UNIVERSITY OF THE WITWATERSRAND, JOHANNESBURG
FACULTY OF COMMERCE, LAW AND MANAGEMENT

Enhancing Global Equity Returns with Trend-Following and Tail Risk Hedging Overlays

Name of Applicant:	João Bruno Meneses Schwalbach
Student Number:	752540
Degree:	Doctor of Philosophy
School:	School of Economics and Finance
Supervisor:	Professor Christo Auret
Date:	12 May 2025

A thesis submitted to the Faculty of Commerce, Law and Management, University of the Witwatersrand, in fulfilment of the requirements for the degree of Doctor of Philosophy

Declaration

I, João Bruno Meneses Schwalbach, declare that this research proposal is my own unaided work. It is submitted in fulfilment of the requirements for the degree of Doctor of Philosophy (PhD) in Finance at the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination in this or any other university.

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Prof. Christo Auret, whose guidance and support have been instrumental throughout this journey.

My sincere thanks go to Prof. Robert Vivian for his valuable contributions and for helping me refine this thesis.

I am grateful to Chris Marais for the mentorship and support over the years, and for believing in me.

I extend my appreciation to Corey Hoffstein and Benn Eifert for setting the standard, and for generously sharing their time and knowledge.

To Katinka Mansvelder-Schwalbach, my wife and best friend, thank you for your patience, motivation, and unwavering support. I love you.

Finally, I dedicate this work to my late parents, Mamia and Pai, who shaped me into the person I am today. I know that you will both be smiling down on me as I walk across the stage.

Glossary of Terms

Acronym	Definition
10Y Australia	10-year Australian government bond future
10Y Canada	10-year Canadian government bond future
10Y Japan	10-year Japanese government bond future
10Y US	10-year US Treasury future
2Y US	2-year US Treasury future
3Y Australia	3-year Australian government bond future
5Y US	5-year US Treasury future
ACWI	MSCI All-Country World Index
AEX	Amsterdam Exchange Index
Alpha	A measure of an investment's performance on a risk-adjusted basis, representing excess returns over a benchmark
ASX SPI 200	Australian Securities Exchange S&P/ASX 200 Index
ATM	At-the-money
AUD/USD	Australian Dollar vs US Dollar
Autocorrelation	The correlation of a signal with a delayed copy of itself over varying time intervals
Behavioural biases	Psychological factors that affect investors' decision-making processes
Beta	A measure of a stock's volatility in relation to the overall market
BMF	Brazilian Mercantile and Futures Exchange
Bovespa	Índice Bovespa - Ibovespa
BRL/USD	Brazilian Real vs US Dollar
BTOP50	Barclays Top 50 Managed Futures Index
CAC 40 10 EUR	Cotation Assistée en Continu
CAD/USD	Canadian Dollar vs US Dollar
CAGR	Compound Annual Growth Rate
CAPM	Capital Asset Pricing Model
CBOT	Chicago Board of Trade
CHF/USD	Swiss Franc vs US Dollar
CME	Chicago Mercantile Exchange
CMX	Commodity Exchange Inc.
Collar strategy	An options strategy that involves holding a long position in an asset, buying protective puts, and selling covered calls to help offset the cost of the puts
CTA	Commodity Trading Advisor
CVaR	Conditional Value at Risk

DAX Index	Deutscher Aktienindex
Delta	Measures the sensitivity of an option's price to a $1 change in the price of the underlying asset
DTE	Days to expiration
Elastic net regression	A statistical method that combines Lasso and Ridge regression techniques for variable selection and regularisation
E-mini DJIA Futures	E-mini–Dow Jones Industrial Average Index
EOE	European Options Exchange
EOP	European Options Exchange
ETF	Exchange-Traded Fund
EUR/USD	Euro vs US Dollar
Euro Bobl	German government bond future based on the 5-year Bund
Euro Bund	German government bond future based on the 10-year Bund
Euro Buxl	German government bond future based on the 30-year Bund
Euro Schatz	German government bond future based on the 2-year Schatz
Euro Stoxx 50	A stock market index representing the 50 largest and most liquid blue-chip companies in the Eurozone
EUX	Eurex
Excess returns	Returns generated above the risk-free rate
FTSE 100	Financial Times Stock Exchange 100 Index
FTSE China A50	Financial Times Stock Exchange China A50 Index
FTSE MIB	A stock market index representing the 40 most liquid and capitalised companies listed on the Borsa Italiana (Italian Stock Exchange)
FTSE/JSE Top 40	A stock market index representing the top 40 companies listed on the Johannesburg Stock Exchange (JSE) based on market capitalisation
Futures contracts	Standardised financial agreements to buy or sell a specific underlying asset at a predetermined price at a specified time in the future
Gamma	Measures the rate of change of delta with respect to changes in the underlying asset's price
Gamma hedging	An options trading strategy that aims to maintain a neutral gamma position to reduce the impact of large price movements
GBP/USD	British Pound vs US Dollar
GFC	Global Financial Crisis
GSCI	S&P GSCI Index
H_0	Null hypothesis
H_1	Alternative hypothesis
Hedge fund	A pooled investment fund that employs various strategies including leveraged and/or short positions with the aim of earning active returns for its investors.

HKG	Hong Kong Futures Exchange
HML	High Minus Low (value factor in asset pricing models)
IBEX 35	Índice Bursátil Español
ICE	Intercontinental Exchange
ICF	ICAP Futures
INR/USD	Indian Rupee vs US Dollar
ITM	In-the-money
JPY/USD	Japanese Yen vs US Dollar
KFE	Korean Futures Exchange
KOSPI200	Korea Composite Stock Price Index 200
KRW/USD	South Korean Won vs US Dollar
Kurtosis	A measure of the "tailedness" of the probability distribution of returns.
Leland's alpha	A performance measure that adjusts for option-based hedging strategies
LIBOR	London Interbank Offered Rate
LME	London Metal Exchange
Long Gilt	UK government bond future based on long-term gilts, typically with maturities ranging from 15 to 50 years
MACD	Moving Average Convergence Divergence
Managed futures	Investment strategies that use futures contracts to gain exposure to various asset classes.
MDX	Mexican Derivatives Exchange
Medium Gilt	UK government bond future based on medium-term gilts
MEX Bolsa	Bolsa Mexicana de Valores (Mexican Stock Exchange)
MFM	Montreal Futures Market
MIL	Milan Stock Exchange
MSCI	Morgan Stanley Capital International
MSCI EAFE Index	Morgan Stanley Capital International Europe, Australasia, and Far East Index
MSCI Emerging Markets Index	Morgan Stanley Capital International Emerging Markets Index
MSE	Moscow Exchange
MXN/USD	Mexican Peso vs US Dollar
NASDAQ 100 E-mini	NASDAQ 100 E-mini-Futures
Nikkei 225	Nikkei Stock Average
NOK/USD	Norwegian Krone vs US Dollar
NSE	National Stock Exchange of India
NSE Nifty 50 Index	A stock market index representing 50 of the largest and most liquid Indian companies listed on the National Stock Exchange of India (NSE)

NYB	New York Board of Trade
NYF	New York Futures Exchange
NYM	New York Mercantile Exchange
NZD/USD	New Zealand Dollar vs US Dollar
Option-based hedging	Using options contracts to protect a portfolio against adverse price movements
OSE	Osaka Securities Exchange
OTM	Out-the-money
Put Option	A financial contract that gives the holder the right, but not the obligation, to sell an underlying asset at a specified price within a specified time period
Quadratic regression	A form of regression analysis where the relationship between the independent variable and the dependent variable is modelled as a second-degree polynomial
Rho	Measures the sensitivity of an option's price to changes in interest rates
Risk parity	An investment strategy that allocations weights to a portfolio’s assets based on their risk contributions rather than capital
Risk-adjusted returns	Returns of an investment adjusted for the amount of risk taken to achieve them, proxied by volatility
Risk-neutral density function	A probability distribution derived from option prices reflecting market expectations of future asset prices
RP	Risk Parity
R-squared	Coefficient of determination indicating the proportion of variance explained by the model
S&P 500	Standard & Poor's 500 Index
S&P 500 E-mini	S&P 500 Index E-mini-Futures
S&P/TSX 60 Index	A stock market index of 60 large companies listed on the Toronto Stock Exchange (TSX) in Canada
SAFEX	South African Futures Exchange
SEK/USD	Swedish Krona vs US Dollar
SFE	Sydney Futures Exchange
SG	Société Generale
SG Trend Index	Société Generale Trend Index
SGX	Singapore Exchange
Sharpe ratio	A measure of risk-adjusted return, calculated as portfolio return minus risk-free rate divided by portfolio volatility
Skewness	A measure of the asymmetry of the probability distribution of returns
SMA	Simple Moving Average
SMB	Small Minus Big (size factor in asset pricing models)
SMI	Swiss Market Index
SPX	S&P 500 Index

Straddle	An options strategy involving the purchase or sale of both a call and a put option with the same strike price and expiration date
Stutzer Index	A performance metric alternative to the Sharpe ratio, designed for non-normal return distributions
Tail risk	The risk of extreme investment losses occurring beyond normal market fluctuations
Theta	Measures the sensitivity of an option's price to the passage of time (time decay)
Time series momentum	A strategy that takes positions based on the trend of past returns over a specified time horizon
TOPIX	Tokyo Stock Price Index
Trend-following	An investment strategy that seeks to capitalise on the momentum of asset prices by taking long positions in rising markets and short positions in falling markets
TSM	Time Series Momentum
t-Stat	t-statistic
UMD	Up Minus Down (momentum factor in asset pricing models)
VaR	Value at Risk
Variance Drag	The phenomenon where high volatility disproportionately erodes compounded returns over time
Variance swap	A financial derivative that allows investors to trade future realised variance of an asset against current implied variance
Vega	Measures the sensitivity of an option's price to changes in the volatility of the underlying asset
VIX	CBOE Volatility Index
Volatility skew	The tendency for implied volatility to vary with the option's strike price, often higher for out-the-money puts than calls
Volatility smile	A pattern where at-the-money options have lower implied volatilities compared to out-the-money and in-the-money options
Volatility swap	A derivative contract where parties exchange a fixed volatility for the realised volatility of an asset over a period
ZAR/USD	South African Rand vs US Dollar

Enhancing Global Equity Returns with Trend-Following and Tail Risk Hedging Overlays

ABSTRACT

This thesis demonstrates that overlaying a combination of trend-following and tail risk hedging strategies onto a global equity portfolio significantly enhances both historical risk-adjusted and absolute returns. Tail risk hedging mitigates equity risk most effectively during sudden market crashes, while trend-following notably supports equity during slower bear markets. These strategies are complementary in that the combination is demonstrated to mitigate severe drawdowns which occur during crises that have manifested both rapidly and gradually. In addition, the application of a trend-following strategy has generated positive returns on average during equity bull markets while remaining uncorrelated under normal conditions. Employing a portable alpha framework, the beta and alpha components of the Combination Overlay portfolio are separated. The performance of a 100% global equity portfolio is compared with a portfolio that embodies the Combination Overlay portfolio, which maintains the same 100% allocation to global equity (representing beta) and overlays it with trend-following and tail risk hedging strategies (representing alpha). The resulting portfolio returns remain largely driven by global equity performance but exhibit a large, positive, and statistically significant alpha of 0.38% per month after controlling for traditional equity factors, global government bond, and commodity excess returns. This alpha stems from a combination of an uncorrelated risk premium earned from the trend-following overlay and a reduction in variance drag due to the superior tail risk characteristics of the Combination Overlay portfolio, which leads to an improved compounding return profile relative to global equity.

Keywords: Trend-Following, Tail Hedging, Portable Alpha, Crisis Alpha, Portfolio Construction

JEL Classification: C22; G01; G11; G13; G15; G17

TABLE OF CONTENTS

Schwalbach_JBM_Thesis_Turnitit_(2024-11-25).pdf

ORIGINALITY REPORT

17%	13%	14%	6%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Bruno Schwalbach, Christo Auret. "Leveraging defence into offence: enhancing absolute and risk-adjusted equity returns with tail risk management overlays", Investment Analysts Journal, 2023 Publication	3%
2	fastercapital.com Internet Source	1%
3	ebin.pub Internet Source	<1%
4	Alex Greyserman, Kathryn Kaminski. "Trend Following with Managed Futures", Wiley, 2014 Publication	<1%
5	wiredspace.wits.ac.za Internet Source	<1%
6	dokumen.pub Internet Source	<1%
7	eprints.lancs.ac.uk Internet Source	<1%

LIST OF TABLES

LIST OF FIGURES

LIST OF FORMULAS

CHAPTER 1: INTRODUCTION

1.1 BACKGROUND OF THE PROBLEM

Equities form the cornerstone of long-term investment portfolios due to their well-established ability to produce returns in excess of the risk-free rate over time. However, equities are particularly vulnerable to variance drag - a phenomenon where high volatility disproportionately erodes compounded returns (Eifert, 2020). This effect is asymmetric: large losses reduce a portfolio's base, thereby compromising compounding potential as future returns are generated from a smaller balance, ultimately reducing terminal wealth. For example, a 10% drawdown requires an 11.11% subsequent gain to recover, while a 50% drawdown necessitates a 100% subsequent gain just to break even. A drawdown is defined as the peak-to-trough loss during a specific period of an investment. This illustrates how extreme volatility spikes coupled with deep drawdowns during market crises can significantly hinder long-term wealth accumulation.

Schwalbach and Auret (2023) show that introducing strategies with positive convexity with respect to extreme equity returns can enhance expected returns, reduce the risk of deep drawdowns and help mitigate the adverse effects of variance drag. Convex strategies exhibit non-linear payoff profiles with respect to the equity return distribution that provides asymmetric returns in extreme events, both in the left tail (large losses) and in the right tail (large gains). Two strategies that have empirically exhibited such convexity in relation to equity returns are trend-following and tail risk hedging.

Trend-following is a well-known hedge fund strategy that capitalises on the momentum of asset prices by buying assets that have been trending upwards and selling assets that have been trending downwards, profiting from a continuation of those trends. The strategy can hold both long and short positions, and its investment universe typically comprises global equity, government bonds, currencies, and commodities. It primarily uses derivatives, such as futures contracts, for implementation. Moskowitz, Ooi and Pedersen (2012), along with a large body of research, have found that trend-following strategies have historically delivered attractively high returns while exhibiting convexity in both the extreme left and right tails of equity return distributions. Schwalbach and Auret (2023) conclude similarly but add nuance by noting that the strategy has historically performed particularly well during prolonged bull and bear markets but less effectively during sudden and severe market crashes.

Tail risk refers to the risk of a rare and unexpected event having a severely adverse impact on investors' portfolios and is associated with market crises. Tail risk events can result in behavioural responses leading to counterproductive investment decisions and permanent investment losses (Bhansali, 2015a). Tail risk hedging primarily serves the purpose of cushioning portfolio losses from sudden and severe market downturns. The strategy uses derivatives, such as put options, to provide non-linear payoffs that tend to accelerate as crises deepen during market crises with extreme volatility spikes. Eifert (2020) finds that these strategies have empirically offered significant left-tail convexity with respect to equity returns. Schwalbach and Auret (2023) further highlight that in contrast to trend-following, tail risk hedging has been particularly effective during sudden and severe market crashes but less effective during extended bear markets.

The complementary nature of the convexity introduced by combining these two strategies forms the foundation of this thesis. Specifically, this thesis investigates whether overlaying trend-following and tail risk hedging strategies onto a global equity portfolio, proxied by the MSCI All Country World Index (ACWI), can improve long-term compounded absolute and risk-adjusted returns, reduce tail risk and ultimately help to mitigate variance drag.

1.2 PROBLEM STATEMENT

The central problem addressed in this study is the susceptibility of equity portfolios to performance deterioration due to deep drawdowns and the most harmful effects of variance drag. While equities are relied upon as a key source of return in long-horizon portfolios, their exposure to deep and long-lasting drawdowns can impair the compounding process, thereby undermining expected returns over time.

Traditional diversification strategies are typically employed to preserve capital and provide stability during market turmoil. This often involves reducing equity allocations in favour of assets that act as safe havens, such as developed market government bonds and gold (Goldwhite, 2009). The rationale is that these assets are expected to perform well when equities underperform, thereby stabilising portfolio returns. However, Eifert (2020) and Bhansali (2008) argue that theory often diverges from practice. During systemic crises, assets commonly used to diversify equity exposure can underperform simultaneously, as correlations among assets tend to increase (Bhansali, 2008). Negative correlation with respect to equity returns is most desirable in the left tail, where the effects of deep drawdowns and variance drag are most severe, making strategies that reliably mitigate risk critical for portfolio construction (Eifert, 2020). Moreover, investors face a compromise between reducing risk and maintaining high return expectations. By reducing equity allocations in favour of lower-yielding assets such as government bonds, investors may decrease the overall risk profile of their portfolios, but also reduce expected returns (Goldwhite, 2009).

In this thesis, it is shown that the trade-off between risk and return can be addressed by broadening the approach equity investors take to diversification. Rather than relying solely on intra-asset diversification through sector and regional allocations or traditional asset classes like government bonds and gold, it is argued that equity investors should consider investment strategies such as trend-following and tail risk hedging, which can be efficiently overlaid onto existing equity exposure. This approach introduces significant convexity with respect to equity returns, reducing the tail risk of equity portfolios and thereby increasing absolute and risk-adjusted returns.

1.3 PURPOSE OF THE STUDY

The primary purpose of this study is to investigate the effectiveness of overlaying trend-following and tail risk hedging strategies onto a global equity portfolio to improve expected returns, reduce portfolio tail risk and mitigate variance drag. Traditionally, equity investors have sought to outperform benchmarks primarily through superior security selection—for example by identifying undervalued stocks or sectors. This approach falls under the mantra of active portfolio management. While this approach can be successful, a large body of research has found that active managers underperform passive equity benchmarks after fees.

An alternative pathway to outperformance is proposed that does not rely solely on security selection as the primary means to deliver alpha but instead through the introduction of uncorrelated risk premia and leveraging the mathematical implications of volatility on compounded returns. By overlaying convex strategies such as trend-following and tail risk hedging onto existing equity exposures, investors can introduce an empirically uncorrelated and attractively high stream of returns with positive convexity with respect to equity returns without having to reallocate away from equities and reduce return expectations.

As previously discussed, variance drag occurs when high volatility disproportionately erodes the compounded returns of a portfolio, even if arithmetic returns are strong. This effect is particularly pronounced in the presence of periods of large drawdowns, where the portfolio's base is significantly reduced, making recovery more challenging and negatively impacting terminal wealth. In line with this perspective, Thapar, Nielsen and Villalon (2019) state that it is considerably better to build a plan for how to act during extreme events rather than being forced to create one during such crises. Bhansali (2013) suggests that a sound risk management programme involves a combination of diversification, trend-following, explicit optionality, and cash. The author suggests that each has a different level of effectiveness depending on the severity of loss, and a practical solution for how to organise and combine these approaches is called for.

This thesis explores each overlay separately and then in combination with a focus on how they collectively synergise with global equity exposure to produce a more robust portfolio that benefits from their complementary strengths. Ultimately, this thesis seeks to broaden the horizon for equity investors by highlighting the importance of considering not just asset allocation but also the incorporation of strategic overlays that can enhance portfolio performance.

1.4 SIGNIFICANCE OF THE STUDY

This thesis makes two significant contributions to the existing literature. First, it demonstrates that investors can expand their universe of equity diversifiers beyond merely diversifying within equities or into other asset classes. Instead, they can consider including investment strategies that introduce reliable positive convexity with respect to extreme equity returns. It is proposed that trend-following and tail risk hedging are highly attractive candidates, showing how they can effectively diversify an equity portfolio and reduce tail risk.

Second, this thesis proposes the utilisation of a portable alpha framework as an efficient method for implementing these convex strategies. This framework allows investors to maintain their full equity exposure, thereby retaining equity beta as the primary driver of returns while generating alpha through strategies that generate their own uncorrelated excess returns. This approach is particularly cash-efficient because tail risk hedging and trend-following strategies are implemented using derivatives with low margin requirements (Asness, 2004). Prime brokers can accept fully collateralised equity portfolios as margin, further enhancing cash efficiency.

While this approach results in a leveraged portfolio that may need to be delivered within a hedge fund vehicle, it is important to note that the leverage is applied using the trend-following and tail risk hedging strategies. These strategies introduce defensive qualities to existing equity exposure rather than simply amplifying profits and losses, as would be the case if leverage were applied to more equity exposure. When viewed through the lens of leveraging positively convex strategies, the application of leverage does not necessarily equate to increased risk; in fact, the opposite is argued. By integrating strategies that are empirically uncorrelated with equity markets, investors can introduce uncorrelated alpha sources without proportionally increasing portfolio tail risk, despite the use of leverage. This offers a practical solution for incorporating advanced risk management techniques into existing portfolios without diluting existing equity exposure.

1.5 PRIMARY RESEARCH QUESTIONS AND HYPOTHESES

The primary research question is represented by the primary hypothesis of this thesis. It will be evaluated following a robust analysis of the secondary hypotheses, which are defined below.

1.5.1 Primary Hypothesis

H_0: Compounded long-term absolute and risk-adjusted equity returns are not improved by overlaying a combination of a trend-following strategy and a tail risk hedging strategy on top of full exposure to global equity.

H_1: Compounded long-term absolute and risk-adjusted equity returns are improved by overlaying a combination of a trend-following strategy and a tail risk hedging strategy on top of full exposure to global equity.

1.5.2 Secondary Hypotheses

$H_{0,A}$: Global markets do not trend.
$H_{1,A}$: Global markets trend.

$H_{0,B}$: SG Trend Index returns do not have a significant relationship with Trend Portfolio returns.
$H_{1,B}$: SG Trend Index returns have a significant relationship with Trend Portfolio returns.

$H_{0,C}$: Trend-following returns are fully explained by traditional equity market factors (global market, SMB, HML, UMD), global government bonds and commodities.
$H_{1,C}$: Trend-following returns are not fully explained by traditional equity market factors (global market, SMB, HML, UMD), global government bonds and commodities.

$H_{0,D}$: Trend-following returns are not convex with respect to global equity returns.
$H_{1,D}$: Trend-following returns are convex with respect to global equity returns.

$H_{0,E}$: Overlaying a trend-following strategy on top of global equity does not enhance risk-adjusted and absolute returns compared to holding global equity alone.
$H_{1,E}$: Overlaying a trend-following strategy on top of global equity does enhance risk-adjusted and absolute returns compared to holding global equity alone.

$H_{0,F}$: S&P 500 Index options are not suitable for hedging sudden and severe drawdowns in MSCI ACWI Index exposure.
$H_{1,F}$: S&P 500 Index options are suitable for hedging sudden and severe drawdowns in MSCI ACWI Index exposure.

$H_{0,G}$: Tail risk hedging does not introduce left-tail convexity with respect to global equity returns when overlaid.

$H_{1,G}$: Tail risk hedging does introduce left-tail convexity with respect to global equity returns when overlaid.

$H_{0,H}$: Overlaying a tail risk hedging strategy on top of global equity does not improve tail risk statistics and leads to a reduction in risk-adjusted returns compared to holding global equity alone.
$H_{1,H}$: Overlaying a tail risk hedging strategy on top of global equity improves tail risk statistics and does not reduce risk-adjusted returns compared to holding global equity alone.

$H_{0,I}$: Overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently.
$H_{1,I}$: Overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity introduces greater convexity than either strategy independently.

1.6 THEORETICAL FRAMEWORK

Markowitz (1952) is credited as the founder of modern portfolio theory and suggested that the expected return of an asset, and therefore its pricing, can be entirely explained through its risk, proxied by variance. Markowitz (1952) popularised the notion of risk considerations thereby laying the groundwork for the benefits of investment diversification. In essence, instead of assessing an asset based on its individual risk and return characteristics, an understanding of how any individual security might impact the overall risk/return characteristics of the entire portfolio is more beneficial. Embedded in this approach is the concept that different assets tend to have different covariances with one another, and thus combining them in a portfolio offers a desirable reduction in overall portfolio variance.

Building upon the foundational work of Markowitz (1952), the Capital Asset Pricing Model (CAPM) was developed by Sharpe (1964) and Lintner (1965). The CAPM derives the expected return of an asset based on its systematic risk relative to the overall market. Black (1972) extended this model by addressing scenarios where investors face borrowing restrictions, introducing the concept of a zero-beta portfolio, which has no correlation with the market, as an alternative benchmark for expected returns.

Black and Scholes (1973) developed an options pricing model that is still widely used today. The model uses five inputs to determine the price of an option: the current asset price, the option strike price, the time to expiration, the risk-free rate and the implied volatility of the underlying asset. The most important underlying assumption of the model is that asset prices follow a lognormal distribution. However, since the market crash of 1987, there has been a tendency for out-the-money put options to be more expensive than out-the-money call options. This is known as the volatility skew and is particularly pronounced in equity options.

Alquist, Frazzini, Ilmanen and Pedersen (2020) explain that geometric returns, which ultimately result in the terminal wealth of an investment portfolio, are lower than arithmetic returns due to variance drag, particularly for riskier portfolios. The effect is most pronounced on the extreme left tail of return distributions (Eifert, 2020). In the context of equity portfolios, variance drag is particularly pernicious due to the asset class's inherent spikes in volatility during market crises. Large losses require disproportionately larger gains to recover, hindering the portfolio's growth over time.

Schwalbach and Auret (2023) show that trend-following strategies have historically offered an attractively high risk premium that can increase equity-dominant portfolio expected returns. The authors show that introducing convexity with respect to extreme equity returns provides asymmetric payoff profiles that mitigate the adverse effects of variance drag. Trend-following strategies exhibit such convexity during prolonged bull and bear markets (Moskowitz et al., 2012). Similarly, tail risk hedging strategies introduce significant left-tail convexity during sudden market crashes (Bhansali, 2014).

Schwalbach and Auret (2023) further demonstrate that the combination of trend-following and tail risk hedging can be efficiently delivered within a portable alpha framework. This framework is built on the concept of separating alpha generation from beta exposure. As outlined by Asness (2004), this approach splits a portfolio between traditional index funds to attain equity beta exposure and hedge funds for alpha generation. Asness (2004) suggests that hedge fund strategies allow for clearer portfolio construction by isolating index exposure from alpha sources, which can be implemented using derivatives with low margin requirements. This structure enables investors to maintain full equity exposure while leveraging uncorrelated alpha strategies, thereby addressing variance drag and improving long-term portfolio performance without significantly increasing overall risk.

1.7 ASSUMPTIONS, LIMITATIONS AND DELIMITATIONS

1.7.1 Assumptions

A fundamental assumption of this study is that investors possess the capability to invest in hedge funds, recognising that the strategies examined herein incorporate embedded leverage and short exposure. This leverage is integral to the implementation of both trend-following and tail risk hedging strategies, given that they are overlaid on top of equity exposure in a portable alpha framework. It is assumed that prime brokers are amenable to accepting fully collateralised equity portfolios as margin behind these overlay strategies. Furthermore, this research operates under specific execution assumptions. It is assumed that the trading costs accounted for are appropriate, that investors can transact options at mid-prices, and that the balance of the instruments used in this thesis are at closing prices.

1.7.2 Limitations

The study is constrained by the availability of options data, which begins on 4 January 1996 and extends to 28 June 2024. Although a more extensive dataset dating back to the 1970s was available for the trend-following strategy, the reliance on options data necessitates this narrower timeframe. The research is limited to using S&P 500 Index (SPX) options to hedge a global equity portfolio proxied by ACWI. While the basis risk introduced by this approach is justified in Chapter 4 due to the high correlation between SPX and ACWI (particularly in volatile conditions), it nonetheless represents a constraint imposed by data availability. Futures data consists of last traded daily prices for both first and second month futures, with second month futures becoming first month on expiration dates rather than first notice dates. This highlights a limitation of the study in that some futures, especially physically delivered commodities, may require earlier roll schedules governed by first notice dates by certain prime brokers.

1.7.3 Delimitations

This study deliberately narrows its focus to specific implementations of tail risk hedging and trend-following strategies, opting for simplicity over complexity. In the realm of tail risk hedging, numerous strategic variations exist and have been demonstrated in Eifert (2020), Krishnan (2017) and Bhansali and Davis (2010b). These include (but are not limited to) early monetisation, delta-hedging to mitigate the drag introduced by holding options over time, sizing option positions based on a percentage of net asset value (NAV) rather than a notional approach and employing more sophisticated option structures such as put ratio spreads. Dynamic switching between different tail hedging approaches has been shown to enhance performance. However, this research opts for the most straightforward method: purchasing long-dated, far out-the-money put options sized to notionally cover the fully collateralised base equity exposure. This choice is made to isolate and demonstrate the convexity benefits introduced by a simple, static tail risk hedging strategy without the confounding effects of optimisation or complex interventions.

Similarly, the trend-following strategy employed in this study adheres to a simple rules-based approach rather than more intricate methodologies that could potentially yield improved outcomes. By maintaining simplicity in both overlay strategies, the study ensures that the observed effects on convexity and portfolio performance can be attributed directly to the fundamental characteristics of trend-following and tail risk hedging rather than to the nuances of strategy implementation.

Finally, the scope of this research is confined to overlaying these strategies on a global equity portfolio represented by ACWI. The study does not extend its analysis to more diversified portfolios, such as a 60/40 equity-bond allocation, in order to maintain a clear focus on the enhancements brought about by the overlays themselves. This delimitation also ensures that the effects of trend-following and tail risk hedging can be assessed in a controlled environment, free from the additional variables introduced by broader asset diversification, however likely to improve performance further still.

1.8 CHAPTER SUMMARY

Equities play an important role in many long-term investment portfolios as the primary driver of absolute returns over time. However, they are particularly vulnerable to variance drag due to extreme volatility spikes and large drawdowns, which detract from compounded returns. Traditional diversification methods may not adequately mitigate these risks, especially during systemic market crises when asset class correlations tend towards one. Introducing convexity with respect to left-tail and right-tail equity returns through strategies like trend-following and tail risk hedging offers a potential solution to enhance portfolio performance. In addition, trend-following strategies have empirically delivered attractively high, uncorrelated returns relative to equity and Schwalbach and Auret (2023) found that overlaying a trend-following strategy on top of an equity portfolio resulted in meaningful historical outperformance both in absolute and relative terms.

Trend-following strategies have demonstrated convexity relative to equity returns in both tails of the distribution, performing particularly well during extended bull and bear markets. Tail risk hedging strategies have demonstrated significant convexity relative to equity returns in the left-tail of the distribution, particularly during sudden and severe market crashes. By overlaying these strategies onto a global equity portfolio within a portable alpha framework, investors can maintain full equity exposure while potentially improving absolute and risk-adjusted returns by effectively reducing the most harmful effects of variance drag.

This study aims to empirically evaluate the effectiveness of these combined overlays in mitigating the worst effects of variance drag, enhancing long-term performance, and improving the risk profile of equity portfolios. The rest of this thesis is organised as follows: Chapter 2 provides a comprehensive review of existing research on the key concepts used in this thesis. Chapter 3 provides an in-depth examination of a trend-following strategy, including an assessment of its standalone performance as well as the risk and return enhancements realised when overlaid onto a global equity portfolio. In Chapter 4, a tail risk hedging strategy is constructed and evaluated with respect to its performance when overlaid onto a global equity portfolio. Chapter 5 investigates the combined effect of both strategies on a global equity portfolio, testing whether the combination enhances convexity and improves risk-adjusted returns beyond what each strategy achieves independently. Finally, Chapter 6 summarises the key findings of this thesis, distils the resulting implications for equity investors, and offers suggestions for future research.

CHAPTER 2: LITERATURE REVIEW

2.1 CHAPTER INTRODUCTION

This literature review explores the existing body of research on the concepts used throughout this study. It begins by exploring the foundational concepts of risk management, touching on the impact of variance drag on portfolio performance and the limitations of traditional diversification methods. It will also consider the behavioural biases that often impair investor judgement, thereby reinforcing the case for a systematic approach to tail risk management rather than one overly reliant on subjective decision-making. A comprehensive literature review on tail risk hedging strategies is then provided. These include explicit option-based strategies, as well as over-the-counter (OTC) derivatives such as variance swaps. For background, it also covers the derivation of risk-neutral density functions from option pricing. It also addresses the practical aspects of rebalancing and rolling options, along with empirical performance analyses and critiques of option-based hedging strategies.

Next, trend-following strategies are discussed in depth. Empirical performance, diversification benefits, and implementation variations are covered. Research which has tested the existence of trend effects in financial markets and the empirical properties of trend-following return distributions is covered. In addition, the application of both trend-following strategies and tail risk hedging strategies, which are typically implemented in hedge funds in a portable alpha framework, is considered. Using this framework, investors can overlay these strategies on top of existing market beta without the need to dilute exposure. Finally, the insights from tail risk hedging and trend-following are synthesised by exploring how these strategies can be combined to create a robust risk management framework that is also geared towards expected return enhancement.

2.2 RISK MANAGEMENT

2.2.1 Variance Drag

Variance drag and the benefits of reducing risk in a portfolio are emphasised by Alquist et al. (2020) in their discussion of low risk investing strategies. The authors point out that variance drag occurs due to the compounding effects of volatility, where large price fluctuations reduce the geometric return compared to the arithmetic return. The implication is that by lowering volatility (without commensurately lowering average returns), investors can mitigate this drag, allowing the portfolio to maintain its compounding power over time. Alquist et al. (2020) also highlight the advantages of how low-risk assets, despite offering lower nominal returns, often outperform on a risk-adjusted basis. Strategies such as "betting against beta," which favour low-beta stocks, deliver returns similar to the market average but with less risk, improving the portfolio's Sharpe ratio. Ultimately, the researchers challenge the notion that higher risk necessarily leads to higher returns.

Bouchey, Nemtchinov, Paulsen and Stein (2012) discuss the concept of volatility harvesting, placing strong emphasis on how reducing portfolio volatility through diversification and rebalancing can enhance long-term portfolio growth. The authors contend that while volatility is traditionally regarded as a risk factor, it also presents an opportunity to generate excess returns when managed systematically. They demonstrate that the excess growth achieved from volatility harvesting stems not from security selection or market forecasting but from the consistent rebalancing of volatile assets that are not perfectly correlated over time. A key concept introduced in the paper is variance drag, which refers to the reduction in an asset's compounded growth rate due to volatility, resulting in geometric returns always being lower than arithmetic returns when volatility is greater than 0. Mathematically, this drag is a function of an asset’s variance, with higher volatility leading to a greater reduction in the compounded growth rate. For instance, Bouchey et al. (2012) illustrate that an asset with 10% volatility experiences a drag of 0.50% on its geometric return relative to its arithmetic return.

In their analysis, Bouchey et al. (2012) advocate for diversification and rebalancing as strategies to manage risk and improve returns. Through rebalancing, the authors argue that investors are able to sell high-performing, riskier assets and reinvest in lower-performing, less volatile ones, capturing the benefits of volatility while managing risk. This systematic approach enables portfolios to take advantage of market fluctuations, enhancing the compounding effect of returns over time. Importantly, Bouchey et al. (2012) emphasise that volatility harvesting is not about increasing returns by taking on more risk but about optimising the existing risk-return profile. By lowering portfolio volatility without altering expected returns, investors reduce the impact of variance drag and create a more stable environment for long-term growth. The researchers conclude that rebalancing uncorrelated return streams not only reduces overall portfolio risk but also enhances compounded returns by mitigating variance drag.

Stein, Nemtchinov and Pittman (2009) focus on the concepts of rebalancing and variance drag in the context of emerging market portfolios. The authors state that emerging market securities have high volatility but also exhibit low correlations and thus argue that a weighting structure that provides equal weighting at a country level with regular rebalancing can reduce variance drag and thereby offer improved performance over the cap-weighted alternative. They emphasise that variance drag reduces the long-term growth of a portfolio due to the impact of volatility. The authors state that for a normally distributed time series of returns, geometric returns are equal to arithmetic returns minus one-half of portfolio variance. This implies that compounded wealth is affected mathematically by high volatility, which drags down the geometric return compared to the arithmetic return. This is particularly relevant in portfolios that are not regularly rebalanced. The authors argue that rebalancing uncorrelated assets, on the other hand, allows portfolios to mitigate the effects of variance drag.

Stein et al. (2009) argue that a rebalanced portfolio has a higher compounded return compared to a cap-weighted portfolio because rebalancing systematically reduces portfolio concentration, ensuring that volatility is better managed. The authors demonstrate through simulations that portfolios using equal-weighted rebalancing strategies experience less concentration and benefit from a reduced volatility drag, which leads to a higher long-term growth rate. Maeso and Martellini (2020) also discuss how portfolio rebalancing in equity markets can significantly improve long-term performance through what the researchers term “volatility pumping effect” or “diversification bonus”, which amounts to the same effect as mitigating variance drag.

Eifert (2020) discusses how volatility reduces long-term compound growth rates through variance drag. The author argues that by thoughtfully reducing volatility, investors can mitigate this drag, which otherwise causes significant reductions in portfolio growth. Eifert (2020) shows how a reduction in volatility from 20% to 10% can increase long-term compound returns by 1.50% per annum. Eifert (2020) ultimately argues that variance drag is a key reason why hedging strategies, even those with negative long-term returns, can still enhance overall portfolio performance by managing volatility and drawdowns.

Eifert (2020) demonstrates that attaching a tail risk hedge to a balanced portfolio can mitigate the negative effects of volatility by cushioning significant drawdowns. Tail risk hedges are designed to perform during market stress, generating outsized returns when traditional risk assets fall sharply. This prevents the portfolio from suffering large, irreversible losses that compound over time, ensuring the long-term geometric return is higher despite the hedge having a negative standalone expected return in normal market conditions.

Moreover, Eifert (2020) emphasises the portfolio effect of rebalancing after a tail risk event. When a hedge offsets losses during a drawdown, rebalancing allows the investor to reinvest hedging profits into risky assets at lower prices. This increases exposure to equities during their recovery phase, effectively capturing gains that would otherwise be lost in a non-hedged portfolio. Eifert (2020) argues that this process strengthens the portfolio's ability to recover after a significant loss, leading to higher compounded returns over time. Lastly, Eifert (2020) points out that even though tail risk hedging strategies lose money in benign markets, their benefit during crises outweighs the long-term costs. By protecting against left-tail events and enhancing rebalancing opportunities, Eifert (2020) argues that tail risk hedges not only reduce volatility but also mitigate the impact of variance drag, allowing a balanced portfolio to achieve a higher overall growth rate in the long run.

2.2.2 Enhancing risk management beyond traditional diversification

Tail risk management is an essential aspect of portfolio construction, particularly when mitigating the impact of extreme market events that could lead to permanent financial losses. Bhansali (2008) emphasises the importance of managing exposure to extreme left-tail events, explaining that the benefits of disaster insurance extend beyond survival. Investors who have appropriate tail-risk hedges can capitalise on reduced liquidity during crises, positioning themselves for favourable returns. While tail risk hedging may appear defensive in the short term, it also enables portfolios to remain resilient and liquid over the long term, especially during periods of systemic distress (Bhansali, 2008). At the portfolio level, tail risks are typically systemic, and their effective management requires macroeconomic models for valuation and hedge construction. Bhansali (2008) argues that systemic risks, such as liquidity shortages, arise when financing pressures affect levered holdings. Traditional investment models, which often overlook rare events, are ill-suited for managing these risks as during systemic crises, correlations between asset classes tend to rise (Bhansali, 2008).

Bhansali (2008) likens tail risk management to disaster insurance. During a market crisis, the foremost benefit of having tail-risk insurance is survival, with the secondary benefit being the ability to redeploy liquidity when others cannot, positioning the investor for future returns. While these advantages come at a cost, it is the responsibility of the tail risk manager to minimise that cost over the long term while maximising protection for unexpected shocks. In a market sell-off, Bhansali (2008) argues that liquidity is significantly depleted and correlations between risky assets increase, reducing their diversification potential.

Bhansali and Davis (2010a) argue that incorporating tail risk hedges can improve the return distribution across various asset mixes, despite the reluctance of investors due to the perceived costs. The researchers suggest that tail risk hedging offers opportunities beyond downside protection and actually enables investors to tilt their portfolios to harvest more risk premia, knowing they are protected from extreme losses. Bhansali and Davis (2010a) consider a 60/40 stock-bond portfolio, where the application of a tail hedge reduces portfolio volatility by truncating the lower end of the return distribution. The comparison between the hedged and unhedged portfolios demonstrates that investors who are comfortable with the risks of a 60/40 portfolio may achieve similar risk profiles in a more aggressive 70/30 portfolio, provided that tail hedging is employed. As such, Bhansali and Davis (2010a) suggest that a portfolio can generate an additional “shadow” return stemming from higher exposure to risky assets made possible due to truncating tail risk from the application of a tail hedge.

Eifert (2020) questions the traditional role of fixed income in long-horizon portfolios as being over-relied upon as a source of diversification, with the paper highlighting that it is unclear how much diversification benefit fixed income instruments can provide going forward. The author proposes the need for investors to seek alternative sources of diversification, absolute return, and convexity.

2.2.3 Tools to manage risk

Bhansali (2013) outlines a toolset for practical risk management by identifying four distinct modes. The author argues that sound risk management requires a dynamic approach as opposed to a one-size-fits-all approach. Specifically, Bhansali (2013) proposes that tools such as diversification, alternative betas, explicit options, and cash reserves are useful to mitigate risk at different levels of severity. These modes span from managing small market fluctuations to protecting against catastrophic events. In Bhansali (2013), the author argues that the first mode, diversification with dynamic rebalancing, deals with smaller market fluctuations ranging from 0% to -5%. For moderate losses, from -5% to -15%, Bhansali (2013) recommends momentum-based approaches that are related to trend-following strategies.

As the potential for larger losses increases (from -15% to -35%), Bhansali (2013) advocates for explicit tail risk hedging using options. In this mode, options provide protection against "jump risks", meaning, sudden and severe market drops that cannot be mitigated through rebalancing or trend-following due to the need for adjusting positioning. Finally, for catastrophic losses beyond -35%, Bhansali (2013) argues that there is no substitute for holding cash that retains its value during crises from a defensive standpoint and even allows investors to take advantage of market dislocations from an offensive standpoint.

Bhansali (2013) suggests that different tools can be used to mitigate various degrees of risk, and the effectiveness of each strategy depends on the magnitude and duration of drawdowns. This notion is in line with the findings of Schwalbach and Auret (2023), who find that for sudden and severe drawdowns, tail risk hedging has historically been particularly effective, as it provides immediate protection during sharp market downturns. However, for more prolonged drawdowns, trend-following strategies may offer better protection, as they can adapt over time to sustained market declines. Since it is impossible to predict tail events ex-ante, Bhansali (2013) advocates for a layered defence approach. This involves incorporating multiple strategies that can address different types of tail events with varying levels of effectiveness, ensuring a robust risk management framework.

In Bhansali (2018), the researcher suggests that research on risk management is mainly focused on downside risk, however, that tail risk also exists on the upside, in the form of under participation in rapid and substantial price surges. Bhansali (2018) thus argues that right tail risk is often overlooked and can also be detrimental if not managed properly. The author presents a case for the strategic use of call options to hedge against this upside risk. Bhansali (2018) notes that while drawdowns tend to be more psychologically impactful, ignoring the possibility of large gains can leave investors unprepared to capitalise on these opportunities. The author further critiques the common assumption that markets only "melt down" and argues that this belief is not consistent with historical data. In environments like the post-2016 US presidential election market rally, where equity markets have exhibited strong upside movements, call options have proven to be valuable tools for portfolio management. Bhansali (2018) broadly shows how convexity stemming from the inherent non-linearity of call options magnifies returns during periods of sharp market gains. As such, the author explains that upside convexity is necessary to manage upside risk.

Harvey, Rattray and Van Hemert (2021) state that when all downside risk is eliminated from an underlying risky asset, for example by using at-the-money put options, the return of the underlying should not exceed the risk-free rate. While the benefits of protection during adverse conditions are valuable, these benefits must be weighed against their associated costs. Instead of employing statistical methods to reduce risk, the researchers prefer to use a 'crisis alpha' framework. This involves utilising more dynamic strategies to reduce portfolio risk, particularly during crises, while preserving excess returns in non-crisis periods.

The authors argue that alpha is most needed during crisis times and focus on researching approaches that primarily deliver benefits during those periods. They find that trend-following and option-based tail hedging fit their established framework well, in addition to volatility targeting strategies, strategic rebalancing, and explicitly reducing exposure through drawdown control approaches. Sound risk management involves combining these tools in such a way as to maximise protection at the minimal possible long-term cost (Harvey et al., 2021). Detailed research has been conducted on each of these approaches individually, but literature demonstrating how to combine these tools into a unified framework is sparse.

2.2.4 Behavioural factors

Bhansali (2015b) explores the behavioural aspects of tail risk hedging, presenting key insights into how investor psychology influences the pricing, application, and effectiveness of these strategies. The author argues that traditional financial models, such as those based on efficient markets and arbitrage-free pricing, often fail to capture the real-world dynamics of how investors perceive and manage tail risk.

Similar to Eifert (2020), Bhansali (2015b) suggests that tail risk hedging should be judged not in isolation, but rather, in the context of the overall portfolio. The reason for this is that even though tail risk hedging might drag total portfolio returns over time, especially during upward markets, its value comes in the interaction effect it offers to a risky asset portfolio: the potential for outsized gains which can be rebalanced into the underlying risky asset portfolio in the midst of crisis in anticipation of subsequent recovery. As such, Bhansali (2015b) states that framing and mental accounting, which are well-established behavioural biases, might result in investors regretting spending money on hedging when tail events do not materialise, which in turn might influence their decision to maintain a tail risk hedging program over time.

