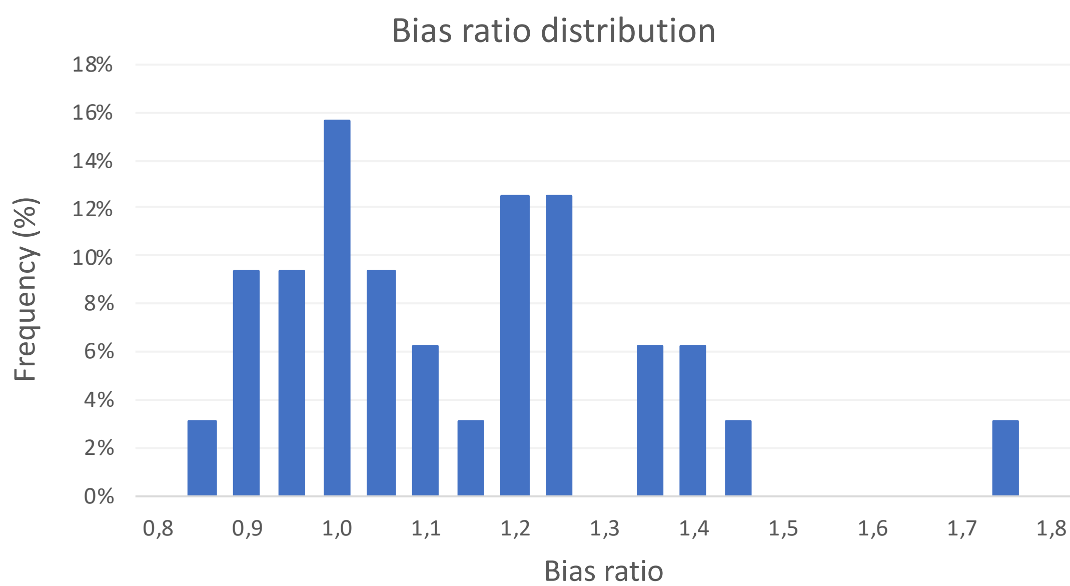

Figure 16: Frequency distribution of the Bias ratio in the SA stock market data.

Like the US stock market, the SA stock market data do not raise significant concerns. The maximum Bias ratio in the SA stock market data is 1.75, which is substantially lower than the 8.09 observed in the Madoff case. Given that the data pertains to the SA market, it would be inappropriate to use the same benchmark applied to the US data. As a result, the JSE all share index has been used as a reference for comparison as the index is closely linked to market movements.

The Bias ratio for the benchmark is 1.13, with the upper fifth and tenth percentiles recorded at 1.80. The average Bias ratio for the SA stock market data is 1.12, which is lower than the benchmark average. The fifth and tenth percentiles for the SA stock market data are 1.41 and 1.37, respectively. These initial results suggest that the SA market reflects a somewhat bearish outlook on future market conditions. While the Bias ratio results do not present any cause for concern, it is still necessary to examine the skewness and kurtosis of the data. Figure 17 provides an illustration of the skewness in the SA stock market data.