Another behavioural factor discussed by Bhansali (2015b) is that some investors are willing to accept negative expected returns on tail hedges because they improve the overall skewness of the portfolio, given that risky asset returns, such as those from equity, are empirically negatively skewed over time. The reason for this is that derivatives such as put options offer non-linear payoffs during rare and extreme downside events. The perceived protection against extreme downside risk, argues Bhansali (2015b), creates a market equilibrium where both buyers and sellers of tail risk hedging can be rational, despite the apparent overpricing of such securities from a traditional risk-neutral perspective as discussed in Ilmanen (2012) and Schwalbach and McClelland (2018).

Bhansali (2015b) suggests that tail risk hedging instruments such as options offer investors an efficient way of ensuring that their behavioural biases do not prevent them from incorporating sound risk management into their investment process. The purchase of options is a precommitment to a tail risk hedging strategy. This makes it less likely for investors to behave counterproductively as a sunk cost is paid before a rare event (Bhansali, 2015b). Interestingly, the fact that investors simultaneously purchase protection while hoping for it to expire without value suggests that the value of the protection to the option buyer exceeds its risk-neutral value.

Bhansali (2015b) also discusses the time inconsistency in tail-hedging decisions. Investors often face the temptation to abandon their hedging strategies prematurely, particularly after experiencing a period of stability where the perceived need for protection decreases. The author suggests that a commitment to a tail hedging program empowers investors to outsource risk management and avoid the behavioural pitfalls of abandoning hedges at the wrong time, ensuring they are still protected when volatility unexpectedly increases.

Baker and Filbeck (2013) explore the behavioural aspects of tail risk events, demonstrating that such setbacks can take years for investors to recover from. The authors argue that during these events, investors are often driven by psychological biases, leading them to sell assets in response to significant losses and re-enter the market only after a substantial recovery has already occurred. This reactionary behaviour reinforces the importance of effective tail risk management, which provides much-needed liquidity during crises when risky assets are undervalued, and risk premiums are elevated.

Bhansali (2015a) investigates tail risk management within the context of retirement funds, highlighting the inherently path-dependent nature of the process. The author emphasises that investors are exposed to the risk that rare but significant events can substantially affect the terminal value of their investments. Consequently, Bhansali (2015a) argues that a systematic approach to managing tail risk is essential to safeguard against large drawdowns, which can have enduring impacts on retirement savings and income.

Bhansali (2015a) proposes that the construction of a retirement plan should satisfy three key criteria. Firstly, it should enable investors to comprehend the likely distribution of investment outcomes. Secondly, it should allow investors to capitalise on market fluctuations, using tools such as tail risk hedging for example. Thirdly, it should incorporate mechanisms to shield investors from behavioural biases that may lead to suboptimal responses during crises, such as liquidating their portfolios in fear at inopportune times. The author explains that such drawdowns are driven by volatility, making the management of volatility crucial in mitigating maximum drawdowns.

To manage these risks, Bhansali (2015a) advocates for a combination of diversification, dynamic risk management, and explicit tail hedging. Diversification reduces exposure to individual asset risks, while dynamic risk management allows investors to adjust their allocations in response to evolving market conditions. Explicit tail hedging, such as purchasing put options, offers protection against extreme market declines and helps manage drawdowns, particularly during periods of heightened volatility, and reduces the risk that behavioural biases influence counterproductive actions. By integrating tail hedging into a long-term strategy, Bhansali (2015a) argues that retirees can reduce the temptation to react emotionally to market fluctuations.

Diversification is the cornerstone of risk management. A significant issue with constrained traditional asset allocation mixes, such as the 60/40 portfolio, is that portfolio risk is often dominated by equity risk. The consequences of such risk concentration become apparent during periods of market turmoil. Investors are sensitive to drawdowns, and Chekhlov, Uryasev and Zabarankin (2005) quantify this by suggesting that the average investor's confidence is first tested when a drawdown exceeds 15% and that investors are unlikely to tolerate a 50% drawdown without opting out of a strategy. Furthermore, the researchers suggest that average investor tolerance is also tested if even mild drawdowns persist over 2 years. This indicates that both the depth and length of drawdowns are useful metrics for assessing the practical usefulness of hedging strategies.

During the 2008 Global Financial Crisis (GFC), Faber (2015) demonstrates that a global 60/40 equity/government bond portfolio suffered a 35% drawdown. Global equity experienced a drawdown of over 50% during the GFC, dominating the 60/40 portfolio performance despite a substantial 40% allocation to safer global government bonds. Qian (2011) examined this phenomenon and found that equity contributes approximately 90% of the overall portfolio volatility in a 60/40 portfolio.

Asness, Frazzini and Pedersen (2012) argue that the application of further gearing can be utilised to approximately match the portfolio's expected return to that of equity while offering lower volatility or to match the portfolio's volatility to that of equity, thereby providing a higher expected return. The same approach can be applied with the inclusion of other asset classes such as commodities, property, gold, corporate bonds, and inflation-linked bonds. The over-reliance on equity as the primary driver of portfolio returns thus represents a vulnerability for many investors should future returns disappoint.

Many forecasts suggest that expected returns over the next decade are approximately half of those generated over the preceding 48 years (Davis et al., 2018). Investigating alternative risk premia from strategies such as trend-following therefore offers a fruitful line of enquiry.

2.3 TAIL RISK HEDGING

2.3.1 Background

Eifert (2020) establishes several key concepts that position tail risk hedging, through explicit option purchases, as a compelling overlay for equity-dominant portfolios. The author challenges the conventional notion that hedging strategies, which often have negative long-term expected returns, should not be considered in the asset allocation mix of long horizon portfolios. Instead, Eifert (2020) argues that, when used in regularly rebalanced asset allocation frameworks, these strategies can add value. Specifically, the author demonstrates that reducing volatility in a risky portfolio can enhance long-term compounded returns when average returns remain constant. This improvement is partly attributed to mitigating the destructive impact of variance drag discussed earlier in this chapter. Secondly, both Eifert (2020) and Bhansali and Davis (2015b) assert that a tail risk hedge should always be evaluated in conjunction with the risky asset portfolio it is intended to protect. While a tail risk hedge may exhibit a negative expected return in isolation over time, it can increase overall expected returns when combined with a risky asset portfolio.

When managed alongside a risky asset portfolio and regularly rebalanced, tail risk hedges can cushion drawdowns and hedging profits can be utilised to increase risky asset exposure during periods of market stress, anticipating potential recoveries and thereby enhancing long-term expected returns (Eifert, 2020). The author further suggests that put options are ideal for hedging tail risk due to their high cash efficiency. Typically, only a small cash outlay as a percentage of the underlying portfolio is required to buffer against portfolio losses in the event of a market crash. Eifert (2020) also highlights delta-hedging as a method to ensure minimal disruption to the underlying risky asset portfolio from a strategic asset allocation perspective, while also helping to monetise hedging gains during crises.

Krishnan (2017) conducts a comprehensive analysis of various systematic buy-and-roll option-based hedging techniques, exploring a diverse array of option durations and target deltas. The researcher concludes that far out-the-money (OTM) options are particularly suitable for managing tail risk, as they provide the highest efficiency and positive convexity during market downturns. The highly convex, non-linear payoff structure of OTM options makes them especially beneficial for tail hedging applications.

Similarly, Bhansali and Davis (2010b) examine the performance of far OTM options during financial crises. They observe that these options can experience substantial value appreciation during market crashes without necessarily transitioning to in-the-money (ITM) status. This phenomenon is attributed to the convexity of option premiums in response to surges in implied volatility. During periods of market instability, implied volatility tends to escalate due to heightened hedging demand and increased realised volatility from significant price declines. This behaviour exemplifies why options are frequently characterised as non-linear hedging instruments, with their convexity allowing option prices to rise significantly as equity prices fall and realised volatility spikes.

Bhansali and Davis (2010b) advocate for the utilisation of far OTM options as a safeguard against extreme market events. The authors highlight the reliably negatively correlated and explosive payoffs that these options provide during such events. However, the authors note that far OTM options have historically been ineffective in mitigating smaller and more prolonged drawdowns. Consequently, these options should be viewed as a form of insurance against extreme events rather than a comprehensive defence against all types of market selloffs.

Eifert (2020) warns against strategies that involve writing calls against existing equity exposure as a potential approach to risk management. The author argues that while these strategies may provide some downside protection in the short term in the form of premium income received for selling equity upside, they fail to protect against severe left tail events. Eifert (2020) states that call writing strategies behave like equity portfolios during significant drawdowns, which negates their utility as true hedges. In contrast, true tail risk hedges, such as those that utilise put options, which are explicitly designed to perform in extreme market conditions, offer much more reliable outcomes.

Traditional long-horizon investment strategies have historically neglected the need for insurance against rare but severe events (Bhansali & Davis, 2010a). The authors suggest that this tendency stems primarily from the perception that tail risk hedging incurs costs that can detract from potential returns. While the primary purpose of portfolio insurance is to protect against substantial downside risks, this defensive view fails to recognise the opportunities that can arise from crises, such as during the GFC in 2007-2008. Bhansali and Davis (2010a) argue that for investors who generate liquidity during such periods using tail risk hedging benefit not only by shielding portfolios from catastrophic losses but also by allowing them to strategically reallocate risk, which can enhance returns during recovery phases. Bhansali and Davis (2010a) advocate for this dual nature of tail risk hedging, coining it as "offensive risk management", where investors can actively take advantage of post-crisis opportunities using liquidity derived from hedging gains during crises.

Options provide exposure to both an underlying asset's price and its volatility. However, options can isolate volatility exposure through a process known as delta hedging. Specifically, the Black-Scholes delta, calculated as the partial derivative of an option's price with respect to its underlying asset, can be neutralised by trading an appropriate amount of the underlying asset to render the overall delta zero (Hull & White, 2017). This parameter is considered the most important option Greek in risk management. In practice, changes in delta can be substantial and challenging to manage. Clewlow and Hodges (1997) demonstrate that large changes in delta result in the delta hedging process becoming expensive due to transaction costs.

Vähämaa (2004) finds that delta hedging can be improved by factoring in the volatility smile, thereby transforming normal Black-Scholes delta into a so-called skew delta. Hedging skew delta accounts for the indirect impact of simultaneous changes in volatility associated with a change in the underlying asset. Jarrow and Turnbull (1994) document that both delta and gamma are required to more effectively hedge movements in underlying interest rate derivatives. Mastinšek (2006) states that hedging errors arise because trading is done discretely, while the Black-Scholes delta adjusts continuously. Raju (2012) highlights additional practical issues related to the delta hedging process. Brière, Burgues, and Signori (2010) suggest that in the short term, the impact of delta hedging can dominate the impact of changes in implied volatility that the hedger aims to isolate.

2.3.2 Tail risk hedging using over-the-counter instruments

An alternative means of hedging tail risk that has historically been used by institutional investors (over-the-counter) is the use of variance and volatility swap contracts. A volatility swap is an agreement to exchange a fixed volatility for the realised volatility of an asset over a specified period. Volatility swaps provide pure exposure to the volatility of an asset alone (Hull, 2003), allowing for direct volatility exposure without the challenges associated with delta hedging options. Variance swaps, where variance is defined as the square of volatility, are more widely used in financial markets. Hull (2003) suggests that variance swaps are easier to value than volatility swaps because a replicating portfolio composed of puts and calls can be constructed. The popular VIX Index is built on this basis (Hobson & Klimmek, 2012), and the fair value of a variance swap is simply the cost of the replicating portfolio. Demeterfi, Derman, Kamal and Zou (1999) suggest that variance swaps can be utilised for various purposes, including isolating the spread between realised and implied volatility, hedging, or speculating on future volatility levels. A variance swap holder receives the agreed notional amount per annualised variance point squared, multiplied by the difference between the realised variance of the underlying and the fixed variance level agreed upon (Demeterfi et al., 1999). Warren (2012) suggests that long positions in variance swaps offer strong hedging properties for equity-centric portfolios.

Konstantinidi and Skiadopoulos (2016) examine the historical performance of variance swaps on the SPX. The authors find that volatility sellers tend to earn positive returns over time, similar to the conclusions of Ilmanen (2012) and Schwalbach and McClelland (2018). However, they also show that short volatility positions are highly vulnerable to abrupt and significant increases in realised volatility during market crises. In contrast, long volatility positions through variance swaps are particularly successful during periods of deteriorating liquidity.

Hafner and Wallmeier (2007) study the historical performance of variance swaps in European markets, specifically examining the DAX and Euro STOXX 50 indices from 1995 to 2004. The authors find a strongly negative volatility risk premium, where implied volatility exceeds realised volatility, rendering a buy-and-hold approach in long variance swaps unattractive. However, they recognise that long variance swap positions yield explosive payoffs during market crashes, not only because they are long volatility but also because they are long the square of volatility (variance). Hafner and Wallmeier (2008) suggest that variance swaps are the most common instruments used to trade volatility. They find strongly negative returns on DAX and Euro STOXX 50 variance swaps from 1995 to 2004, consistent with similar findings on the SPX and indicative of a negative volatility risk premium.

2.3.3 Background on risk-neutral density functions

Options markets provide valuable insights into the risk-neutral return distribution of underlying assets, particularly through the analysis of risk-neutral densities (RNDs). Doran, Peterson and Tarrant (2007) find that the risk-neutral density function of equities priced by option markets is generally fat-tailed and negatively skewed. Using the framework established by Black and Scholes (1973), many have attributed the higher pricing of out-the-money puts compared to out-the-money calls to the tendency for markets to experience sudden and extreme drawdowns but slow and steady upward movements.

Bates (2000) examines the 1987 market crash and finds that strong negative skewness developed in options' implied distributions after the crash as investors rushed to purchase put protection. Similarly, Bakshi, Cao and Chen (1997) develop an option pricing model that allows for random jumps, stochastic volatility, and interest rates. The authors find evidence of excess kurtosis and negative skewness in the implied probability distribution of option prices.

Pan (2002) examines the SPX and short-dated, near at-the-money options, finding that jump-risk premia tend to respond quickly to market volatility. The author suggests that this finding significantly explains the volatility smirk observed particularly in equity options. RNDs offer a glimpse into the implied probability of specific outcomes over specific horizons (Tian, 2011). Breeden and Litzenberger (1978) demonstrate how volatility smiles can be used to estimate RNDs. Specifically, given a range of strike prices spanning future payoffs, the probability priced into each band, and therefore the entire implied probability distribution, can be derived from the second derivative of option prices. The benefit of this approach is that it does not require a model with restrictive assumptions. Insights into market participants' expectations and risk preferences can be gleaned from deriving RNDs, making the procedure valuable to investors.

Hull (2003) builds on the work of Breeden and Litzenberger (1978) and uses a series of butterfly spreads to estimate an RND. A butterfly spread involves buying one call at a lower strike price, selling two calls at a higher strike price, and buying one call at an even higher strike price, all with the same expiration date and equidistant strike prices. Birru and Figlewski (2012) expand on the methodologies of Breeden and Litzenberger (1978) and Hull (2003), suggesting practical steps to calibrate the resulting RND. First, given that trading is sporadic for many strikes, using transaction prices is problematic. Therefore, it is preferable to use bid/ask prices established by market makers, which are more reliably available across strikes. The authors also eliminate far out-the-money data points where the option value is too small relative to its bid/ask spread. Second, they construct a smooth implied volatility curve across strikes. Since options are traded at discrete intervals (strikes), smoothing and interpolating between them is necessary to obtain an RND.

The authors maintain that it is more accurate to smooth implied volatilities rather than option prices, converting these back to premiums afterward. Third, the authors interpolate implied volatilities using a fourth-degree smoothing spline rather than the more commonly used cubic spline method. This approach captures less noise because the smoothing mechanism does not require passing through every data point. Fourth, the spline is fitted using least squares to the midpoint implied volatilities, with adjustments for wider bid/ask spreads. Fifth, to construct the RND, the authors recommend using out-the-money puts and calls for the left and right sides of the distribution, respectively, and using a blend for at-the-money options. This is due to the wider bid/ask spreads for deep in-the-money options. Sixth, they convert the interpolated implied volatilities back to option prices and apply a similar method as proposed by Hull (2003) to arrive at the body of the RND. Finally, to complete the density, the left and right tails of the RND need to be determined. The authors use generalised extreme value distributions to model the tails of the RND.

Birru and Figlewski (2012) examine intraday changes in the RND for the SPX during the GFC. They quantify changes in investor risk preferences and expectations of future returns to the minute, observing a fivefold increase in short-term, intraday volatility from October 2006 compared to October 2008. Expectations around future returns, overall levels of uncertainty, and price volatility are clearly reflected in option prices and can be observed from RNDs. However, the literature generally concludes that RNDs are not useful for predicting financial crises. Bates (1991) concludes that index options did not forecast the crash of 1987 to any significant degree. Bhabra, Gonzalez, Kim and Powell (2001) find that options react to crises rather than predict them; index option implied volatilities did a poor job of predicting the Korean crisis of 1997. This suggests that implied volatilities used to price options are more lagging than leading variables. Dennis and Mayhew (2002) show that periods of high market volatility tend to lead to more negative levels of implied skewness in option contracts, concurring with the notion that implied volatilities are more lagging than leading parameters.

2.3.4 Rebalancing and rolling options

The concept of rebalance timing luck was introduced by Hoffstein, Sibears and Faber (2019). The authors explore the extent to which choices around the timing and frequency of rebalancing influence investment outcomes. The authors find that different protocols create substantial performance differences between portfolios that are otherwise managed identically, highlighting the stochastic nature of how returns compound when portfolios are periodically rebalanced. The authors emphasise that rebalancing is a fundamental aspect of both passive and active investment strategies. However, the choice of rebalancing intervals, often determined by convenience or fixed schedules, can have long-lasting impacts on portfolio performance. For instance, the timing of rebalancing a 60/40 equity/bond portfolio was shown to lead to a performance difference of approximately 0.22% per annum between portfolios rebalanced in May and those rebalanced in October.

Blitz, van der Grient and van Vliet (2010) further validate the findings presented in Hoffstein et al. (2019), showing that rebalance timing can result in non-trivial deviations in realised performance, even among portfolios rebalanced under identical methodologies. Hoffstein et al. (2019) provide statistical models that quantify the variance in returns attributable to rebalance timing luck. Their analysis suggests that this phenomenon can be minimised by staggering the action of rebalancing rather than conducting it all at once. Furthermore, Hoffstein et al. (2019) highlight that rebalance timing luck can be exacerbated during periods of heightened market volatility, as these conditions amplify the differences in asset class performance, resulting in the potential for performance dispersion.

In line with Hoffstein et al. (2019), Eifert (2020) states that because tail hedges are path-dependent instruments, the timing of portfolio rebalancing can significantly impact performance. To mitigate this, the paper utilises a staggered rebalancing protocol on options, reducing the likelihood of understated or overstated performance due to unlucky or lucky rebalance timing. Schwalbach and Auret (2023) follow suit, employing three separate tranches of put options for tail risk hedging, with each tranche rolled quarterly in distinct, non-overlapping months.

2.3.5 Empirical performance analysis

Bhansali and Davis (2010b) provide a compelling case for incorporating tail risk hedging strategies into traditional portfolios, particularly for long-horizon investors who avoid such hedges due to concerns about cost and underperformance during stable market conditions. The authors argue that this perception has led many to forgo tail risk strategies altogether. The authors highlight the potential to enhance overall portfolio returns through the introduction of tail risk hedging by enabling investors to take on greater risk with confidence in extreme downside protection. Bhansali and Davis (2010b) suggest that tail hedging, when properly applied, can allow investors to increase their allocation to riskier assets, such as equities, without incurring the full brunt of downside risk. This approach enables investors to capture more risk premium, which can significantly enhance long-term returns.

Bhansali and Davis (2010b) illustrate the empirical performance of tail risk hedging through a comparison of portfolios that employ such strategies with those that do not. In their analysis, they construct a 60/40 equity-bond portfolio and compare it to an equity-heavy portfolio with a tail risk hedge. The authors demonstrate that the introduction of tail hedging not only empirically mitigated downside risk but also improved the portfolio's overall performance by allowing for a more aggressive asset allocation by reducing fixed income allocation in favour of equity. Bhansali and Davis (2010b) also focus on the tail-hedged portfolio’s performance during the 2007-2008 financial crisis. Their results reveal that portfolios with tail hedges significantly outperformed unhedged portfolios during this period.

Bhansali and Davis (2010b) reemphasise that investors can be exposed to more risk premia by attaching tail risk hedges to their risky asset portfolios. The authors concede that naïve passive applications of tail hedges empirically drag portfolio returns, suggesting that more active intervention is required for an effective tail risk hedging strategy to enhance the expected returns of a risky asset portfolio. In their analysis, Bhansali and Davis (2010b) discuss early monetisation of tail risk hedges, suggesting that there can be opportunities to monetise hedges early, especially when markets experience sudden spikes in volatility. The reason for this is that when realised volatility in the underlying market spikes, implied volatilities across option surfaces tend to spike in response, factoring in higher uncertainty.

Bhansali and Davis (2010b) show that using a systematic method in which monetisation is triggered as a function of original option premiums increasing n-fold in size in response to a systemic shock. The authors use monetisation multiples of 2.5x, 5x, 7.5x and 10x. The authors suggest that monetisation decisions should be event driven. For example, if a 25% out-the-money put option with one-year to expiration is purchased, and the market falls by 35% during the year only to recover to a new high at the expiration of the contract, the contract would expire without value. If, however, an active monetisation rule resulted in selling the put option for a profit during the market drawdown, then proceeds could be immediately reinvested into the market, benefitting from the subsequent recovery.

Bhansali and Davis (2010b) also suggest that investors should not necessarily wait for tail hedges to become in-the-money to monetise them. An option's price consists of intrinsic value (the maximum of the amount that the option's strike price exceeds the cost of the underlying and zero) as well as extrinsic value (time value). Thus, even while the option holds zero intrinsic value, a fall in the market that is severe enough will result in the extrinsic value component of an option being worth many multiples of its original cost. This highly convex (non-linear) payoff profile is the main attraction of using deep out-the-money options for tail hedging purposes. Bhansali and Davis (2010b) also use a threshold factor augmented vector autoregression model to backfill implied volatilities on US equity back to 1928. Given that most option data can only reliably be sourced since the early-2000s, applying such a model allows researchers to glean useful findings.

Bhansali and Davis (2010b) find that the larger the monetisation multiple, the longer the average time it historically took to earn a payout. As such, the authors argue that investor patience can be tested without applying tail risk hedging with a permanent, asset allocation-based approach in which premium allocations are considered as a portfolio allocation. The authors compare the performance of three distinct portfolios, each of which invested in a passive US equity portfolio. The first strategy involves a straightforward equity buy-and-hold approach without any tail risk hedge. The second strategy introduces a naïve, passive tail risk management approach, where the investor spends 1% of the portfolio value annually on put options (following an asset allocation-based approach), which are held to expiration and rolled into a new tranche of options purchased with 1% of portfolio NAV, and so on. Finally, the third strategy adopts an active tail risk management approach, where put options are purchase in the same way as the naïve portfolio but are monetised whenever their value increases fivefold, and the gains are reinvested into the underlying equity portfolio.

Bhansali and Davis (2010b) find that the naïve tail hedge portfolio provided protection during periods of severe market decline but also imposed a consistent cost on the portfolio. As a result, the passive strategy tended to underperform the unhedged portfolio during stable or rising markets, since the put options expire worthless in the absence of tail events. The authors find that the portfolio with the active monetisation rule did best over the period in absolute and risk-adjusted terms. It should be noted that the period under investigation was from January 1988 to October 2009. The fact that the study used data with an end date within 12 months of the 2008 GFC must be acknowledged.

When Bhansali and Davis (2010b) extend their study back to January 1950 (with the same end date) by backfilling implied volatility data using a model with assumptions around how options would have been priced, both the naïve and active tail hedged portfolios outperformed naked equity in both the absolute and relative terms. Finally, when Bhansali and Davis (2010b) extend their study back to April 1928 (again, with the same end date) using the same modelling assumptions for pricing options, they find that the naïve portfolio significantly outperformed in both absolute and risk-adjusted terms, followed by the active tail-hedged portfolio and then naked equity. The authors explain that during the Great Depression, equity fell 85% over a 5-year period, and the active approach would have sold too early, having locked in some hedging gains but not nearly the same amount as the naïve portfolio, which did not realise hedging gains until its annual roll date thus pointing out a flaw in a monetisation strategy. Specifically, the risk of monetising all optionality too early and leaving investors without protection should conditions worsen. Ultimately, Bhansali and Davis (2010b) argue that the value gained from a Great Depression-level worst-case scenario, which is a rare but severe event, outweighs the carrying cost of attaching a tail risk hedging program to a pre-existing equity portfolio.

Despite the substantial shift in the asset management industry from active to passive management, as documented by Prondzinski and Miller (2018), Bhansali and Davis (2010b) demonstrate that tail risk managers can still add significant value. A considerable body of research has shown that most active managers underperform passive benchmarks after accounting for fees, which accounts for this migration. However, in the context of risk management, it remains less clear whether a passive approach is more effective than an active one.

Baker and Filbeck (2013) argue that even long-term investors can benefit from employing tail risk management strategies. The authors assert that such strategies position investors to redeploy liquidity into risky assets during the depths of crises when risk premiums are at their highest. This perspective challenges the notion that tail hedges should be viewed solely as a sunk cost with occasional payoffs. Instead, tail risk hedging can be considered a strategic component of a larger portfolio, enabling a more opportunistic and proactive investment approach, particularly given the fortuitous timing of reinvestment opportunities that arise during market downturns.

Eifert (2020) compares the performance of a 60/40 equity/bond portfolio to that of the same 60/40 portfolio with an overlay of a tail risk hedging strategy. The author utilises 30-delta, one-year puts, divided into three separate tranches, which were delta-hedged to neutralise the directional risk embedded in the options. The tail risk hedging strategy targeted 1% of the portfolio's NAV in an asset allocation approach, meaning the portfolio was rebalanced to a 60/40/1 equity/bond/option premium allocation, where the option premium was funded at the risk-free rate.

Eifert (2020) finds that the portfolio with the tail hedge overlay outperformed the unhedged portfolio in both absolute and relative terms over the period from 2002 to 2020. Specifically, the compound annual growth rate (CAGR) of the hedged portfolio was 8.04% compared to 7.79% for the unhedged portfolio. This outperformance is attributed to the hedge's ability to protect the portfolio during significant market drawdowns, thus allowing the rebalanced portfolio to take advantage of subsequent recoveries without having to reduce equity exposure. The delta-hedged puts reduced the drag caused by volatility, thereby mitigating the negative impact of variance drag on long-term portfolio compounding. Furthermore, Eifert (2020) finds that the inclusion of the tail risk hedge not only reduced volatility but also significantly mitigated portfolio drawdowns during major market events, such as the 2008 GFC and the March 2020 Covid Crash. The hedged portfolio experienced smaller drawdowns, allowing for quicker recoveries post-crisis.

Evidence of tail risk hedging outperformance during sudden and severe market crises is also supported by industry-level outcomes. Herbst-Bayliss and Delevingne (2020) report that the CBOE Eurekahedge Tail Risk Hedge Fund Index, which is a composite of tail risk hedging funds, gained 14% in February 2020 and between 32% and 41% for the year up until March 2020.

Szado and Kazemi (2009) find that a long collar strategy outperforms a protective put strategy on the Invesco QQQ Trust Series 1 ETF from 1999 to 2008. Partially offsetting the cost of a put option by selling an upside call option can result in improved performance (Szado and Kazemi, 2009). Similarly, Szado and Schneeweis (2010) examine a more active approach to hedging the QQQ and find that it generally outperforms a more passive strategy. Goldwhite (2009) examines the performance of diversification strategies during episodes of market stress. The author concludes that changes in the VIX offer a useful way to categorise investments based on the prevailing risk environment.

Fernandes, Medeiros and Scharth (2014) show that the VIX Index exhibits long-range dependence, aligning with similar research demonstrating volatility clustering. Daigler and Rossi (2006) measure the performance of a portfolio comprising a long position in the SPX hedged with a long position in VIX Index futures. The authors find a significant reduction in portfolio volatility with the addition of VIX futures. This outcome is explainable through the highly negative correlation between the VIX and the SPX (Hsu and Murray, 2007). Dash and Moran (2005) find similar risk management benefits when adding long volatility exposure to hedge funds. They observe that the negative correlation offered by long volatility is especially useful when hedge funds deliver their worst returns, suggesting significant tail risk benefits.

Whaley (2013) evaluates VIX Exchange-Traded Products (ETPs) as a buy-and-hold investment and finds similarly disappointing long-run returns as observed in VIX futures, but with additional product-related peculiarities. Warren (2012) suggests that the unsatisfactory performance of investing in short-dated VIX futures can be somewhat improved by investing in longer-dated VIX futures. Szado (2009) considers the impact of long VIX futures exposure and finds that asset class correlation profiles, which held for many years before the GFC, broke down during the tail event. Thus, while suffering slow and steady long-run losses, the explosive nature of VIX futures returns during episodes such as the GFC offers a desirable source of diversification to long-horizon portfolios. Berkowitz and DeLisle (2018) find that VIX products offer attractive portfolio diversification in times of crisis but suffer negative long-term returns.

2.3.6 Arguments against the use of option-based tail hedging

Critics of tail risk hedging primarily argue that put options are prohibitively expensive, thereby questioning the economic viability of such strategies. From a theoretical perspective, Ilmanen (2012) posits that option writers require compensation for the inherent risk of insuring option holders against extreme market downturns. This compensation is embedded in the form of a volatility risk premium, analogous to the equity risk premium sought by equity investors.

Essentially, while equity investors demand a premium for bearing market risk, option writers seek a premium to assume the risk associated with low-probability, high-impact adverse movements in asset prices (Ilmanen, 2012). This perspective is further supported by Schwalbach and McClelland (2018) and Israelov and Nielsen (2015), who emphasise that the elevated cost of tail risk hedging is significantly influenced by the volatility risk premium inherent in option prices. Schwalbach and McClelland (2018) argue that externalising the risk of catastrophe is a highly valuable action and justifies the persistence of the volatility risk premium. Consequently, holders of put options should expect to generate negative returns in the long term in standalone terms.

Short volatility strategies capitalise on the inherent richness in option pricing, particularly the tendency for implied volatility to surpass subsequent realised volatility. Bhansali and Harris (2018) delve into the increasing prevalence of short volatility strategies, arguing that these strategies effectively involve the sale of financial insurance. Investors earn modest premiums in exchange for assuming substantial risks during extreme market events. The authors caution that such strategies can contribute to market destabilisation, especially if there is a coordinated unwinding of short volatility positions, potentially precipitating a market crash.

Bhansali and Harris (2018) note that short volatility strategies, once predominantly utilised by hedge funds, have become accessible to a broader spectrum of investors, including retail participants through exchange-traded volatility products. A significant concern highlighted is the lack of awareness among many investors regarding the extent of correlation within their strategies, which could lead to a feedback loop where escalating volatility induces widespread deleveraging, thereby exacerbating market declines.

Options inherently truncate the return distribution of portfolios, rendering traditional performance metrics such as the Sharpe ratio and Jensen's alpha less effective for evaluating their risk and return profiles (Kapadia & Szado, 2012). To address this limitation, Kapadia and Szado (2012) propose alternative metrics, namely Leland's alpha and the Stutzer Index, which are specifically designed to account for the non-normal distribution of returns resulting from option-based strategies.

From an empirical performance perspective, Figlewski et al. (1993) provide some of the earliest assessments of tail risk hedging strategies by investigating the effectiveness of a protective put approach. The authors' findings indicate that the protective put strategy tended to underperform in practice. The study stresses the limitations associated with the strategy of buying and holding put options until expiration without implementing early monetisation, a weakness further described by Bhansali and Davis (2010b).

Asvanunt, Nielsen and Villalon (2015) identify equities as the primary source of tail risk impacting most investors. The researchers argue that, from a risk perspective, a traditional 60/40 equity/fixed income portfolio resembles a 90/10 equity/fixed income allocation. Their study explores both direct and indirect methods for hedging tail risk. The direct method employs a collar strategy, which partially finances the purchase of 3-month, 7.50% out-the-money (OTM) put options through the sale of one-month, 10% OTM call options. In contrast, the indirect approach comprises three distinct investment strategies: allocating to low-beta equities to mitigate equity risk, adjusting the equity and fixed income mix via a risk parity methodology, and implementing a trend-following strategy for portfolio rebalancing.

The low-beta equity strategy exploits the low-beta anomaly, which posits that low-beta stocks tend to outperform the returns predicted by the Black et al. (1972) one-factor model, the Fama and French (1992) three-factor model, and the extended Carhart (1997) four-factor asset pricing models. By overweighting low-beta equities and underweighting high-beta equities based on a trailing 12-month beta assessment conducted monthly, the authors aim to reduce the overall portfolio beta without entirely diminishing the equity allocation. This approach is grounded in the hypothesis that during tail events, investors disproportionately liquidate high-beta equities in favour of low-beta equities. Consequently, this strategy enhanced average returns to 10.2% compared to the benchmark's 9.7%. The Sharpe ratio improved to +0.72 from the benchmark's +0.55, with returns exhibiting an 84% correlation and a 5.1% tracking error. Although low-beta equities slightly underperformed during four of the five worst equity months, they outperformed during five of the seven most severe multi-month drawdown periods identified by the authors.

When integrating the collar strategy with the 60/40 benchmark, Asvanunt et al. (2015) observed a reduction in average returns from 9.7% to 8.1%, accompanied by a decline in Sharpe ratios from +0.55 to +0.48. This outcome occurred despite the 60/40 with collar portfolio achieving lower volatility, decreasing from 9.5% to 7.5%. The returns between the two approaches were highly correlated at 96%, with a modest tracking error of 3.2%. From a risk management perspective, the collar-enhanced portfolio proved effective during severe, short-term drawdowns but was less satisfactory during prolonged drawdown periods.

Regarding the modification of the 60/40 equity/bond portfolio through a risk parity approach, Asvanunt et al. (2015) aimed to achieve a more balanced risk distribution by equalising the volatility contributions of equities and bonds. This was accomplished by allocating a greater capital proportion to bonds rather than equities and subsequently leveraging the portfolio to align its risk with that of the traditional 60/40 mix.

The authors argue that this methodology diminishes the portfolio's reliance on equities, thereby reducing its exposure to tail risk. Upon implementation, the risk parity adjustment yielded the highest outperformance relative to the benchmark, generating a statistically significant alpha of 4% and a Sharpe ratio of +0.79, compared to the benchmark's +0.55. However, the returns were only 78% correlated with the 60/40 benchmark and exhibited a higher tracking error of 7.0%, indicating that investors prioritising benchmark-like performance may experience periods of substantial deviation. Nevertheless, the risk parity portfolio outperformed the 60/40 benchmark during the most challenging equity periods, attributed to the consistent overperformance of bonds during these times.

Ilmanen, Thaper, Tummala and Villalon (2021) explore the efficacy of integrating trend-following strategies with regular index put purchasing as a method for portfolio tail hedging. Specifically, the researchers assess the long-term average costs associated with each approach alongside their reliability. The study employs a strategy that involves buying and rolling 5% out-the-money (OTM) put options with a 1-month tenor as a proxy for tail hedging. This approach is considerably different to those presented in the previous section. Specifically, advocates of trend-following generally suggest puts that are far OTM and cite significantly improved convexity as their primary reason. 5% OTM is thus arguably not a tail risk hedging strategy. Nevertheless, the authors' findings reveal that the option-based strategy consistently mitigates equity drawdowns, particularly during shorter periods, albeit at the expense of negative long-term geometric returns.

Ilmanen et al. (2021) argue that, akin to insurance, put option strategies are expected to incur capital losses when considered in isolation over extended periods. The authors further assert that option pricing incorporates risk-neutral expectations, which encapsulate both market anticipations and the required risk premiums. The presence of the implied volatility smile indicates that markets do not discount Gaussian returns, and the persistent elevation of implied volatilities above realised volatilities may result from biased expectations or the existence of a volatility risk premium, or a combination thereof. However, the authors contend that this discrepancy is inconsequential, as a model-free approach can effectively capture the real-world dynamics inherent in long-term realised option returns. Ilmanen et al. (2021) conclude that during the most severe drawdown episodes, both option-based protection and trend-following strategies have proven to be reliable forms of portfolio hedging, outperforming traditional assets such as US Treasuries and gold.

2.4 TREND-FOLLOWING

2.4.1 Background

Hurst, Ooi and Pedersen (2017) define trend-following as an investment strategy that involves going long markets that have been rising and short markets that have been falling, with the expectation that those trends will continue. Trend-following is also referred to as time-series momentum. In contrast to cross-sectional momentum, which compares the performance of different assets at a specific point in time, trend-following focuses on the performance of individual assets over time. The conceptual foundations of trend-following are deeply rooted in early financial thought. Hurst et al. (2017) reference David Ricardo's 200-year-old adage to "cut short your losses and let your profits run" as a seminal principle underpinning trend-following strategies.

Historically, trend-following has been a widely employed strategy among diverse investor groups, including hedge funds, institutional investors, and individual traders. Its longevity and widespread adoption can be attributed to its simplicity and adaptability across various asset classes and market environments, as well as substantial empirical evidence demonstrating an attractive risk and return profile (Hurst et al., 2017). Empirical research has consistently demonstrated the efficacy of trend-following strategies in generating significant alpha, particularly during periods characterised by sustained market momentum or heightened volatility (Hurst et al., 2017).

Thapar et al. (2019) argue that trend-following presents an attractive strategy for diversifying equity holdings and mitigating downside risk. The researchers' findings indicate that trend-following has historically effectively reduced portfolio volatility and enhances risk-adjusted returns by systematically adjusting exposures in response to prevailing market trends. This diversification benefit was particularly valuable in volatile or turbulent market environments, where traditional asset classes empirically exhibited high correlations and thus reduced diversification efficacy.

Babu et al. (2019) examines the historical performance of trend-following across various major asset classes and trend horizons. The authors' findings demonstrate that trend-following has been effective across a diverse range of market conditions and asset classes, demonstrating its versatility and robustness as a risk management tool. The authors highlight that trend-following strategies have historically delivered consistently positive returns, particularly during market turmoil.

Hurst et al. (2017) further explore the performance of trend-following across different economic environments, revealing significant diversification benefits when integrated into traditional asset allocation mixes. The authors find that trend-following generated positive returns in the vast majority of historical episodes of severe market stress, including the 2008 GFC, during which the conventional 60/40 equity/government bond portfolio experienced substantial drawdowns.

The authors partly attribute trend-following's impressive crisis-time performance to the strategies positioning themselves short falling assets. However, the researchers note that trend-following tends to underperform during rapid crash environments, such as the 1987 stock market crash, highlighting a limitation of the strategy in swiftly deteriorating market conditions. The findings of Till (2016) are similar, suggesting that trend-following strategies require sufficient time to readjust positioning to fully unlock the risk mitigation benefits of a trend-following approach.

Babu et al. (2020) observes that trend-following, over the 2010s, experienced a decade of lower performance relative to previous decades. The authors attribute this underperformance to low volatility and the absence of large positive and negative market moves, which historically drive the success of trend-following strategies. Despite the decade of relative underperformance, the researchers assert that trend-following remains an effective strategy for diversification and can offer attractive prospective returns, provided that the magnitude of market movements reverts to historical levels.

Trend-following also complements optionality as a tool for portfolio risk management. After severe crises, such as the GFC, liquidity recovery can take extended periods (Anand, Irvine, Puckett & Venkataraman, 2013). During these periods, tail risk hedging using optionality remains expensive due to elevated volatility levels (Ilmanen, 2012). In contrast, trend-following offers an alternative, enabling continued risk management without reliance on expensive tail risk options. Brush (1997) suggests that the primary advantage of long-short strategies, including trend-following, lies in their potential to reduce portfolio volatility while still delivering competitive absolute returns.

2.4.2 Testing for trend effects in financial markets

In financial literature, pooled panel regressions are a commonly used statistical method for examining whether financial markets exhibit trending behaviour. This methodology is thoroughly described in Baltas and Kosowski (2013) and Moskowitz et al. (2012). Baltas and Kosowski (2013) investigate the performance of trend-following strategies by analysing the historical returns of a diverse array of lookback and holding period combinations, employing daily, weekly, and monthly return data.

The researchers utilise futures contracts across various asset classes, including equities, bonds, commodities, and currencies, spanning multiple decades. To ensure comparability, the authors normalise returns by dividing each periodic return for every ticker by its lagged ex-ante volatility. This normalisation process standardises the returns, allowing for a consistent scale across different assets and time periods. To determine whether markets trend, the authors conduct pooled panel regressions to assess the predictive power of lagged returns on subsequent returns. Recognising the potential biases in standard errors due to autocorrelation and cross-sectional dependence, they implement robust standard error techniques. Specifically, the authors cluster standard errors by both time and asset. This dual clustering approach accounts for correlations within the same futures contract over time and between different futures contracts within the same time period, thereby addressing both time-series and cross-sectional dependencies.

Petersen (2008) provides detailed insight into the necessity of these adjustments. The researcher finds that if independent variables and residuals are correlated, standard errors in the resulting regression model can be biased downwards, leading to an inflated sense of confidence in the model's reliability. When residuals (defined as the differences between observed and predicted values) are correlated across time or cross-sectionally within the same period, errors are not independently distributed.

Given that the predictive power of lagged returns against subsequent returns is tested by stacking all futures contracts and dates, it is important to control for both time series and cross-sectional dependence in returns. This means addressing autocorrelation in the return series per futures contract and correlations between returns across different futures contracts within any given period, especially since markets are influenced by the same economic environment and may respond similarly to shocks. When executing pooled panel regressions, Baltas and Kosowski (2013) and Moskowitz et al. (2012) calculate the t-statistics of the coefficients using standard errors that are clustered by both time and asset.

2.4.3 The empirical return distribution properties of trend-following strategies

Trend-following strategies have garnered significant attention in financial literature due to their potential to generate crisis alpha and enhance portfolio diversification. Bhansali (2008) analyses various hedge fund strategies and their correlation to tail risk in equities, as proxied by the CBOE Volatility Index (VIX). Bhansali (2008) finds that trend-following strategies exhibit a positive correlation with the VIX while remaining uncorrelated to equity returns. Furthermore, Bhansali (2008) posits that a well-executed trend-following strategy behaves similarly to a series of straddles, thereby replicating option payoff profiles without the explicit use of options. Fung and Hsieh (2001), along with Darius et al. (2002), employ a lookback straddle approach to evaluate the performance of trend-following strategies. Their research demonstrates that lookback straddles better explain the returns generated by trend-following hedge funds compared to standard asset indices.

Hurst et al. (2017) also identify that the return distribution of trend-following exhibits a “smile” with respect to equity returns, akin to a straddle payoff profile. This implies that trend-following performs optimally during periods of consistent market uptrends and downtrends, capitalising on sustained movements and suffering from abrupt reversals.

Kulp, Djupsjöbacka and Estlander (2005) investigate the historical risk-return profile of the S&P Managed Futures Index (of which trend-following strategies represent a key part) and a long volatility portfolio constructed using option straddles from 1 January 2000 to 31 December 2003. Their findings indicate a 0.78 correlation between the two approaches, consistent with prior research. The study concludes that, over the examined period, trend-following outperformed explicit long volatility exposure by 9.50%, while providing comparable risk management qualities. However, the authors acknowledge that sudden and significant changes in implied volatility, such as those experienced during short-term crises, are better captured by explicit long volatility strategies than by trend-following, which requires more time to adjust positions.

Brooks (2017) extends the analysis of trend-following by incorporating security fundamentals alongside price-based signals. The study finds that using fundamentals to establish trends also results in an attractive historical risk-return profile. Brooks (2017) notes that the 'smile' shaped return distribution of trend-following strategies, which signifies strong returns during extreme market movements, is particularly pronounced during down markets when following fundamental trend signals. This is attributed to the rapid deterioration of fundamentals preceding significant equity crises. Moreover, the author suggests that combining price and fundamental trend signals offers complementary benefits, as each approach tends to experience drawdowns during different periods while maintaining effective risk hedging properties.

Harvey et al. (2021) conducts a comprehensive study on trend-following strategies, examining their efficacy as a source of crisis alpha over a multi-decade horizon beginning in 1960. Given the absence of a benchmark extending this far back, the researchers construct a proxy by replicating a trend-following portfolio that closely resembles the BTOP50 Managed Futures Index, which is an industry benchmark with data going back to 1987. This replication involves building a systematic approach backdated to 1960, utilising data from broad-based asset classes including fixed income instruments, equities, commodities, and currencies up to 2015.

The methodology employed by Harvey et al. (2021) involves using monthly returns and scaling them by their volatilities for signal generation and targeting an ex-ante annualised volatility of 10%. As is typical in trend-following strategies, risk is evenly allocated across global equities, bonds, commodities, and currencies, with equal risk distribution within each asset class. Despite using monthly data (whereas the underlying managers of the BTOP50 Index likely utilised daily or even tick data for signals and risk management), the replicated portfolio achieved a correlation coefficient of +0.62 with the BTOP50 Index.

Over the full period from 1960 to 2015, the replicating portfolio delivered a Sharpe ratio of +1.56 with an ex-post volatility of 10.5% per annum, closely aligning with the ex-ante target of 10%. It is important to note that this performance does not account for transaction costs. The authors argue that while historical performance might be affected by such costs, the fundamental dynamics of trend-following returns remain intact. They estimate that a 0.02% portfolio-level transaction cost per trade would result in a 0.42% annualised reduction in portfolio performance, consistent with their real-world trading experiences.

A notable finding from Harvey et al. (2021) is the positive skewness of trend-following returns. Over the full period, skewness measured 1.04 for 3-month overlapping returns and 1.48 for 12-month overlapping returns. This positive skewness is a rare trait, as strategies that have historically delivered attractively high positive absolute returns, such as broad-based equity investments, typically exhibit negative skewness.

The researchers observed that positive skewness became more pronounced over longer measurement horizons (quarterly and annually versus daily or weekly), suggesting that trend-following is particularly suitable for long-horizon investors, such as institutional investors. The tendency for trend-following returns to behave similarly to long straddles, where the strategy increases investment in winning positions and reduces exposure to losing ones, partially explains this positive skewness. This characteristic was robust across different sub-periods, including before and after 1985, covering varying long-term interest rate regimes, and remained evident even when excluding 2008, the best single year for the strategy under investigation.

Daniel and Moskowitz (2016) delve into the risks and characteristics of momentum strategies, specifically their propensity for sudden drawdowns, termed 'momentum crashes.' These crashes often occur after extended periods of trending prices that reverse abruptly due to catalyst events. The researchers highlight the importance of understanding these risks when employing momentum-based strategies, including trend-following, to manage expectations and mitigate potential losses.

Shepherd et al. (2018) discusses the advantages of strategies characterised by long/short exposures to various asset classes, emphasising their capacity to enhance the likelihood of achieving long-term return objectives, such as the challenging 5% annualised real return target often sought in retirement planning. The researchers suggest that alternative risk premia (ARP) strategies, of which trend-following is a key component, have historically exhibited low correlation with traditional assets like equities and bonds. Their findings indicate that a well-constructed ARP strategy can contribute to portfolio diversification, reduce volatility, and improve risk-adjusted returns. Shepherd et al. (2018) demonstrates that portfolios allocating 20% to ARP can more than triple the probability of meeting the 5% real return target over the next decade. The systematic, transparent, and liquid nature of ARP strategies enhances their appeal compared to more complex and less accessible hedge fund structures (Brightman & Shepherd, 2016).

Beck, Hsu, Kalesnik and Kostka (2016) provides a comprehensive analysis of momentum strategies within the context of factor investing. They highlight that momentum strategies, initially documented by Jegadeesh and Titman (1993), are designed to exploit the continuation of stock price trends, where past winners continue to outperform, and past losers underperform. The researchers suggest that high turnover costs due to frequent rebalancing and the risk of momentum crashes, as documented by Daniel and Moskowitz (2016), are key factors contributing to the sustained profitability of momentum and trend-following strategies.

Ilmanen et al. (2021) examine trend-following strategies as diversifiers for equity portfolios, employing the same trend-following portfolio described in Hurst et al. (2017). In contrast to the protective put strategy analysed alongside, the authors find that the trend-following approach not only adds value during significant equity drawdowns but also generates attractive absolute returns over time, despite muted performance during the 2010s. This reinforces the role of trend-following as a valuable component in portfolio construction, offering both defensive qualities during market downturns and growth potential during varying market conditions.

2.4.4 Trend-following strategies diversification potential with respect to other assets

While substantial evidence exists regarding the protective qualities of trend-following strategies on equity returns, the literature on whether this has empirically held true for government bonds is now reviewed. Together with equities, government bonds form the core of many long-horizon investment strategies. Therefore, a strategy like trend-following offering downside protection for both bonds and equities, rather than equities alone, holds significant value.

Harvey et al. (2021) explores this proposition by examining the BTOP50 Managed Futures replication strategy's performance during adverse market conditions. The authors find that their replication strategy delivers significant absolute returns during the worst equity and bond market environments, reinforcing the notion of trend-following as a reliable source of crisis alpha. While numerous studies have highlighted the diversification benefits of trend-following in relation to equity returns, this research extends the potential benefits to bond returns as well. This suggests that trend-following can offer diversification advantages not only for equity-centric portfolios but also for balanced strategies comprising equities and bonds, including risk parity strategies that often have substantial bond exposures.

In addition to outperforming during adverse market conditions, Harvey et al. (2021) note that trend-following strategies have also generated positive absolute returns during strong equity and bond market environments. By applying a quadratic fit to the observations when comparing trend-following returns with matching period equity returns, the authors find a straddle-like “smile” payoff profile, consistent with other trend-following literature. This “smile” pattern is even more pronounced when trend-following returns are compared with matching period bond returns, indicating the strategy's ability to diversify other risky assets beyond equities.

To further demonstrate the diversification benefits of trend-following during crises, Harvey et al. (2021) compute rolling 3-month returns for the SPX and 10-year US Treasury bonds, sorting them into quintiles to compare corresponding trend-following returns. The authors discover that trend-following returns are highest in the lowest quintile of both SPX and 10-year US Treasury bond returns, and second highest in the highest quintile of these returns.

Moreover, Harvey et al. (2021) performs a return attribution of trend-following across sub-asset classes, including equities, bonds, currencies, and commodities. The authors find that equities, bonds, and currencies have generated fairly consistent performance in both the lowest and highest quintiles of SPX and 10-year US Treasury bond returns. However, commodities display more of a left skew, delivering outsized performance during the lowest equity and bond quintiles. These findings remain robust when applying 12-month rolling returns, reinforcing the strategy's effectiveness across different time horizons.

Butler and Butler (2023) analyse performance data from the Barclay BTOP50 Index, a well-known industry benchmark, from 1 August 1990 to 31 January 2023. The authors note that during this period, the BTOP50 Index generated competitive absolute returns comparable to benchmark equity and fixed income indices, with volatilities roughly twice that of bonds and half that of equities, and with low correlation to both. Consequently, the authors state that trend-following is considered an ideal strategy to complement a traditional strategic asset allocation comprising stocks and bonds.

Complementing this research, Hurst et al. (2017) conduct an extensive study of trend-following performance by constructing a time-series momentum strategy spanning from 1880 to 2013, covering over a century of data. The investment universe consists of 67 markets, including 29 commodities, 11 equity indices, 15 bond markets, and 12 currency pairs. The authors splice data from various sources to backfill historical information. The researchers implement a simple time-series momentum rule, establishing positions based on an equal weight of 1-month, 3-month, and 12-month lookback period returns for each market. Long positions are taken in markets with positive returns, and short positions in those with negative returns, with position sizing targeting a 10% ex-ante volatility to maintain consistent risk over time.

After accounting for estimated transaction fees and applying both a management fee of 2% and a performance fee of 20% subject to a high-water mark, Hurst et al. (2017) find that trend-following has historically delivered positive absolute returns in every non-overlapping decade examined. Over the entire period, the strategy achieved an annualised return of 11.2% net of fees and costs, with realised volatility of 9.7%, aligning with the ex-ante volatility target of 10% and a Sharpe ratio of +0.77. The correlation with US equities was 0.0, and with US Treasury bonds was –0.04, indicating minimal correlation and thus significant diversification benefits.

Furthermore, Hurst et al. (2017) identify a time-series momentum "smile" when plotting annual trend-following returns against US equity market returns, indicated by a polynomial line of best fit. This "smile" is clear evidence of trend-following delivering convexity with respect to equity returns. This pattern suggests that trend-following has tended to outperform when equities have both outperformed and suffered their worst episodes, while exhibiting low correlation during periods of average equity returns. The authors note that this "smile" pattern is even more pronounced when comparing trend-following returns with matching period bond returns, indicating the strategy's ability to diversify other risky assets beyond equities and to be a source of convexity for bond returns as well.

Hurst et al. (2017) also focus on trend-following outcomes during the 10 most significant peak-to-trough drawdown events in their sample period, contrasting performance against a US 60/40 equity/Treasury bond portfolio. The authors conclude that trend-following generated positive returns in 8 of the 10 observed drawdowns. Notable periods of valuable hedging contribution include the 2008 GFC, the Dot-Com bubble burst of the early 2000s, the oil crisis in 1973–1974, the Great Depression, and World War I. The authors attribute trend-following outperformance during bear markets in US equities to the tendency for gradual losses to accumulate over time before crises become fully entrenched, providing trend-followers with sufficient time to position appropriately and profit from trend continuations. This theory is supported by their finding that the average duration of the 10 deepest drawdowns in the study was 15 months.

2.4.5 Allocations to Trend-Following as part of a Portfolio

Asvanunt et al. (2015) incorporate a 20% allocation to a trend-following strategy into the 60/40 equity/bond portfolio. The authors contend that bear markets typically develop gradually, resulting from extended economic downturns. The trend-following portfolio is constructed as an equal-weighted blend of 1-, 3-, and 12-month time series momentum rules, where positive returns during formation periods lead to long positions and negative returns result in short positions, with position sizes targeting a 10% ex-ante volatility. This strategy yielded marginally higher average returns than the 60/40 benchmark, enhancing the Sharpe ratio from +0.55 to +0.68. The portfolio achieved a statistically significant alpha of 1.3% over the benchmark, with returns being 96% correlated and a modest tracking error of 2.7%. Consistent with previous findings, the trend-following strategy bolstered the portfolio's performance during the worst equity months.

Overall, Asvanunt et al. (2015) conclude that trend-following successfully generated positive and significant alpha relative to the 60/40 equity/fixed income benchmark portfolio. However, it did not consistently add value during sudden and short-lived market crashes. Specifically, the collar strategy provided dependable defensive characteristics during abrupt crises, but the associated hedging gains were diminished during subsequent equity recoveries. The authors suggest that exceptional timing capabilities are essential to optimally implement direct hedges and effectively unwind them post-tail risk events to maximise the benefits of such strategies.

Krishnan (2017) demonstrates that trend-following, as a hedge fund style, has historically generated strong performance when equities and other hedge fund styles have performed poorly. This suggests its usefulness as a source of crisis alpha and as a general diversifier for an equity-dominant portfolio. To compute an optimal allocation to trend-following, the author uses the SPX and the Barclays CTA Index as proxies for equity and trend-following, respectively. By making simplifying assumptions regarding prospective expected returns, such as 5% for equities and 0% for trend-following (a conservative estimate given the historically significant positive returns of trend-following strategies), Krishnan (2017) concludes that a 22% allocation to trend-following maximises the Sharpe ratio of an equity/CTA portfolio.

Hurst et al. (2017) examine the impact of incorporating trend-following strategies into a traditional 60/40 equity/bond benchmark portfolio. The authors adjust the benchmark pro-rata to accommodate a 20% capital allocation to their trend-following strategy. The authors find that from 1880 to 2013, the resulting portfolio delivered an annualised net return of 8.3% with an ex-post volatility of 8.8%, compared to the benchmark's performance of 7.8% with an ex-post volatility of 10.8%. The maximum drawdown experienced by the benchmark was reduced from –62.3% to –50.2% when trend-following was included. The Sharpe ratio improved from +0.38 for the benchmark to +0.54 with the inclusion of trend-following, indicating enhanced risk-adjusted returns.

Similarly, Clare, Seaton, Smith and Thomas (2016) finds that integrating trend-following into an equal-weighted strategic asset allocation mix enhanced risk-adjusted returns. Clare et al. (2016) demonstrates that including trend-following significantly reduced historical portfolio volatility and negative skewness, as well as maximum drawdowns. The researchers highlight these risk reduction benefits as particularly relevant for investors nearing retirement, where purchasing an annuity is a consideration and large drawdowns could materially affect future living standards.

Harvey et al. (2021) investigates whether constraining trend-following strategies by prohibiting long positions in equities and bonds enhances crisis alpha. Specifically, the authors restrict long positions in equities and bonds to zero and compare the resulting trend-following returns against corresponding rolling 3-month SPX returns and 10-year US Treasury bond returns, sorted into quintiles. Their findings indicate that, for the SPX, constraining long equity positions significantly increases performance in the lowest quintile of returns, suggesting that crisis alpha potential is enhanced when long positions in equities are prohibited. However, this enhancement comes at the expense of reduced performance in quintiles 3, 4, and 5 relative to the unconstrained trend-following strategy. Similar results are observed when examining the 10-year US Treasury bond return quintiles. The authors conclude that while crisis alpha was historically larger when long positions in equities and bonds were disallowed, the inability to harvest risk premia during other market conditions led to reduced average returns.

2.4.6 Variations in systematic trend-following protocols

Harvey et al. (2021) investigates how trend-following outcomes vary when different lag periods are used to determine strategy signals. The authors find that historical performance is particularly strong when trends are measured using relatively short-term lags of 1 to 4 months. Similar positive results are observed when trends are established using longer-term lags of 9 to 11 months. However, consistent underperformance is detected when lags of 5 to 8 months are used, a finding that is robust across all assets studied. In an attempt to replicate the BTOP50 Managed Futures Index (a widely used benchmark with historical data spanning multiple decades) the researchers use monthly returns with lags of 1 to 4 months on broad-based equities, fixed income instruments, commodities, and currencies. The authors find that this approach explains the majority of the benchmark's returns. This finding differs somewhat from existing literature regarding 12-month lags, indicating that the 12-month lag has historically offered only marginally positive returns and significantly less predictive power than the shorter lags mentioned.

The authors argue that this discrepancy arises from the use of monthly data, as opposed to daily data, which allows for more nimble positioning and more frequent rebalancing. For example, when using a 12-month lag with monthly data, the position for January 2000 would be determined using the return from January 1999 and would only be adjusted in February 2000. In contrast, daily data would permit the position on 15 January 2000 to include additional data up to that date, providing a more current assessment. This distinction is particularly important for a 12-month lag, considering that Baltas and Kosowski (2013) found that historical weekly returns are positively predictive for lags up to 52 weeks and negatively predictive for lags of 53 and 54 weeks. Therefore, it is conceivable that some of this adverse predictive power affected the 12-month lag in Harvey et al. (2021).

To replicate the excess returns of the BTOP50 Managed Futures Index, Harvey et al. (2021) compute returns for portfolios constructed using trend signals formed from lags of 1 to 11 months. The researchers find a U-shaped quadratic relationship (commonly referred to as a "smile"), with index returns best explained by the shorter lags and, to a lesser extent, the longer lags, while intermediate lags exhibit low explanatory power. The authors find that 76% of the optimal quadratic weights are attributed to the first 4 lags. Furthermore, there is a correlation coefficient of +0.92 between the more elaborate trend-following portfolio computed using the U-shaped quadratic fit and a simpler portfolio employing equal weights across lags 1 to 4.

The authors suggest that this finding indicates that trend-followers predominantly utilise fairly short-term signals to establish positions. The researchers also observe that the quality of the signals diminishes at lags 5 to 8, which they suggest is consistent with markets initially underreacting to new information, followed by overreactions in subsequent periods. They are less certain about the reasons for the strong signal quality at lags 9 to 11, positing that seasonal effects and annual reporting cycles in financial data may offer an explanation. Harvey et al. (2021) find that shorter lags contribute the majority of the positive skewness exhibited by trend-following strategies.

According to Zakamulin and Giner (2020), the two most popular methods used to establish trend-following signals are momentum rules and moving average rules. Momentum rules generate a buy signal when the current price of an instrument exceeds its price *n* periods ago. In contrast, moving average rules establish a buy signal when the current price is higher than its average price over the past *n* periods. The simple moving average is most commonly used (Zakamulin and Giner, 2020).

Marshall, Nguyen and Visaltanachoti (2017) also compares momentum rules with moving average rules and find significant overlap between the approaches. The researchers report a correlation of 0.78 between the return series of long-only strategies developed using momentum and moving average rules, respectively. However, across many strategies investigated, the authors find that moving average rules consistently generate higher returns and higher Sharpe ratios.

The researchers attribute this to moving average methodologies resulting in earlier signal generation than momentum rules, leading to improved return statistics. Specifically, with moving average rules, a long position is established when the current price exceeds the average price over the past *n* periods (e.g., 200 days), which occurs more frequently than the corresponding holding period return turning positive in a momentum rule. This allows moving average rules to establish positions sooner and more frequently, and to exit losing positions earlier than momentum rules. The authors also find that these results hold even after accounting for higher transaction costs associated with moving average rules. Neely, Rapach, Tu and Zhou (2014) find similar results, concluding that moving average rules tend to outperform momentum rules.

Hong and Satchell (2015) explore the reasons behind the popularity of moving average rules, deriving the autocorrelation structures of various markets. The authors propose that moving average rules have become dominant because they can identify price trends with minimal computational effort and provide a straightforward method of utilising price autocorrelation structures without requiring a deep understanding.

The authors suggest that market participants following this approach may also amplify price autocorrelation structures, demonstrating how stochastic processes can have their autocorrelation properties transformed in line with the investment horizons of trend-followers. Hong and Satchell (2015) recommend using short-term moving average rules when asset price, earnings, and return autocorrelations are low, and long-term moving average rules when autocorrelations are high. This, the authors theorise, explains why many practitioners employ long-term simple moving averages to capture price momentum, as autocorrelations tend to be higher over longer horizons.

2.4.7 Outlook for trend-following expected returns

Harvey et al. (2021) note that US government bonds benefited from a prolonged bull market from 1985 to the end of their study period. This strong performance was accompanied by relatively low levels of realised volatility, with US government bonds (proxied by 10-year Treasury bonds) seldom experiencing drawdowns exceeding 10%. The authors attribute this multi-decade period of robust performance to a compression in interest rates, with yields declining from close to 16% in the early 1980s to below 2% during the first quarter of 2016. The authors further observe that while many countries have experienced negative real interest rates over this period, some have even faced negative nominal yields across the yield curve. The authors indicate that continued yield compression from such low levels is unlikely, as interest rates cannot remain significantly negative for extended periods.

Against this backdrop, Harvey et al. (2021) conduct research to determine the likelihood of sustained strength in trend-following performance in the likely absence of a bond market tailwind, given its significant contribution to trend-following returns. Specifically, they investigate whether expected returns on trend-following strategies remain positive and economically significant during periods when bond yields rise. The authors highlight a stark contrast in bond returns between the periods 1960–1985 and 1985–2015. In the former, bond excess returns were negative on average, whereas in the latter, they were strongly positive on average. Consequently, the earlier period offers fertile ground to examine trend-following performance during rising yield environments.

Niederhoffer and Weddepohl (2014) suggest that much of the returns generated by trend-following strategies since the 1980s can be attributed to the bond bull market, casting doubts on prospective return expectations given the current low starting point for US Treasury bond yields. The researchers argue that trend-following strategies have delivered disappointing returns during periods of rising interest rates and conclude that the approach is highly dependent on the interest rate regime.

In contrast, Hurst et al. (2017) conduct a study covering over a century of empirical trend-following returns from 1880 to 2013 and find the opposite to be true. Hurst et al. (2017) report that trend-following has generated positive returns net of transaction costs and fees in every non-overlapping decade within their sample period. Notably, they highlight the 1970s as the best-performing decade for trend-following, during which the 10-year US Treasury bond yield rose from 7.8% to 11.1%, accompanied by high volatility.

Despite this challenging environment, their trend-following strategy yielded a 21.3% annualised return net of transaction costs and fees, with a 9% ex-post volatility (compared to a 10% ex-ante target) and a Sharpe ratio of +1.70. The portfolio exhibited a correlation of –0.24 with US equity returns and –0.25 with US Treasury bond returns during a decade characterised by rising bond yields. Hurst et al. (2017) also note that trend-following experienced a mixed period of performance following the 2008 GFC up to the end of their sample period in 2013.

The authors examine popular explanations for this lacklustre performance. First, they consider the substantial growth in assets under management (AUM) employing the strategy over time, which could potentially impact return profiles. Second, they suggest that the economic environment, marked by a lack of clear trends and frequent trend reversals, has not been conducive to the strategy's success. Regarding the AUM hypothesis, the authors cite BarclayHedge estimates indicating that total AUM employing trend-following strategies grew from $22 billion in 1999 to $280 billion in 2014. While acknowledging this significant growth, they argue that the underlying markets, already vast, have also expanded over the period, resulting in trend-following strategies holding a relatively stable percentage of the underlying markets.

The researchers do not consider the mixed performance in the latter years of their study to be anomalous, noting that strategy drawdowns were not unusually large or persistent relative to historical norms. However, they acknowledge an increase in pairwise correlations within and across the markets comprising the trend-following strategy. The authors also state that transaction costs have decreased considerably over time due to increased competition in market-making.

2.4.8 Replicating trend-following indices

Butler and Butler (2023) propose an alternative method to generate trend-following returns other than by means of constructing a strategy using similar approaches to those covered in this literature review so far. The authors detail a methodology that allows investors to broadly replicate the returns of the SG Trend Index, which is a well-established industry benchmark

made up of the top 10 trend-following funds (by AUM) that are open to new investment. The index is equal-weighted and reconstituted at the beginning of every year.
The authors propose the use of statistical methods to systematically replicate the beta of SG Trend Index returns. By doing so, the authors argue that investors can harvest the majority of returns stemming from market inefficiencies while foregoing attempts at earning alpha, thereby saving on the high fees associated with alpha pursuit. However, the authors caution against allocating to a single trend-following fund or a small sample of funds due to the wide historical dispersion in fund performance resulting from diverse methodologies. The reason attributed to this is that their research indicates that individual trend-following funds deviated from the average of all funds by 9% on average from 2018 to 2022.

Highlighting the high degree of dispersion among trend-following mutual funds, Butler and Butler (2023) set out to replicate the SG Trend Index, which has published daily returns since 2000, in contrast to the Barclay BTOP50 Index. Although the Barclay BTOP50 Index has provided returns since 1987, these were published monthly until 2010, after which daily returns became available. The SG Trend Index, however, has offered daily returns from its inception in 2000. The researchers note that there are two main methods for replicating benchmark indices. The first is a "top-down" replication, which uses a similar approach to return-based style analysis by regressing short-term index returns against broad asset class returns to estimate the average holdings of the managers that constitute the index at a single point in time. The second is a "bottom-up" approach, wherein the researcher attempts to uncover the strategy that best reproduces the index returns.

To conduct their analysis, Butler and Butler (2023) utilise 27 futures contracts across fixed income, currencies, energy, equities, and metals. They also create a subset of nine futures in the same sectors, referred to as the "small universe." For the top-down replication approach, the authors compute a daily fit in which short-term index returns are used as the dependent variable, and the universe of futures returns serves as independent variables. They scale daily returns to match the same rolling 40-day volatility, computed as an exponentially weighted moving average, and then regress these returns on the returns of the SG Trend Index. An Elastic Net regression is employed for this task, as it can simultaneously identify the most influential futures at any given point in time while controlling for over-allocations in any single market.

The researchers fit the Elastic Net regression using rolling returns over the past 20 to 40 days in 5-day increments and apply a linear model to the output to compute the weights of the individual models in the final top-down model. The results presented in Butler and Butler (2023) are impressive, with the replication accurately capturing the performance of the underlying index. As such, index replication through the use of robust regression techniques offers a viable approach to capturing trend-following beta.

2.5 PORTABLE ALPHA

Portable alpha strategies have gained prominence as a means to enhance portfolio performance by separating market returns (beta) from active returns (alpha). Asness (2004) discusses the concept of portable alpha in detail, highlighting its ability to improve traditional portfolios by using derivatives, such as futures or swaps, to maintain exposure to market beta while simultaneously deploying capital into alpha-generating strategies like hedge funds. Traditionally, both tail risk hedging and trend-following strategies have been implemented in hedge funds, making the approach relevant.

This approach enables investors to "port" alpha onto the base portfolio, capturing both passive market returns and potential outperformance through active management. The key advantage lies in the efficient use of capital, as the same capital base supports both beta exposure and the alpha overlay.

Asness (2004) identifies several benefits of portable alpha, particularly in environments where traditional portfolios composed of equities and bonds struggle to generate sufficient returns. With equity markets becoming more efficient, the potential for outperformance via traditional active management diminishes. Portable alpha addresses this issue by combining passive market exposure with alternative alpha sources.

One core finding from Asness (2004) is that portable alpha can significantly improve risk-adjusted returns. By adding an uncorrelated alpha source, the portfolio retains broad market exposure while benefiting from excess returns generated by skilled managers or uncorrelated investment strategies, leading to potential improvements in performance. A significant advantage of portable alpha is its ability to diversify return sources. While the beta component remains tied to traditional market movements, the alpha component can be derived from strategies like market-neutral hedge funds or alternative investments that exploit inefficiencies.

This separation of return drivers mitigates risk and provides downside protection during market corrections. For instance, when the beta portion underperforms, the alpha component may still produce positive returns, stabilising overall portfolio performance during periods of market stress (Asness, 2004). Adding hedge fund alpha to a base beta portfolio allows investors to earn total return from traditional market exposure while also benefiting from the excess returns generated by hedge fund strategies.

However, the implementation of portable alpha is not without challenges. Asness (2004) concedes that the strategy depends heavily on manager skill in generating consistent alpha. Portable alpha strategies can introduce complexity, particularly with the use of derivatives, and may require sophisticated risk management and governance structures. Despite these considerations, Asness (2004) concludes that the long-term benefits of portable alpha make it a valuable strategy for investors to consider.

More recently, AQR Capital Management (2024) reiterate that portable alpha strategies are an effective solution for improving risk and return profiles of long-horizon investors. The research emphasises the importance of diversifying alpha sources, suggesting that relying on multiple alpha strategies (such as equity market-neutral or trend-following) can further reduce the risk of underperformance. Ensuring that the strategies that are overlayed generate alpha that is independent of market beta (and of one another) provides meaningful diversification (AQR Capital Management, 2024).

Cavaglia, Fan and Wang (2022) provide a detailed examination of portable alpha, emphasising its effectiveness in enhancing the return profile of a traditional beta portfolio by overlaying hedge fund strategies. The core premise, according to the researchers, is that investors can retain market exposure while simultaneously generating additional returns by overlaying alpha strategies on top of a base beta portfolio. This structure allows for portfolio outcomes to remain largely driven by the market, but with a possibility of outperformance stemming from the hedge fund strategies being overlaid. The key benefit is that equity dilution is not necessary to overlay potentially valuable overlay strategies such as trend-following and tail risk hedging.

2.6 COMBINING TREND-FOLLOWING AND TAIL RISK HEDGING

Bhansali (2013) proposes four distinct approaches to effective risk management: cash, diversification, trend-following strategies, and tail risk hedging. The author states that these methods are designed to mitigate portfolio risks of varying severities. While most investors are familiar with the use of cash and diversification to manage risk, it is notable that trend-following and tail risk hedging represent the other two modes of effective risk management according to Bhansali (2013) and are less widely known.

Ilmanen et al. (2021) examine the effectiveness of combining trend-following strategies with systematic index put purchases for tail risk hedging in portfolios. The authors evaluate the long-term costs associated with each method and analyse their reliability in mitigating tail risk. Ilmanen et al. (2021) find that there are significant complementary benefits of using these approaches in combination, with trend-following effectively mitigating extended market declines and put options providing protection during sharp market crashes.

Ilmanen et al. (2021) conclude that both option-based protection and trend-following strategies have consistently served as effective portfolio hedges during the most significant market downturns. Moreover, the authors find that each approach was more effective than traditional assets, such as US Treasuries and gold, in providing downside protection during periods of extreme volatility, proving to be more reliable forms of risk mitigation in times of severe financial stress.

McQuinn, Thapar and Villalon (2021) further investigate the roles of trend-following and option-based hedging as portfolio safeguards against drawdowns. The authors posit that slower, more prolonged drawdowns are more detrimental to long-term investors compared to rapid, sharp declines of equivalent magnitude. To illustrate this, they present a hypothetical scenario involving two portfolios, each generating a 6% annual return. The first portfolio experiences a swift 20% crash within a single month before resuming a 6% annual growth rate, whereas the second portfolio undergoes a gradual 20% decline over the course of a year.

This example demonstrates that spreading a 20% loss over a longer duration results in a more adverse terminal outcome, as the prolonged drawdown negates an entire year's worth of earnings, unlike the swift loss which only affects a single month's returns. Empirical data from the 2008 GFC and the Tech Burst of the early 2000s corroborate the robustness of this illustrative example.

McQuinn et al. (2021) then evaluate two option-based hedging strategies that involve buying and rolling unhedged put options at 10% and 20% OTM levels, respectively. Both strategies are rebalanced quarterly and feature a one-year tenor. The authors highlight that passive option-based hedging strategies can vary based on specifications such as rebalancing protocols, option tenors, levels of moneyness, and the implementation of delta hedging.

This analysis reveals that option-based hedges perform exceptionally well during short, sharp crashes but their effectiveness diminishes as drawdowns extend over longer periods. The study identifies defensive equity styles, including low-beta and risk parity approaches, as effective means to reduce tail risk within a 60/40 equity/bond portfolio while simultaneously enhancing both risk-adjusted and absolute returns. Furthermore, McQuinn et al. (2021) find that trend-following strategies serve as a valuable source of crisis alpha, particularly during extended bear markets.

Krishnan (2017) investigates trend-following strategies as a portfolio hedge following significant market declines. The author notes that implied volatility levels are generally elevated after large market drops due to surges in the realised volatility of risky assets. Increases in implied volatility raise option prices, making tail hedging through explicit option purchases prohibitively expensive. Krishnan (2017) asserts that volatile sideways markets are unfavourable environments for trend-following strategies, whereas persistently upward or downward markets are ideal. The author argues that post-sell-off market environments tend to be highly volatile with a strong likelihood of directional moves, making trend-following particularly effective.

Krishnan (2017) highlights the significant benefit of capacity offered by trend-following strategies, as the underlying futures markets are highly liquid and traded in large volumes. The largest trend-following managers manage in excess of $100 billion in total gross notional exposure. Another notable feature is that trend-following strategies are fully committed to hedging in down markets without profit-taking strategies typical of tail risk hedging using options. The author contends that trend-following complements tail hedging approaches using options, which become progressively less potent as implied volatility rises.

Malek and Dobrovolsky (2009) investigate trend-following performance as a function of underlying market volatility. While acknowledging that trend-following returns have historically been strongest during the worst periods for equities, the authors challenge the notion that this implies that trend-following is a long volatility strategy that benefits from rising volatility levels and underperforms when volatility falls. The authors differentiate between changes in volatility and general levels of volatility, focusing on disassociating trend-following returns from volatility changes, which is more appropriate for option-based hedging.

The authors' findings suggest that trend-following is not so much a 'long volatility' strategy as a 'long gamma' strategy. Specifically, in a flat market environment where a trend-following strategy has no exposure, its directional market exposure (strategy delta) is zero. If the underlying market moves upwards, the trend-following strategy adopts a long position, making the strategy delta positive relative to the underlying market. Conversely, if the market moves downwards, the strategy takes a short position, rendering the strategy delta negative. Malek and Dobrovolsky (2009) argue that this behaviour is akin to a straddle position, which is long gamma, with the strategy delta dependent on underlying market price movements rather than volatility movements.

Krishnan (2017) adds that while a trend-following strategy eventually positions itself favourably in risk-off markets, this occurs in response to adverse price changes that have already transpired. This suggests that trend-following strategies are not direct replacements for long option-based approaches, which directly benefit from surges in volatility.

Moreover, Krishnan (2017) suggests that trend-following offers investors the ability to diversify across time. Standard risk models focus on diversifying across assets at a single point in time, where assumptions about covariance between assets are crucial. The author posits that investing in trend-following strategies constitutes a departure from traditional frameworks, as covariances become dynamic due to the rebalancing of positions based on underlying market movements. This view is demonstrated by showing that the pairwise correlation between different moving average crossover strategies decreases as lookback periods increase, creating a diversifying effect.

2.7 CHAPTER SUMMARY

This chapter conducted a comprehensive literature review on the essential concepts upon which this study is based, conducting an extensive review of risk management strategies for long-term investment portfolios. It began by discussing the concept of variance drag and the importance of reducing portfolio volatility through diversification and rebalancing. The limitations of traditional diversification were highlighted and the need for enhanced risk management tools beyond conventional methods was emphasised.

The literature on tail risk hedging was thoroughly explored, with an emphasis on option-based strategies and the use of over-the-counter instruments such as variance and volatility swaps. Studies on the derivation of risk-neutral density functions and their role in interpreting market expectations were examined. Empirical analyses illustrated the effectiveness of tail risk hedging strategies during extreme market events, despite criticisms regarding their cost and implementation complexity.

The chapter also examined trend-following strategies, outlining their historical efficacy across various market conditions. Different methodologies for implementing trend-following strategies were assessed, including variations in signal generation and the impact of different lookback periods. The diversification potential of trend-following was highlighted, particularly its ability to provide crisis alpha and complement other risk management strategies.

The concept of portable alpha was also explored, with a focus on how overlaying alpha-generating strategies like trend-following and tail risk hedging onto a beta portfolio can enhance overall performance without diluting market exposure. Finally, literature covering the synergistic benefits of combining tail risk hedging and trend-following to manage risk and enhance expected return was examined.

CHAPTER 3: TREND-FOLLOWING

3.1 CHAPTER INTRODUCTION

This chapter introduces the trend-following strategy employed throughout this thesis and explores its impact when overlaid on global equity exposure. Trend-following is a hedge fund strategy that invests across multiple markets, including equities, government bonds, currencies, and commodities. The strategy seeks to exploit trending behaviour by buying markets that have been rising, selling markets that have been falling, and profiting from continuing trends. With respect to equity returns, trend-following strategies have historically exhibited negative correlation in the left tail, positive correlation in the right tail, and little to no correlation in the body of the distribution (Moskowitz et al., 2012). The strategy generally uses futures contracts for implementation, which offer cash efficiency at the portfolio level. This is because futures typically require only a small percentage of the total exposure to be posted as margin, while prime brokers can partially utilise existing cash positions as collateral for margin requirements.

Given the attractive historical payoff profile and the cash efficiency of using futures, this thesis contends that the strategy is well-suited for a portable alpha framework, where portfolio beta (global equity) and alpha (trend-following that is overlaid) are separated. This chapter examines whether the combined portfolio-level results are enhanced due to diversification benefits stemming from the low correlation between trend-following and the underlying equity portfolio. Improvements in risk-adjusted returns from trend-following are partly attributable to the introduction of convexity in relation to both left-tail and right-tail equity returns when overlaid on a global equity portfolio. Testing for convexity will form a central focus of this chapter.

The study begins by examining the core phenomenon that underpins trend-following strategy returns through hypothesis $H_{0,A}$: Global markets do not trend. It then constructs a simple trend-following strategy based on systematic rules commonly found in the literature. To ensure that the strategy is representative of a typical trend-following approach, the performance of the resulting Trend Portfolio is compared to the SG Trend Index, a trend-following industry benchmark, by testing hypothesis $H_{0,B}$: SG Trend Index returns do not have a significant relationship with Trend Portfolio returns.

Consistent with existing research, the study also demonstrates that the returns generated by the trend-following strategy are uncorrelated with traditional asset classes and equity market factors, making the strategy a compelling overlay candidate for global equity exposure in a portable alpha application. To support this, it tests $H_{0,\,C}$: Trend-following returns are fully explained by traditional equity market factors (global market, SMB, HML, UMD), global government bonds, and commodities. The study evaluates the diversification benefits that trend-following provides by assessing whether the strategy introduces convexity with respect to both left-tail and right-tail equity returns through $H_{0,\,D}$: Trend-following returns are not convex with respect to global equity returns.

It evaluates $H_{0,\,E}$: Overlaying a trend-following strategy on top of global equity does not enhance risk-adjusted and absolute returns compared to holding global equity alone. To do this, the study analyses the long-term risk and return characteristics of ACWI and the Trend Overlay portfolio, formed by overlaying 50% of its trend-following strategy on top of 100% ACWI exposure.

3.2 DATA DESCRIPTION

In this section, the various sources of data used to construct the trend-following strategy are discussed. The data used comprise futures contracts, exchange rates and interest rates. It also outlines the data used for regression analyses and performance comparisons. Given the global exposure of trend-following strategies, which utilise futures contracts traded on exchanges worldwide, public holidays do not perfectly coincide across different markets. Consequently, there may be instances where some futures contracts are traded on a given day while others are not due to market closures for public holidays. To address this, the data are standardised to include all weekdays in a typical non-leap year, amounting to 261 trading days. If an exchange is closed on a particular day due to a holiday, it assigns a 0% return to the instruments traded on that exchange for that day.

3.2.1 Futures contract data

The data used in the trend-following strategy comprise 79 futures contracts, as detailed in Table 3.1. These data include 23 equity index futures (18 from developed markets and five from emerging markets), 13 developed market government bond futures, 14 currency futures (five of which are emerging market currencies against the US dollar), and 29 commodity futures. The futures data, sourced from Bloomberg, consist of last traded daily prices, which are transformed into daily returns. These daily returns are excess returns given that they are derived from futures contracts.

Table 3.1 reports the summary statistics given by annualised mean (Mean), annualised standard deviation (Vol.), skewness (Skew.) and excess kurtosis (Kurt.) of monthly returns per contract. It also summarises the exchange where the futures are traded, the base currency (Cur.) of each contract, and the first date on which data was available (Start Date). The starting date of the study is 4 January 1996, which coincides with the first date for which options data are available (as described in the next chapter). Given that returns are scaled by volatility, a 'warm up' period of one-year is used where data are available. As such, many of the contract starting dates are the beginning of 1995. Contracts that have a later starting date are treated similarly, being incorporated into the study following a 'warm up' period used to compute volatility, which is then used to scale returns.

The futures contracts used in this study are traded on various exchanges with different operating hours. On any given date, the last traded price for futures contracts on US exchanges will be captured later in the day in the United States than the last traded prices recorded for the same day on Asian and European exchanges due to time zone differences. Lookahead bias is guarded against by forming and executing trade decisions using lagged return information. This process is described in detail in the methodology section of this chapter.

Table 3.1 Summary statistics for futures contracts

Asset	**Start Date**	**Cur.**	**Exch.**	**Mean**	**Vol.**	**Skew.**	**Kurt.**
Commodities							
Wheat	1995-01-04	USD	CBT	-4.11	29.72	0.54	1.41
Kansas City HRW Wheat	1995-01-04	USD	CBT	2.49	29.35	0.58	1.55
Corn	1995-01-04	USD	CBT	2.10	26.58	0.30	1.16
Soybean	1995-01-04	USD	CBT	11.81	24.35	-0.08	0.72
Coffee	1995-01-04	USD	NYB	3.30	34.80	0.78	1.45
Sugar	1995-01-04	USD	NYB	5.99	31.10	0.22	0.60
Cocoa	1995-01-04	USD	NYB	11.11	32.56	1.23	5.95
Cotton	1995-01-04	USD	NYB	-0.11	28.51	0.15	0.66
Lean Hogs	1995-01-04	USD	CME	4.42	31.78	0.26	1.60
Live Cattle	1995-01-04	USD	CME	9.47	15.91	-0.38	2.93
Feeder Cattle	1995-01-04	USD	CME	2.96	14.63	-0.10	1.60
WTI Crude Oil	1995-01-04	USD	NYM	10.44	38.81	0.84	11.23
NY Harbour ULSD	1995-01-04	USD	NYM	14.31	35.06	0.39	1.66
Brent Crude Oil	1995-01-04	USD	ICE	14.03	34.23	-0.34	3.30
Low Sulphur Gasoil	1995-01-04	USD	ICE	14.82	34.62	0.15	1.00
Natural Gas	1995-01-04	USD	NYM	-6.82	55.88	0.76	2.09
Copper	1995-01-04	USD	CMX	8.88	25.46	0.10	3.02
Gold	1995-01-04	USD	CMX	5.90	15.60	0.17	0.90
Silver	1995-01-04	USD	CMX	8.54	29.39	0.24	0.79
Soybean Meal	1995-01-04	USD	CBT	19.99	29.01	0.54	1.45
Platinum	1995-01-04	USD	NYM	6.55	21.68	-0.48	2.19
Frozen Orange Juice	1995-01-04	USD	NYB	13.45	32.10	0.32	0.34
Oats	1995-01-04	USD	CBT	17.23	33.85	0.64	1.81
Soybean Oil	1995-01-05	USD	CBT	3.90	27.06	0.69	3.25
Aluminium	1997-07-24	USD	LME	-1.11	19.00	0.17	0.34
Nickel	1997-07-24	USD	LME	9.52	33.50	0.33	0.60
Lead	1997-07-24	USD	LME	5.47	26.06	0.09	1.59

Zinc	1997-07-24	USD	LME	2.73	25.49	0.01	1.54
RBOB Gasoline	2005-10-04	USD	NYM	9.23	29.70	-1.01	10.00
Currencies							
AUD/USD	1995-01-04	USD	CME	1.60	11.66	-0.21	1.25
CAD/USD	1995-01-04	USD	CME	0.37	8.06	-0.27	2.34
JPY/USD	1995-01-04	USD	CME	-3.68	10.58	0.67	3.42
CHF/USD	1995-01-04	USD	CME	-0.28	10.06	0.35	1.47
GBP/USD	1995-01-04	USD	CME	0.00	8.31	-0.27	1.18
MXN/USD	1995-04-26	USD	CME	4.96	11.13	-0.97	3.51
BRL/USD	1995-11-09	USD	CME	5.06	16.83	-1.36	10.70
NZD/USD	1997-05-08	USD	CME	2.09	12.14	-0.08	1.44
ZAR/USD	1997-07-01	USD	CME	1.50	15.14	-0.23	0.63
EUR/USD	1998-05-20	USD	CME	-0.44	8.78	0.05	1.68
NOK/USD	2002-05-17	USD	CME	0.00	10.33	-0.13	1.41
SEK/USD	2002-05-17	USD	CME	-0.06	9.64	0.12	1.27
KRW/USD	2006-09-19	USD	CME	-0.85	8.61	0.13	8.04
INR/USD	2013-01-29	USD	CME	0.41	4.00	-0.06	16.99
Equities							
DAX Index	1995-01-04	EUR	EUX	8.47	22.25	-0.39	1.48
IBEX 35	1995-01-04	EUR	MFM	8.47	21.35	-0.23	1.29
CAC 40 10 EUR	1995-01-04	EUR	EOP	8.50	20.31	-0.22	0.66
TOPIX	1995-01-04	JPY	OSE	8.65	23.94	0.25	2.98
AEX	1995-01-04	EUR	EOE	9.03	20.89	-0.44	0.99
FTSE 100	1995-01-04	GBP	ICF	5.61	15.53	-0.27	0.63
FTSE/JSE Top 40	1995-01-04	ZAR	SAF	11.94	23.10	0.25	2.67
Nikkei 225	1995-01-04	JPY	OSE	7.83	24.67	-0.25	0.92
KOSPI200	1996-05-06	KRW	KFE	9.97	27.46	1.46	13.56
Bovespa	1997-02-17	BRL	BMF	10.90	28.66	4.67	63.30
S&P 500 E-mini	1997-09-10	USD	CME	6.45	14.89	-0.55	1.23
Euro Stoxx 50	1998-06-23	EUR	EUX	5.22	18.93	-0.32	1.12
Swiss Market	1998-09-17	CHF	EUX	3.43	17.43	-0.50	0.87
NASDAQ 100 E-mini	1999-06-22	USD	CME	8.52	22.00	-0.39	2.34
S&P/TSX 60 Index	1999-11-24	CAD	MSE	4.97	11.28	-0.30	1.53
ASX SPI 200	2000-05-03	AUD	SFE	5.00	12.86	-0.24	0.88
NSE Nifty 50 Index	2000-06-13	INR	NSE	13.17	17.93	-0.34	2.97
E-mini–Dow Jones Industrial Average Index	2002-04-24	USD	CBT	5.71	12.88	-0.30	2.42
FTSE MIB	2004-04-05	EUR	MIL	4.74	17.23	-0.03	2.67

FTSE China A50	2006-12-28	USD	SGX	3.72	21.85	0.29	5.85
MEX Bolsa	2009-03-24	MXN	MDX	7.07	12.63	0.45	1.99
MSCI EAFE Index	2009-09-09	USD	NYF	2.72	11.26	0.05	3.83
MSCI Emerging Markets Index	2009-09-09	USD	NYF	1.81	13.00	-0.03	4.99
Government Bonds							
3Y Australia	1995-01-04	AUD	SFE	1.63	12.11	0.62	2.55
10Y Australia	1995-01-04	AUD	SFE	1.51	11.98	0.59	2.41
Euro Bobl	1995-01-04	EUR	EUX	3.13	10.19	0.47	1.20
Euro Bund	1995-01-04	EUR	EUX	4.39	11.37	0.39	0.53
Long Gilt	1995-01-04	GBP	ICF	3.96	11.37	0.50	1.08
10Y Canada	1995-01-04	CAD	MSE	3.33	10.61	0.59	2.70
2Y US	1995-01-04	USD	CBT	0.91	1.55	0.42	1.88
5Y US	1995-01-04	USD	CBT	1.96	3.85	0.17	1.45
10Y US	1995-01-04	USD	CBT	2.89	5.79	0.19	1.75
Euro Schatz	1997-03-10	EUR	EUX	1.37	9.52	0.41	1.54
Euro Buxl	1998-10-05	EUR	EUX	5.85	18.84	0.48	7.39
10Y Japan	1999-03-15	JPY	OSE	3.59	10.39	-0.36	2.41
Medium Gilt	2009-11-24	GBP	ICF	1.45	8.88	0.54	1.57

3.2.2 Exchange rate data

Although the majority of futures contracts used in this study are denominated in US dollars, some are traded on non-US exchanges. For those contracts, the study transforms daily returns from local currency to the base currency of the study, US dollars, using daily spot exchange rate data sourced from Bloomberg. All exchange rate data are available from the starting date of the study, 4 January 1996.

3.2.3 Interest rate data

The trend-following strategy constructed in this study uses exchange-traded futures contract data, and thus, the resulting portfolio returns are excess returns. To compute daily total returns for the portfolio, the study uses daily 3-month US LIBOR sourced from Bloomberg, which is transformed into a time series of daily returns. The same yields are also used to convert the total returns of benchmark indices discussed below into excess returns for comparison purposes. Interest rate data are available from the starting date of the study, 4 January 1996.