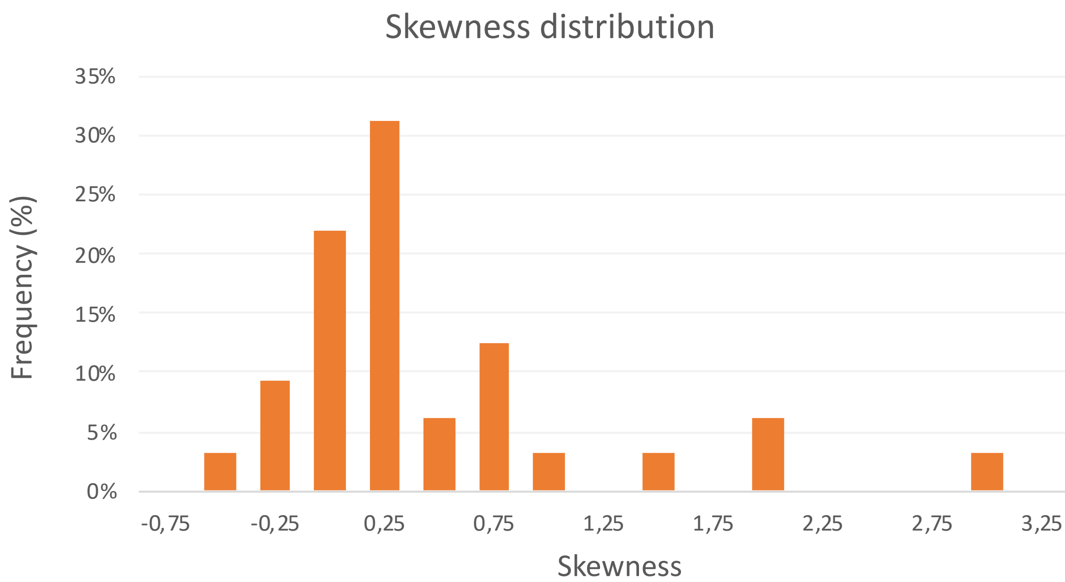

Figure 17: Frequency distribution of the skewness of returns in the SA stock market data.

Skewness in SA stock market data cluster around zero, like the skewness patterns observed in the US stock market data. The presence of mild positive skewness in the data may explain why the Bias ratio deviates slightly from 1. Low skewness values suggest that the data are symmetrical: differences between the numerator and denominator of the Bias ratio are smaller compared to the hedge and index fund datasets.

This symmetry helps explain why the Bias ratio for the SA stock market is lower than those observed in the hedge and index fund data. These results are consistent with the findings from the US stock market data. Given this, it is reasonable to expect that the kurtosis values for the SA stock market data will be smaller than those observed in the hedge and index fund data. Figure 18 confirms this hypothesis, demonstrating that the kurtosis values are relatively modest, with most of the data exhibiting kurtosis levels < 4.

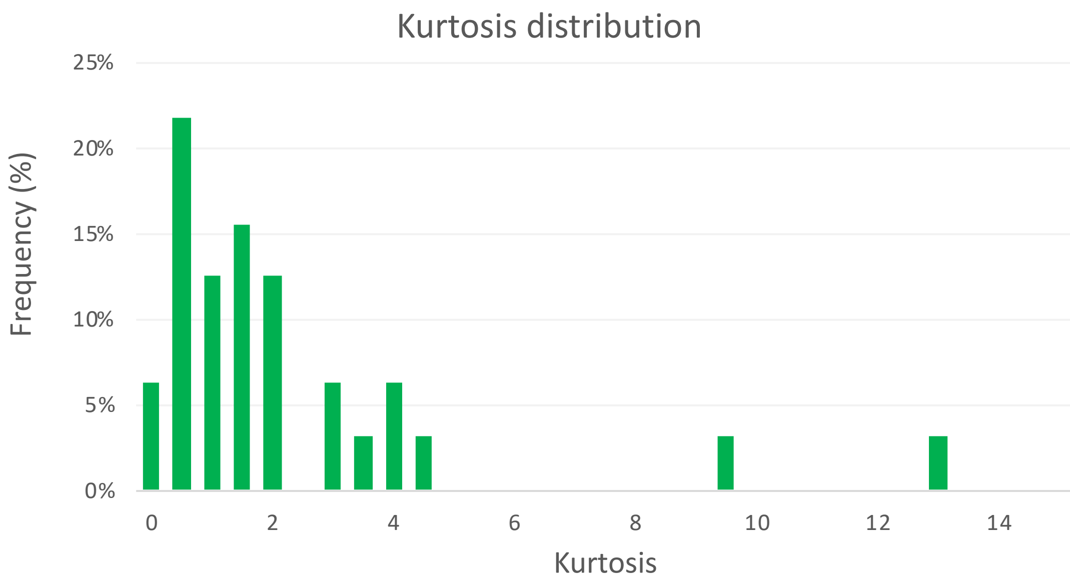

Figure 18: Frequency distribution of the kurtosis of returns in the SA stock market data.

Lower kurtosis values suggest that the return distribution in the SA stock market data exhibit a shallower peak and narrower tails. These findings are consistent with the patterns observed in the US stock market data. It is important to note that stock market data and index fund data exhibit distinct distributions, with stock market data typically characterised by lower skewness and kurtosis scores, whereas index fund data tend to show higher skewness and kurtosis values. Consequently, the Bias ratio is smaller (larger) for stock market (index) data. Kurtosis demonstrates explanatory power over the Bias ratio in both index fund and stock market data. Supporting these observations, the R^2 value indicates a positive relationship between kurtosis and the Bias ratio in the SA stock market data, with an R^2 value of 46%.