3.2.4 Data used for regression analysis and performance comparison

The study employs the daily data summarised in Table 3.2 below to evaluate the returns of the trend-following strategy. For standard asset class performance, ACWI serves as a proxy for global equity; the Bloomberg Global-Aggregate Total Return Index represents global bonds; and the S&P GSCI Index is utilised for commodities. All asset class data are sourced from Bloomberg.

Factor return data, including small minus big (SMB) and high minus low (HML) are used to proxy for the size and value factors respectively. These data sets are obtained from Kenneth French's website (French, 2024). The momentum factor dataset from the same source is incorporated into the analysis. To validate the trend-following strategy, the study compares it against an industry benchmark—the Société Générale Trend Index (SG Trend Index). The SG Trend Index is an equal-weighted index comprised of the largest 10 trend-following managers by AUM open to new investment. This index, sourced from Société Générale, is reconstituted annually and reported net of fees including both fixed and performance fees. Notably, the SG Trend Index does not have daily data dating back to 4 January 1996, instead, it starts on 3 January 2000. In all cases, daily excess returns are considered, denominated in US dollars.

Table 3.2 Standard asset pricing benchmarks used

Asset Pricing Benchmark	Asset Class/Factor	Start Date
ACWI	Global Equity	1996-01-04
Bloomberg Global-Aggregate Total Return Index Value Unhedged USD	Global Bonds	1996-01-04
S&P GSCI Index	Commodities	1996-01-04
Société Générale Trend Index	Trend-Following	2000-01-03
Fama and French (1992) Size Factor	Size Factor	1996-01-04
Fama and French (1992) Value Factor	Value Factor	1996-01-04
Carhart (1997) Momentum Factor	Momentum Factor	1996-01-04

3.3 METHODOLOGY

In the previous section, the investment universe was selected for the trend-following strategy. This section details the construction of the trend-following portfolio. At each stage, simplicity is prioritised over optimisation to emphasise the robustness of the methodology. The process begins by illustrating how daily returns are computed from futures prices and then describes how returns are normalised using their ex-ante volatility estimates. Next, the generation of the trend signal and the sizing of positions are discussed. Finally, the incorporation of transaction costs and the rebalancing methodology are described.

3.3.1 Transforming futures prices into daily returns

To transform futures prices into a time series of daily returns, daily returns for the front month contract of each ticker are computed as described in (3.1).

$$R_{\mathrm{t,i}} = \frac{\mathrm{P}_{t,i}^{(1)}}{\mathrm{P}_{t-1,i}^{(1)}} - 1 \qquad \textbf{3.1}$$

Where $R_{t,i}$ is the daily return for ticker *i* at time *t*, $P_{t,i}^{(1)}$ is the price of the front month contract for ticker *i* at time *t*, $P_{t-1,i}^{(1)}$ is the front month contract price for ticker *i* at time *t-1*.

To manage contract rollovers (rolls), each contract gets rolled forward the day before its expiration. On the expiration date of each contract, daily returns are computed using the second month contract prices, as shown in (3.2).

$$R_{\mathrm{t,i}} = \frac{\mathrm{P}_{t,i}^{(2)}}{\mathrm{P}_{t-1,i}^{(2)}} - 1 \qquad \textbf{3.2}$$

Where $R_{t,i}$ is the daily return for ticker *i* at time *t*, $P_{t,i}^{(2)}$ is the price of the second month contract for ticker *i* at time *t*, $P_{t-1,i}^{(2)}$ is the second month contract price for ticker *i* at time *t-1*.

For example, the front month E-mini SPX futures contract from 18 March 2024 to its expiration on 21 June 2024 was the June 2024 contract, and the second month contract was the September 2024 contract. Daily returns are computed for this ticker using (3.1) up to the day before expiration, 20 June 2024. On this day, at market close, the contract was rolled into the second month contract, which was the September 2024 contract. Thus, on the expiration day, 21 June 2024, (3.2) was employed. After 21 June 2024, the September contract became the front month contract, and (3.1) was once again applied.

Table 3.1 presents the summary statistics on all futures contracts used in the study. Daily returns are compounded into monthly returns for each futures contract and the first four moments of each contract's historical return distribution are presented. Similar to Baltas and Kosowski (2013) and Moskowitz et al. (2012), considerable differences in mean returns and annualised volatility between futures contracts over the period under investigation are found.

As an asset class, bonds generated the lowest volatility, while commodities generated the highest. The median bond volatility was about nine times less than the median commodity (3.25% and 29.43%, respectively). Notably, all equity and bond mean returns were positive, which is in line with the findings of Moskowitz et al. (2012) and Baltas and Kosowski (2013), and some currency and commodity mean returns were negative. Equity returns were mostly negatively skewed, and the same can also be said of currency returns. Commodity returns tended to be positively skewed.

3.3.2 Computing ex-ante volatility estimates per ticker

The meaningful differences in the distributional properties across and within asset classes illustrated in Table 3.1 suggest that returns need to be normalised to appropriately construct a diversified trend-following portfolio per Baltas and Kosowski (2013) and Moskowitz et al. (2012). To do this, returns are scaled by their recent exponentially weighted moving average volatility.

The implication of using an exponentially weighted moving average volatility is that more recent observations are weighted more heavily in the calculation. As a result, recent events play a more significant role in the signal generation process compared to historical events. For each ticker, a rolling window of 65 trading days (corresponding to three months) is used and the squared difference of each return within the window and the mean return of the whole window is computed. These squared differences are then weighted by the normalised exponential decay weights, and the sum of the weighted squared differences is the resulting ex-ante variance, which is finally transformed into annualised volatility. This process is distilled in (3.3).

$$\sigma_{\mathrm{t,i}} = \sqrt{261} \times \sqrt{\sum_{j=0}^{L-1} \delta_j (R_{\mathrm{t-j,i}} - \bar{R}_{\mathrm{t,i}})^2} \tag{3.3}$$

Where $\sigma_{t,i}$ is the ex-ante volatility for ticker *i* at time *t*, 261 is the number of trading days in a year, L is the lookback period of 65 trading days, $R_{t-j,i}$ is the return of ticker *i* at time *t-j*, $\bar{R}_{t,i}$ is the mean return over the lookback period of ticker *i* at time *t* and δ_j is the decay factor calculated using (3.4).

$$\delta_j = \frac{\lambda^j}{\sum_{k=0}^{L-1} \lambda^k} \tag{3.4}$$

Where δ_j is the decay factor, λ^j represents the weights for day *j*, $\sum_{k=0}^{L-1} \lambda^k$ represents the sum of the exponential decay weights over the entire lookback period *L*, and *k* is the variable that iterates through each day within the lookback period to calculate the total sum of the weights. This sum is then used to normalise the weight for each specific day *j* so that the weights sum to one. It is important that weights are normalised to ensure they sum to one, ensuring that the resulting volatility measure is not biased by the magnitude of the weights.

3.3.3 Pooled panel regressions

Having computed rolling ex-ante volatility estimates per ticker, they can be used to normalise returns across the investment universe by dividing them by their own volatility. In so doing, all tickers are put on roughly the same scale. The resulting normalised returns are used to test the null hypothesis $H_{0,A}$: markets do not trend.

To evaluate whether trends exist in global markets, a similar methodology to Moskowitz et al. (2012) and Baltas and Kosowski (2013) is followed. Data are transformed from daily to monthly returns, and then, over 36 monthly lookback periods per ticker, pooled panel regressions are run to determine whether lagged monthly returns have predictive power over subsequent month returns.

Using pooled panel regressions, the resulting t-statistics are computed that account for group-wise clustering by time (monthly level) per asset class and across the entire 79 futures contract universe in accordance with (3.5).

$$\frac{R_{t,i}}{\sigma_{\mathrm{t-1,i}}} = \alpha + \frac{\beta_h R_{t-h,i}}{\sigma_{\mathrm{t-h-1,i}}} + \varepsilon_{t,i} \qquad \textbf{3.5}$$

Where $R_{t,i}$ is the monthly return for ticker *i* at time *t*, $\sigma_{\mathrm{t-1,i}}$ is the ex-ante volatility estimate computed using (3.4) return for ticker *i* at time *t-1*, α is the intercept term, β_h is the coefficient for the lagged normalised return, *h* is the lag monthly period, from *h* = 1, 2, …, 36 months, and $\varepsilon_{t,i}$ is the error term for ticker *i* at time *t*.

A significant and large t-statistic for the coefficient β_h would suggest that the null hypothesis, $H_{0,\,A}$: markets do not trend, should be rejected in favour of the alternative hypothesis, $H_{1,\,A}$: markets do trend. Per Petersen (2008), to account for return dependence resulting from potential autocorrelation at an individual futures contract level or cross-sectional correlation among the different futures contracts per period across the data set, t-statistics are computed using standard errors that are clustered by time and by ticker.

3.3.4 Trend signals

Moskowitz et al. (2012), Baltas and Kosowski (2013) and Harvey et al. (2021) find that markets exhibit trend effects, where lagged returns up to one-year have predictive power on subsequent returns. For simplicity, normalised returns per ticker are computed over four lookback periods, *h*: 65, 131, 196, and 261 trading days, corresponding to approximately 3, 6, 9, and 12 months. Signals are assigned to each lookback period: either a one to represent a long position resulting from a positive return for ticker *i* from time *t-h* to time *t*, or -1 to represent a short position resulting from a negative return for ticker *i* from time *t-h* to time *t*.

Then, the resulting signals are equally weighted to establish a final position for each ticker *i* at time *t*. For example, should three lookback signals be long, and 1 lookback signal be short for a given ticker on a given day, then the resulting final position will be 0.50, which is the average of the 4 positions. This process is defined in (3.6) and (3.7).

$$Signal_{t,i}^{h} = \begin{cases} 1 \; if \; \dfrac{R_{t-h,i}}{\sigma_{\mathrm{t-h-1,i}}} \geq 0 \\ \hline -1 \; if \; \dfrac{R_{t-h,i}}{\sigma_{\mathrm{t-h-1,i}}} < 0 \end{cases} \tag{3.6}$$

Where $Signal_{t,i}^{h}$ is the signal for ticker *i* at time *t* for lookback period *h*. $\frac{R_{t-h,i}}{\sigma_{\mathrm{t-h-1,i}}}$ is the normalised return for ticker *i* at time *t* for lookback period *h* as described earlier in the chapter.

$$Final\ Signal_{t,i} = \frac{\sum_{h \in \{65,131,196,261\}} Signal_{t,i}^{h}}{4} \tag{3.7}$$

Where $Final\ Signal_{t,i}$ is the final signal for ticker *i* at time *t* which is an average of all 4 signals derived from (3.6) across the lookbacks discussed.

3.3.5 Position sizing

Maximising diversification of the trend-following portfolio is a key objective. Signal diversification per ticker has already been achieved using (3.7) by employing four different lookback periods. Next, to compute weights per ticker, this study ensures that the portfolio roughly receives an equal risk contribution over time across all asset classes (equities, commodities, bonds and currencies), as well as within each asset class. This objective is met using (3.8).

$$w_{t,i} = Final\ Signal_{t,i} \times \frac{40\%}{\sigma_{\mathrm{t-1,i}}} \times \frac{1}{N_t^{Asset\ Classes}} \times \frac{1}{N_t^{Tickers\ in\ Asset\ Class}} \tag{3.8}$$

Where $Final\ Signal_{t,i}$ is the long or short position established in ticker *i* at time *t* determined using (3.7) and can result in either a fully long, fully short, neutral or partially (50%) long or short position in ticker *i* at time *t*. 40% is a scaling factor used by Moskowitz et al. (2012) to achieve an ex-ante volatility roughly equal to 40% per ticker that results in an ex-post annualised portfolio volatility of approximately 12% (with the reduction in volatility attributed to diversification) in their trend-following strategy. The significance of the 12% ex-post volatility figure is that it is similar to the volatility of the risk factors used in their study.

Baltas and Kosowski (2013) incorporate the same 40% scaling factor in their study. $\sigma_{\mathrm{t-1,i}}$ is the 1-trading-day lagged ex-ante exponentially weighted moving average volatility for ticker *i* at time *t* computed in (3.2) and (3.3). The term $\frac{40\%}{\sigma_{\mathrm{t-1,i}}}$ ensures that tickers with higher volatility receive smaller weights and those with lower volatility receive larger weights, normalising the risk contribution from each position. To ensure equal risk contribution across and within asset classes, the formula multiplies the factors $\frac{1}{N_t^{Asset\ Classes}}$ and $\frac{1}{N_t^{Tickers\ in\ Asset\ Class}}$ respectively, using the number of asset classes as well as tickers within each asset class that ticker *i* belongs to and is available to trade at time *t*.

This approach ensures that no single asset class or ticker disproportionately influences the overall portfolio risk. For instance, all four asset classes are present throughout the period under investigation, implying that each class should contribute roughly 25% to the overall portfolio risk.

3.3.6 Portfolio returns

To compute the unscaled daily return of the portfolio, the term on the left-hand side of (3.9) below is used. A final scaling factor $\frac{12\%}{\sigma_{t-1}^{PU}}$ on the right-hand side of (3.9) is also applied and explained below.

$$R_t^P = \left(\sum_i \left(w_{t-1,i} \times R_{t,i}\right) - \sum_i \left(0.02\% \times \left|w_{t,i} - w_{t-1,i}\right|\right)\right) \times \frac{12\%}{\sigma_{t-1}^{PU}} \qquad \textbf{3.9}$$

Where R_t^P is the daily return of the portfolio at time *t*, $w_{t-1,i}$ is the weight of ticker *i* at time *t-1* as calculated in (3.8), $R_{t,i}$ is the daily excess return of ticker *i* at time *t*. The second term of (3.9) addresses the impact of transaction costs. Harvey, Rattray and Van Hemert (2021) state that a 0.02% transaction cost is appropriate for medium-term trend-following strategy. Baltas and Kosowski (2013) argue that exchange-traded futures have relatively low transaction costs when compared to cash equity or bond markets due to their liquidity, and as such, do not incorporate transaction costs into their study. A conservative approach is taken and a 0.02% transaction cost is applied for trades per Harvey et al. (2021), which explains the constant in (3.9).

The term $\left|w_{t,i} - w_{t-1,i}\right|$ refers to the absolute change in the position held in ticker *i* from time *t-1* to time *t*. For example, if the position in gold changed from -10% to 0% from one day to the next, the portfolio will incur a 0.002% transaction cost at the portfolio level to execute the trade. For simplicity, the trend-following portfolio is rebalanced daily.

As a final measure to ensure that the trend-following portfolio delivers as constant a volatility profile as possible, an overall portfolio volatility target of 12% is applied, in line with Moskowitz et al. (2012). In testing, the unscaled portfolio delivered an overall portfolio volatility in line with the 12% target over the entire period, but there were periods in which rolling one-year volatility was slightly higher and slightly lower than the target.

As such, the final term $\frac{12\%}{\sigma_{t-1}^{PU}}$ is applied to (3.9), where 12% is a constant annualised volatility target and σ_{t-1}^{PU} is the volatility of the modelled unscaled portfolio returns computed using (3.3). This final term evenly increases exposure across the entire portfolio should the realised volatility of the portfolio be lower than the target and evenly reduces exposure across the entire portfolio should the realised volatility of the portfolio be higher than the target. Applying this final adjustment factor helps to bring returns closer in line with the 12% volatility target over time.

3.4 FINDINGS: TREND-FOLLOWING STRATEGY

3.4.1 Pooled panel regression analysis

To test the secondary null hypothesis $H_{0,\ A}$: markets do not trend, multiple pooled panel regressions are run to evaluate the predictability of lagged monthly futures returns. Figure 3.1 illustrates the resulting t-statistics from pooled panel regressions for all asset classes by monthly lags *h*. Positive t-statistics can be interpreted as evidence of trend effects in financial market returns, and negative t-statistics can be interpreted as evidence of reversals in financial market returns.

Figure 3.1 Pooled panel regressions: t-statistics by monthly lag *h* for all asset classes

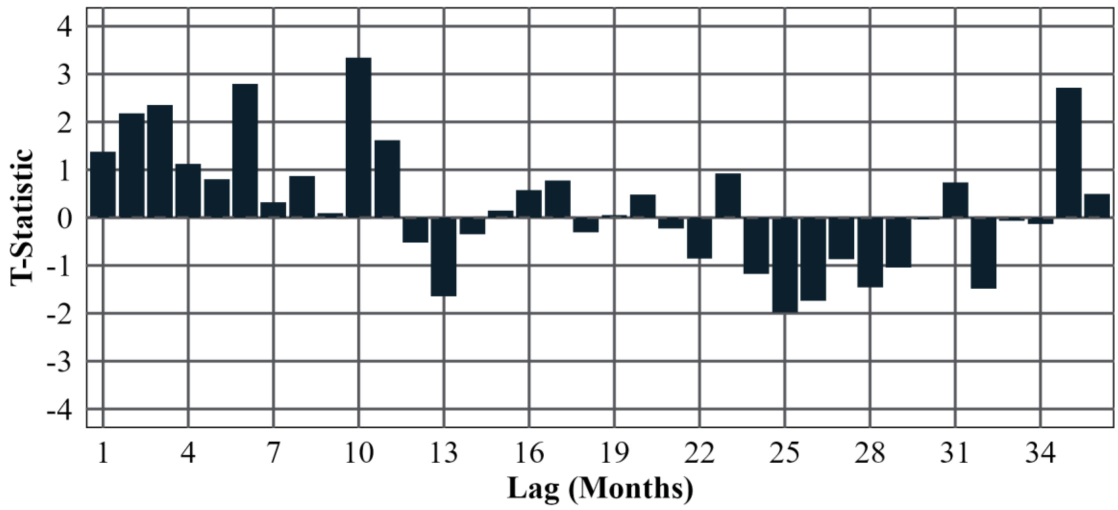

Consistent with the literature, positive t-statistics are found for all monthly lags up to 11 months, indicating the presence of trend effects. The analysis reveals that several lags are statistically significant at various confidence levels. Specifically, lags 2 and 3 are statistically significant at the 5% level, with t-statistics of +2.18 and +2.35, respectively. Lag 6 is statistically significant at the 1% level, with a t-statistic of +2.79. Lag 10 is statistically significant at the 0.1% level, with a t-statistic of +3.35. These results suggest that past returns have predictive power for future returns over these horizons, supporting the alternative hypothesis $H_{1,\ A}$: markets do trend. Consistent with previous studies, a reversal effect is observed around the 12-month mark. This phenomenon indicates that the momentum effect tends to weaken and eventually reverse after one-year.

Next, the same process is repeated and tests for the presence of trend effects are conducted per asset class. Figure 3.2 illustrates the resulting t-statistics from pooled panel regressions for the 23 equity futures contracts by monthly lags *h*. All t-statistics are positive up to the 10-month lag except for the 5-month lag. The 10-month lag is statistically significant at the 5% level, with a t-statistic of +1.99. Reversal effects are found for equity from the 11-month lag.

Figure 3.2 Pooled panel regressions: t-statistics by monthly lag *h* for equity

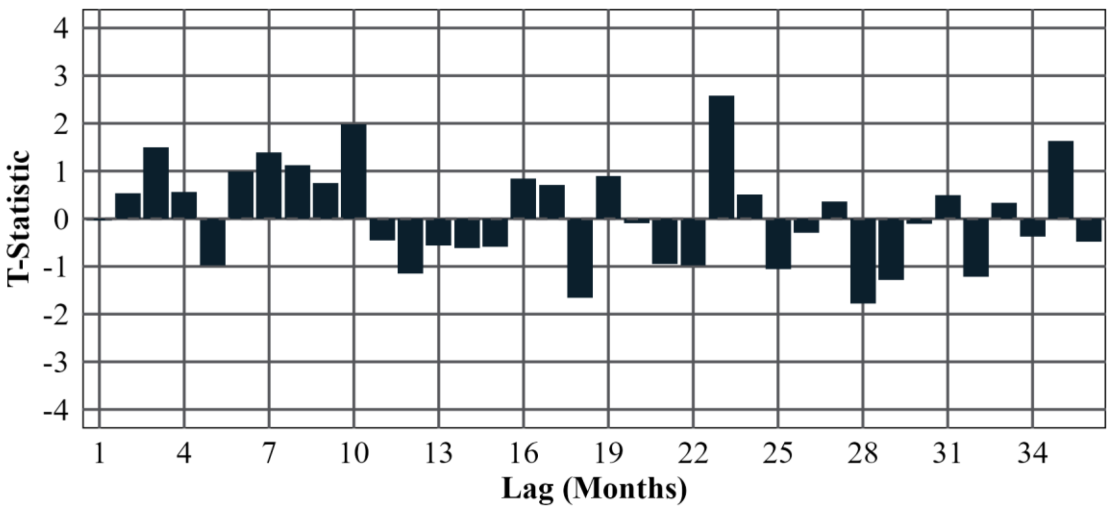

Figure 3.3 illustrates the resulting t-statistics from pooled panel regressions for the 13 bond futures contracts by monthly lags *h*. The t-statistics for bonds were positive for eight out of the first 11 lags, with mixed results thereafter. Both the 3-month and the 10-month lags were statistically significant at the 5% level, with t-statistics of +2.30 and +1.98, respectively.

Figure 3.3 Pooled panel regressions: t-statistics by monthly lag *h* for bonds

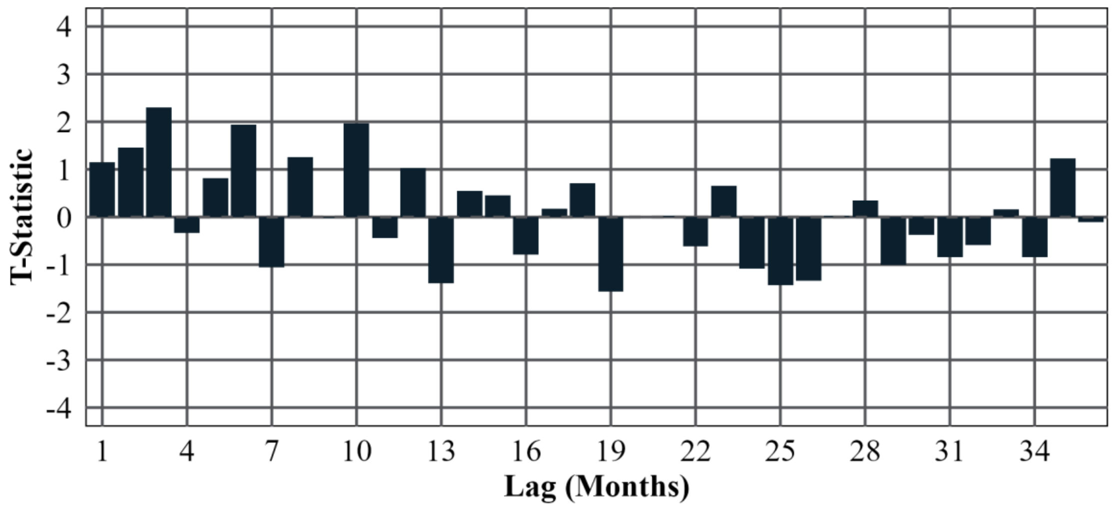

Figure 3.4 illustrates the resulting t-statistics from pooled panel regressions for the 14 currency futures contracts by monthly lags *h*. Only 2 t-statistics for the first 12-months were negative, and a noticeable reversal effect was apparent after the 12-month mark. The 6-month lag was statistically significant at the 5% level, with a t-statistic of +2.47.

Figure 3.4 Pooled panel regressions: t-statistics by monthly lag *h* for currencies

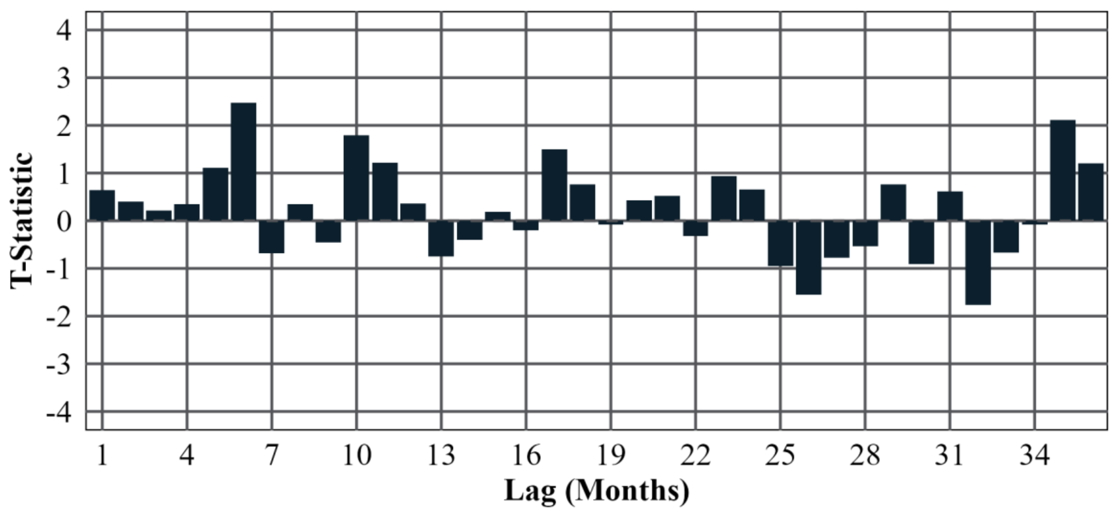

Figure 3.5 illustrates the resulting t-statistics from pooled panel regressions for the 29 commodity futures contracts by monthly lags *h*. Among all the asset classes tested,

commodities had the strongest evidence of trend effects. Only lags 8 and 9 were negative for the first 11 lags, albeit only marginally. Lag 6 was statistically significant at the 5% level, with a t-statistic of +2.90. Both the 2-month and the 10-month lags were statistically significant at the 1% level, with t-statistics of +2.90 and +3.08, respectively. Lag 11 was statistically significant at the 0.1% level, with a t-statistic of +3.44. After lag 11, a sharp reversal effect, present in lags 12 and 13, is observed.

Figure 3.5 Pooled panel regressions: t-statistics by monthly lag *h* for commodities

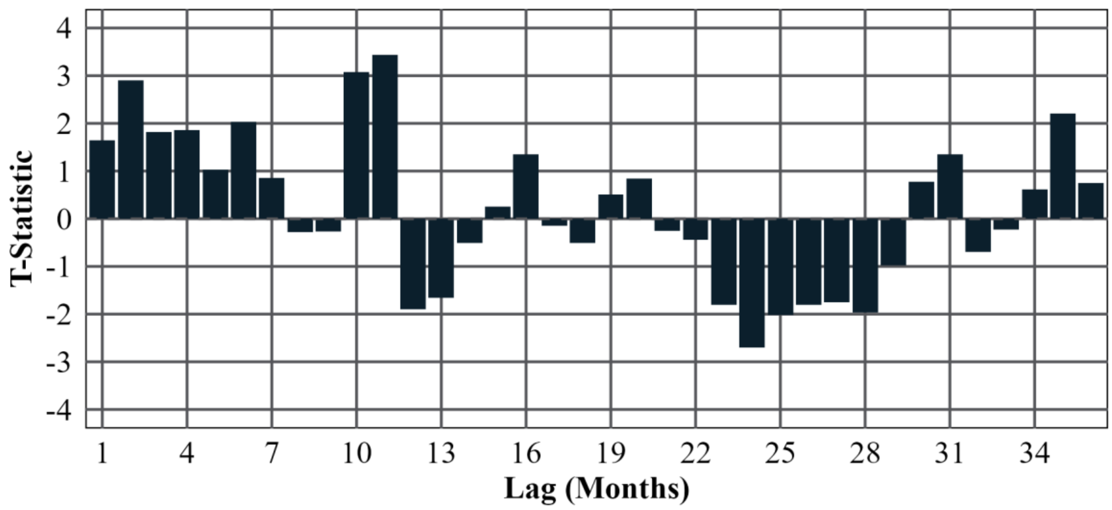

The comprehensive analysis depicted in Figures 3.1 to 3.5, which is largely in line with past research, provides compelling evidence that trend effects do exist in financial market returns. The statistically significant t-statistics at various lags across different asset classes, both individually and in combination, suggest the presence of some predictive power in lagged returns. The secondary null hypothesis $H_{0,A}$: markets do not trend is rejected in favour of the alternative hypothesis $H_{1,A}$: markets do trend.

3.4.2 Trend-following strategy vs. SG Trend Index

Having rejected $H_{0,A}$: markets do not trend in favour of the alternative hypothesis $H_{1,A}$: markets do trend, the trend-following strategy is now constructed as described in the methodology section of this chapter. In the following section, the performance of the resulting trend-following portfolio, referred to as the *Trend Portfolio*, is compared and contrasted against the performance of the SG Trend Index, an industry benchmark for trend-following strategies. The SG Trend Index is an equal-weighted composite of the 10 largest trend-following funds by AUM open to new investment. The index is net of all fees, which include transaction costs, fixed asset management fees and performance fees. The Trend Portfolio is net of transaction costs. Excess returns of both the Trend Portfolio and the SG Trend Index are compared from 3 January 2000 to 28 June 2024 to match the dates for which SG Trend Index data exists. This is the only deviation from the full period under analysis, which is 4 January 1996 to 28 June 2024.

3.4.2.1 Excess performance comparison

In Figure 3.6, the cumulative excess returns of the Trend Portfolio are illustrated against the SG Trend Index, and Table 3.2 reports the performance statistics of both strategies. The Trend Portfolio has outperformed the SG Trend Index over time. Table 3.3 shows that the Trend Portfolio generated a CAGR of 9.7% compared to 3.6% for the SG Trend Index. The Sharpe ratio of the Trend Portfolio was +0.71, which was meaningfully higher than the +0.27 delivered by the index.

Figure 3.6 Cumulative excess returns: SG Trend Index vs Trend Portfolio

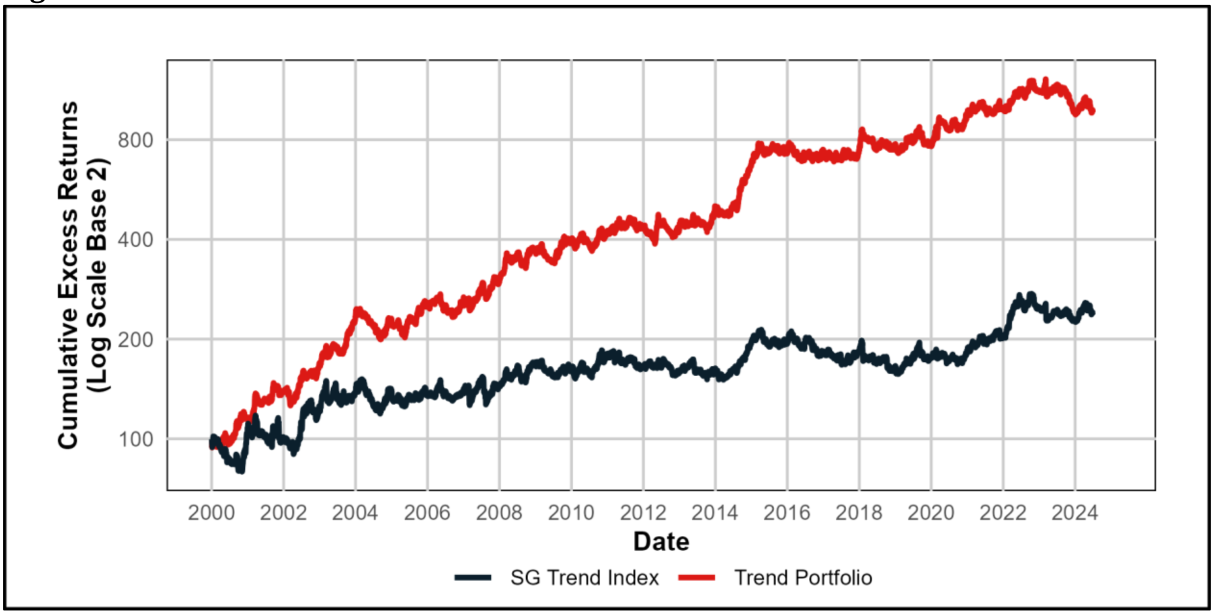

In Figure 3.7, the drawdowns of both strategies are illustrated, showing the peak-to-trough performance of each approach from inception up to each point in time. While the drawdowns of both strategies were positively correlated, the SG Trend Index tended to suffer slightly deeper drawdowns over the period. Annualised realised volatility was similar across both strategies: 13.7% for the Trend Portfolio and 13.5% for the SG Trend Index, as summarised in Table 3.2.

Both strategies' return distributions were positively skewed over the period, which is a key feature of trend-following portfolios that is well documented in the literature. The Trend Portfolio exhibited a positive skew of 0.85 compared to 0.11 for the SG Trend Index. The excess kurtosis of the SG Trend Index was -2.31 and was more negative than the -1.32 of the Trend Portfolio.

When analysing the tail risk statistics of the strategies, the Trend Portfolio had lower 95% and 99% value-at-risk (VaR) and conditional value-at-risk (CVaR) statistics. The worst 3-month performance on the Trend Portfolio of -11.86% was slightly better than the -14.35% suffered by the SG Trend Index. However, the best 3-month performance of the Trend Portfolio of 20.99% lagged the SG Trend Index, which delivered 30.88%.

Figure 3.7 Drawdowns: SG Trend Index vs Trend Portfolio

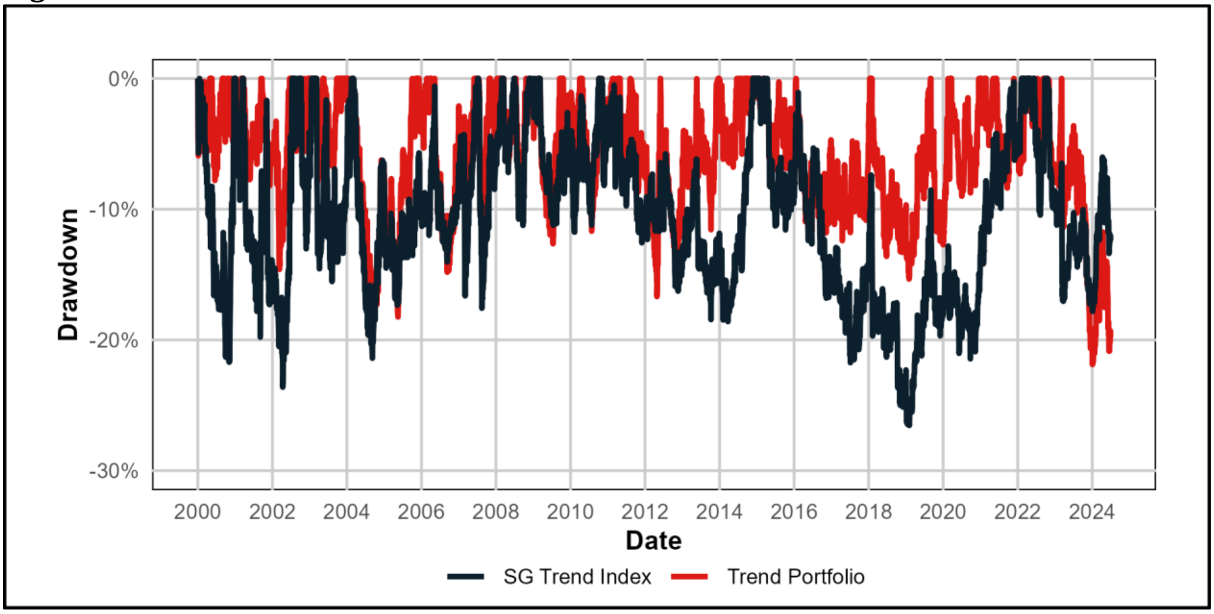

In Figure 3.8, the rolling 12-month correlation between the SG Trend Index and the Trend Portfolio is illustrated. Rolling correlations were high over the period and tended to oscillate between 60% and 90% for most of the time.

Figure 3.8 Rolling one-year correlation of daily returns: SG Trend Index vs Trend Portfolio

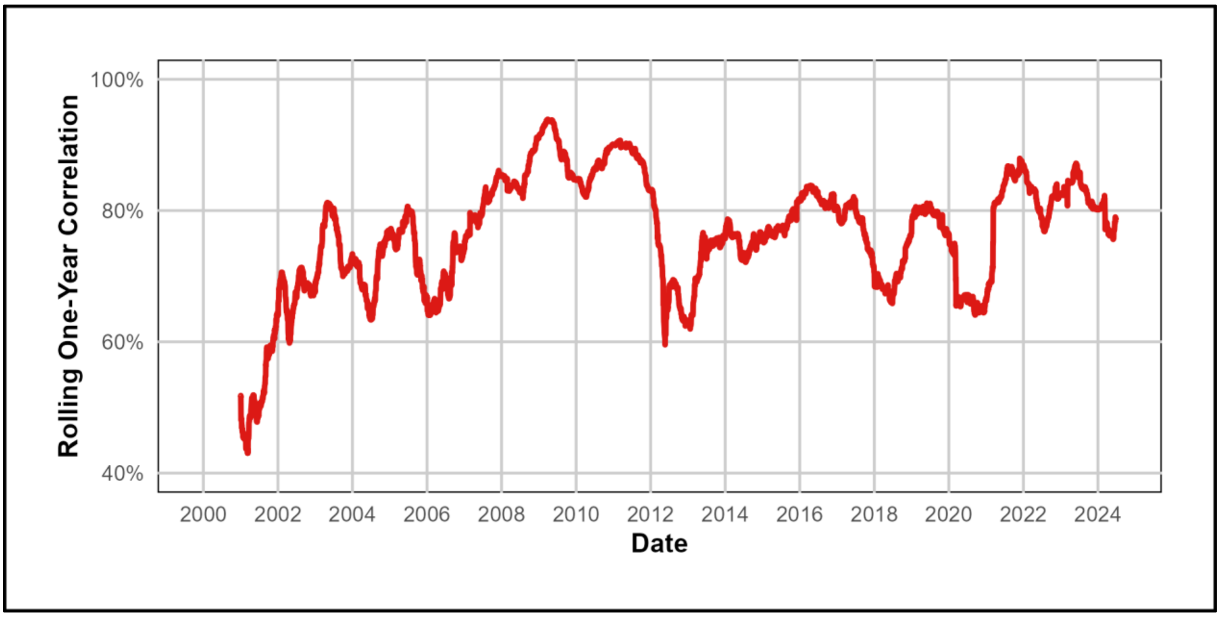

Table 3.3 Summary statistics for the SG Trend Index vs Trend Portfolio

Statistic	**Trend Portfolio**	**SG Trend Index**
CAGR	9.7%	3.6%
Annualised Volatility	13.7%	13.5%
Sharpe Ratio	0.71	0.27
Skewness	0.85	0.11
Excess Kurtosis	-1.32	-2.31
VaR (95%)	-4.49%	-5.29%
VaR (99%)	-6.39%	-8.74%
CVaR (95%)	-5.56%	-7.50%
CVaR (99%)	-7.03%	-10.68%
Best 3M Performance	20.99%	30.88%
Worst 3M Performance	-11.86%	-14.35%

In Figure 3.9, the regression analysis conducted on the excess monthly returns of the SG Trend Index, the independent variable, against the monthly excess returns of the Trend Portfolio, the dependent variable, is illustrated. Table 3.4 summarises the regression output. This regression analysis is used to test the null hypothesis $H_{0,\,B}$: SG Trend Index returns do not have a significant relationship with Trend Portfolio returns. The coefficient of +0.66 indicates that for every 1% excess monthly return in the SG Trend Index, the Trend Portfolio tended to deliver 0.66% on average, excluding the alpha component. The t-statistic of the coefficient was +14.58, statistically significant at a 1% level. This finding, in conjunction with the high degree of rolling 12-month correlations presented in Figure 3.8 suggests that the Trend Portfolio captured a substantial portion of the performance of the SG Trend Index.

Figure 3.9 Regression analysis of monthly returns: Trend Portfolio vs. SG Trend Index

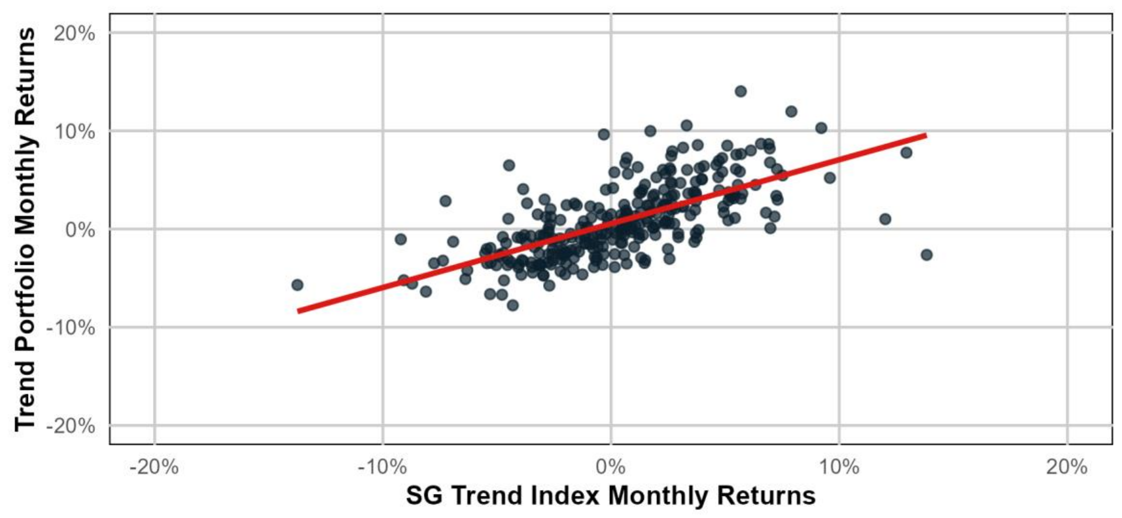

After adjusting for the SG Trend Index, the results demonstrate that the Trend Portfolio generates a statistically significant monthly alpha of 0.61% with a t-statistic of +3.44, significant at the 1% level. Furthermore, 42% of the variance in monthly Trend Portfolio returns can be explained by monthly SG Trend Index returns, which is further evidence of a positive and significant relationship.

The regression analysis presented above suggests that the null hypothesis $H_{0, B}$: SG Trend Index returns do not have a significant relationship with Trend Portfolio returns should be rejected in favour of the alternative hypothesis, $H_{1, B}$: SG Trend Index returns have a significant relationship with Trend Portfolio returns. This finding implies that there is a statistically significant relationship between a trend-following industry benchmark SG Trend Index and the Trend Portfolio returns, meaning that the Trend Portfolio returns are a fair proxy for the broad strategy that captures typical trend-following return distribution dynamics effectively.

Table 3.4 Summary of regression results: Trend Portfolio and SG Trend Index

Statistic	**SG Trend Index**	**Intercept**	**R-Squared**
Coefficient	0.66	0.61%	42%
(t-Stat)	(14.58)***	(3.44)***	

3.4.3 Performance Attribution in Trend Portfolio Returns

Having established that the Trend Portfolio is a suitable proxy for a trend-following strategy, attention is now turned to understanding the drivers of portfolio returns through a performance attribution analysis. The purpose of this analysis is to decompose the strategy's return into the individual contributions from underlying asset class returns. In the previous section, the monthly returns of the Trend Portfolio were examined from January 2000 to June 2024, as this period was consistent with the available data for the SG Trend Index. However, to gain a more comprehensive understanding of the Trend Portfolio's performance over time, the strategy's performance will now be extended backwards in time to match the starting date of the period under investigation, 4 January 1996. To achieve this, the daily excess returns of the Trend Portfolio for the 4-year period not investigated in the previous section are backdated, so that daily returns are available for the full period from 4 January 1996 to 28 June 2024 using the same methodology as applied in the initial analysis.

In Figure 3.10, a performance attribution is illustrated by analysing the contributions to Trend Portfolio excess return from the four key asset classes that comprise it: global equities, government bonds, currencies, and commodities. The results indicate a roughly equal contribution from each of the asset classes, consistent with the underlying methodology designed to allocate equal risk across these classes. This balanced contribution reflects the portfolio's construction principles, emphasising risk parity across asset classes and aligning with the existing literature on trend-following strategies.

Figure 3.10 Performance attribution by asset class for the Trend Portfolio

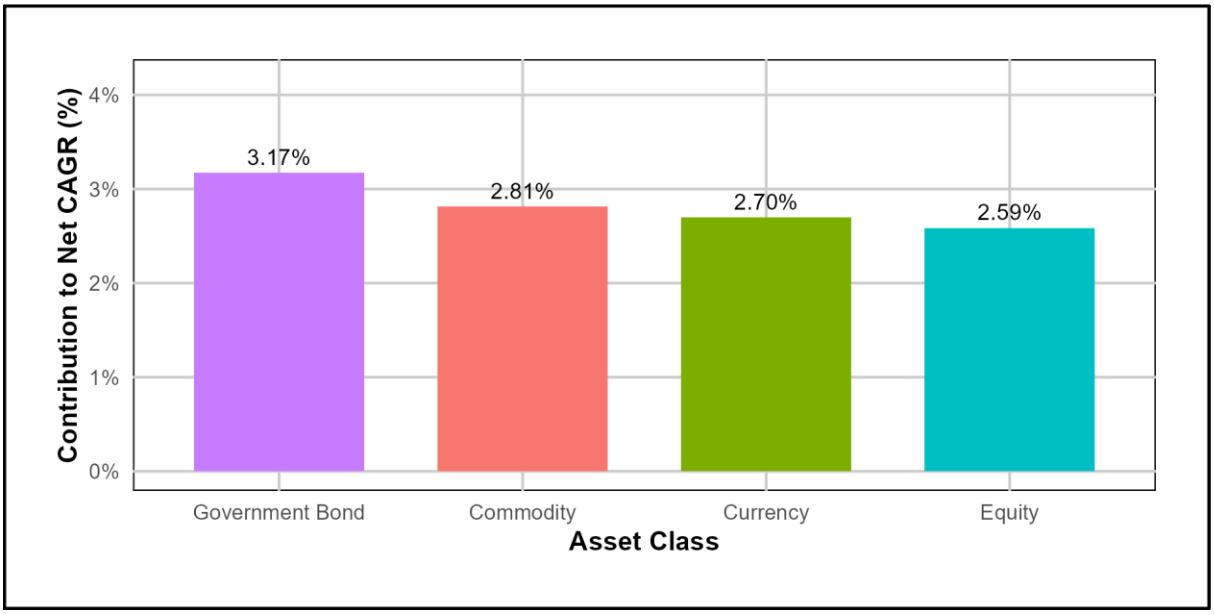

Government bonds contributed the most of all asset classes that made up the Trend Portfolio, adding 3.17% to the CAGR of the Trend Portfolio, followed by commodities, which contributed 2.81%, currencies, which contributed 2.70% and equities, which contributed 2.59%. This outcome is particularly noteworthy given the inherent differences in return distributions and volatilities across asset classes, as highlighted in Table 3.1. The methodology used ensured that no single asset class dominated the portfolio's risk profile, thereby achieving the intended diversification benefits.

3.4.4 Alpha and loading on risk factors in Trend Portfolio excess returns

Having established that the Trend Portfolio is a suitable proxy for a trend-following portfolio, its risk-adjusted performance is next examined. Hypothesis $H_{0,C}$ is tested, that trend-following returns are fully explained by traditional equity market factors, including the global market (ACWI), size (SMB), value (HML), and momentum (UMD), as well as by global government bonds and commodities.

Table 3.5 summarises the regression output of the monthly excess Trend Portfolio returns on the excess returns of the Fama-French three-factor model, extended to a four-factor model by adding a cross-sectional momentum factor. ACWI represents the global market portfolio, SMB represents the size factor, HML represents the value factor and UMD represents the cross-sectional momentum factor. The asterisks in the t-statistics denote the levels of statistical significance for each coefficient in the regression analysis. Specifically, a single asterisk indicates significance at the 10% level, two asterisks signify significance at the 5% level, and three asterisks represent significance at the 1% level.

Table 3.5 Summary of regression results: Trend Portfolio vs. traditional market factors

Statistic	**ACWI**	**SMB**	**HML**	**UMD**	**Intercept**	**R-Squared**
Coefficient	0.01	-0.04	0.06	0.25	0.87%	9%
(t-Stat)	(0.24)	(-0.59)	(0.93)	(5.35)***	(4.09)***	

Consistent with the literature, the results indicate that the Trend Portfolio delivers a positive alpha of 0.87% per month after risk-adjusting for the factors, statistically significant at the 1% level. The Trend Portfolio shows a near-zero loading on the global equity market factor (ACWI), and this coefficient is not statistically significant (as indicated by the t-statistic of +0.24). The lack of significance suggests that the Trend Portfolio's returns are not strongly driven by overall market movements. This implies that the strategy returns are equity-market-neutral, consistent with the existing trend-following literature. As such, the strategy offers the potential for meaningful diversification benefits to a global market portfolio.

Similarly, both SMB and HML have near-zero coefficients. The very slightly negative coefficient on the SMB factor indicates that the Trend Portfolio tends to load negatively on the size factor, meaning it may have a tilt towards larger-cap equities rather than small-cap equities. However, this relationship is not statistically significant (t-statistic of -0.59), implying that the size effect does not play a crucial role in the returns of the Trend Portfolio. The Trend Portfolio shows a very slightly positive loading on the value factor (HML), suggesting a slightly greater loading on value stocks over growth stocks. However, this coefficient is also not statistically significant (t-statistic of +0.93), which indicates that the value exposure does not significantly influence the Trend Portfolio's returns.

The Trend Portfolio has a positive and statistically significant loading on the cross-sectional momentum factor (UMD) at the 1% significance level. This is in line with the findings in Moskowitz et al. (2012). It is also intuitive that a trend-following strategy and the momentum factor are somewhat related in that they each attempt to exploit the tendency for markets that are rising (falling) to continue rising (falling). The key difference between the approaches is that trend-following positions can be both long and short and utilise multiple asset classes. However, the intercept term shown in Table 3.5 indicates a 0.87% alpha per month, which is positive and statistically significant at the 1% level. Therefore, despite the loading on UMD, the large, positive and statistically significant alpha of 0.87% illustrated in Table 3.5 suggests that the momentum factor plays a small role at best in trend-following outperformance over time.

Next, the regression analysis is expanded to include global government bond (BONDS) and commodity index (GSCI) returns, proxied by the Bloomberg Global-Aggregate Total Return Index (Value Unhedged) and the S&P GSCI Index, respectively. The regression output is summarised in Table 3.6.

Table 3.6 Summary of regression results: Trend Portfolio vs. traditional market factors, global government bonds and commodities

Statistic	**ACWI**	**SMB**	**HML**	**UMD**	**BONDS**	**GSCI**	**Intercept**	**R-Squared**
Coefficient	0.05	-0.03	0.08	0.26	-0.07	-0.05	0.85%	10%
(t-Stat)	(0.88)	(-0.43)	(1.2)	(5.50)***	(-0.53)	(-1.35)	(4.00)***	

Table 3.6 expands the model to include global government bonds and commodities. The inclusion of these additional factors allows us to assess whether exposure to these markets influences Trend Portfolio returns. The Trend Portfolio maintains a statistically significant and positive alpha of 0.85% per month, even after accounting for the expanded set of factors. This alpha is statistically significant at the 1% level.

There was not much change in the coefficients and t-statistics in the ACWI, SMB, HML and UMD factors, and the interpretation of Table 3.5 holds for the regression summarised in Table 3.6. The introduction of the global government bond factor (BONDS) reveals a slightly negative coefficient of -0.07 that is not statistically significant (the t-statistic is -0.53). Similarly, the coefficient on the commodity factor (GSCI) is slightly negative and statistically insignificant, with a t-statistic of -1.35. This indicates that the Trend Portfolio does not exhibit a strong relationship with bond and commodity returns. These findings are in line with those presented in the literature review and emphasise the value that trend-following offers as a diversifier in traditional asset allocation mixes.

Based on the regression analyses presented in Tables 3.5 and 3.6, the null hypothesis $H_{0, C}$ that trend-following returns are fully explained by traditional equity market factors (global market, SMB, HML, UMD), global government bonds, and commodities is rejected. The findings demonstrate that the Trend Portfolio consistently produces a statistically significant positive alpha of 0.85% per month, even after accounting for these factors.

This significant alpha indicates that there are additional sources of return that are not captured by the traditional equity factors, global government bonds, or commodities. Furthermore, while the momentum factor (UMD) plays a role in explaining the returns, the economically and statistically significant alpha of the Trend Portfolio suggests that its performance cannot be explained by the factors controlled for. The evidence implies that trend-following returns, as proxied by the Trend Portfolio, may offer meaningful diversification benefits to portfolios with global equity, bond, and commodity exposure.

3.4.5 Trend-following convexity with respect to global equity returns

In the next section, the potential usefulness of diversifying global equity portfolios with trend-following is investigated by testing $H_{0, D}$: Trend-following returns are not convex with respect to global equity returns. To do this, several scatterplots of Trend Portfolio returns relative to global equity returns, represented by the ACWI, are illustrated to better understand the relationship between the strategies.

Subsequently, a regression of Trend Portfolio returns against both ACWI and squared ACWI monthly returns is conducted to statistically test whether trend-following returns are convex with respect to extreme equity market performance. Finally, the performance of the Trend Portfolio and ACWI is stress-tested on 9 historical crisis episodes defined in Harvey et al. (2021) and Schwalbach and Auret (2023).

3.4.5.1 Scatterplots illustrating trend-following convexity with respect to global equity returns

Scatterplots of non-overlapping Trend Portfolio excess returns versus ACWI excess returns over 3 different reporting periods: monthly, quarterly, and annually, are shown, and findings are presented in Figures 3.11 to 3.13, respectively. The quadratic regression fit is then illustrated using a curved red line in each figure. The scatter plots in these figures illustrate a key characteristic of trend-following portfolios: the “smile”.

The “smile” can be described as the tendency for trend-following returns to be convex with respect to equity returns, and where this convexity becomes more pronounced as the measurement horizon lengthens. This "smile" reflects how the Trend Portfolio tends to generate its highest periodic returns during periods of extreme market movements, whether bullish or bearish, while showing little or no correlation with equities during more normal market conditions, which was demonstrated in the previous section.

Figure 3.11 Convexity of Trend Portfolio monthly excess returns relative to ACWI monthly returns

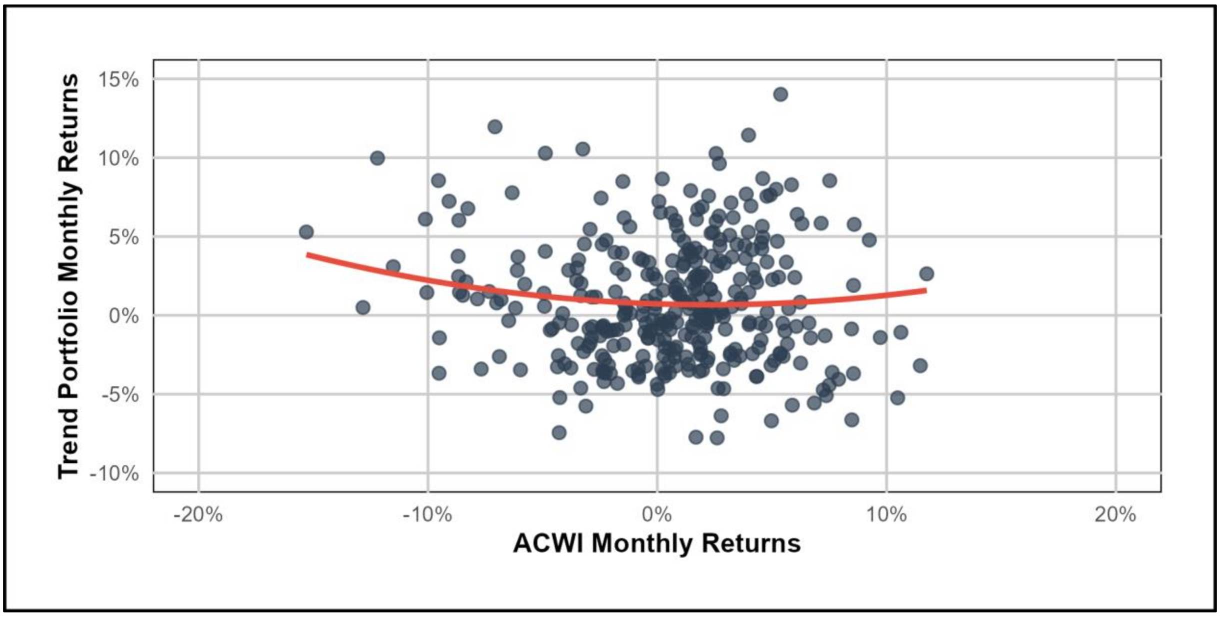

Figure 3.12 Convexity of Trend Portfolio quarterly excess returns relative to ACWI quarterly excess returns

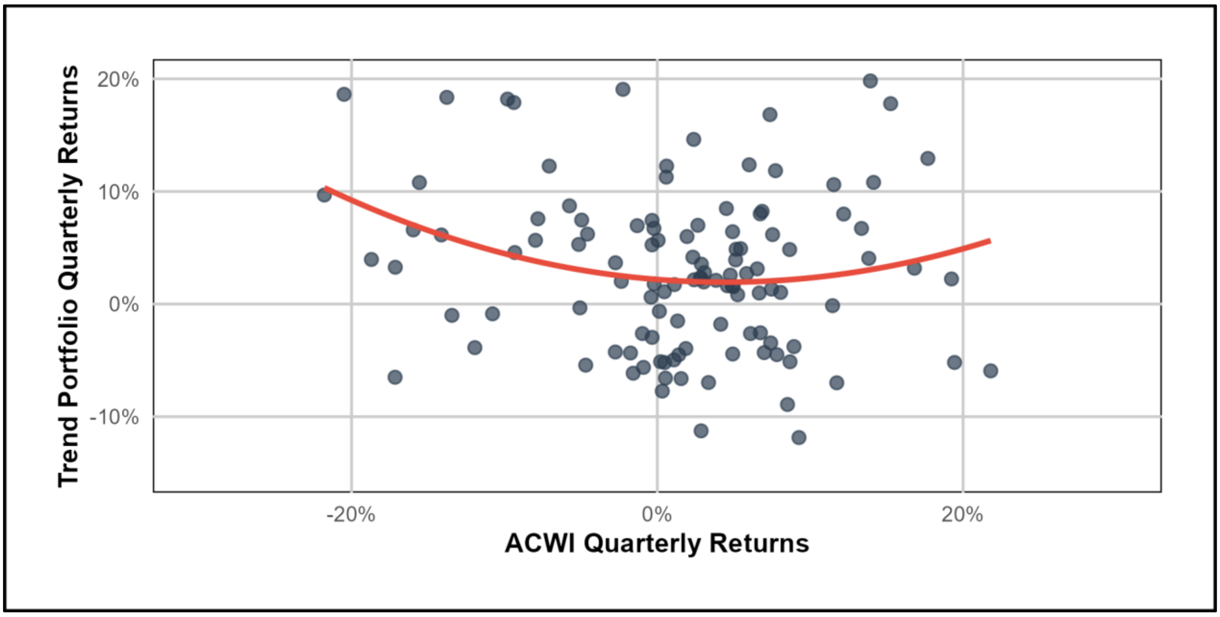

Figure 3.13 Convexity of Trend Portfolio annual excess returns relative to ACWI annual excess returns

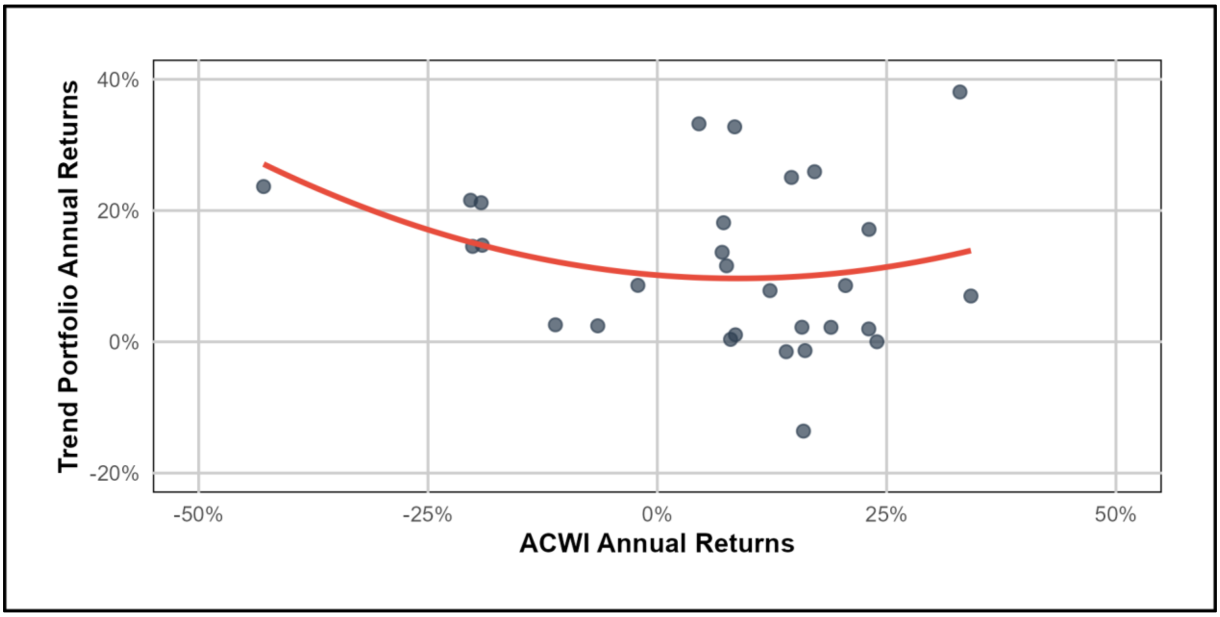

The findings demonstrated in Figures 3.11 to 3.13 are consistent with the existing literature on trend-following strategies. During bear markets, the Trend Portfolio typically shifts toward a greater allocation of short positions, enabling it to capture profits from the continuation of downward price trends. Conversely, during bull markets, the portfolio adjusts its positioning to a higher proportion of long positions, allowing it to benefit from sustained upward price movements. This dynamic rebalancing, automatically responding to the prevailing market conditions, results in the portfolio capitalising on directional trends, regardless of the market's overall direction.

This ability to dynamically adjust between long and short positions creates a convex return profile with respect to ACWI returns across multiple time horizons. This behaviour is reminiscent of a lookback straddle, as described by Fung and Hsieh (2002). However, unlike a lookback straddle, the trend-following “smile” does not cross into negative territory in any of Figures 3.11 to 3.13, indicating that the Trend Portfolio exhibits a consistently positive return bias across the distribution of equity returns.

3.4.5.2 Regression summary: Trend Portfolio monthly returns against ACWI and squared ACWI monthly returns

To statistically test whether the convexity observed in the Trend Portfolio's performance is significant, particularly during extreme upward and downward movements in the equity markets, the analysis follows Moskowitz et al. (2012) and reports, in Table 3.7, the coefficients from a regression of Trend Portfolio monthly returns on ACWI monthly returns as well as squared ACWI monthly returns.

Table 3.7 Summary of regression results: Trend Portfolio, ACWI and squared ACWI excess returns

Statistic	ACWI	ACWI Squared	Intercept	R-Squared
Coefficient	-0.07	1.41	0.72%	3%
(t-Stat)	(-1.46)	(2.30)**	(2.85)***	

Consistent with the findings presented in Tables 3.5 and 3.6, the loading on ACWI is not statistically significant, indicating that trend-following returns remain uncorrelated with global equity returns, which is in line with the existing literature. However, Table 3.7 shows that the coefficient on squared monthly ACWI returns is +1.41, with a t-statistic of +2.30, which is statistically significant at the 5% level. This positive and significant coefficient provides evidence of convexity in the Trend Portfolio's excess returns with respect to global equity returns.

The results are consistent with Moskowitz et al. (2012), who also found a positive and statistically significant relationship between trend-following and squared equity returns. Moreover, these findings align with the trend-following “smile”, the characteristic convexity of trend-following returns with respect to global equity returns, as illustrated in Figures 3.11 to 3.13 and documented in the literature.

3.4.6 Historical crisis episode performance comparison

Figures 3.14 and Table 3.8 provide a comprehensive overview of the performance of both the ACWI and the Trend Portfolio during several notable market crises. The starting and ending dates for these crises are defined by Harvey et al. (2021), as well as Schwalbach and Auret (2023). The crises span from the Asian Crisis in 1998 to the High Inflation period in 2022.

In Figure 3.14, the full duration of each crisis episode is highlighted (in grey) and the cumulative excess returns of both the Trend Portfolio and the ACWI are annotated from the crisis start date to its conclusion. In absolute terms, it is observed that the Trend Portfolio's returns have generally been largest during longer-lived crises. For example, during the Tech Bubble Burst, which spanned 548 weekdays, the Trend Portfolio generated a cumulative return of 51.42%, while the ACWI fell by 51.02%. Similarly, during the first leg of the GFC in 2007-2008, which lasted 292 weekdays, the Trend Portfolio delivered a return of 32.34%, while the ACWI suffered a loss of 55.20%.

During the High Inflation episode of 2022, the Trend Portfolio returned 19.69%, while the ACWI declined by 27.06%. The direction and magnitude of the Trend Portfolio returns in these periods are roughly in line with those of the SG Trend Index.

The literature establishes that trend-following strategies have historically generated outsized positive performance during prolonged crises. These strategies typically position themselves short falling markets, profiting as these trends continue. They take long positions in markets that have been rising during the crisis, benefiting from a continuation in those positive trends. Given the relatively long-term lookback periods used to form trend positions (as discussed in the methodology section), the longer the crisis, the more time is afforded to trend-following managers to adjust their positioning. This allows them to better capitalise on sustained market movements.

Figure 3.14 Cumulative Excess Returns of Trend Portfolio and ACWI with Crises Highlighted

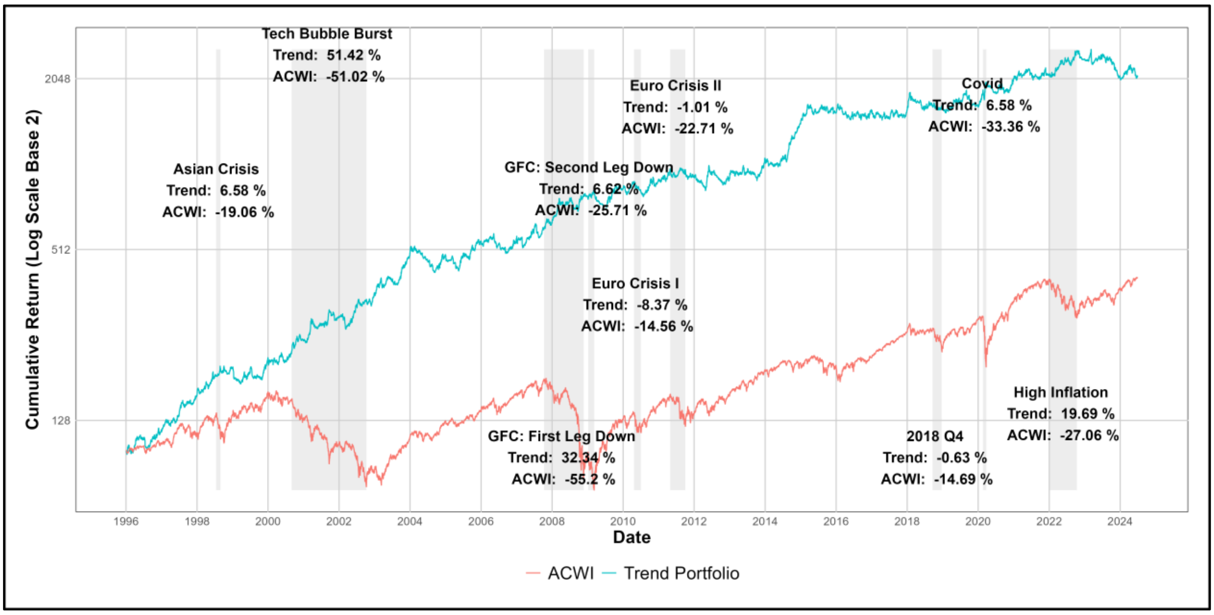

In Table 3.8, the starting and ending dates for each crisis episode as defined by Harvey et al. (2021), and Schwalbach and Auret (2023) are presented. The longest period of stress examined was the Tech Bubble Burst from 2000 to 2002, while the shortest was the Covid Crash of 2020, which spanned only 23 weekdays. The evidence for the Trend Portfolio's ability to deliver convex returns during shorter-lived crises is mixed. During the Asian Crisis, which lasted 31 weekdays, the Trend Portfolio achieved a cumulative return of 6.58%, while ACWI posted a loss of 19.06%. Similarly, during the second leg of the GFC, the Trend Portfolio returned 6.62%, while ACWI dropped 25.71%.

During the 50 weekdays of the first Euro Crisis, the Trend Portfolio declined by 8.37%, while ACWI fell by 14.56%. In the second Euro Crisis, which spanned 111 weekdays, the Trend Portfolio remained relatively stable, with a slight loss of 1.01%, compared to a 22.71% fall in ACWI. During the final quarter of 2018, the Trend Portfolio was largely unchanged, losing just 0.63%, while ACWI dropped by 14.69%. Conversely, the Trend Portfolio gained 6.58% during the Covid Crash, outperforming ACWI, which dropped by 33.36%. During the High Inflation episode of 2022, the Trend Portfolio gained 19.69%, while ACWI experienced a loss of 27.06%.

In summary, across the nine crisis episodes where ACWI experienced significant losses, the Trend Portfolio posted flat returns in two episodes, negative returns in two, and positive returns in the remaining five. These findings suggest that trend-following strategies have empirically offered meaningful diversification benefits to buy-and-hold global equity strategies during extreme periods of market stress.

Table 3.8 Crisis episode performance: Trend Portfolio and ACWI

Period Name	**Peak Day**	**Trough Day**	**Number of Weekdays**	**Trend Portfolio Cumulative Return**	**ACWI Cumulative Return**
Asian Crisis	1998-07-17	1998-08-31	31	6.58%	-19.06%
Tech Bubble Burst	2000-09-01	2002-10-09	548	51.42%	-51.02%
GFC: First Leg Down	2007-10-09	2008-11-20	292	32.34%	-55.20%
GFC: Second Leg Down	2009-01-06	2009-03-09	44	6.62%	-25.71%
Euro Crisis I	2010-04-23	2010-07-02	50	-8.37%	-14.56%
Euro Crisis II	2011-04-29	2011-10-03	111	-1.01%	-22.71%
2018 Q4	2018-09-20	2018-12-20	65	-0.63%	-14.69%
Covid	2020-02-19	2020-03-23	23	6.58%	-33.36%
High Inflation	2022-01-03	2022-10-12	202	19.69%	-27.06%

3.4.7 Summary statistics of performance and the distribution of returns

Next, summary statistics of excess monthly Trend Portfolio returns are presented and compared against ACWI in Table 3.9. The Trend Portfolio's CAGR is meaningfully higher than ACWI at 11.26% compared to 5.05%, respectively. The Trend Portfolio also exhibited lower volatility than ACWI. The Sharpe ratio of +0.81 for the Trend Portfolio compared to +0.32 for ACWI indicates that Trend Portfolio returns were also higher on a risk-adjusted basis.

Trend Portfolio returns exhibit positive skewness throughout the period. This property is well established in trend-following literature and is consistent with the findings of Fung and Hsieh (2002), who note the similarity of trend-following returns and a long lookback straddle option strategy. In a long straddle strategy, an investor buys both a call option and a put option with the same strike price and expiration date. This strategy benefits from significant price movements in either direction, leading to outsized gains when the underlying asset moves significantly. Similarly, trend-following strategies tend to maintain and even increase positions in profitable trends while cutting losses quickly when trends reverse. This approach naturally leads to positively skewed returns, as the strategy captures large gains during strong trending markets. In contrast, ACWI returns exhibit negative skewness, meaning that downside risk can be substantial during periods of market stress, as established in the previous section.

Table 3.9 Summary Statistics: Trend Portfolio and ACWI

Statistic	Trend Portfolio	ACWI
CAGR	11.26%	5.05%
Annualised Volatility	13.91%	15.59%
Sharpe Ratio	0.81	0.32
Skewness	0.73	-0.65
Excess Kurtosis	-1.74	-1.54
VaR (95%)	-4.43%	-8.16%
VaR (99%)	-6.67%	-12.04%
CVaR (95%)	-5.72%	-10.44%
CVaR (99%)	-7.41%	-15.18%
Best 3M Performance	19.86%	33.06%
Worst 3M Performance	-12.30%	-34.86%

The visual difference in downside risk are evident in Figure 3.15, which shows the peak-to-trough drawdowns of the Trend Portfolio and ACWI. The Trend Portfolio experiences frequent but comparatively shallow and short-lived drawdowns when compared to ACWI. In addition, Trend Portfolio drawdowns rarely exceeded 20%, suggesting that the strategy cuts loss-making positions quickly, either because the volatility of the position spikes leading to a smaller position or because the actual signal generated from the return over the lookback period changes.

This finding is further supported by the tail risk statistics presented in Table 3.9. VaR and CVaR, at both the 95% and 99% confidence levels show that the Trend Portfolio is less exposed to extreme losses than ACWI. At the 95% confidence level, the VaR for the Trend Portfolio is -4.43%, compared to -8.16% for ACWI, and the CVaR shows a similar pattern, -5.72% for the

Trend Portfolio and -10.44% for ACWI. At the 99% level, the differences are even more pronounced: the Trend Portfolio's VaR is -6.67% compared to -12.04% for ACWI, and its CVaR is -7.41% versus -15.18% for ACWI. The worst 3-month performance of the Trend Portfolio was -12.30%, compared to -34.86% for ACWI.

Figure 3.15 Drawdowns of the Trend Portfolio and ACWI

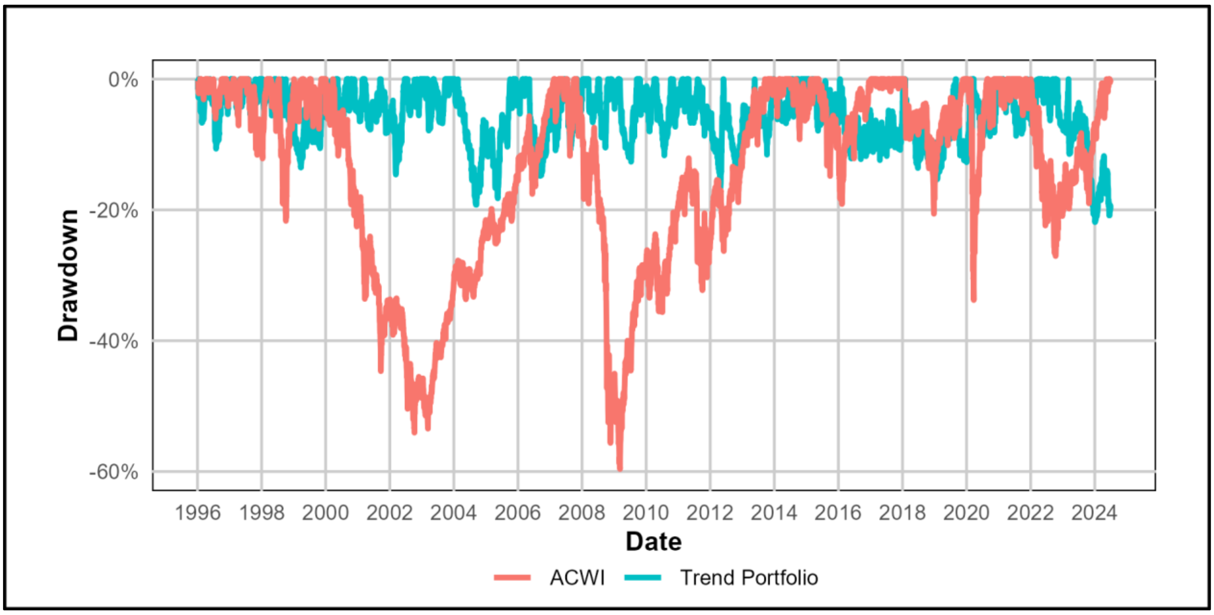

The findings presented in Figures 3.11 to 3.15 and in Tables 3.7 to 3.9 suggest that the null hypothesis $H_{0,\,D}$: Trend-following returns are not convex with respect to global equity returns should be rejected in favour of the alternative hypothesis $H_{1,\,D}$: Trend-following returns are convex with respect to global equity returns. In addition, the dichotomy in skewness, and the superior tail risk profile of the Trend Portfolio add further substance to the argument for the diversification potential of trend-following strategies.

3.5 FINDINGS: TREND OVERLAY PORTFOLIO

The results presented in the previous section support the notion that trend-following strategies have the potential to offer meaningful diversification benefits to equity-dominant portfolios, especially during extreme market conditions. In this next section, the notion is put to the test and the final hypothesis of the chapter, $H_{0, E}$: Overlaying a trend-following strategy on top of global equity does not enhance risk-adjusted and absolute returns compared to holding global equity alone, is explored. To test the hypothesis, the same procedure as in Schwalbach and Auret (2023) is followed, where 50% of the Trend Portfolio strategy is overlaid on top of a full 100% ACWI portfolio (Trend Overlay) and compared against a 100% ACWI portfolio (ACWI). It is assumed that the underlying ACWI portfolio serves as collateral for the margin required to overlay 50% of the trend-following strategy. The daily returns of the Trend Overlay portfolio thus reflect the contributions of the ACWI equity exposure and the performance of 50% of the trend-following strategy. All findings are based on the excess returns of the resulting portfolios.

3.5.1 Excess performance comparison

In Figure 3.16, the cumulative excess return of the Trend Overlay portfolio is illustrated and compared against ACWI. Figure 3.16 shows that cumulative returns for the Trend Overlay portfolio were meaningfully higher than ACWI while still largely resembling the contours of an equity-dominant portfolio. According to Table 3.10, the CAGR of the Trend Overlay portfolio was 11.12% compared to 5.05% for ACWI. This finding is consistent with Schwalbach and Auret (2023), who showed that long-term ACWI cumulative returns were empirically higher when overlaid with 50% exposure in the trend-following industry benchmark, SG Trend Index. The difference between the findings of this study and the findings and Schwalbach and Auret (2023) is the magnitude of the outperformance. This could be partially explained by manager fees and performance fees. The SG Trend Index is net of these fees, whereas the Trend Overlay portfolio is gross.

Figure 3.16 Cumulative excess returns of the Trend Overlay portfolio and ACWI

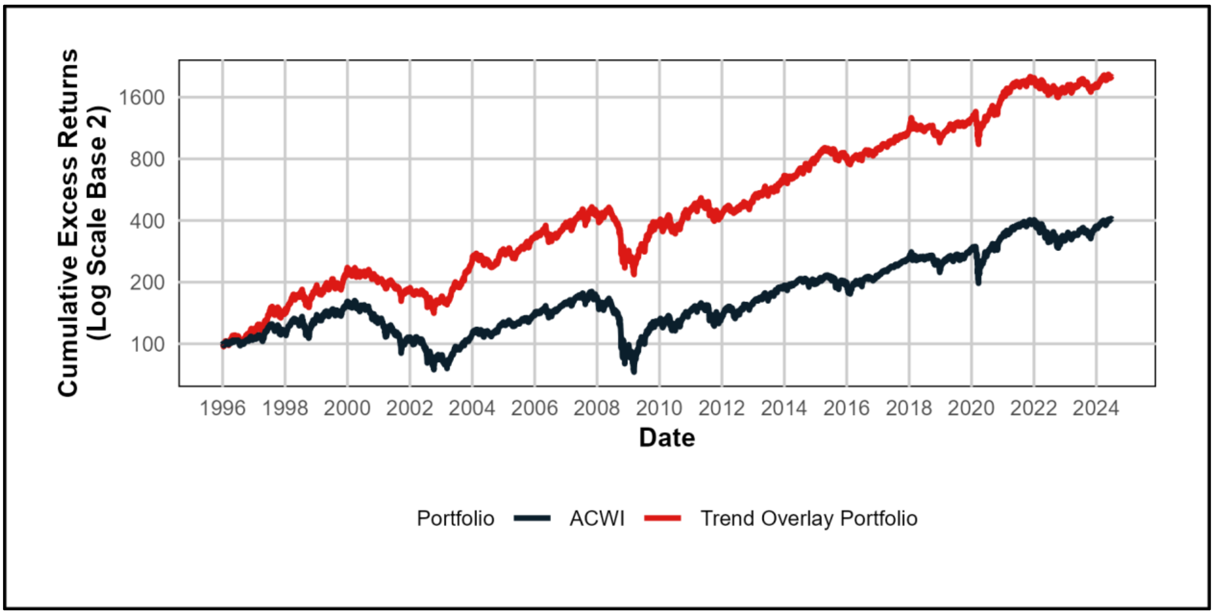

In terms of drawdowns, the findings are in line with those in Schwalbach and Auret (2023) in that overlaying a trend-following strategy on top of a global equity portfolio has empirically led to outperformance on average during periods of market stress, particularly during prolonged drawdown episodes. The Trend Overlay portfolio experienced drawdowns at roughly the same times as ACWI as shown in Figure 3.17. However, when ACWI suffered deep drawdowns that exceeded 20%, Trend Overlay drawdowns were shallower on average. Recoveries tended to be much faster for the Trend Overlay portfolio after prolonged market crises such as the Dot-Com Bubble Burst episode in the early 2000s and the full GFC in 2007-2009.

Despite the standalone Trend Portfolio volatility of 13.91% and ACWI volatility of 15.62%, when the two were combined to form the Trend Overlay portfolio, the resulting volatility was only slightly higher than ACWI at 16.40% per Table 3.10. This finding demonstrates the diversification benefits offered by trend-following with respect to long equity portfolios. Overlaying a 50% allocation to the Trend Portfolio on top of 100% in ACWI empirically added meaningfully to CAGR while increasing volatility by less than a percentage point. This resulted in a large improvement in the Sharpe ratio of the Trend Overlay portfolio of +0.68 compared to +0.32 delivered by ACWI.

Figure 3.17 Drawdowns of Trend Overlay and ACWI

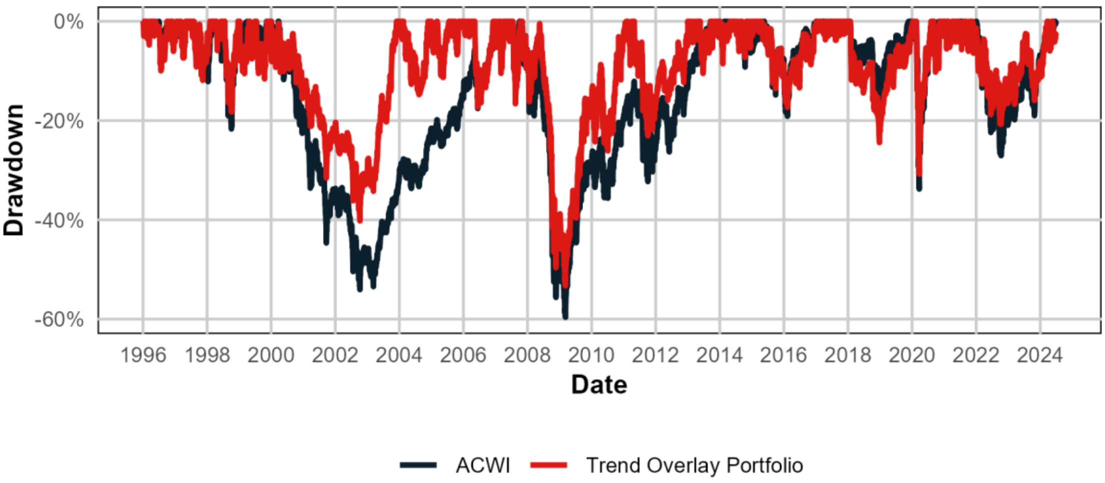

Overlaying Trend Portfolio excess returns on top of global equity exposure also reduced the negative skewness of ACWI from -0.65 to -0.22, as shown in Table 3.10. This is in line with the findings presented in the previous section, where it is demonstrated that standalone trend-following returns tend to be positively skewed in contrast with equity returns, which tend to be negatively skewed. When examining the tail risk statistics, it is notable that the Trend Overlay portfolio meaningfully improves on those of ACWI.

At the 95% confidence level, the VaR for the Trend Overlay is -6.93%, noticeably smaller than ACWI's -8.16%. The 99% VaR also reveals a similar pattern: -10.66% for the Trend Overlay versus -12.04% for ACWI. When looking at CVaR at the 95% confidence level, the Trend Overlay portfolio's -9.22% outcome was less severe than that of ACWI, at -10.44%. The same is true for the 99% CVaR: -12.98% for the Trend Overlay and -15.18% for ACWI.

These results suggest that the trend-following strategy overlaid on top of equity empirically resulted in dampening extreme losses relative to global equities, despite having slightly higher overall volatility. This finding is notable as the overlay effectively adds leverage to an equity portfolio, and yet, reduces tail risk statistics due to the diversifying nature of a trend-following strategy with respect to equity. The best-performing 3-month period for the Trend Overlay portfolio is 28.03%, and 33.06% for ACWI. However, the diversifying potential of trend-following is evident in the worst 3-month performance for the Trend Overlay portfolio being -30.66%, compared to ACWI of -34.86%.

Table 3.10 Summary statistics: Trend Overlay portfolio and ACWI

Statistic	Trend Overlay Portfolio	ACWI
CAGR	11.12%	5.05%
Annualised Volatility	16.40%	15.59%
Sharpe Ratio	0.68	0.32
Skewness	-0.22	- 0.65
Excess Kurtosis	-2.75	-1.54
VaR (95%)	-6.93%	-8.16%
VaR (99%)	-10.66%	-12.04%
CVaR (95%)	-9.22%	-10.44%
CVaR (99%)	-12.98%	-15.18%
Best 3M Performance	28.03%	33.06%
Worst 3M Performance	-30.66%	-34.86%

In Figure 3.18, the correlation between the rolling annual returns of the Trend Overlay portfolio and ACWI is shown. Rolling annual returns are chosen because they represent a common reporting frequency. Correlation oscillates over time between 80% and 100%, indicating a high degree of correlation. Correlations have tended to make large stepwise changes following extreme events.

Figure 3.18 Rolling one-year correlations of the Trend Overlay portfolio and ACWI

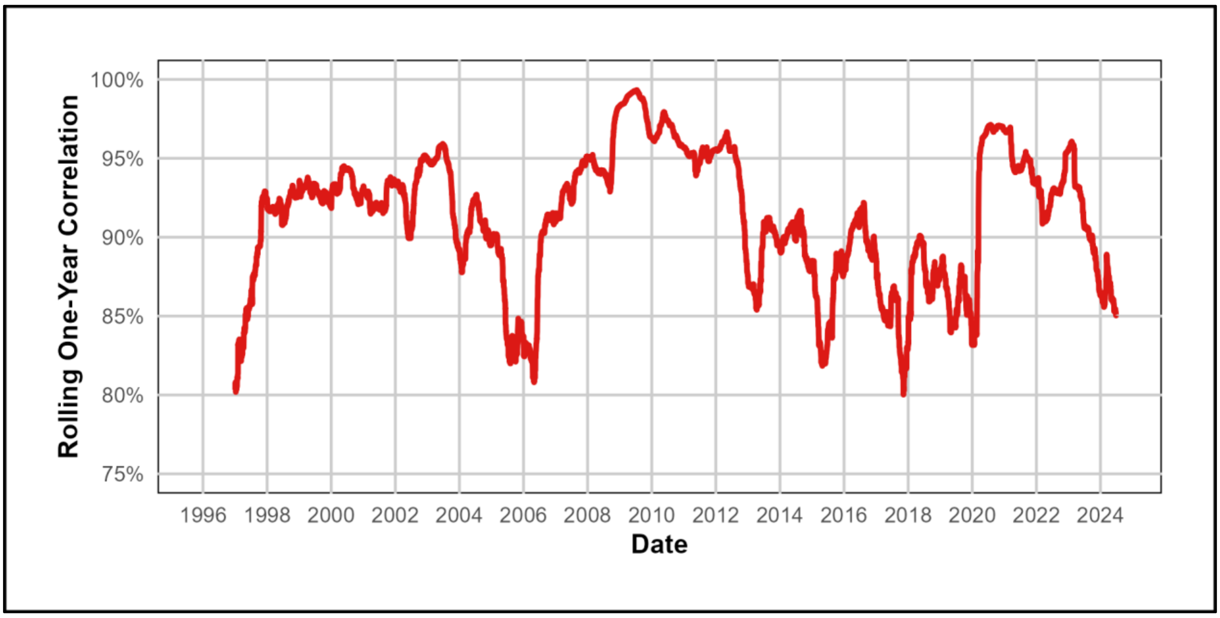

3.5.2 Alpha and loading on risk factors in Trend Overlay portfolio excess returns

In Figure 3.19, a plot displaying the monthly excess returns of ACWI on the x-axis against matching monthly excess returns of the Trend Overlay portfolio on the y-axis is illustrated. A linear regression is run, and the output is presented in Table 3.11. The monthly excess returns of the Trend Overlay portfolio exhibit a strong and positive coefficient of +0.96 with ACWI monthly excess returns, statistically significant at the 1% level. After adjusting for ACWI, the Trend Overlay portfolio generates a statistically significant monthly alpha of 0.50% with a t-statistic of +4.64, also significant at the 1% level. This alpha indicates that the Trend Overlay portfolio has empirically delivered excess returns above what would be expected from exposure to ACWI alone, demonstrating the diversification benefits achieved through overlaying the trend-following strategy. A strong and positive loading on ACWI, combined with a high R-squared of 83%, suggests that the Trend Overlay portfolio has empirically generated alpha while remaining suitable for investors who seek the majority of portfolio returns to stem from global equity beta.