At the benchmark level of 1.13, there is no significant relationship between skewness and the Bias ratio. This result aligns with expectations based on the analysis conducted. Further, the R^2 analysis reveals that the degree of kurtosis has a significant impact on the Bias ratio when kurtosis is low. However, as kurtosis increases, the explanatory power of kurtosis over the Bias ratio diminishes, as observed in the hedge fund data. As seen above when kurtosis is able to explain the Bias ratio in context, higher levels of Bias ratios are not concerning. Hence there is not fraudulent activity observed in the above dataset.

4.7 Summary of results

When looking at the US market as a whole there was no indication of fraudulent activity. That is across the hedge fund, index fund and stock market data. In order to ensure the Bias ratio does flag fraud sufficiently both Madoff and Steinhoff cases were analysed. In both cases it is clear to see that the Bias ratio does indeed flag fraudulent activity. Similar to the US markets, the SA stock market did not flag any fraudulent activity. The results presented do indeed substantiate H1, in the sense that the Bias ratio does effectively identify fraudulent activity. In order to substantiate these findings, statistical measures such as skewness and kurtosis were introduced. Both measures helped provide context on the shape and distribution of the data. This is important as it will help contextualise the Bias ratio and substantiate the results. Once the correlation between both skewness and kurtosis to the Bias ratio have been evaluated results can be verified. Both skewness and kurtosis do substantiate the results in all of the observed datasets. Hence H2 can be substantiated. An additional function of skewness and kurtosis in validating the results is that, these statistical measures help eliminate false positives and negatives. The particular distribution of the results leads the user of the Bias ratio to expect a certain level of the Bias ratio.

Interestingly, the results for both the US and SA stock markets yield similar results. The Bias ratios for both markets fall in a similar band and can both be explained through the kurtosis distribution. The Bias ratio can be universally adapted and compared. However, Bias ratios across asset classes differ. This is likely due to the systematic risks and characteristics of these markets. For example, the hedge fun market is less regulated than the stock market, and the stock market adjusts quicker to new information than hedge funds. A result of this is that hedge funds will have inherently higher Bias ratios than the stock market. This is substantiated by the comparison between the Madoff and Steinhoff Bias ratios.

5. Conclusions, limitations and future work

5.1 Conclusions

There is no universally acceptable value for the Bias ratio. Rather, its validity is contingent upon the specific instrument, as well as the skewness and kurtosis characteristics of the data. Higher Bias ratios associated with p-values > 0.01, can be used to identify potentially fraudulent performance. Through the use of the Bias ratio, p-values, skewness and kurtosis, it is possible to identify potentially fraudulent activities. Bias ratios however need to be considered in the context of the asset class and benchmark and specific asset classes are susceptible to higher outperformance of the benchmark. Thus the inclusion of p-values assists in verifying that the observed Bias ratio's are indeed statistically significant. Hedge funds, on average, exhibit higher Bias ratios, which can present challenges when assessing the risk of fraudulent activity. Index funds display more varied Bias ratios, with the potential for higher values under certain conditions. Stock market data, both from the US and SA, tend to have lower Bias ratios, which reduces the likelihood of fraud in these contexts.

Practical implications of the uses of the Bias ratio are twofold. The Bias ratio can be used as a performance metric and as a fraud detection tool. Investors can use the Bias to evaluate the returns of a stock or index. The use of the Bias ratio is easy to interpret as the higher the Bias ratio the better the stock is performing. This however cannot be used in isolation and will need a benchmark to give context into the actual performance of the stock. A Bias ratio >1 gives insight similar to the Sharpe ratio. What this means is that there are more positive returns than negative. If the Bias ratio gets higher this then means that there are more larger positive values than negative. As a fraud detection tool the Bias ratio can be used by investors to reliably detect fraud at a p-value 0.01 level. As the stock market is a zero sum game, there is an expected asymmetry of returns. Hence the Bias ratio is expected to be 0. Due to the nature of the stock market and the abnormal returns, a Bias ratio of 1 is thus realistically expected. In terms of fraud detection, the higher the Bias ratio the more cause for concern. This Bias ratio however cannot be used in isolation when detecting fraud. In the statistical context of returns, a high Bias ratio may actually be correct. Hence other data analysis metrics are needed to validate the Bias ratio.

The presence of higher Bias ratios in hedge and index funds warrants concern, but this issue can be mitigated through further analysis of skewness and kurtosis. Skewed data and high kurtosis are associated with higher Bias ratios, which can, in turn, alleviate initial concerns regarding fraudulent activity. When comparing the Madoff case to the hedge fund data, it is evident that skewness and kurtosis play a crucial role in substantiating the likelihood of fraud. The Bias ratio exhibits a strong correlation with kurtosis at lower levels of kurtosis, though this correlation weakens as kurtosis increases. By incorporating supporting metrics such as skewness and kurtosis, it becomes feasible to verify and assess potential fraudulent activity.

5.2 Limitations & Future Work

The study is subject to several limitations, primarily arising from the availability of data. As a result, the timeframes of the datasets used do not align perfectly. The hedge fund data utilised in this study span January 2000 to December 2011. These data are particularly difficult to obtain, and they represent the only available hedge fund data for the purposes of this thesis. In contrast, the index fund data cover the period from 1997 to 2021. The timeframes for the US and SA stock market data extend from January 2000 to September 2024. Benchmark data also present timing discrepancies. The S&P 500 data cover the period from January 2000 to April 2016, while the Dow Jones data span August 2004 to December 2024. The study faced the limitation of a lack of SA index or hedge fund data for comparison to the US hedge and index fund data.

Future research could address these limitations by acquiring more recent hedge fund data for testing the Bias ratio. The inclusion of SA hedge and index fund data would provide a more robust comparative basis, enhancing the study's ability to contrast SA and US market dynamics. The study would also benefit from exploring different markets across different countries. Similarly, investigating different sectors in a single country may help gain a greater understanding of how inter-sector behaviours. Another potential avenue for future research would involve exploring the efficacy of the Bias ratio as a performance metric.

7. References

Abdulali, A. (2006). The Bias ratio: Measuring the shape of fraud. *Protégé Partners Quarterly Letter, 3rd Quarter.*

Association for saving and investment South Africa (ASISA). (2023, 27 February). *Hedge fund industry sees assets under management jump by an unprecedented 30% in 2022* [Press release].

Auret, C. & Vivian, R. (2014). A comparative analysis of returns of various financial asset classes in South Africa: a triumph of bonds? *South African business review, 18*(3), 175-195.

Blythe, S. E. (2020). Financial statement fraud: Lessons learned from selected U.S. legal cases in the past twenty years. *Journal of modern accounting and auditing, 16*(1), 1-18.

Braun, M., Riutort, J., & Roche, H. (2024). Hedge fund fee structure and risk exposure. *Economic modelling, 132*.

Chen, H., Yang, C., & Peng, Y. (2014). The trading on the mutual funds by gene expression programming with Sortino ratio. *Applied software computing, 15*(1), 219-230.

Di Miceli da Silveira, A. (2013). The Enron scandal a decade later: Lessons learned? *Homo oeconomicus, 30*(3), 315-347.

Douady, R., Abdulali, A., & Adlerberg, I. (2009). The Madoff case: Quantitative bests qualitative. Retrieved 11 Aug. 2024 from http://www.riskdata.com/resources/hedge_funds_risk_and_replication.html.

Felixson, K. & Pelli, A. (1999). Day end returns-stock price manipulation. Journal of multinational financial management, 9(2), 95-127.

Fung, W. & Hsieh, D. A. (1999). A primer on hedge funds. *Journal of empirical finance, 6*(3), 309-331.

Govan. B. (2024). An analysis of fraud detection using Benford's law and the Bias Ratio (publication No. 31979320) [Doctoral dissertation, University of Witwatersrand]. ProQuest Dissertations & Theses Global.

Granschow, K., Kulemann, L., Ripp, F., & Zastrow, L. (2020). What do we know about the financial consequences of the WorldCom accounting scandal. *Management studies, 10*(2), 39-49.

Homer, E. M., Jalain, C. I., & Hoover, K. B. (2023). Hearing from the forgotten victims: A content analysis of the consequences of Bernard L. Madoff's ponzi scheme. *International journal of evidence-based research, policy and practice, 18*(7), 1335-1353.

Huang, Y., Liu, W., Li, S., Guo, Y., & Chen, W. (2024). Interpretable Single-dimension Outlier Detection (ISOD): An unsupervised outlier detection method based on quantiles and skewness coefficients. *Applied sciences, 14*(1), 1-11.

Jory, S. R. & Perry, M. J. (2011). Ponzi schemes: A critical analysis. *SSRN Electronic Journal*, Available at SSRN: https://ssrn.com/abstract=1894206.

Kasipo, M. (2016). Governing white-collar crime in South Africa: The challenges of a state-centric approach. *AML Journal of Africa, 2*(3), 5-14.