Figure 3.19 Linear regression analysis of monthly excess returns: Trend Overlay portfolio vs. ACWI

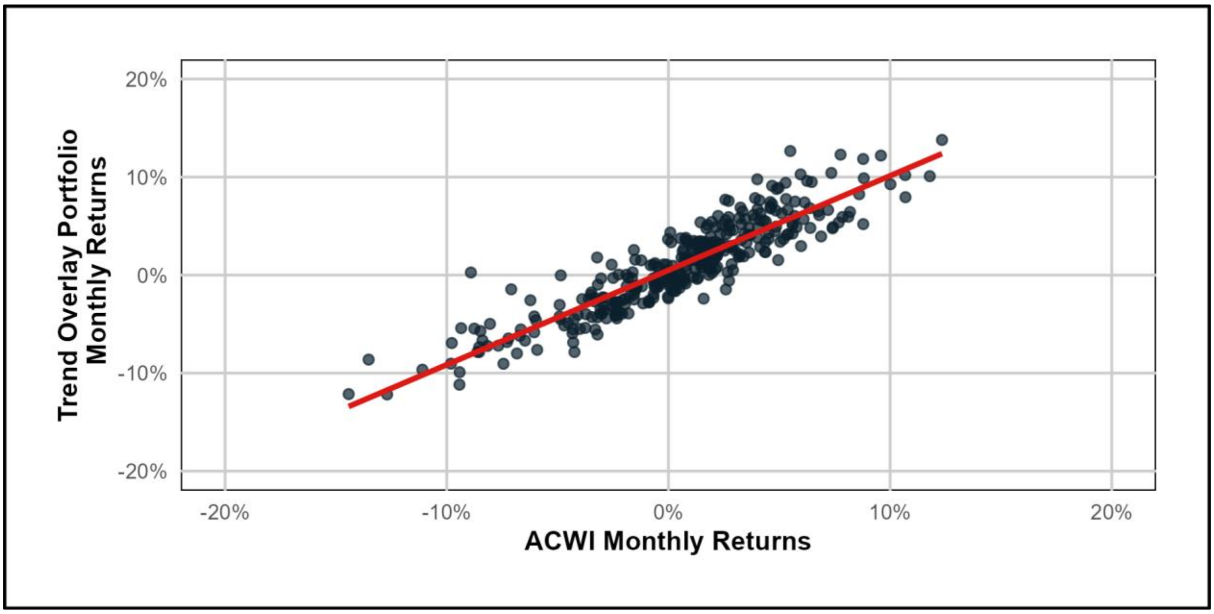

Table 3.11 Summary of regression results: Trend Overlay portfolio and ACWI

Statistic	ACWI	Intercept	R-Squared
Coefficient	0.96	0.50%	83%
(t-Stat)	(40.08)***	(4.64)***	

Next, the analysis assesses whether the Trend Overlay portfolio returns can be explained by traditional equity market factors, namely the global market (ACWI), size (SMB), value (HML), and momentum (UMD), as well as by global government bonds and commodities. Table 3.12 summarises the regression output of the monthly excess Trend Overlay portfolio returns on the excess returns of the Fama-French three-factor model, extended to a four-factor model by including the cross-sectional momentum factor.

Table 3.12 Summary of regression results: Trend Overlay portfolio vs. traditional market factors

Statistic	ACWI	SMB	HML	UMD	Intercept	R-Squared
Coefficient	1.01	-0.01	0.04	0.12	0.42%	84%
(t-Stat)	(39.03)***	(-0.39)	(1.15)	(5.27)***	(3.96)***	

The evidence suggests that the Trend Overlay portfolio returns continue to load significantly on ACWI returns, with a coefficient of +1.01, which is statistically significant at the 1% level. Consistent with the literature, the findings indicate that the Trend Overlay Portfolio also loads positively on the momentum factor, UMD, with a coefficient of +0.12, which is statistically significant at the 1% level, as indicated by a t-statistic of +5.27. Notably, the strategy still delivers a statistically significant and positive alpha of 0.42% per month after risk adjusting for the factors. The Trend Overlay portfolio shows a near-zero loading on HML and SMB, and the coefficients are not statistically significant.

In Table 3.13, the regression analysis is extended to include global government bond and commodity index returns, proxied by the Bloomberg Global-Aggregate Total Return Index (Value Unhedged) and the S&P GSCI Index, respectively. Both bonds and commodities exhibit a small but negative coefficient that is not statistically significant.

These results indicate that the Trend Overlay portfolio's performance is largely independent of the size factor, the value factor, bond returns and commodity returns. Trend Overlay returns load positively and significantly on both global equity and the momentum factor. However, the strategy maintains a statistically significant and positive alpha of 0.41% per month after controlling for the full complement of factors.

Table 3.13 Summary of regression results: Trend Overlay portfolio against traditional market factors, global bonds and commodities

Statistic	ACWI	SMB	HML	UMD	BONDS	GSCI	Intercept	R-Squared
Coefficient	1.03	-0.01	0.04	0.13	-0.04	-0.02	0.41%	84%
(t-Stat)	(34.88)***	(-0.26)	(1.37)	(5.39)***	(-0.56)	(-1.13)	(3.88)***	

3.5.3 Trend Overlay portfolio return convexity with respect to equity returns

Having rejected the null hypothesis $H_{0,D}$ (trend-following returns are not convex with respect to global equity returns) earlier in the chapter, attention is now turned to whether overlaying a trend-following strategy on top of a global equity portfolio has resulted in introducing convexity from the embedded trend-following strategy.

3.5.3.1 Scatterplots illustrating Trend Overlay portfolio convexity with respect to ACWI

Figures 3.20 to 3.22 illustrate how overlaying a trend-following strategy on top of a long global equity portfolio has empirically introduced convexity. Convexity refers to the non-linear relationship between the returns of the Trend Overlay portfolio and ACWI, indicating that the overlay provided value during extreme market conditions on both the left and right tail of the equity return distribution.

The monthly, quarterly, and annual ACWI excess returns are plotted on the x-axis against the corresponding Trend Overlay portfolio returns on the y-axis. For each return frequency, a regression model with a quadratic fit is applied, illustrated with a red curved line. Across all time frequencies, there is clear evidence of convexity in Trend Overlay portfolio returns with respect to ACWI returns.

In terms of left-tail ACWI returns, the Trend Overlay portfolio has, on average, delivered smaller losses, meaning that the trend-following strategy embedded has mitigated risk. In terms of right-tail ACWI returns, the Trend Overlay portfolio has, on average, delivered slightly higher returns, meaning that the trend-following strategy embedded has enhanced upside participation. The tendency for trend-following strategies to deliver a U-shaped payoff profile, as found earlier in the chapter and reminiscent of a lookback straddle strategy, has effectively been translated into an equity-dominant portfolio. This finding helps to explain the sustained outperformance of the Trend Overlay portfolio relative to ACWI.

Figure 3.20 Convexity of Trend Overlay portfolio monthly excess returns relative to ACWI monthly excess returns

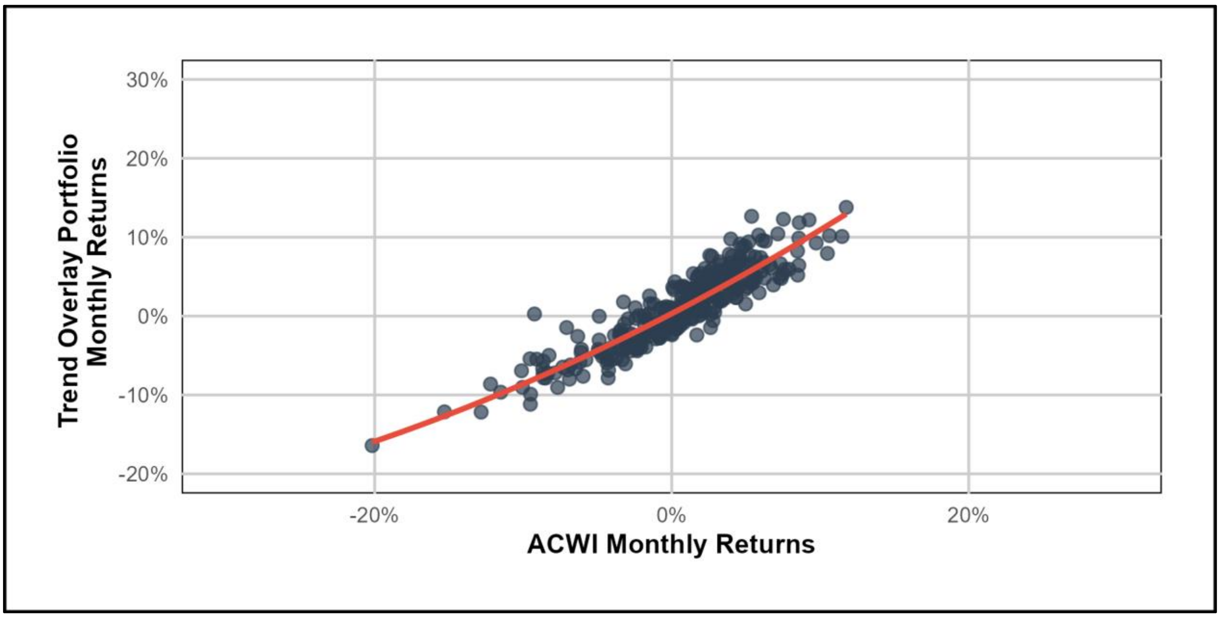

Figure 3.21 Convexity of Trend Overlay portfolio quarterly excess returns relative to ACWI quarterly excess returns

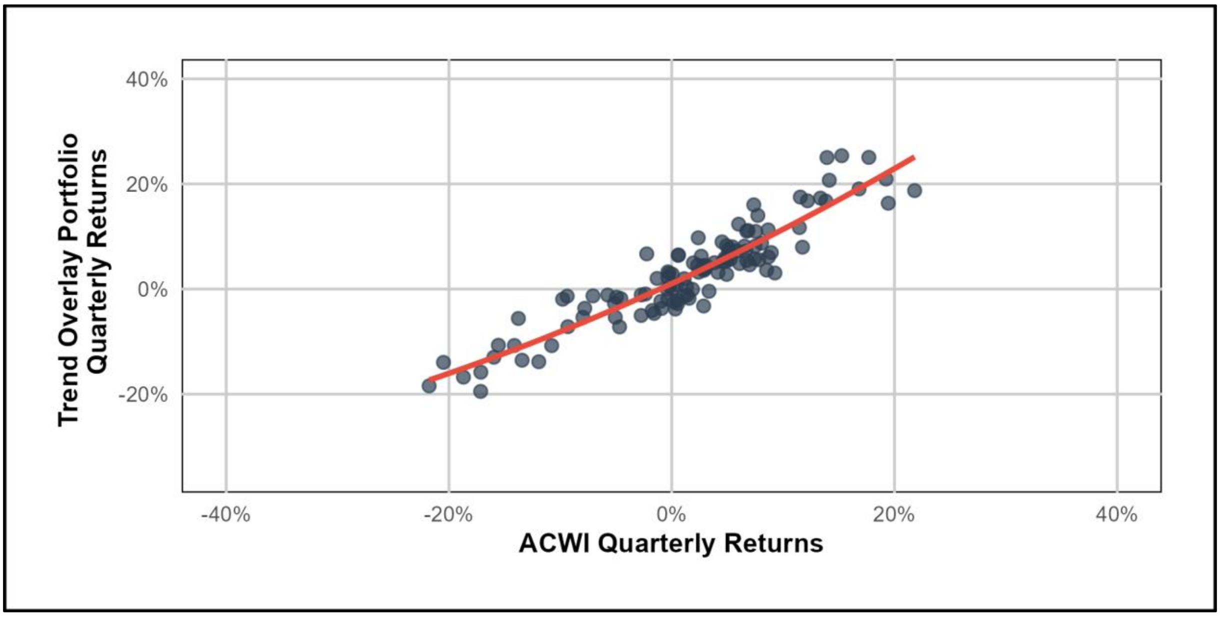

Figure 3.22 Convexity of Trend Overlay portfolio annual excess returns relative to ACWI annual excess returns

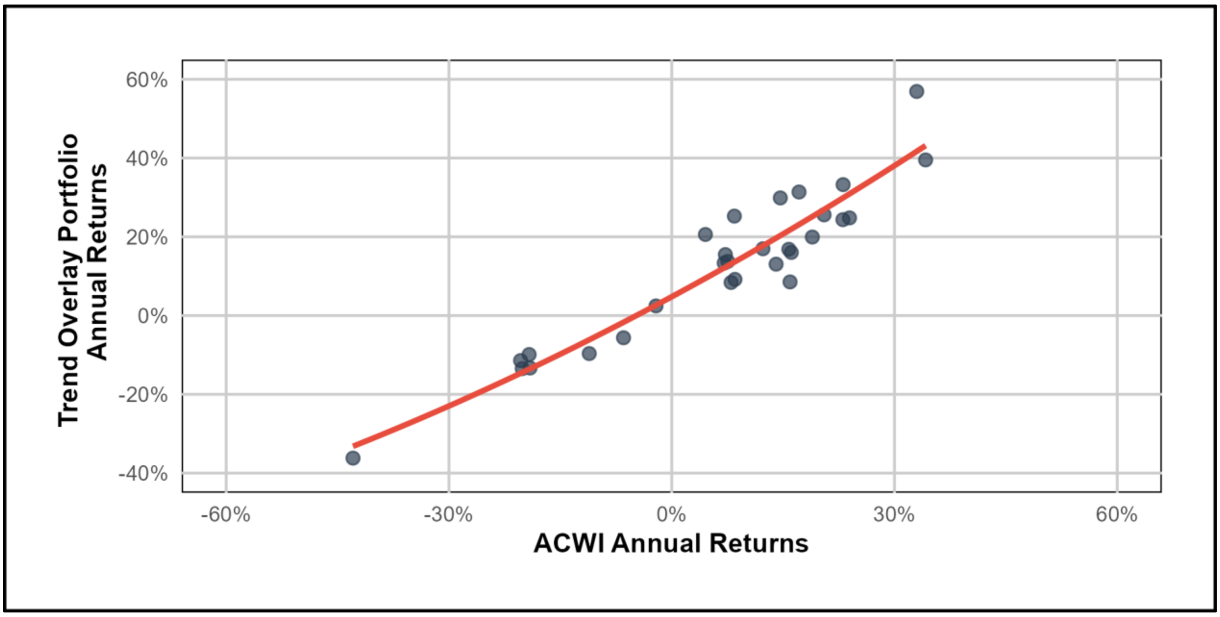

3.5.3.2 Regression summary: Trend Overlay portfolio monthly excess returns against ACWI and squared ACWI monthly excess returns

Next, a regression of Trend Overlay portfolio excess returns against both ACWI and squared ACWI monthly excess returns is conducted to statistically test whether overlaying trend-following returns on top of ACWI has empirically resulted in portfolio returns that are convex with respect to extreme equity market performance. A large and significant coefficient on squared ACWI returns would suggest that the Trend Overlay portfolio exhibits convexity.

Table 3.14 summarises the regression output. The findings suggest that the Trend Overlay portfolio continues to load positively and statistically significantly on ACWI monthly excess returns. The coefficient on squared monthly ACWI returns is +0.83, with a t-statistic of +2.71, which is statistically significant at the 1% level. This highly positive and significant coefficient, coupled with the visual evidence from Figures 3.20 to 3.22, indicates the presence of convexity in Trend Overlay portfolio excess returns with respect to ACWI.

Table 3.14 Summary of regression results: Trend Overlay portfolio, ACWI and Squared ACWI Returns

Statistic	**ACWI**	**ACWI Squared**	**Intercept**	**R-Squared**
Coefficient	0.98	0.83	0.32%	83%
(t-Stat)	(39.25)***	(2.71)***	(2.54)**	

3.5.4 Historical crisis episode performance comparison

Figure 3.23 presents the same cumulative excess return graph as presented earlier in Figure 3.16 but with a focus on several notable market downturns. Using the crisis periods defined in Harvey et al. (2021) and Schwalbach and Auret (2023), the cumulative excess performance of the Trend Overlay portfolio and ACWI is examined throughout the entire period under investigation, highlighting notable crisis periods ranging from the Asian Crisis in 1998 to the High Inflation era of 2022.

Figure 3.23 Cumulative excess returns of the Trend Overlay portfolio and ACWI with crises highlighted

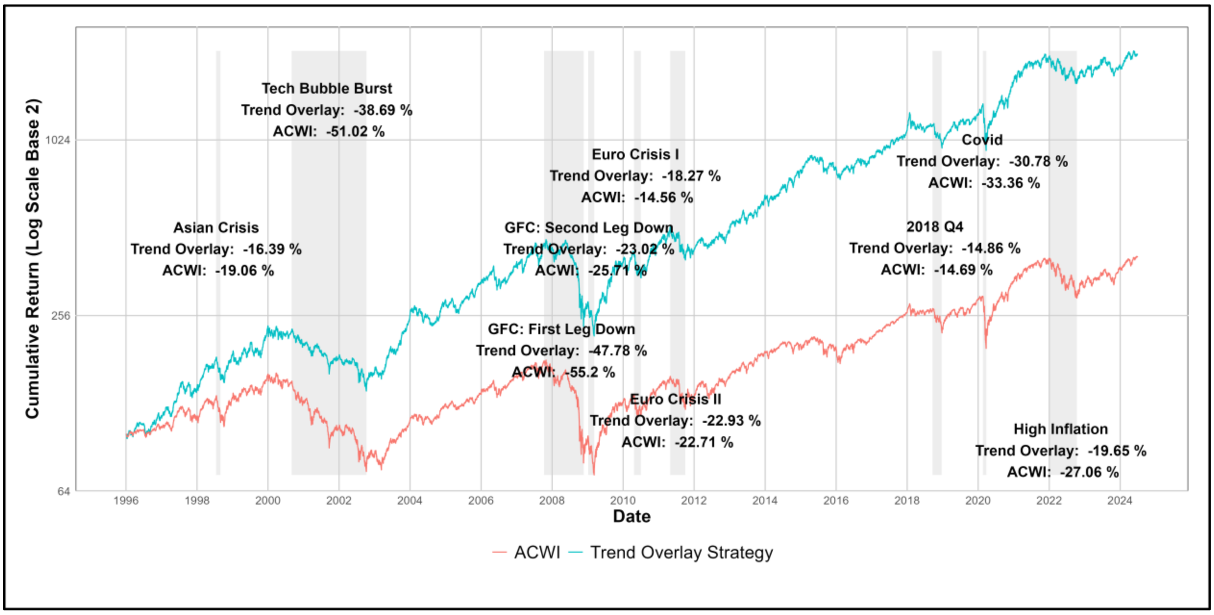

Throughout these episodes, it is observed that while the Trend Overlay portfolio tended to fall when ACWI fell, its drawdowns were mostly smaller than those of ACWI. Consistent with the literature and earlier findings, trend-following has empirically outperformed during prolonged bear markets, and this is reflected in Trend Overlay portfolio returns. For example, during the Tech Bubble burst (lasting 548 weekdays), the Trend Overlay Portfolio fell by 38.69%, a relatively milder loss compared to ACWI's 51.02% drawdown.

Similarly, during the inflationary 2022 episode, the Trend Overlay portfolio declined by 19.65% compared to ACWI's 27.06% drawdown. In Table 3.15, a more granular view of each episode is provided, listing the start and end dates for each crisis, along with the number of weekdays between the peak and trough of each event. Across all nine crises, the Trend Overlay Portfolio experienced losses, underscoring that ACWI beta is still largely responsible for strategy returns. This is consistent with regression findings presented earlier. In longer-lived episodes, the Trend Overlay portfolio tended to fall meaningfully less than ACWI on account of trend-following buffering losses. In shorter-lived episodes, the Trend Overlay portfolio drawdowns tended to be more similar to ACWI.

Table 3.15 Cumulative excess returns of the Trend Overlay portfolio and ACWI during specific crises

Period Name	Peak Day	Trough Day	Number of Weekdays	Trend Overlay portfolio Cumulative Return	ACWI Cumulative Return
Asian Crisis	1998-07-17	1998-08-31	31	-16.39%	-19.06%
Tech Bubble Burst	2000-09-01	2002-10-09	548	-38.69%	-51.02%
GFC: First Leg Down	2007-10-09	2008-11-20	292	-47.78%	-55.20%
GFC: Second Leg Down	2009-01-06	2009-03-09	44	-23.02%	-25.71%
Euro Crisis I	2010-04-23	2010-07-02	50	-18.27%	-14.56%
Euro Crisis II	2011-04-29	2011-10-03	111	-22.93%	-22.71%
2018 Q4	2018-09-20	2018-12-20	65	-14.86%	-14.69%
Covid	2020-02-19	2020-03-23	23	-30.78%	-33.36%
High Inflation	2022-01-03	2022-10-12	202	-19.65%	-27.06%

3.5.5 Relative return comparison

Thus far, a large and significant empirical relationship has been established between ACWI and the Trend Overlay portfolio. To further assess whether the Trend Overlay portfolio is a viable alternative to ACWI for investors focused on relative returns and sensitive to tracking error, the relative return dynamics between the two strategies is explored.

As shown in Table 3.16, the Trend Overlay portfolio outperformed ACWI in approximately two-thirds of observed quarters when considering quarterly excess returns. This consistent outperformance suggests that the strategy has the potential to add value over time for relative return-oriented investors.

It is important to note that in the remaining one-third of the periods, the Trend Overlay underperformed ACWI, which could challenge investor confidence, particularly for investors with shorter time horizons or those more concerned with sustained periods of underperformance. This is especially true during prolonged negative streaks, where the longest consecutive number of quarters of underperformance was seven.

While this finding represents a significant challenge for investors who are sensitive to tracking error, it is worth noting that the longest streak of underperformance was much shorter than the longest streak of outperformance, which spanned 13 quarters. In terms of the magnitude of relative performance, the maximum difference observed was a 10.82% positive spread, which marked the highest quarterly excess return difference between the Trend Overlay portfolio and ACWI. Conversely, the lowest quarterly excess return difference between the Trend Overlay portfolio was -6.59%, less in absolute terms than the highest difference but material, nonetheless.

Table 3.16 Relative return comparison of the Trend Overlay portfolio and ACWI

Statistic	**Value**
Percentage >= 0%	65.50%
Percentage < 0%	34.50%
Maximum Difference	10.82%
Minimum Difference	-6.59%
Longest Positive Streak	13
Longest Negative Streak	7

Figure 3.24 presents the difference between quarterly excess returns of the Trend Overlay portfolio and ACWI over the period under investigation and offers deeper insight into the summary statistics presented in Table 3.16.

A key factor for investors who are tracking-error sensitive is the variability in short-term relative performance. Investors may overreact to periods of short-term underperformance, even when the long-term performance statistics indicate strong evidence of alpha. In Figure 3.24, both periods of strategy outperformance (65.50% of quarters) and underperformance (34.50% of quarters) in relation to ACWI are found. The red dotted line, representing the average relative return, indicates that over the full period examined, the average quarterly relative return is positive, which is in line with the regression findings presented earlier in this chapter.

Figure 3.24 Quarterly excess return difference between the Trend Overlay portfolio and ACWI

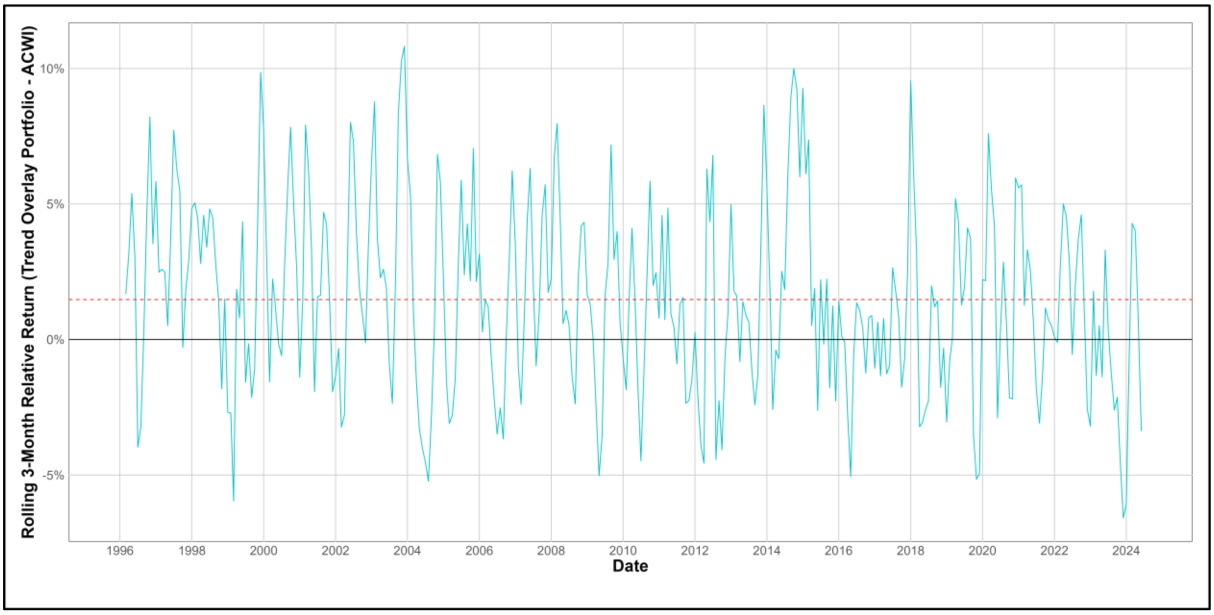

3.5.6 Summary of Trend Overlay portfolio findings

To test hypothesis $H_{0,E}$: overlaying a trend-following strategy on top of global equity does not enhance risk-adjusted and absolute returns compared to holding global equity alone, the Trend Overlay portfolio was constructed. This was done by combining a 100% allocation to ACWI with a 50% allocation to the Trend Portfolio, constructed earlier in this chapter, effectively overlaying a trend-following component on top of the equity portfolio. The performance of the Trend Overlay portfolio was then compared against 100% ACWI.

The results confirm that the Trend Overlay portfolio consistently outperformed ACWI. Specifically, the Trend Overlay portfolio generated a significantly higher CAGR than ACWI, achieving 11.12% compared to 5.05%. Although the strategy exhibited slightly higher volatility, its Sharpe ratio was +0.68, markedly higher than the +0.32 observed for ACWI. Furthermore, overlaying a trend-following strategy introduced convexity to the global equity portfolio, leading to outperformance during extreme market conditions in both upward and downward movements.

Despite the Trend Overlay portfolio employing leverage due to the nature of the embedded trend-following strategy, it experienced shallower drawdowns during key crisis episodes compared to ACWI, particularly during prolonged bear markets. It also exhibited lower tail risk statistics than ACWI, further reinforcing the diversification benefits introduced by trend-following.

Finally, across all regression analyses presented, the Trend Overlay portfolio empirically delivered a statistically significant and positive alpha ranging between 0.41% and 0.50% per month after controlling for traditional risk factors. Therefore, the null hypothesis $H_{0,\ E}$ is rejected in favour of the alternative hypothesis $H_{1,\ E}$: overlaying a trend-following strategy on top of global equity does enhance risk-adjusted and absolute returns compared to holding global equity alone.

3.6 CHAPTER SUMMARY

In this chapter, the impact of overlaying a trend-following strategy on a global equity portfolio represented by ACWI was examined. A trend-following strategy, called the Trend Portfolio, was constructed, utilising an investment universe of 79 futures contracts across four key asset classes: global equities, government bonds, commodities, and currencies.

The futures contracts were transformed into daily US dollar-denominated returns and normalised using ex-ante volatility estimates to standardise risk across and within all asset classes. Employing pooled panel regressions, robust evidence of trend effects across markets was found, leading to the rejection of the null hypothesis $H_{0, A}$: Markets do not trend.

The Trend Portfolio was built using an ensemble of simple long/short trend-following signals, determined daily for each market. Position sizes were scaled based on ex-ante volatility estimates to prevent any single market or asset class from dominating portfolio outcomes.

The resulting Trend Portfolio exhibited key characteristics of classic trend-following strategies documented in a large body of existing research. These findings include positively skewed returns, high absolute returns, and a convex return profile with respect to equity returns. When compared to the SG Trend Index, a trend-following industry benchmark, a significant and positive relationship between the returns of the Trend Portfolio and the SG Trend Index was found, prompting the rejection of the null hypothesis $H_{0, B}$ that SG Trend Index returns do not have a significant relationship with Trend Portfolio returns.

Further investigation was conducted to determine whether traditional equity market factors, namely, the global market (ACWI), size (SMB), value (HML), and momentum (UMD) as well as global government bonds and commodities, could explain the returns of the Trend Portfolio. The regression analysis revealed near-zero and statistically insignificant loadings on these factors, except for the momentum factor, which showed a significant and positive loading.

After adjusting for these factors, the Trend Portfolio still generated a statistically significant monthly alpha of 0.85%, leading to the rejection of the null hypothesis $H_{0, C}$: Trend-following returns are explained by traditional factors.

A key focus of the chapter was the diversification potential of trend-following strategies for equity-dominant portfolios. The convexity of the Trend Portfolio's returns relative to ACWI was analysed across different time horizons, including monthly, quarterly and annually, and the characteristic "smile" was observed in the resulting payoff profiles, indicating that trend-following returns have been empirically strongest during periods of extreme equity market movements, both in the left-tail and in the right-tail of the return distribution.

A statistically significant positive coefficient on squared ACWI returns provided further evidence of convexity, resulting in the rejection of the null hypothesis $H_{0, D}$: Trend-following returns are not convex with respect to global equity returns.

Finally, following the methodology used by Schwalbach and Auret (2023), the performance enhancements from overlaying 50% of the Trend Portfolio on top of a 100% global equity portfolio (ACWI), forming the Trend Overlay portfolio, were investigated. The Trend Overlay portfolio empirically outperformed ACWI, delivering a higher CAGR of 11.12% compared to 5.05%, and a higher Sharpe ratio of +0.68 versus +0.32.

The Trend Overlay portfolio also exhibited meaningfully lower tail risk statistics and less negative skewness while suffering shallower drawdowns during prolonged bear markets. While the risk reduction during short-lived crises was less pronounced, the diversification benefits introduced by incorporating trend-following into an equity portfolio were evident.

The regression analysis revealed a statistically significant monthly alpha of 0.41%, even after controlling for traditional equity market factors, global government bonds, and commodities. These findings led to the rejection of the null hypothesis $H_{0, E}$: Overlaying a trend-following strategy on top of global equity does not enhance risk-adjusted and absolute returns.

This chapter demonstrated the substantial diversification and risk-adjusted return benefits that trend-following strategies can provide when overlaid on global equity portfolios. The empirical evidence suggests that incorporating a trend-following overlay can significantly enhance portfolio performance, both in absolute terms and on a risk-adjusted basis, while also introducing desirable convexity properties that reduce tail risk during periods of market stress.

CHAPTER 4: TAIL RISK HEDGING

4.1 CHAPTER INTRODUCTION

In this chapter, a simple tail risk hedging strategy is introduced and its efficacy in introducing left-tail convexity when overlaid on a global equity portfolio is examined. Building on the foundation established in the previous chapter, ACWI continues to be used to represent global equity exposure. For the tail risk hedging component, exchange-traded options on the SPX) are utilised.

The period under investigation is 4 January 1996 to 28 June 2024, with all results denominated in US dollars. The analysis begins by addressing the null hypothesis $H_{0,F}$: SPX options are not suitable to hedge sudden and severe drawdowns in ACWI exposure. Despite the basis risk introduced by using SPX options to hedge ACWI exposure, it is argued that the significant liquidity of SPX options and the strong correlation between SPX and ACWI (particularly during left-tail events when tail risk hedging is most critical) justify this approach.

Hypothesis $H_{0,G}$: Tail risk hedging does not introduce left-tail convexity with respect to global equity returns when overlaid is then tested. To address this, the impact of compound returns being negatively impacted by volatility through a phenomenon known as variance drag is discussed, explaining how large losses and extreme volatility can erode long-term compounded returns. It is contended that the most critical area in the distribution of returns to hedge is the extreme left tail. By studying how different option specifications have historically provided positive convexity during sharp market crises, it is demonstrated that overlaying such options introduces significant left-tail convexity at the portfolio level with respect to ACWI returns.

Hypothesis $H_{0,H}$: Overlaying a tail risk hedging strategy on top of global equity does not improve tail risk statistics and leads to a reduction in risk-adjusted returns compared to holding global equity alone is evaluated. The long-term empirical risk and return profiles of ACWI and the Tail Overlay portfolio, formed by overlaying the tail risk hedging strategy on a fully invested ACWI position, are analysed. The examination assesses whether the strategy enhances tail risk statistics without compromising risk-adjusted returns over time.

The structure of this chapter is as follows: a detailed description of the data used in this chapter as well as the methodology used in constructing the Tail Overlay portfolio is first provided. The findings are then presented, comparing the historical risk and returns characteristics of the Tail Overlay portfolio and ACWI, which are used to evaluate the hypotheses described above.

4.2 DATA DESCRIPTION

In this section, the various sources of data used to construct the tail-hedging strategy and the resulting tail-hedged overlay strategy created when overlaying the tail risk hedging strategy on top of a global equity portfolio are described. Consistent with the previous chapter, ACWI is used to represent global equity exposure and 3-month US LIBOR is used to represent the risk-free rate (both sourced from Bloomberg), aligning the date range with that of the trend-following strategy. The data are standardised to include all weekdays in a non-leap year, totalling 261 trading days. If no data were available for a particular exchange due to a holiday, a 0% return is assigned to the instruments traded on that exchange for the day.

Exchange-traded SPX options data are used for the tail risk hedging strategy, covering the period from 4 January 1996 to 28 June 2024, with all data denominated in US dollars. The options data, sourced from OptionMetrics, include end-of-day bid and ask prices for each SPX option, along with underlying SPX Index closing prices, total returns, and end-of-day values for delta, gamma, vega, rho, and theta. Mid-prices are calculated by averaging the end-of-day bid and ask prices for each option. For hedging, long-dated, far out-the-money (OTM) put options are used. Options which most closely match the criteria of being 10-delta and approximately one-year to expiration at each roll date are selected. This process is described below.

4.2.1 Motivating the decision to hedge the tail risk of a global equity portfolio using SPX Options

The underlying equity portfolio in this study is represented by the ACWI, encompassing a broad basket of global equities. However, options on the SPX are used for the tail risk hedging strategy, which introduces basis risk since the derivative does not precisely match the underlying exposure. There are three reasons for this choice.

The first reason is due to data limitations. OptionMetrics provides SPX options data starting from 4 January 1996, allowing the analysis of over 28 years of returns, including several key crises during which both the SPX and ACWI experienced sudden and severe drawdowns. Secondly, liquidity is of utmost importance when considering a tail hedging strategy. Krishnan (2017) discusses how liquidity levels in financial markets tend to diminish during crises. Using highly liquid exchange-traded SPX put options is attractive not only because of how sought after they become in the midst of systemic crisis, but also due to the existence of a centralised clearinghouse. This contrasts with OTC derivatives, which involve counterparty risk while offering customised contracts between the hedger and the financial institution counterparty. Such risk becomes a concern during sudden and severe market crises when the financial health of counterparties can become uncertain.

Finally, global equity markets have historically exhibited high correlations between one another during systemic crises. This high correlation reduces the effectiveness of regional diversification and diminishes the significance of basis risk during a systemic shock. US equities, as proxied by SPX, have historically constituted the majority of the ACWI's exposure, which is reflected in historical monthly total returns. In Table 4.1, a linear regression is run between ACWI monthly returns (dependent variable) and corresponding SPX monthly returns (independent variable). The resulting coefficient is close to one and statistically significant at the 1% level, suggesting a strong, nearly one-to-one relationship between the monthly returns of SPX and ACWI. However, the intercept of -0.16% per month, statistically significant at the 5% level, indicates that SPX has historically generated alpha over ACWI. It is found that 91% of the variability in ACWI monthly returns can be explained by SPX monthly returns.

Table 4.1 Summary of regression results: monthly total returns of ACWI (dependent variable) vs. SPX (independent variable)

Statistic	**SPX**	**Intercept**	**R-Squared**
Coefficient	0.96	-0.16%	91%
(t-Stat)	(59.54)***	(-2.23)**	

It is found that changes in the realised volatility of SPX are closely related to changes in the realised volatility of ACWI. In Table 4.2, the regression output is summarised where monthly changes in realised trailing 12-month SPX volatility are treated as the independent variable, while monthly changes in realised trailing 12-month ACWI volatility serve as the dependent variable. A large and positive coefficient of +0.75 that is statistically significant at the 1% level is found. The R-squared value implies that changes in SPX's realised volatility can explain 89% of the variation in ACWI's realised volatility. The strong linkage between the volatilities of the SPX and ACWI can be attributed to the interconnectedness of global financial markets.

Table 4.2 Summary of regression results: monthly changes in realised volatility for ACWI (dependent variable) vs. SPX (independent variable)

Statistic	Change In SPX Volatility	Intercept	R-Squared
Coefficient	0.75	-0.00%	89%
(t-Stat)	(53.98)***	(-0.10)	

In summary, the decision to use SPX options to hedge an underlying ACWI portfolio is empirically supported by Tables 4.1 and 4.2, which demonstrate a strong and significant relationship between the return distributions of SPX and ACWI.

4.2.2 Motivating the decision to use far-out-the-money put options

Having settled on the use of SPX options in this research, the decision to use far out-the-money (OTM) put options as the instrument to implement the tail risk hedging strategy used in this study is now justified. The primary rationale stems from the disproportionate impact that large, discrete negative returns have on long-term compounded performance. This is illustrated in Table 4.3, which compares the CAGR of ACWI to those when the best and worst 10 months are excluded from the calculation. Missing the top 10 months of ACWI performance would have reduced buy-and-hold ACWI performance by 3.59% per annum.

This finding is often cited within the investment industry to discourage market timing, emphasising the risk of missing out on significant positive returns by attempting to time the market. However, what is seldom discussed is the opposite perspective: excluding the worst 10 months of ACWI performance would have improved performance by 4.88% per annum relative to buy-and-hold ACWI, which is larger in absolute terms and is illustrative of the greater consequence of severe negative returns on compounded returns.

Table 4.3 CAGR (total returns) of ACWI and ACWI excluding the top 10 best and worst months from January 1996 to June 2024

Strategy	CAGR
ACWI	7.81%
ACWI Excluding Bottom 10 Months	12.69%
ACWI Excluding Top 10 Months	4.22%

The contrasting effects observed in Table 4.3 can be attributed to the phenomenon known as variance drag. Specifically, Eifert (2020), Bouchey et al. (2012) and Stein et al. (2009) show that to approximate the CAGR of a time series, one can reduce annualised arithmetic returns by half of the annualised variance of those returns. It is for this reason that Alquist et al. (2020) describe variance drag as the tendency for geometric returns to be less than arithmetic returns due to the impact of volatility.

This disparity arises because geometric returns account for the compounding effect of portfolio growth over time, while arithmetic returns represent a simple average that does not consider the sequencing of returns. To illustrate how variance drag impacts compounded returns, consider the stylised example presented in Figure 4.1. If a portfolio falls by 10% in one period, it requires an 11.11% gain in the subsequent period to break even (achieving a geometric return of 0% over both periods). However, a portfolio that falls by 50% in one period requires a 100% subsequent gain to break even. As shown in Figure 4.1, this relationship generalises into an exponential curve, demonstrating that large losses disproportionately affect long-term compounded returns because they demand disproportionately large gains for recovery. This effect is a direct result of the mathematical nature of compounding, where significant losses shrink the base from which future returns must grow.

Figure 4.1 Stylised example of the gain required to breakeven following losses of varying magnitudes

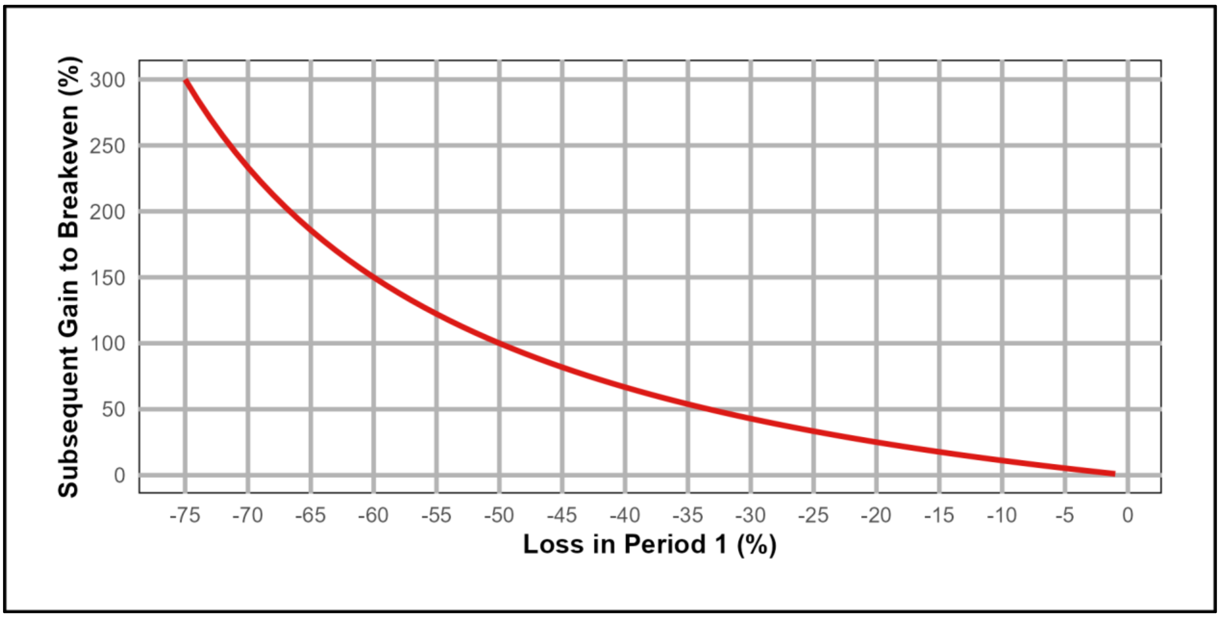

To further demonstrate the impact of variance drag on long-term returns, this chapter presents a stylised example of the relationship between compound returns and volatility. Figure 4.2 illustrates the resulting CAGR of a portfolio with an average (arithmetic) annual return of 10% as a function of portfolio volatility. In Figure 4.2, it is observed that increasing volatility tends to reduce long-term compounded returns at an exponential rate. At sufficiently high levels of volatility, even a portfolio that consistently delivers an average return of 10% can result in negative compounded return over time, emphasising the impact of variance drag.

Figure 4.2 Stylised example of the relationship between CAGR and a portfolio with 10% average returns as a function of annualised volatility

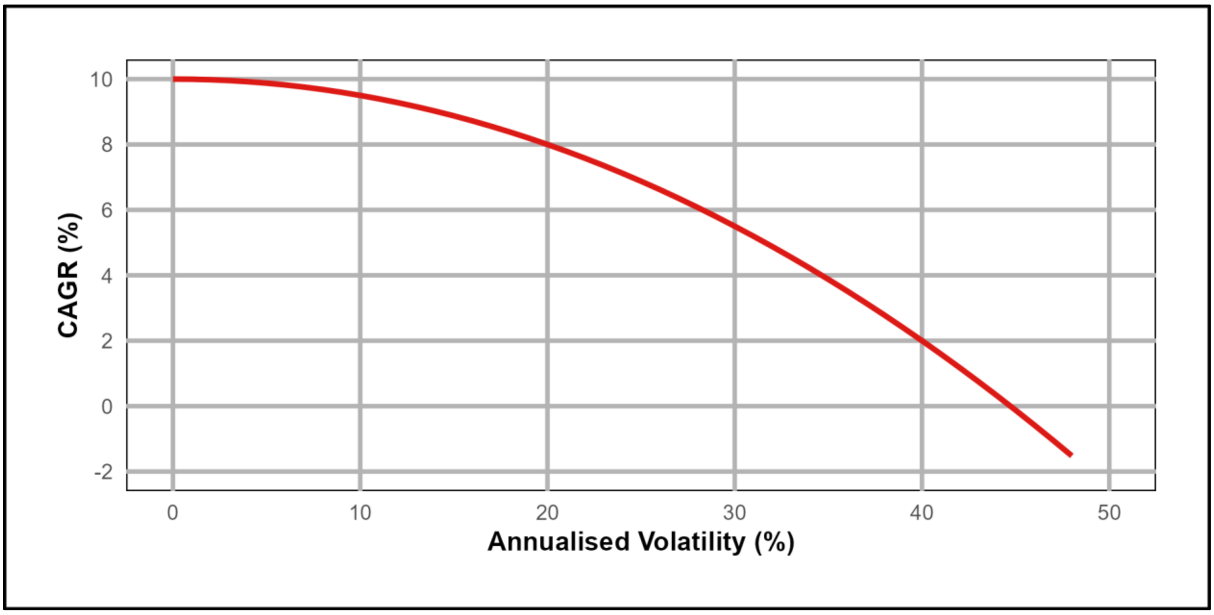

Table 4.3 and Figures 4.1 and 4.2 support the notion that the impact of variance drag is most noticeable at high levels of volatility and large levels of portfolio losses (where the exponential effects become significantly pronounced), while it has a comparatively modest effect at lower levels of volatility and smaller losses. Given that the systematic purchasing of put options can drag on portfolio returns, particularly in the absence of market crises, investors should aim to balance the conflicting objectives of minimising premium costs over time while maximising potential hedging gains during crises. Table 4.3 and Figures 4.1 and 4.2 suggest that focusing on mitigating extreme events offers the biggest 'bang-for-the-buck' in reducing the most harmful effects of variance drag.

To better understand how different option prices have responded to fluctuations in market volatility over time, the study next examines how changes in option implied volatilities, and therefore, option prices, have been related to changes in the realised volatility of the underlying market using the SPX. It is found that option market participants have tended to react to prevailing market conditions by increasing option implied volatilities (and thus option prices) in response to rising levels of realised volatility in underlying price movements and vice versa. This is notable because volatility tends to spike due to prices falling in the underlying market. Table 4.4 provides a summary of standalone regression results comparing monthly changes in implied volatility (the dependent variable) for SPX put options across different constant deltas and days to expiration (DTE), with monthly changes in realised 12-month SPX volatility (the independent variable).