Ketenci, U. G., Kurt, T., Önal, S., Erbil, C., Aktürkoglu, S., & Ilhan, H. S. (2021). A Time-Frequency Based Suspicious Activity Detection for Anti-Money Laundering. *IEEE Access, 9*(1), 59957-59967.

Kircher, F. & Rösch, D. (2021). A shrinkage approach for Sharpe ratio optimal portfolios with estimation risks. *Journal of banking & Finance, 133*.

Lassance, N. (2022). Maximizing the out-of-sample Sharpe ratio. *SSRN Electronic Journal*, Available at SSRN: https://ssrn.com/abstract=3959708.

Ojah, K., Muhanji, S., & Kodongo, O. (2020). Insider trading laws and price informativeness in emerging stock markets: The South African case. *Emerging markets review, 43*.

Quisenberry, W. L. (2017). Ponzi of all ponzis: Critical analysis of the Bernie Madoff scheme. *International journal of economics and financial management, 5*(1), 1-6.

Rossouw, J. & Styan, J. (2019). Steinhoff collapse: A failure of corporate governance. *International review of applied economics, 33*(1), 163-170.

Sen, J. (2022). *A comparative study on the Sharpe ratio, Sortino ratio and Calmar ratio in portfolio optimization*. PsyArXiv. https://www.researchgate.net/publication/366517929_A_Comparative_Study_on_the_Sharpe_Ratio_Sortino_Ratio_and_Calmar_Ratio_in_Portfolio_Optimization.

Sharpe, W. F. (1966). Mutual fund performance. *The Journal of business, 39*(1), 119-138.

Sharpe, W. F. (1994). The Sharpe ratio. *The Journal of portfolio management, 21*(1), 49-58.

Shechtman, D., Wilensky, M., & Fusfeld, L. (2009). Someone made off with my money, now what? Tax issues affecting ponzi scheme victims. *Journal of taxation of investments,* 1-26.

Sortino, F. A. & Van Der Meer, R. (1991). Downside risk. *Journal of portfolio management, 17*(4), 27-31.

Spalding, A. D. (2012). Auditing due diligence in law and ethics: The ponzi ‘feeder fund’ cases. *Review of business and finance case studies, 3*(1), 1-12.

Srivastava, P. & Mazhar, S. S. (2018). Comparative analysis of Sharpe and Sortino ratio with reference to top ten banking and finance sector mutual funds. *International journal of management studies, 4*(2), 93-100.

Steenkamp, P. (2007). *Fidentia: A strategic and corporate governance analysis* [Unpublished masters thesis]. University of Stellenbosch.

U.S. securities and exchange commission (SEC). (2008, 11 December). *SEC charges Bernard L. Madoff for multi-billion dollar ponzi scheme* [Press release].

U.S. securities and exchange commission. (2005, 29 September). *Samuel Israel III; Daniel E. Marino; Bayou Management, LLC; Bayou Accredited Fund, LLC; Bayou Affiliates Fund, LLC; Bayou No Leverage Fund, LLC; and Bayou Superfund, LLC* [Press release].

Van Der Linde, K. E. (2022). The Steinhoff corporate scandal and the protection of Investors who purchased shares on the secondary market. *PER, 25*(1), 1-23.

Van Heerden, C. (2020). Establishing the risk denominator in a Sharpe ratio framework for share selection from a momentum investment strategy approach. *South African journal of economic and management sciences, 23*(1), 1-19.

Van Rensburg, P. & Robertson, M. (2003). Size, price-to-earnings and beta on the JSE Securities Exchange. *Investment analyst journal, 58*, 7-16.

Van Vuuren, J. & Van Vuuren, G. (2022). Detecting investment fraud using the Bias ratio. Sage Journals, 12(2).

Wang, C. D., Chen, Z., Lian, Y., & Chen, M. (2022). Asset selection based on high frequency Sharpe ratio. *Journal of econometrics, 227*(1), 168-188.

Woker, T. (2013). Property-syndication investment v property scam: who decides? A critical examination of the regulation of public-property syndication schemes in South Africa. *Obiter, 34*(2), 233-251.

Xiong, J. X., Idzorek, T. M. & Ibbotson, R. G. (2014). Volatility versus tail risk: Which one is compensated in equity funds? *Journal of portfolio management, 40*(2), 112-121.

Zask, E. (2020). Hedge funds. *The Journal of alternative investments, 3*(3), 33-42.