For each combination of DTE and delta, separate regressions are run to examine how sensitive changes in implied volatility (and thus SPX option prices) are to changes in realised volatility. Isolating for changes in implied volatilities for combinations of options with terms of 1-month, 3-month and 12-months that are approximately 50-delta (to proxy for near-the-money puts), 25-delta (to proxy for out-the-money puts) and 10-delta (to proxy for far out-the-money puts) allows us to offer insight into the relationship between changes in the realised volatility of SPX and matching-period changes in the SPX option volatility surface.

Across the surface, a statistically significant (at the 1% level) and positive relationship is found between changes in the realised volatility of SPX and changes in the implied volatility of the SPX option surface. Closer-to-expiration options tended to have the highest coefficients, indicating that while changes in implied volatility still responded positively to changes in realised volatility, the impact was lower for options with longer tenors. These options tended to be less sensitive to short-term market fluctuations.

Table 4.4 Summary of standalone regression results: monthly changes in implied volatility (dependent variable) for constant delta and tenor SPX put options vs. monthly changes in realised SPX return volatility (independent variable)

DTE	Delta	Coefficient	t-Stat
30	-50	0.29	(10.92)***
30	-25	0.34	(10.95)***
30	-10	0.37	(10.72)***
91	-50	0.23	(12.42)***
91	-25	0.27	(12.22)***
91	-10	0.29	(11.49)***
365	-50	0.14	(11.62)***
365	-25	0.15	(11.45)***
365	-10	0.16	(11.29)***

Table 4.5 extends this analysis by examining the squared monthly changes in realised volatility as the independent variable against monthly changes in implied volatility as the dependent variable. This allows for the magnification of the importance of spikes in realised volatility, which are associated with periods of market stress. It is found that across all tenors, far-OTM, proxied by 10-delta put options, exhibit the strongest response to squared changes in realised volatility, suggesting that far-OTM options are most convex with respect to large spikes in equity volatility. In addition, the short-term tenor had lower coefficients than the medium-term and long-term tenors, suggesting that longer-term tenors are more convex with respect to large changes in realised volatility.

This is notable because, as has been discussed, variance drag is most pernicious at extreme levels of volatility. Krishnan (2017), Eifert (2020), and Bhansali (2015b) also demonstrate that far-OTM options with low deltas offer an attractive balance between positive convexity during market crashes and reduced theta decay during normal market conditions. For these reasons, one-year to expiration, 10-delta SPX put options are selected for the tail risk hedging strategy.

A useful by-product of the choice of option type is that it impacts effective market participation less than closer-to-the-money options. A deliberate approach is taken for tail risk hedging, avoiding interventions that could potentially improve strategy performance in order to underscore the robustness of the strategy. These interventions include early monetisation rules, discussed in Bhansali (2010b), and delta-hedging, which is typically employed to offset the negative delta exposure inherent in a long-put option position, as discussed in Eifert (2020).

Regarding delta-hedging, when implied volatility is low, an at-the-money (ATM) option has a delta of approximately -0.50 for a put, often referred to as "50-delta". When attaching a 50-delta put, notionally matched, to an equity index portfolio with the same underlying index as the option, the initial exposure to the equity market is effectively reduced to 50%. This stems from the full exposure to the underlying index being offset by the -50% delta of the put.

Although the effective exposure to the underlying market fluctuates as conditions change, holding puts can impede equity market participation at a portfolio level. A key advantage of using 10-delta puts for tail risk hedging relative to closer-to-the-money deltas (when not delta-hedging) is that they impact effective exposure to the underlying market less, while still offering meaningful hedging potential in extreme events. The choice of delta is of particular importance due to the SPX empirically outperforming ACWI, as shown in Table 4.1.

The asterisks presented next to the t-statistics in Table 4.5 denote statistical significance levels. One asterisk indicates significance at the 10% level, two at the 5% level, and three at the 1% level.

Table 4.5 Summary of standalone regression results: monthly changes in implied volatility (dependent variable) for constant delta and tenor SPX put options vs. squared monthly changes in realised SPX realised volatility (independent variable)

DTE	Delta	Coefficient	t-Stat
30	-50	0.12	(1.47)
30	-25	0.15	(1.62)
30	-10	0.19	(1.90)*
91	-50	0.20	(3.54)***
91	-25	0.26	(3.81)***
91	-10	0.30	(4.08)***
365	-50	0.18	(5.09)***
365	-25	0.22	(5.70)***
365	-10	0.26	(6.09)***

4.3 METHODOLOGY

The tail risk hedging strategy used in this study is built to overlay on top of a 100% long buy-and-hold ACWI portfolio, creating the Tail Overlay portfolio. A similar approach as in Schwalbach and Auret (2023) is used. The hedge is constructed by systematically purchasing 10-delta put options on the SPX with one-year to expiration, executed in three separate tranches. Each tranche is rolled quarterly and notionally sized to represent one-third of the fully invested equity position. The use of three tranches, rather than a single position, is employed to reduce rebalance timing risk, following Hoffstein et al. (2019). This method helps smooth the hedge performance by avoiding large, abrupt adjustments to the entire hedge position at a single point in time. Option rebalancing occurs monthly, coinciding with the standard monthly SPX options expiration cycle, which is typically the third Friday of each month.

At each roll date, the tranche closest to expiration is rolled into new 10-delta put options with a fresh one-year-to-expiration horizon. At each monthly roll date, the portfolio holds one tranche with 12 months to expiration, another with 11 months, and a final tranche with 10 months to expiration. To maintain simplicity in the tail risk hedging strategy, any further interventions, such as monetisation, tactical premium size adjustments, more complex multi-legged option structures and delta-hedging are deliberately avoided.

Options are not held to expiration. Each option is rolled using a staggered approach every three months as the intention is for them to provide a significant price appreciation during periods of market distress that then gets partly realised at roll dates and reinvested in equity at the lower prevailing prices. This approach is the same as in Schwalbach and Auret (2023). To illustrate this concept by way of example, on 19 December 2019, one-third of the puts were rolled into one-year, 10-delta options that cost $39.25. Approximately three months later, on 19 March 2020, in the midst of crisis, these puts (now with approximately nine months to expiration) were sold for $364.25 and rolled into new one-year, 10-delta puts, with some of the profit being rebalanced into underlying equity exposure.

Similar to Eifert (2020), the cost of purchasing the options is financed through a short position in 3-month US LIBOR, which is a proxy for the risk-free rate. This financing mechanism allows the strategy to remain fully invested in the underlying equity index. The daily returns of the Tail Overlay portfolio reflect the contributions of the ACWI equity exposure, the performance of the 3-tranches of puts and the interest charges from funding.

Long puts provide the holder with the right, but not the obligation, to sell an underlying security at a defined price at (or before, in the case of an American-style put option) the expiration date. The payoff profile for a long put is illustrated by Figure 4.3 below. K1 indicates the strike price, the depth of the payoff line that is in loss represents the option premium, which is the maximum loss. The breakeven price is the strike price minus the premium paid. When enough long puts are purchased to cover the notional exposure of a position in the underlying security, the resulting combination is called a protective put. The attractiveness of a protective put is that the long-put option places a clear floor on the maximum potential loss represented by the degree to which the option is OTM and the premium.

Figure 4.3 Payoff profile of a long-put option

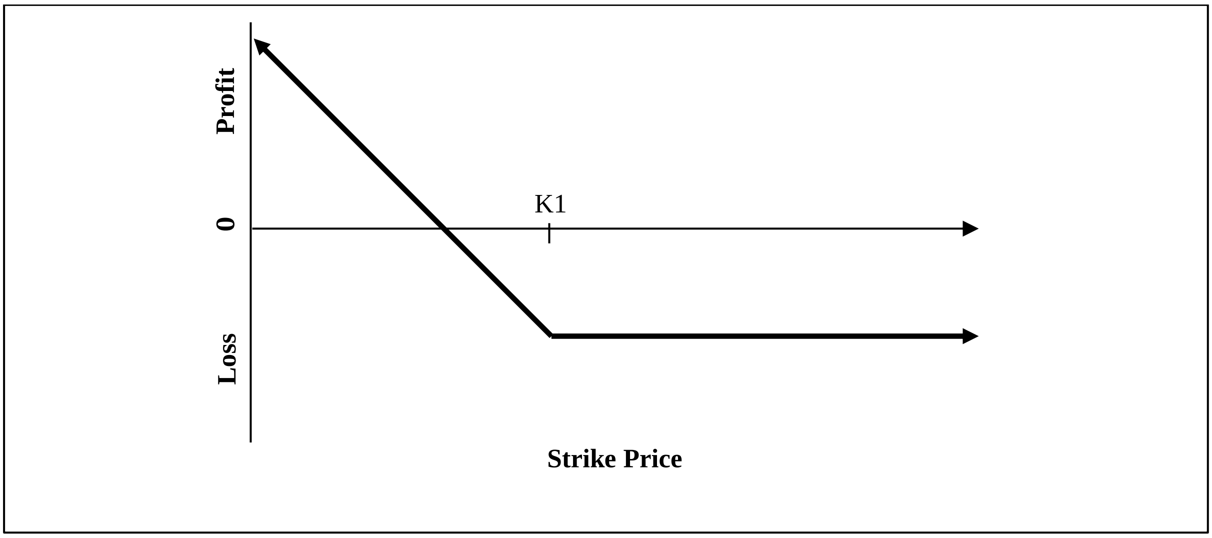

4.4 FINDINGS

4.4.1 Excess performance comparison

The results presented in this section compare the excess returns of ACWI to those of the Tail Overlay portfolio, which is comprised of 100% ACWI overlaid with the tail risk hedging strategy from 4 January 1996 to 28 June 2024. In Figure 4.4, the cumulative excess return of the Tail Overlay portfolio is illustrated and compared against ACWI. In contrast to the Trend Overlay portfolio discussed in the previous chapter, the cumulative returns of the Tail Overlay portfolio were lower than ACWI over the sample period in absolute terms.

Table 4.6 shows that the CAGR of the Tail Overlay portfolio was 4.11% over the period compared to 5.05% for ACWI. This indicates that the Tail Overlay portfolio captured approximately 80% of ACWI's long-term returns. Systematically holding puts that are not delta-hedged reduces the effective equity exposure of the portfolio due to the negative delta exposure embedded within the puts. A meaningful portion of the remaining underperformance can be attributed to the well-documented tendency for implied volatility to exceed realised volatility in the underlying asset, as demonstrated in Schwalbach and McClelland (2018). This discrepancy results in hedgers paying a volatility risk premium on average over time to truncate downside risk.

Figure 4.4 Cumulative excess returns of the Tail Overlay portfolio and ACWI

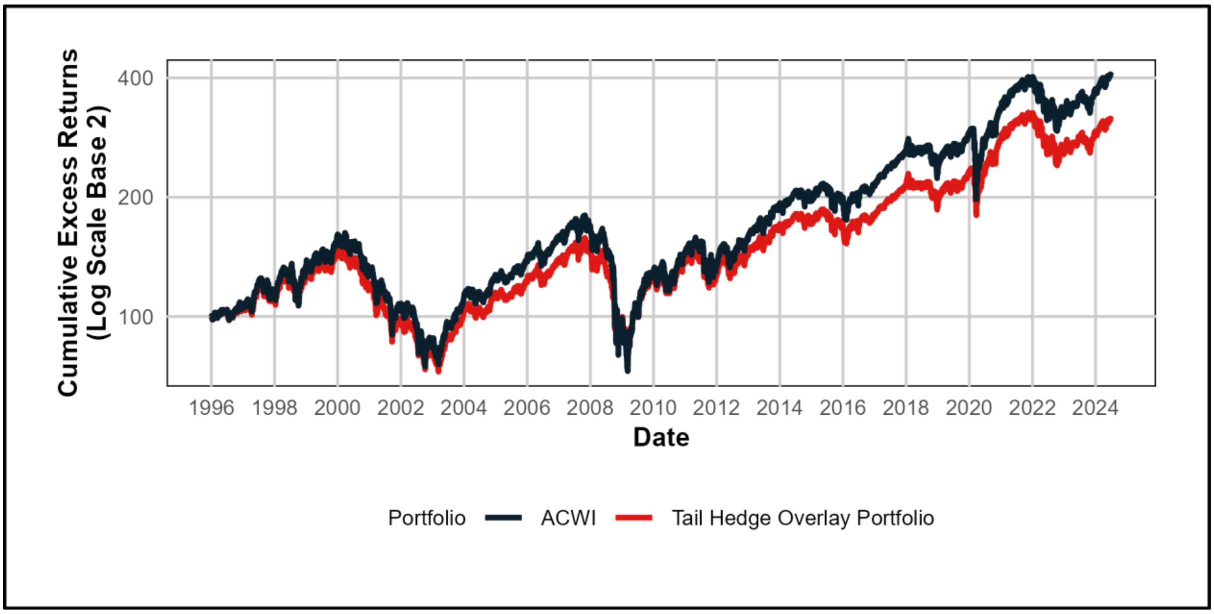

The benefits brought about by the put options can be seen in Figure 4.5, which illustrates the peak-to-trough drawdowns of both strategies over the period under investigation. Visually, the Tail Overlay portfolio experienced shallower drawdowns compared to ACWI, particularly during episodes of sudden and severe market stress. The most notable improvements in drawdowns occurred during the 2008 GFC and the 2020 Covid Crisis. Following these episodes, the Tail Overlay portfolio exhibited faster recoveries than ACWI, owing to a reduction in variance drag, which helped mitigate the compounding impact of losses over time.

These events, namely the 2008 GFC and the 2020 Covid Crisis, are widely considered as crashes, characterised by sudden and severe market declines. In contrast, more prolonged downturns, typically termed bear markets, such as the Dot-Com Bubble Burst in the early 2000s and the inflationary episodes of 2022, presented less favourable conditions for tail risk hedging. During these extended drawdowns, the effectiveness of the Tail Overlay portfolio was more muted, as the slow, grinding nature of these declines did not trigger the kind of sharp volatility spikes that benefit the systematic use of put options.

Figure 4.5 Drawdowns of Tail Overlay portfolio and ACWI

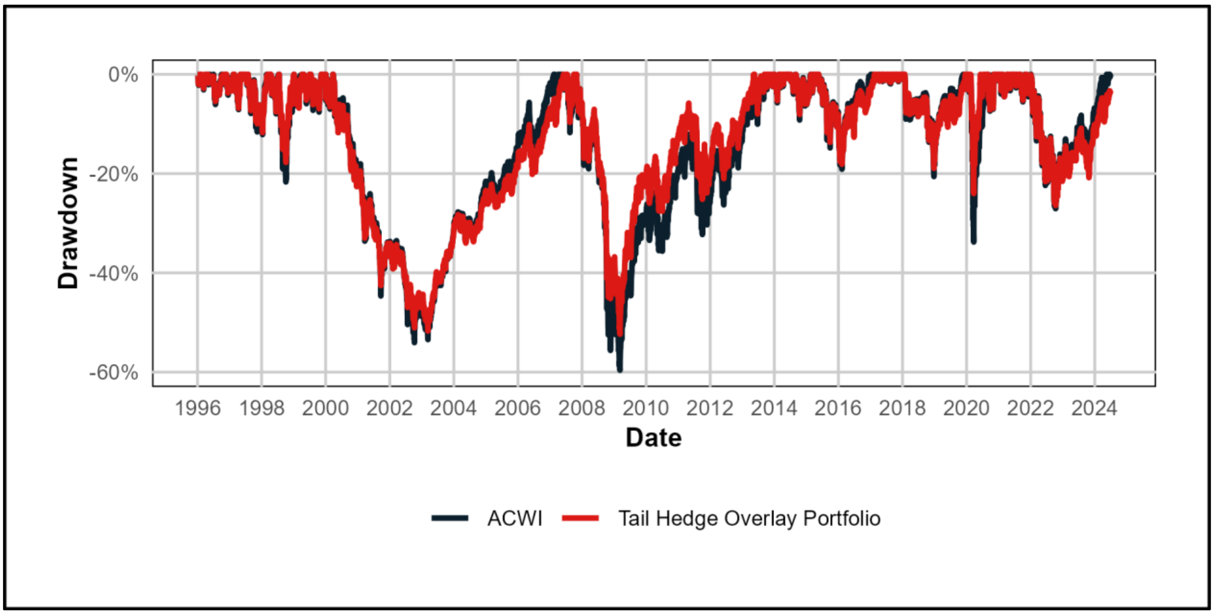

The tail risk hedging strategy performed best during periods characterised by substantial increases in realised volatility, which, in turn, drove implied volatility higher, as evidenced by Schwalbach and Auret (2023). In contrast, drawdown periods accompanied by moderate but not extreme elevations in implied volatility resulted in suboptimal hedging outcomes. For example, while the ACWI drawdown during the Dot-Com Bubble Burst was notably deeper than that experienced during the Covid Crash, the Tail Overlay portfolio demonstrated a greater degree of relative outperformance during the short-lived Covid Crash.

This highlights that implied volatility surfaces tend to react differently to crises of varying durations. By way of example, consider Figure 4.6, which contrasts constant one-year to expiration implied volatilities for 10-delta puts observed daily over the course of the Dot-Com Bubble Burst (1 September 2000 to 9 October 2002) and the more abrupt Covid Crash (19 February 2020 to 23 March 2020). Although the drawdown during the Dot-Com Bubble was more severe and prolonged, the implied volatility for put options saw a much more modest rise over the course of the event. The reason for this is that option markets tend to react to rising levels of realised volatility by adjusting implied volatilities (and thus option prices) in response to rising levels of uncertainty.

When drawdowns are suffered over prolonged periods, realised volatility may only rise modestly relative to when drawdowns of similar magnitudes are suffered abruptly. This is seen illustrated empirically in Figure 4.6. Figure 4.6 demonstrates how implied volatility spiked sharply during the relatively brief Covid Crash, peaking at levels above 60%, despite the underlying market falling by just over 30%.

Implied volatility levels remained elevated for several months during the Dot-Com Bubble episode, leading to sustained high premium costs being incurred throughout the episode. Elevated premium costs contributed to the tail risk hedging strategy underperforming relative to the Covid Crash episode, where implied volatility spiked sharply but briefly. Notably, the same Dot-Com Bubble period, during which the tail risk hedging strategy underperformed, was a time when the trend-following strategy discussed in the previous chapter generated significantly positive performance.

In contrast, during the shorter and more acute Covid Crisis, the tail risk hedging strategy produced outsized returns due to the pronounced volatility spike, while the trend-following strategy delivered comparatively weaker results. This dichotomy in performance across two of the most severe ACWI drawdowns indicates that tail risk hedging and trend-following may offer superior diversification when combined. The differing performance profiles suggest that overlaying these strategies in combination could introduce greater convexity into an ACWI portfolio than either would achieve on their own.

Figure 4.6 Implied volatilities of constant one-year to expiration, 10-Delta SPX Index put options

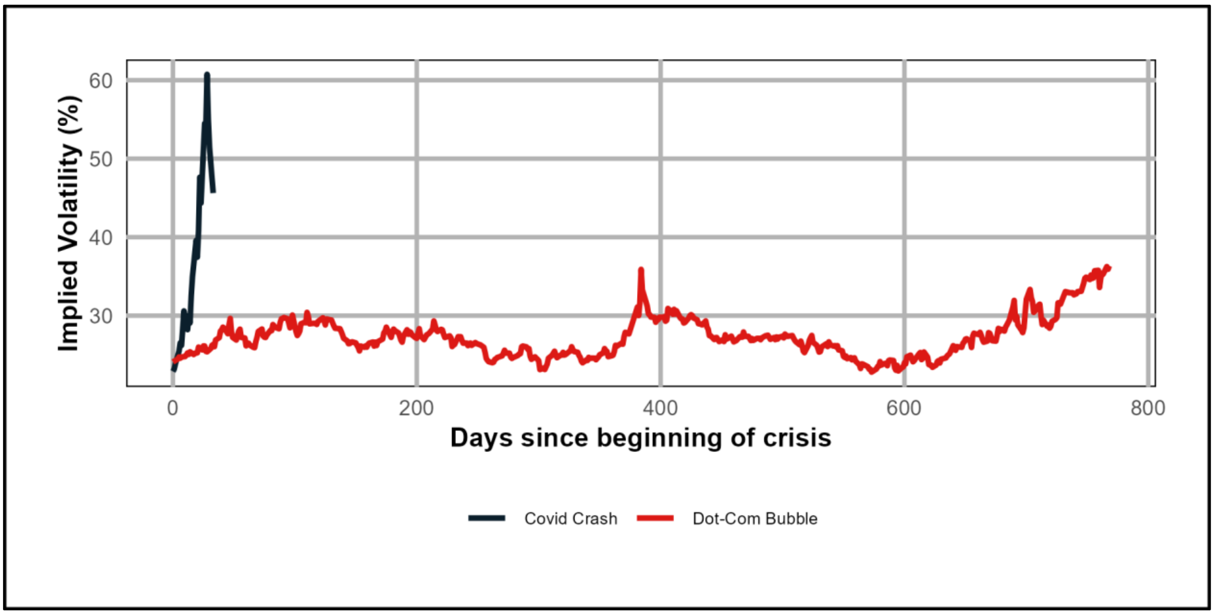

Continuing the comparison of the risk profiles of ACWI and the Tail Overlay portfolio, the evidence suggests that the Tail Overlay portfolio generated lower volatility than ACWI (12.72% vs. 15.59%, respectively). The Tail Overlay portfolio also produced considerably less negative skewness (-0.43 vs. -0.65). This finding is in line with Bhansali (2015b), who noted that from a behavioural perspective, some investors are willing to accept negative returns from tail hedging in standalone terms because they can improve the overall skewness of a risky asset portfolio.

Further improvements in tail risk could be seen by comparing the Tail Overlay portfolio with ACWI in terms of VaR at both the 95% and 99% confidence levels (-6.91% vs. -8.16% and -9.15% vs. -12.04%, respectively). Similarly, the CVaR for the Tail Overlay was lower, suggesting better downside protection in extreme market scenarios (-8.65% vs. -10.44% at the 95% confidence level, and -11.25% vs. -15.18% at the 99% confidence level). While ACWI had a better 3-month best performance result than the Tail Overlay portfolio (33.06% vs. 29.05%, respectively), the Tail Overlay portfolio experienced a meaningfully less severe drawdown in its worst 3-month period (-25.27% vs. -34.86%, respectively), further demonstrating its effectiveness in cushioning extreme market declines. Despite the lower excess compounded return over time, the Tail Overlay portfolio exhibited a comparable Sharpe ratio to ACWI, with significantly superior tail risk statistics.

Table 4.6 Summary statistics: Tail Overlay portfolio and ACWI

Statistic	Tail Overlay portfolio	ACWI
CAGR	4.11%	5.05%
Annualised Volatility	12.72%	15.59%
Sharpe Ratio	0.32	0.32
Skewness	-0.43	-0.65
Excess Kurtosis	-2.50	-1.54
VaR (95%)	-6.91%	-8.16%
VaR (99%)	-9.15%	-12.04%
CVaR (95%)	-8.65%	-10.44%
CVaR (99%)	-11.25%	-15.18%
Best 3M Performance	29.05%	33.06%
Worst 3M Performance	-25.27%	-34.86%

Figure 4.7 illustrates the rolling one-year correlation between the Tail Overlay portfolio and ACWI. Overall, the correlations remain high over time, fluctuating between 93% and 100%. Compared to the Trend Overlay portfolio discussed in the previous chapter, the Tail Overlay portfolio demonstrates a higher average correlation with ACWI. However, during periods of market turmoil, such as the 2008 GFC and the 2020 Covid Crisis, the correlations tended to drop. This suggests that while the Tail Overlay portfolio closely tracks ACWI under normal market conditions, it diverges during extreme market events, reflecting its role in mitigating downside risk during periods of heightened volatility.

Figure 4.7 Rolling one-year correlations of the Tail Overlay and ACWI

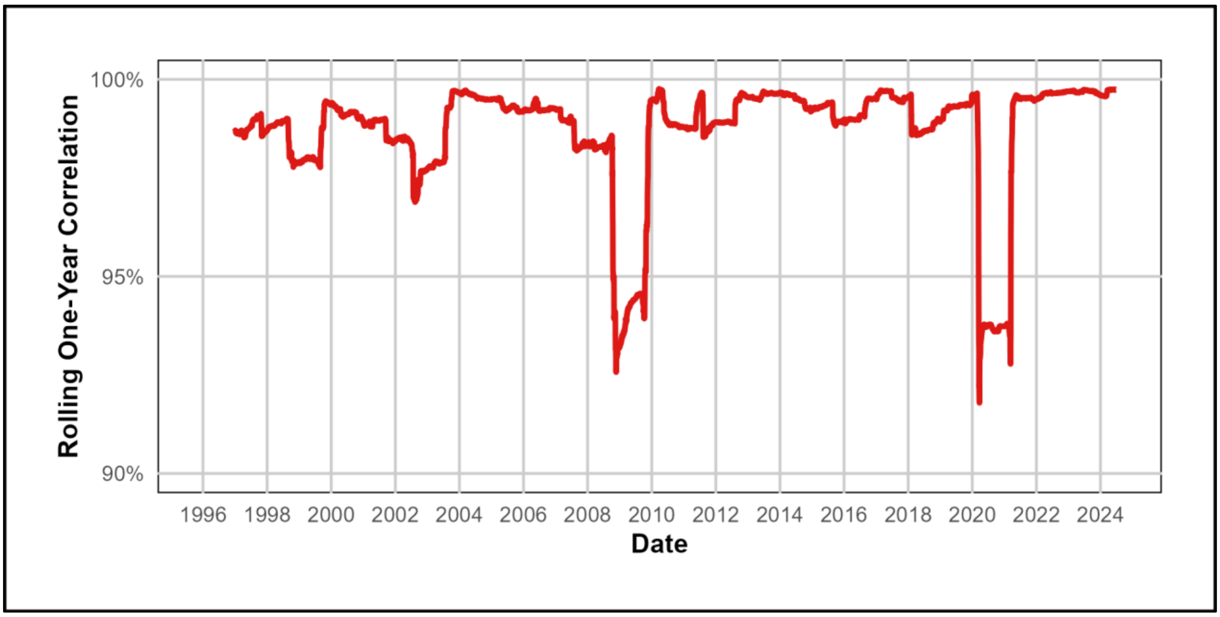

4.4.2 Alpha and loading on risk factors in Tail Overlay portfolio returns

In Figure 4.8, a plot of the monthly excess returns of ACWI on the x-axis against the matching monthly excess returns of the Tail Overlay portfolio on the y-axis is presented. A linear regression is fitted, with the results shown in Table 4.7. The regression yields a coefficient of +0.85, indicating a strong positive relationship between the Tail Overlay portfolio and ACWI's monthly excess returns, statistically significant at the 1% level. The Tail Overlay portfolio produces an insignificant intercept of -0.03%, with a t-statistic of -0.94, suggesting no significant monthly alpha generation beyond its exposure to ACWI. This is supported by the high R-squared result of 98%, indicating that ACWI returns explain nearly all of the variation in the Tail Overlay portfolio's returns. This suggests that the portfolio has closely tracked the performance of global equity over time while improving upon its tail risk hedging characteristics.

Figure 4.8 Linear regression analysis of excess monthly returns: Tail Overlay portfolio vs. ACWI

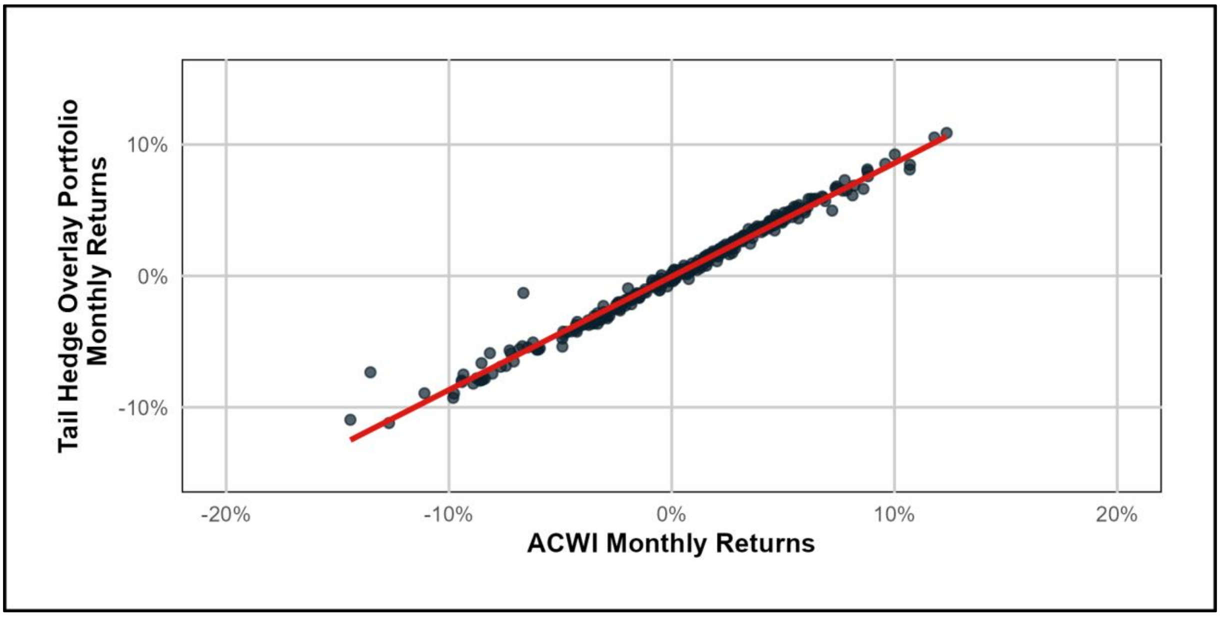

Table 4.7 Summary of regression results: Tail Overlay portfolio and ACWI

Statistic	**ACWI**	**Intercept**	**R-Squared**
Coefficient	0.85	-0.03%	98%
(t-Stat)	(135.20)***	(-0.94)	

Table 4.8 presents the regression output for the Tail Overlay portfolio's monthly excess returns against traditional equity market factors, specifically the global market factor (ACWI), size (SMB), value (HML), and momentum (UMD). The coefficient for ACWI is +0.86, indicating a strong positive relationship with global equity, which is statistically significant at the 1% level. This high beta to the market factor suggests that the Tail Overlay portfolio captures a substantial portion of the equity market's movements, consistent with its construction as an overlay on a fully invested ACWI portfolio.

Unlike the Trend Overlay portfolio discussed in the previous chapter, the Tail Overlay portfolio does not exhibit a positive loading on the momentum factor (UMD). The coefficient for UMD is effectively zero, with an insignificant t-statistic of -0.30, indicating no meaningful exposure to momentum-driven returns. However, the Tail Overlay portfolio does exhibit a statistically significant, albeit near-zero, loading on the value factor (HML), with a small positive coefficient. While the magnitude of this coefficient is minimal, its statistical significance suggests a slight tilt towards value stocks within the strategy's returns. The R-squared value of 98% indicates that the regression model explains nearly all the variation in the Tail Overlay portfolio's returns.

Table 4.8 Summary of regression results: Tail Overlay portfolio against traditional market factors

Statistic	ACWI	SMB	HML	UMD	Intercept	R-Squared
Coefficient	0.86	-0.01	-0.02	0.00	-0.03%	98%
(t-Stat)	(124.33)***	(-1.49)	(-2.42)**	(0.30)	(-0.95)	

Table 4.9 extends the analysis by incorporating global government bonds and commodities, proxied by the Bloomberg Global Aggregate Total Return Index and the S&P GSCI Index, respectively. The inclusion of these additional factors reveals statistically significant coefficients for the value factor (HML), global government bonds, and commodities. However, these coefficients are near zero, indicating that the Tail Overlay portfolio's performance is still driven primarily by its exposure to global equities. The findings presented so far in this chapter, along with those in Tables 4.1 and 4.2 of the Data Description section, support the rejection of hypothesis $H_{0,\,F}$**:** SPX options are not suitable for hedging sudden and severe drawdowns in ACWI exposure, in favour of the alternative hypothesis $H_{1,\,F}$**:** SPX options are suitable for hedging sudden and severe drawdowns in ACWI exposure.

Table 4.9 Summary of regression results: Tail Overlay portfolio against traditional market factors, global government bonds and commodities

Statistic	ACWI	SMB	HML	UMD	BONDS	GSCI	Intercept	R-Squared
Coefficient	0.85	-0.01	-0.02	0.00	0.08	-0.01	-0.03%	98%
(t-Stat)	(112.17)***	(-0.87)	(-2.07)**	(0.46)	(4.62)***	(-1.99)**	(-1.08)	

4.4.3 Tail Overlay portfolio Convexity with respect to Equity Portfolios

In this section, the null hypothesis $H_{0,\,G}$: Tail risk hedging does not introduce left-tail convexity with respect to global equity returns when overlaid is evaluated.

4.4.3.1 Scatterplots illustrating Tail Overlay portfolio convexity with respect to ACWI

While overlaying a trend-following strategy on top of ACWI (discussed in the previous chapter) has empirically introduced convexity on both the left and right tails of the equity return distribution, overlaying a tail risk hedging strategy on top of ACWI tends to only introduce left-tail convexity with respect to equity returns. This is illustrated in Figures 4.9 to 4.11. These figures plot the monthly, quarterly, and annual ACWI excess returns on the x-axis against the corresponding Tail Overlay portfolio returns on the y-axis. For each frequency, a quadratic regression is fitted to the observations.

It is found that the Tail Overlay portfolio exhibits stronger convexity on the left tail than the Trend Overlay portfolio, providing more robust downside protection during extreme market declines. This finding is consistent with the Tail Overlay portfolio's superior tail risk statistics compared to the Trend Overlay portfolio, including lower volatility, VaR, and CVaR metrics. However, the Tail Overlay portfolio does not display the same right-tail enhancement observed in the Trend Overlay, as its primary function is risk mitigation rather than return enhancement.

Convexity is more noticeable for monthly excess returns than for annual excess returns. This observation aligns with our earlier finding that the tail risk hedging strategy has empirically delivered outsized returns during sudden and severe market crises in which volatility spiked. The empirical evidence presented supports the rejection of the null hypothesis $H_{0, G}$ in favour of the alternative hypothesis $H_{1, G}$: Tail risk hedging introduces left-tail convexity with respect to global equity returns when overlaid.

Figure 4.9 Convexity of Tail Overlay portfolio monthly excess returns relative to ACWI monthly excess returns

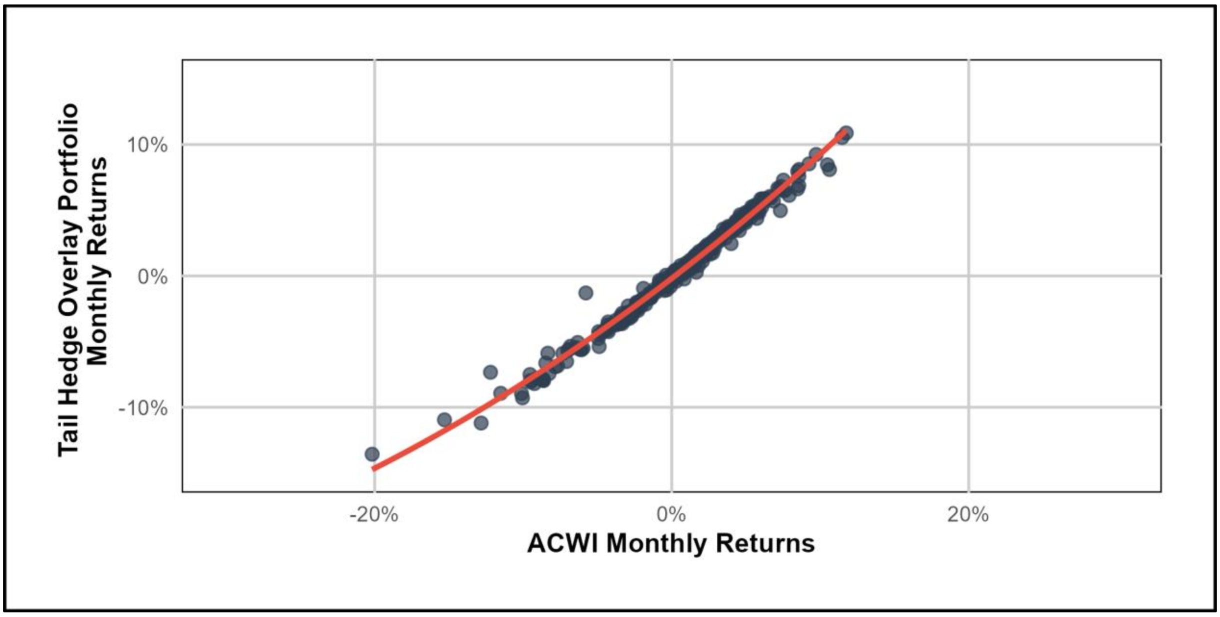

Figure 4.10 Convexity of Tail Overlay portfolio quarterly excess returns relative to ACWI quarterly excess returns

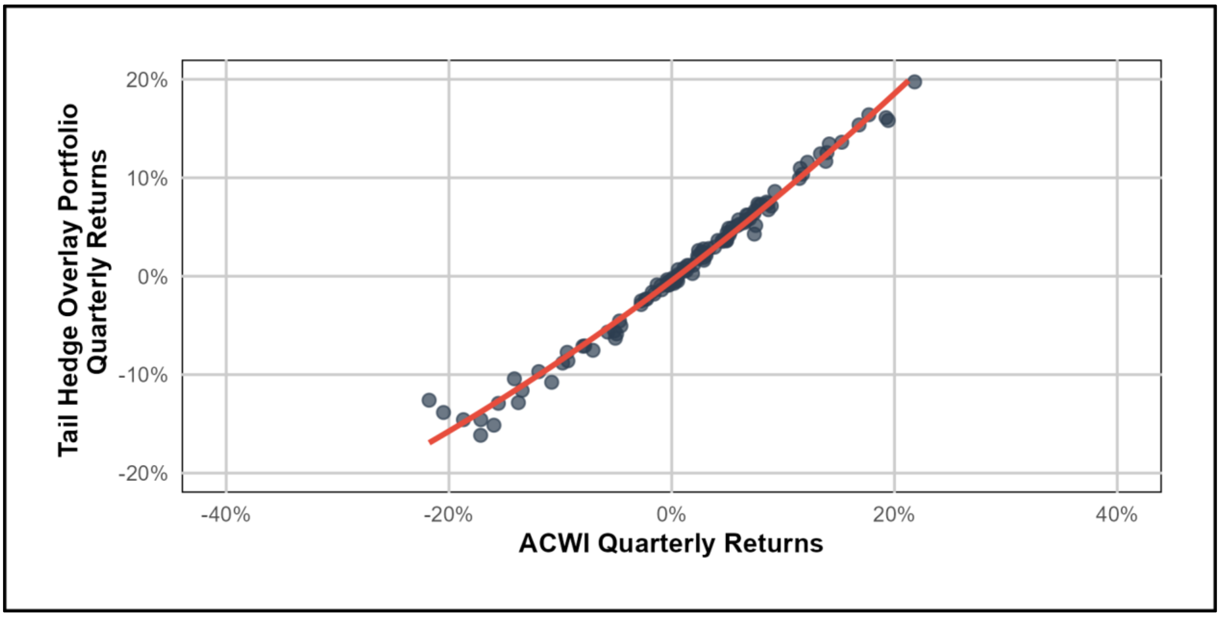

Figure 4.11 Convexity of Tail Overlay portfolio annual excess returns relative to ACWI annual excess returns

4.4.3.2 Regression summary: Tail Overlay portfolio monthly excess returns against ACWI and squared ACWI monthly excess returns

To continue the analysis of the convexity introduced by the tail risk hedging strategy, a linear regression model of the Tail Overlay portfolio's excess monthly returns against both ACWI's monthly excess returns and the squared ACWI monthly excess returns is fitted. Squaring ACWI returns allows for the overweighting of extreme observations and thus enables the statistical testing for the presence of convexity with respect to extreme equity market performance. Table 4.10 summarises the regression output.

The Tail Overlay portfolio continues to show a strong positive loading on ACWI, with a coefficient of +0.87, which is statistically significant at the 1% level. Importantly, the coefficient on the squared ACWI returns is +0.75, with a t-statistic of +10.80, which is also statistically significant at the 1% level. This positive and significant coefficient on the squared term suggests the presence of convexity in the Tail Overlay portfolio's performance relative to ACWI. This finding is in line with the improved tail risk statistics and drawdown profile relative to ACWI presented earlier in the chapter, further showing that the Tail Overlay portfolio empirically provided asymmetrical responses to extreme losses on ACWI.

The significance of the squared ACWI term, combined with the visual evidence provided in Figures 4.9 to 4.11, clearly illustrates that the Tail Overlay portfolio's performance is convex with respect to extreme left-tail global equity returns. These findings support the rejection of hypothesis $H_{0, G}$: Tail risk hedging does not introduce left-tail convexity with respect to global equity returns when overlaid. Instead, the alternative hypothesis $H_{1, G}$: Tail risk hedging introduces left-tail convexity with respect to global equity returns when overlaid is accepted.

Table 4.10 Summary of regression results: Trend Overlay portfolio, ACWI and squared ACWI returns

Statistic	ACWI	ACWI Squared	Intercept	R-Squared
Coefficient	0.87	0.75	-0.19%	99%
(t-Stat)	(155.67)***	(10.80)***	(-6.64)***	

4.4.4 Historical crisis episode performance comparison

The summary statistics presented have indicated that the Tail Overlay portfolio has empirically delivered significantly improved tail risk metrics compared to both ACWI and the Trend Overlay portfolio discussed in Chapter 3. To expand the analysis, specific episodes of market turmoil are now isolated to compare the portfolios' performance during these critical periods. Figure 4.12 illustrates the cumulative excess returns of the Tail Overlay portfolio and ACWI over the entire period under investigation, highlighting notable crisis periods as defined in Harvey et al. (2021) and Schwalbach and Auret (2023). Table 4.11 provides detailed data on each episode.

Figure 4.12 Cumulative excess returns of the Tail Overlay portfolio and ACWI with crises highlighted

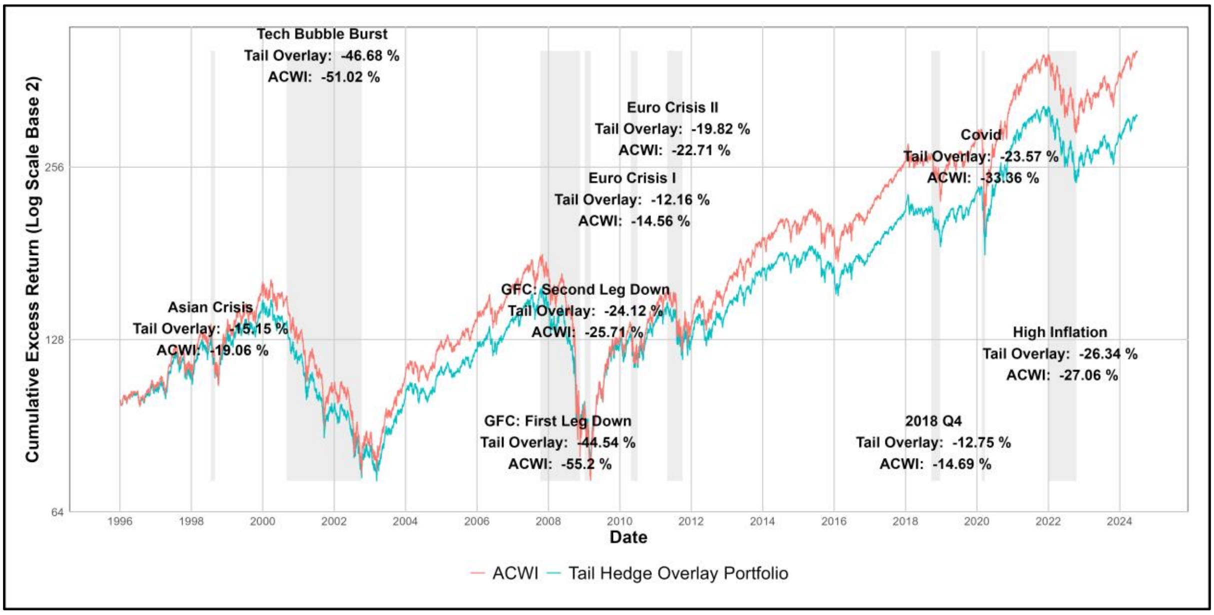

The Tail Overlay portfolio has empirically outperformed ACWI during sudden and severe market crashes. For example, during the Covid Crash, the Tail Overlay portfolio declined by -23.57% compared to ACWI's -33.36%. Similarly, during the first leg of the 2008 GFC, the Tail Overlay portfolio provided significant downside protection, falling -44.54% vs. ACWI's -55.20%, an outperformance of more than 10 percentage points in a short period. In contrast, the Tail Overlay portfolio was less effective during more prolonged drawdowns, such as the Dot-Com Bubble Burst. Over this extended period, the Tail Overlay portfolio declined by -46.68%, only moderately better than ACWI's decline of -51.02%.

This reduced efficacy aligns with the nature of tail risk hedging, which is ideally utilised to protect against sudden and severe market downturns characterised by sharp spikes in volatility. Prolonged market declines without significant volatility spikes have not triggered the same level of protective benefits from the tail risk hedging strategy. Conversely, the Trend Overlay portfolio, discussed in Chapter 3, has historically performed better relative to ACWI during extended drawdowns like the Dot-Com Bubble Burst. Trend-following strategies are adept at capturing sustained market trends, allowing them to benefit from gradual market declines by adjusting positions in response to persistent negative momentum.

Table 4.11 Cumulative excess returns of the Tail Overlay portfolio and ACWI during specific crises

Period Name	Peak Day	Trough Day	Number of Weekdays	ACWI Cumulative Return	Tail Portfolio Cumulative Return
Asian Crisis	1998-07-17	1998-08-31	31	-19.06%	-15.15%
Tech Bubble Burst	2000-09-01	2002-10-09	548	-51.02%	-46.68%
GFC: First Leg Down	2007-10-09	2008-11-20	292	-55.20%	-44.54%
GFC: Second Leg Down	2009-01-06	2009-03-09	44	-25.71%	-24.12%
Euro Crisis I	2010-04-23	2010-07-02	50	-14.56%	-12.16%
Euro Crisis II	2011-04-29	2011-10-03	111	-22.71%	-19.82%
2018 Q4	2018-09-20	2018-12-20	65	-14.69%	-12.75%
Covid	2020-02-19	2020-03-23	23	-33.36%	-23.57%
High Inflation	2022-01-03	2022-10-12	202	-27.06%	-26.34%

4.4.5 Relative Return Comparison

Investors focused on tracking error may find the Tail Overlay portfolio challenging to hold despite its improved tail risk profile relative to ACWI. As shown in Table 4.12, the Tail Overlay portfolio underperformed ACWI in most quarters due to the performance drag during normal or rising market conditions brought about by the overlaid tail risk hedging strategy. While the strategy has provided significant tail risk mitigation during severe market downturns, its benefits are realised only in a few extreme episodes. This persistent underperformance during benign periods makes it particularly challenging for tracking-error-sensitive investors to maintain confidence over the long term. This is in contrast to the Trend Overlay portfolio, which outperforms ACWI in approximately two-thirds of the observed quarters. However, the Tail Overlay portfolio's superior performance during market crashes remains an important point of differentiation.

Table 4.12 Relative return comparison of the Trend Overlay portfolio and ACWI

Statistic	Value
Percentage >= 0%	23.10%
Percentage < 0%	76.90%
Maximum Difference	10.61%
Minimum Difference	-4.01%

Longest Positive Streak	9
Longest Negative Streak	31

Figure 4.13 presents the difference between the quarterly excess returns of the Tail Overlay portfolio and ACWI over the period under investigation. The red dotted line, representing the average relative return, is slightly negative, reflecting the performance drag introduced by the embedded tail risk hedging strategy. Large spikes in relative outperformance occurred during periods of market stress, particularly in the 2008 and 2020 crashes. Underperformance, while persistent, rarely exceeded 3%. When it did exceed 3%, it tended to follow periods of outperformance, aligning with market recoveries.

When comparing the relative return statistics of the Trend Overlay and Tail Overlay portfolios, the Tail Overlay portfolio tended to exhibit significantly lower variability in its quarterly excess return differences with ACWI. While the Trend Overlay portfolio experiences both large outperformance and underperformance during different market regimes, the Tail Overlay portfolio exhibits less pronounced variation, reflective of its primary focus on left-tail risk management.

Figure 4.13 Quarterly excess return difference between the Tail Overlay portfolio and ACWI

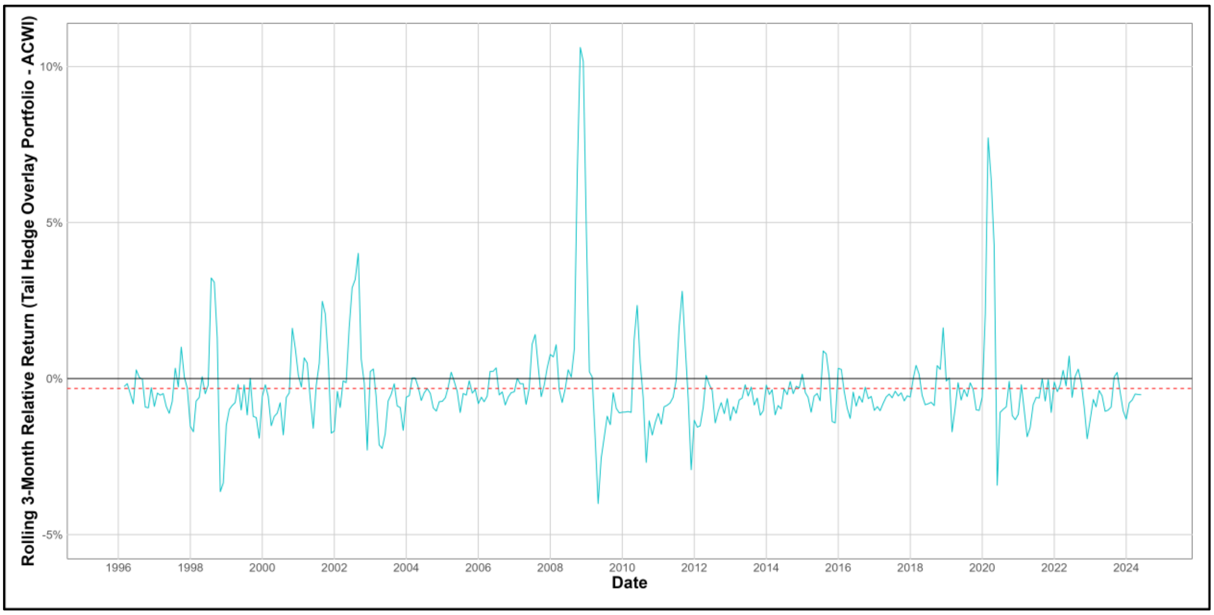

The findings presented in this chapter support the rejection of $H_{0,\ H}$: Overlaying a tail risk hedging strategy on top of global equity does not improve tail risk statistics and leads to a reduction in risk-adjusted returns compared to holding global equity alone, in favour of the alternative hypothesis $H_{1,\ H}$: Overlaying a tail risk hedging strategy on top of global equity improves tail risk statistics and does not reduce risk-adjusted returns compared to holding global equity alone.

4.5 CHAPTER SUMMARY

In this chapter, the risk and return profile of the ACWI was compared to that of the Tail Overlay portfolio, which comprises of the same allocation to ACWI except that it also incorporates a tail risk hedging strategy. The primary goal was to assess how effectively the tail risk hedging strategy mitigates extreme left-tail risk, and whether this comes at the expense of reduced risk-adjusted performance.

The analysis began by testing hypothesis $H_{0,\,F}$: SPX options are not suitable to hedge sudden and severe drawdowns in ACWI exposure. The findings suggests that overlaying SPX options mitigated the tail risk profile of an ACWI portfolio during major market crashes, such as the 2008 GFC and the 2020 Covid Crash. These results demonstrate the suitability of SPX options in protecting ACWI against sharp drawdowns. By systematically using 10-delta, one-year to expiration put options, the Tail Overlay portfolio cushioned severe market drawdowns relative to ACWI, leading to the rejection of hypothesis $H_{0,\,F}$.

Next, the hypothesis $H_{0,\,G}$: Tail risk hedging does not introduce left-tail convexity with respect to global equity returns when overlaid was explored. It was found that the Tail Overlay portfolio displayed clear left-tail convexity with respect to ACWI returns, meaning that during sharp equity market declines, the strategy helped mitigate losses significantly. This was particularly evident during sudden and severe market crises when implied volatility spiked, triggering large gains in the portfolio's option positions. The rejection of $H_{0,\,G}$ highlights the enhanced left-tail risk reduction benefits provided by the tail risk management strategy embedded with the Tail Overlay portfolio.

In contrast, the Trend Overlay portfolio discussed in Chapter 3 demonstrated more balanced convexity with respect to ACWI returns, offering improved performance across both the left-tail (albeit less in magnitude than the Tail Overlay portfolio) and the right-tail of the ACWI return distribution.

Finally, the hypothesis $H_{0,\,H}$: Overlaying a tail risk hedging strategy on top of global equity does not improve tail risk statistics and reduces risk-adjusted returns compared to holding global equity alone was tested. Despite a consistent performance drag brought about by the embedded tail risk hedging strategy during normal market conditions, the Tail Overlay portfolio provided significant improvements in key tail risk metrics, such as lower VaR and CVaR, while delivering the same Sharpe ratio as ACWI. Although the Tail Overlay portfolio generated lower absolute returns than ACWI, it offered superior protection in the worst market environments, particularly during market crashes.

This chapter demonstrates that overlaying a tail risk hedging strategy on a global equity portfolio has historically improved tail risk statistics without compromising risk-adjusted returns. The tail risk hedging strategy embedded within the Tail Overlay portfolio resulted in consistent outperformance relative to ACWI during sharp and severe market drawdowns.

The trend-following strategy embedded within the Trend Overlay portfolio discussed in Chapter 3 has historically delivered stronger results during prolonged market drawdowns, where the tail risk hedging strategy was empirically less effective. Thus, it is found that while both strategies have empirically offered valuable diversification to an equity portfolio, the Tail Overlay portfolio is particularly effective in providing left-tail convexity during acute market crises, while the Trend Overlay portfolio offers more balanced convexity and performs better in extended drawdowns. These complementary characteristics point to the potential benefits of combining our trend-following strategy and tail risk hedging strategy as overlays within a portfolio.

This leads to the exploration of hypothesis $H_{0,\,I}$: Overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently. This will be examined, along with the primary hypothesis of our study, H_0: Compounded long-term absolute and risk-adjusted equity returns are not improved by overlaying a combination of a trend-following strategy and a tail risk hedging strategy on top of full exposure to global equity, in the following chapter.

CHAPTER 5: COMBINING TREND-FOLLOWING AND TAIL RISK HEDGING OVERLAYS

5.1 CHAPTER INTRODUCTION

This chapter introduces the Combination Overlay portfolio, which integrates both trend-following and tail risk hedging strategies as overlays on top of an ACWI portfolio, aiming to use their complementary strengths. As discussed in previous chapters, trend-following is a hedge fund strategy that invests in multiple globally diversified markets—including equities, government bonds, currencies, and commodities—to capitalise on trending behaviour by buying markets that have been rising and selling those that have been falling. Tail risk hedging, on the other hand, involves strategies designed to truncate extreme portfolio losses from large market crashes through the use of optionality.

This chapter builds on earlier findings that trend-following strategies have historically delivered significant positive returns during prolonged bear markets, while tail risk hedging has generated substantial profits during sudden and severe market crashes. This chapter combines these approaches in order to introduce desirable convexity into an equity portfolio, with the aim of outperforming during extreme events, regardless of whether they manifest rapidly or gradually. Should this be achieved, the findings will align with those of Schwalbach and Auret (2023) in that the combination of trend-following and tail risk hedging offers investors a powerful tail risk mitigation solution with the potential to enhance returns as well.

This thesis has thus far demonstrated that, with respect to equity returns, trend-following strategies have exhibited significant convexity. That is, trend-following returns have historically been negatively correlated in the left tail of the equity return distribution, positively correlated in the right tail, and exhibited little to no correlation in the body of the distribution. The analysis has shown that tail risk hedging strategies have offered significant left-tail convexity with respect to equity returns by cushioning portfolio losses during market crashes.

Given these attractive and complementary historical characteristics, coupled with the efficiency with which these overlays can be applied to a fully collateralised global equity portfolio (since they are implemented using cash-efficient futures and options with low margin requirements) it is contended that the strategy is ideal for application in a portable alpha framework. This framework is one in which portfolio beta (equity exposure) and alpha (from the overlay strategies) are separated. Furthermore, this chapter examines whether the combined portfolio-level results are improved due to the diversification benefits arising from the low correlation between the overlay strategies and the underlying equity portfolio.

Enhancement in risk-adjusted returns occurs in part due to the introduction of greater convexity with respect to both left-tail and right-tail equity returns when both strategies are overlaid on a global equity portfolio. Testing for the presence of this increased convexity will thus be central to this chapter, the hypothesis $H_{0, I}$: Overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently will be evaluated. This chapter will conclude with the evaluation of the primary hypothesis of this thesis, H_0: Compounded long-term absolute and risk-adjusted equity returns are not improved by overlaying a combination of a trend-following strategy and a tail risk hedging strategy on top of full exposure to global equity.

5.2 DATA AND METHODOLOGY

5.2.1 Data description

The data used in this chapter are a combination of the data employed in the previous chapters, encompassing both the trend-following and tail risk hedging strategies. The analysis covers the period from 4 January 1996 to 28 June 2024, and all returns are excess returns and in US dollars (USD). The data include futures contracts, exchange rates, interest rates, and option contracts. Given the global nature of the trend-following strategy, which utilises futures contracts traded across various international exchanges, public holidays do not align perfectly across markets. As a result, the data are standardised to include all weekdays in a non-leap year, amounting to 261 trading days. On days when an exchange is closed due to a public holiday, 0% returns are assigned to the affected instruments. This standardised approach is also applied to the data used for the tail risk hedging strategy to ensure consistency across the combined analysis.

5.2.1.1 Trend-following strategy

The data used in the trend-following strategy include the same 79 futures contracts detailed in Table 3.1. This data set comprises 23 equity index futures (18 from developed markets and 5 from emerging markets), 13 developed market government bond futures, 14 currency futures (of which 5 represent emerging market currencies against the USD), and 29 commodity futures. All futures data are sourced from Bloomberg and comprise last traded daily prices, which are converted into daily returns. These returns represent excess returns, as they are derived from futures contracts. For non-USD denominated contracts, daily returns are converted into USD using spot exchange rates, also sourced from Bloomberg. As the trend-following strategy employs exchange-traded futures contracts, the resulting portfolio return is considered an excess return. To proxy for the risk-free rate, 3-month US LIBOR is used throughout this analysis.

5.2.1.2 Tail Risk Hedging Strategy

For the tail risk hedging strategy employed in this study, exchange-traded SPX options data are utilised. The analysis spans the period from 4 January 1996 to 28 June 2024, with all data denominated in US dollars. The options data are sourced from OptionMetrics, which provides end-of-day bid and ask prices for each SPX option; mid-prices are computed by averaging bid and ask prices. The tail risk hedging strategy uses the same 10-delta, one-year to expiration SPX put options used in the same way as detailed in Chapter 4.

5.2.1.3 Data used for performance comparison and regression analysis

ACWI is utilised as a proxy for global equity and for the performance comparison of the Combination Overlay portfolio, which is constructed by overlaying both the tail risk hedging strategy and 50% of the trend-following strategy, both presented earlier in this research, on top of a fully invested ACWI position. This represents the same approach as in Schwalbach and Auret (2023). For the broader asset class performance comparisons, ACWI is also employed to represent the global equity market. The Bloomberg Global Aggregate Total Return Index is used as a proxy for global bond performance, while the S&P GSCI Index serves as the benchmark for commodities. All data are sourced from Bloomberg. Factor return data are obtained from Kenneth French's website (French, 2024). Small minus big (SMB) and high minus low (HML) factors to proxy for size and value, respectively. Momentum factor data are also sourced from Kenneth French's database (French, 2024).

5.2.2 Methodology

To assess the potential risk and return enhancements achieved by combining trend-following and tail risk hedging overlays on a conventional equity portfolio, this chapter focuses on the Combination Overlay portfolio, comparing its performance to that of ACWI. The methodology used in this chapter is similar to that of Schwalbach and Auret (2023). The Combination Overlay portfolio consists of a 100% long position in ACWI, overlaid with a 50% allocation to the trend-following strategy discussed in Chapter 3 and the full tail risk hedging strategy introduced in Chapter 4. This construction is designed to capitalise on the convexity introduced by these overlays while maintaining full equity exposure throughout the analysis period.

The same methodologies and rebalancing protocols for the trend-following and tail risk hedging overlays as established in the previous chapters are employed in this chapter. The daily returns of the Combination Overlay portfolio are calculated by maintaining a 100% exposure to ACWI, funded by cash (proxied by 3-month US LIBOR), and overlaying the aforementioned strategies. The daily portfolio return is computed as the sum of the return contributions from each component. Sufficient portfolio margin is assumed to be available throughout the entire analysis period. This methodology is consistent with that of Schwalbach and Auret (2023), except in this thesis, trend-following exposure is attained using a fully-fledged trend-following strategy (defined in Chapter 3), where Schwalbach and Auret (2023) proxied for trend-following returns using an industry benchmark.

The risk and return profile of the Combination Overlay portfolio is compared not only to ACWI but also to the individual overlay portfolios, namely, the Trend Overlay and Tail Overlay portfolios discussed in the previous chapters. This comparison involves analysing different risk and return statistics, including CAGR, volatility, Sharpe ratio, skewness, and kurtosis. In addition, VaR and CVaR statistics are computed, as well as drawdown metrics. Linear regression analysis is performed to evaluate the Combination Overlay portfolio's excess return loadings on several factors.

5.3 FINDINGS

In this section, the empirical performance characteristics of the Combination Overlay portfolio are summarised. This portfolio combines the tail risk hedging strategy with a 50% allocation to the trend-following strategy, overlaid on a 100% allocation to the ACWI. Its performance is compared against that of ACWI and key distinctions relative to the Tail Overlay and Trend Overlay strategies examined in previous chapters are highlighted. This analysis aims to test Hypothesis $H_{0,I}$: overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently. Finally, the primary hypothesis of this thesis, H_0, which posits that compounded long-term absolute and risk-adjusted equity returns are not improved by overlaying this combination on full exposure to global equity is evaluated.

5.3.1 Excess performance comparison

Figure 5.1 illustrates the cumulative excess returns of the Combination Overlay portfolio and ACWI. The results are presented on a logarithmic scale for ease of comparison. Table 5.1 presents the summary statistics of both portfolios. It is found that the CAGR for the Combination Overlay portfolio was 10.06% over the period under investigation, which is significantly higher in absolute terms than both the Tail Overlay portfolio, at 4.11%, and ACWI, which delivered 5.05%. The Trend Overlay portfolio exhibited the highest CAGR of 11.12%. This result suggests that the tail risk hedging strategy embedded in the Combination Overlay portfolio exerted a mild drag on absolute returns. These findings are consistent with those of Schwalbach and Auret (2023), who observed the same hierarchy in empirical performance, with the Trend Overlay portfolio generating the highest returns and the Tail Overlay portfolio delivering the lowest.

Figure 5.1 Cumulative excess returns of the Combination Overlay portfolio and ACWI

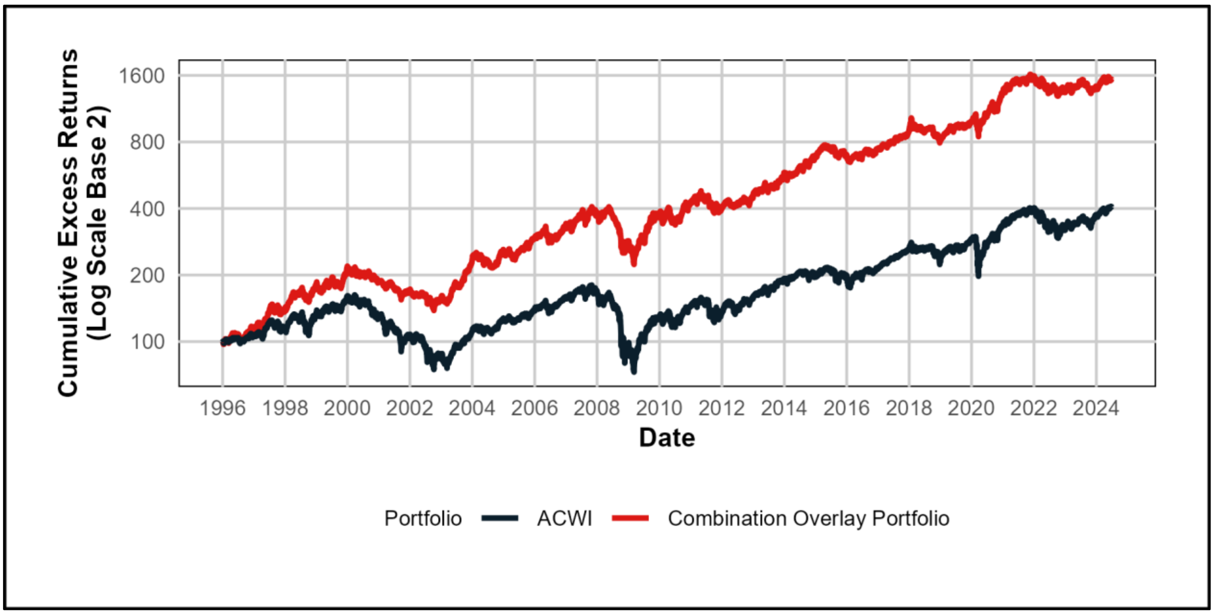

The Combination Overlay portfolio suffered meaningfully shallower drawdowns across the period than ACWI, as depicted in Figure 5.2. While similar conclusions were reached for both the Trend Overlay and Tail Overlay portfolios, the Combination Overlay portfolio uniquely benefited from the combined strengths of the two strategies. The embedded tail risk hedging strategy empirically cushioned sharp market crashes, while the embedded trend-following strategy empirically mitigated losses during extended bear markets.

This convergence of complementary risk mitigation historically led to markedly shallower drawdowns and faster recoveries. This is evidenced by the portfolio's worst 3-month performance of -22.92%, which is substantially less severe than ACWI's -34.86%. This metric was also lower than that of either the Trend or Tail Overlay portfolios individually. The volatility of the Combination Overlay portfolio was 13.95%, lower than that of ACWI (15.59%) and the Trend Overlay portfolio (16.40%) and was only surpassed by the Tail Overlay portfolio (12.72%). As a result, the Combination Overlay portfolio delivered the highest Sharpe ratio of all the strategies investigated in this paper, at +0.72. This figure was more than double that of ACWI, suggesting a significant improvement in risk-adjusted returns over the period under investigation brought about by the diversification improvements of the tail risk hedging and trend-following overlays.

The characteristic negative skewness typically associated with equity strategies, as demonstrated by ACWI's skewness of -0.65, was almost eliminated by the application of both overlays. While the Tail Overlay portfolio reduced negative skewness to -0.43 and the Trend Overlay portfolio further improved it to -0.22, the Combination Overlay portfolio delivered an almost neutral skewness of -0.01. This effectively eliminated the downside bias typically seen in equity portfolios, highlighting the superior diversification benefits of combining trend-following and tail risk hedging strategies, rather than employing them independently or not at all.

Figure 5.2 Drawdowns of Combination Overlay and ACWI

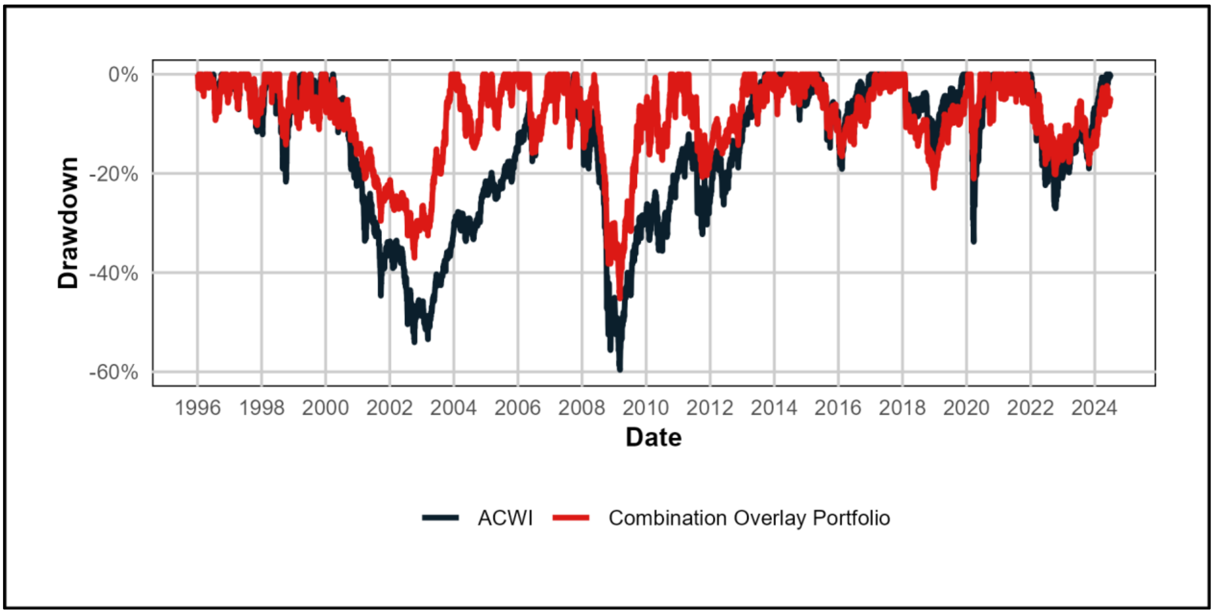

When considering tail risk statistics presented in Table 5.1, the Combination Overlay portfolio exhibited impressive protection against extreme market events relative to ACWI and the individual overlay strategies. Specifically, the 95% and 99% monthly VaR metrics for the Combination Overlay portfolio were -6.11% and -8.65%, respectively. This was notably lower than ACWI's VaR figures of -8.16% (95%) and -12.04% (99%). In terms of monthly CVaR, the Combination Overlay portfolio again performed best, with a 95% CVaR of -7.69% and a 99% CVaR of -9.71%. These figures are meaningfully lower than ACWI's CVaR values of -10.44% (95%) and -15.18% (99%). The Combination Overlay's tail risk metrics were also an improvement on those of the Trend Overlay portfolio, which had a 95% CVaR of -9.22% and a 99% CVaR of -12.98%, as well as the Tail Overlay portfolio, which showed a 95% CVaR of -8.65% and a 99% CVaR of -11.25%.

Table 5.1 Summary statistics: Combination Overlay portfolio and ACWI

Statistic	Combination Overlay Portfolio	ACWI
CAGR	10.06%	5.05%
Annualised Volatility	13.95%	15.59%
Sharpe Ratio	0.72	0.32
Skewness	-0.01	-0.65
Excess Kurtosis	-3.10	-1.54
VaR (95%)	-6.11%	-8.16%
VaR (99%)	-8.65%	-12.04%
CVaR (95%)	-7.69%	-10.44%
CVaR (99%)	-9.71%	-15.18%
Best 3M Performance	24.14%	33.06%
Worst 3M Performance	-22.92%	-34.86%

Figure 5.3 illustrates the rolling 12-month correlation between the returns of the Combination Overlay portfolio and ACWI. Over the sample period, correlations have oscillated between the mid-70% and mid-90% range, reflecting strong co-movement between the two portfolios during normal market conditions. However, the Combination Overlay portfolio exhibited lower correlations with ACWI compared to the individual overlay portfolios. Similar to the individual overlay strategies, correlations tended to break down slightly during periods of market stress. This was primarily due to the Combination Overlay portfolio outperforming ACWI in extreme events.

Figure 5.3 Rolling one-year correlations of the Combination Overlay portfolio and ACWI

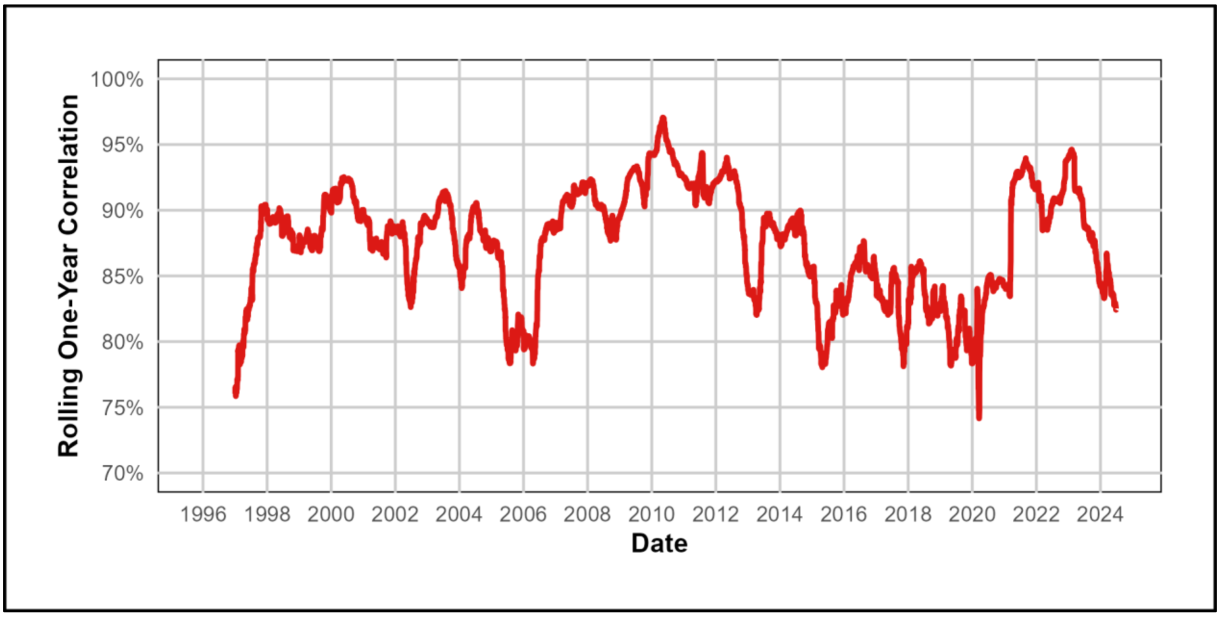

5.3.2 Alpha and loading on risk factors in Combination Overlay portfolio returns

Figure 5.4 presents a scatterplot of the monthly excess returns of ACWI on the x-axis against those of the Combination Overlay portfolio on the y-axis. A linear regression is fit to the data, and the results are summarised in Table 5.2. The regression output shows that the excess monthly returns of the Combination Overlay portfolio have a coefficient of +0.81 with ACWI's monthly returns, which is statistically significant at the 1% level. This coefficient is lower than that of the Trend Overlay portfolio (+0.96) and slightly lower than that of the Tail Overlay portfolio (+0.85). Table 5.2 shows that the Combination Overlay portfolio generated a positive monthly alpha of 0.47%, statistically significant at the 1% level and close to the 0.50% monthly alpha generated by the Trend Overlay portfolio, as shown in Table 3.11. Approximately 75% of the variability in the Combination Overlay portfolio's monthly excess returns can be explained by ACWI's monthly excess returns.

Figure 5.4 Linear regression analysis of excess monthly returns: Combination Overlay portfolio vs. ACWI

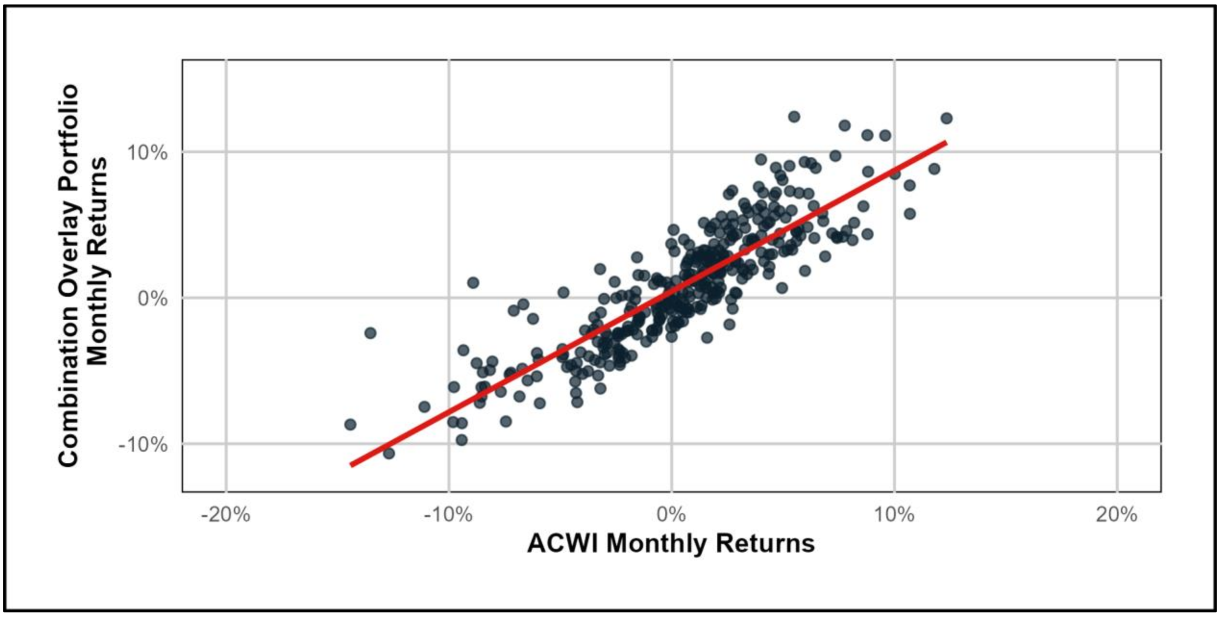

Table 5.2 Summary of regression results: Combination Overlay portfolio and ACWI

Statistic	ACWI	Intercept	R-Squared
Coefficient	0.81	0.47%	75%
(t-Stat)	(31.58)***	(4.09)***	

Table 5.3 highlights the Combination Overlay portfolio's loadings on traditional equity market factors. A large, positive and statistically significant loading on ACWI is found, with a coefficient of +0.87. This coefficient is slightly larger than that of the Tail Overlay portfolio presented in Table 4.8 but lower than the coefficient of +1.01 for the Trend Overlay portfolio in Table 3.12. The difference in these coefficients shows the impact of the negative delta embedded in the tail risk hedging strategy, which reduces effective equity exposure over time primarily because delta hedging was not performed.

The Combination Overlay portfolio also loads positively and significantly on the momentum factor (UMD), with a coefficient of +0.13. This momentum loading reflects the influence of the embedded trend-following strategy, as the Trend Overlay portfolio's loading on UMD was 0.12, also statistically significant at the 1% level. However, the size (SMB) and value (HML) factors show no statistically significant loading, with coefficients of -0.03 and +0.01, respectively. This suggests that the Combination Overlay portfolio does not exhibit a bias towards small-cap or value stocks, reinforcing the notion that its performance is primarily driven by broader market beta and, to a much lesser extent, the momentum factor.

A positive monthly alpha of 0.40% is observed, significant at the 1% level. This indicates that the portfolio delivered excess returns beyond what is explained by its significant loadings on ACWI and the momentum factor.

Table 5.3 Summary of regression results: Combination Overlay portfolio against traditional market factors

Statistic	**ACWI**	**SMB**	**HML**	**UMD**	**Intercept**	**R-Squared**
Coefficient	0.87	-0.03	0.01	0.13	0.40%	76%
(t-Stat)	(31.29)***	(-0.89)	(0.23)	(5.02)***	(3.47)***	

Next, the model presented in Table 5.3 is extended by incorporating global government bond excess returns and commodity excess returns, proxied by the Bloomberg Global Aggregate Total Return Index and the S&P GSCI Index, respectively. The findings, presented in Table 5.4, show that the key results relating to traditional equity factors remain largely unchanged.

The Combination Overlay portfolio continues to exhibit a strong, positive and statistically significant loading on ACWI, with a coefficient of +0.89 (slightly higher than in the previous regression). The portfolio also maintains its significant positive loading on the momentum factor (UMD), with a coefficient of +0.13, reinforcing the influence of the embedded trend-following strategy. The loadings on size (SMB) and value (HML) factors remain statistically insignificant, consistent with the results from Table 5.3. With the inclusion of global government bonds (BONDS) and commodities (GSCI) in the regression, it is observed that the portfolio displays a near-zero, positive, and insignificant loading on global government bonds, with a coefficient of +0.04. The coefficient on the commodity index is -0.03, statistically significant at the 10% level. This slight negative loading suggests a minor inverse relationship between commodity prices and the Combination Overlay portfolio's performance, though very small in magnitude.

Notably, the portfolio still delivers a positive monthly alpha of 0.38%, statistically significant at the 1% level. This demonstrates that even after accounting for the additional global government bond and commodity factors, the Combination Overlay portfolio continued to empirically produce excess returns beyond what can be explained by its loadings on equity factors, global government bonds, and commodity exposures. The R-squared value marginally increases to 77%, indicating that the extended model explains a slightly larger proportion of the portfolio's return variation, although the impact of the added factors remains minimal in driving performance.

Table 5.4 Summary of regression results: Combination Overlay portfolio against traditional market factors, global government bonds and commodities

Statistic	**ACWI**	**SMB**	**HML**	**UMD**	**BONDS**	**GSCI**	**Intercept**	**R-Squared**
Coefficient	0.89	-0.02	0.02	0.13	0.04	-0.03	0.38%	77%
(t-Stat)	(28.02)***	(-0.59)	(0.58)	(5.21)***	(0.54)	(-1.71)*	(3.37)***	

5.3.3 Combination Overlay portfolio convexity with respect to ACWI returns

The results presented thus far suggest that the Combination Overlay portfolio has empirically delivered a statistically significant alpha of similar magnitude to that of the Trend Overlay portfolio after accounting for traditional equity factors, global government bonds, and commodities. What is different, however, is that due to the incorporation of the tail risk hedging strategy, the Combination Overlay portfolio exhibits a significantly improved risk profile relative to the Trend Overlay portfolio. To further illustrate this finding, the analysis is extended by exploring the convexity introduced by combining the trend-following and tail risk hedging overlays on top of ACWI.

5.3.3.1 Scatterplots illustrating Combination Overlay portfolio convexity with respect to ACWI

Figures 5.5 to 5.7 present scatterplots of the monthly, quarterly, and annual ACWI excess returns on the x-axis against the Combination Overlay portfolio returns on the y-axis. These scatterplots include quadratic fits to visually demonstrate the convexity of the Combination Overlay portfolio relative to ACWI excess returns. At each frequency, a non-linear relationship is observed between the Combination Overlay portfolio and ACWI, characterised by shallower losses and larger gains at the extremes. This suggests that the combination of trend-following and tail risk hedging strategies provides convexity on both the left and right tails of the ACWI return distribution.

The resulting quadratic fits are more convex than those of the Trend Overlay and Tail Overlay portfolios when considered individually. This increased convexity indicates that the characteristic U-shaped payoff profile typically associated with trend-following strategies has not only been retained but also amplified on the left tail (where convexity is most prized) due to the inclusion of the tail risk hedging strategy.

Figure 5.5 Convexity of Combination Overlay portfolio monthly excess returns relative to ACWI monthly excess returns

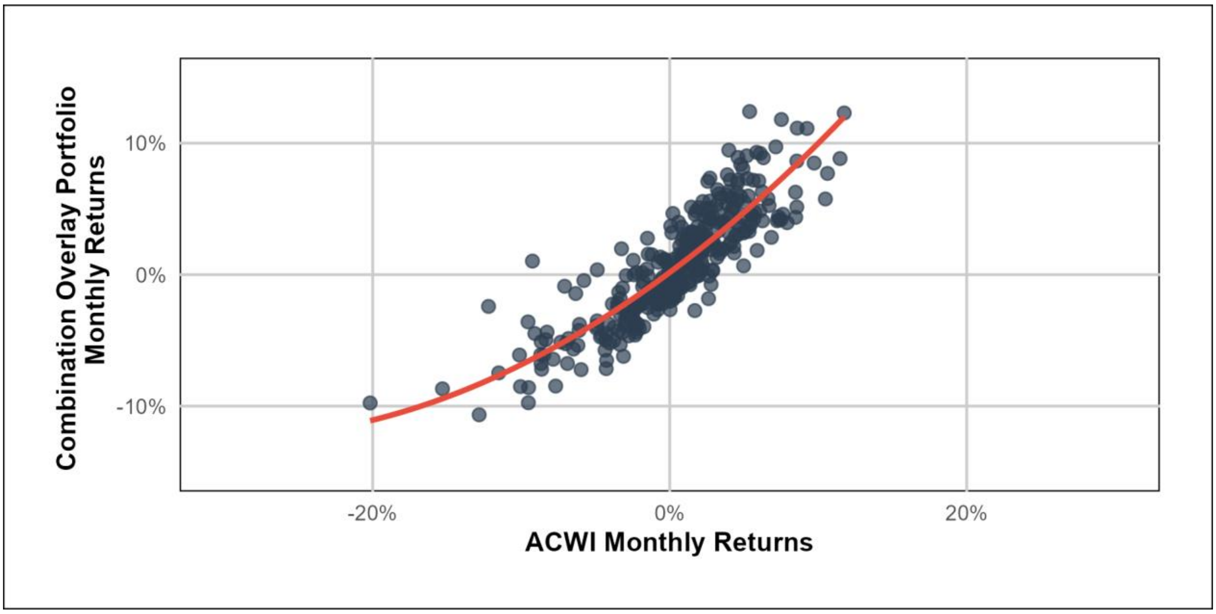

Figure 5.6 Convexity of Combination Overlay portfolio quarterly excess returns relative to ACWI quarterly excess returns

Figure 5.7 Convexity of Combination Overlay portfolio annual excess returns relative to ACWI annual excess returns

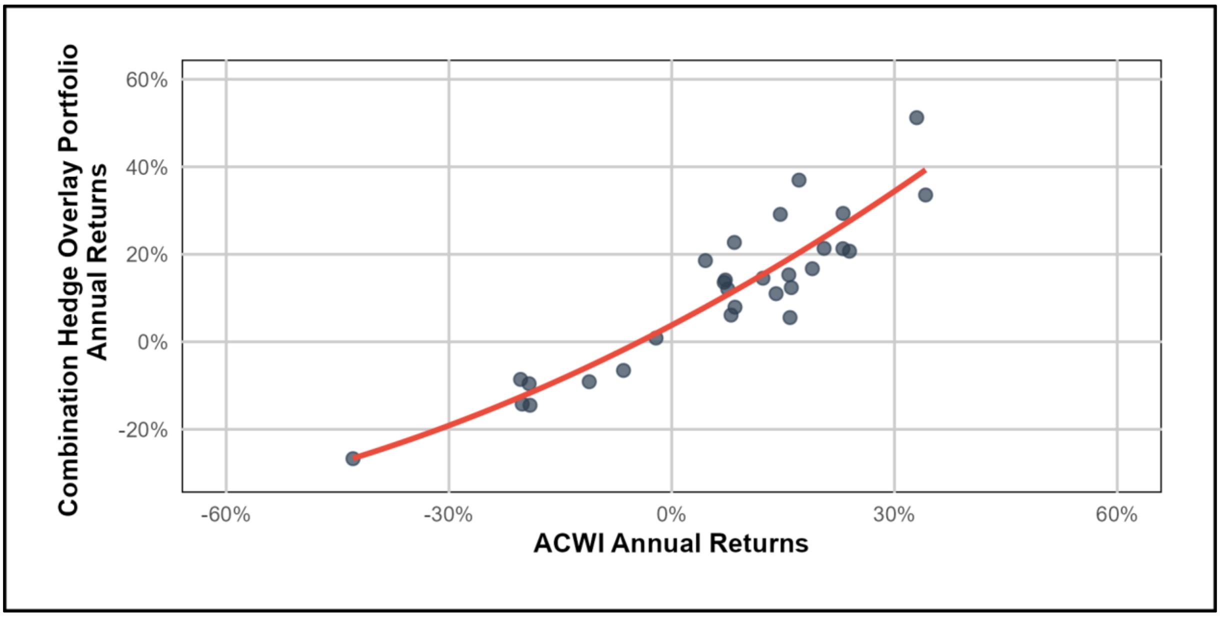

5.3.3.2 Regression summary: Combination Overlay portfolio monthly excess returns against ACWI and squared ACWI monthly excess returns

Next, the presence of convexity is formally tested by regressing the Combination Overlay portfolio's excess monthly returns against both ACWI and squared ACWI monthly excess returns. The inclusion of a squared ACWI term allows for overweighting extreme returns relative to more benign returns in the model. The findings are presented in Table 5.5, which summarises the regression output.

The coefficient on the squared ACWI term is large and positive at +1.40, statistically significant at the 1% level. This coefficient is substantially larger than both the Trend Overlay portfolio’s coefficient of +0.83 and the Tail Overlay portfolio’s coefficient of +0.75, presented in previous chapters. This indicates that the Combination Overlay portfolio has empirically delivered a significantly greater degree of convexity with respect to extreme ACWI returns than the individual overlay portfolios, consistent with the visual evidence presented in Figures 5.5 to 5.7. These findings suggest that, while the Combination Overlay portfolio tends to be less correlated with the month-to-month linear movements of ACWI compared to the individual overlay portfolios, it offers significantly more convexity with respect to extreme equity returns, both positive and negative.

The Combination Overlay portfolio demonstrates that positive convexity serves as a highly effective remedy for mitigating the most harmful effects of variance drag. Improving the compounding profile of arithmetic returns by mitigating the impact of extreme portfolio volatility and large losses can result in enhanced expected geometric returns over time.

The implication is that investors can effectively improve tail risk metrics of equity-dominant portfolios without sacrificing the potential for attractively high long-term absolute returns. As explored in earlier chapters, variance drag refers to the disproportionate impact that large portfolio losses have on compounded returns over time. Viewed through this lens, it becomes evident why the Combination Overlay portfolio has empirically generated alpha comparable to that of the Trend Overlay portfolio while simultaneously offering markedly improved risk metrics. These results support the rejection of $H_{0,\ I}$: overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently.

Table 5.5 Summary of regression results: Combination Overlay portfolio, ACWI and squared ACWI Returns

Statistic	**ACWI**	**ACWI Squared**	**Intercept**	**R-Squared**
Coefficient	0.84	1.40	0.17%	76%
(t-Stat)	(32.22)***	(4.35)***	(1.29)	

5.3.4 Historical crisis episode performance comparison

To highlight the Combination Overlay portfolio's performance during isolated market crises, Figure 5.8 illustrates the cumulative excess returns of the strategy and ACWI. Specific crisis periods, as defined in Harvey et al. (2021) and Schwalbach and Auret (2023), are highlighted in grey. More detail on each episode is provided in Table 5.6.

Figure 5.8 Cumulative excess returns of the Combination Overlay portfolio and ACWI with crises highlighted

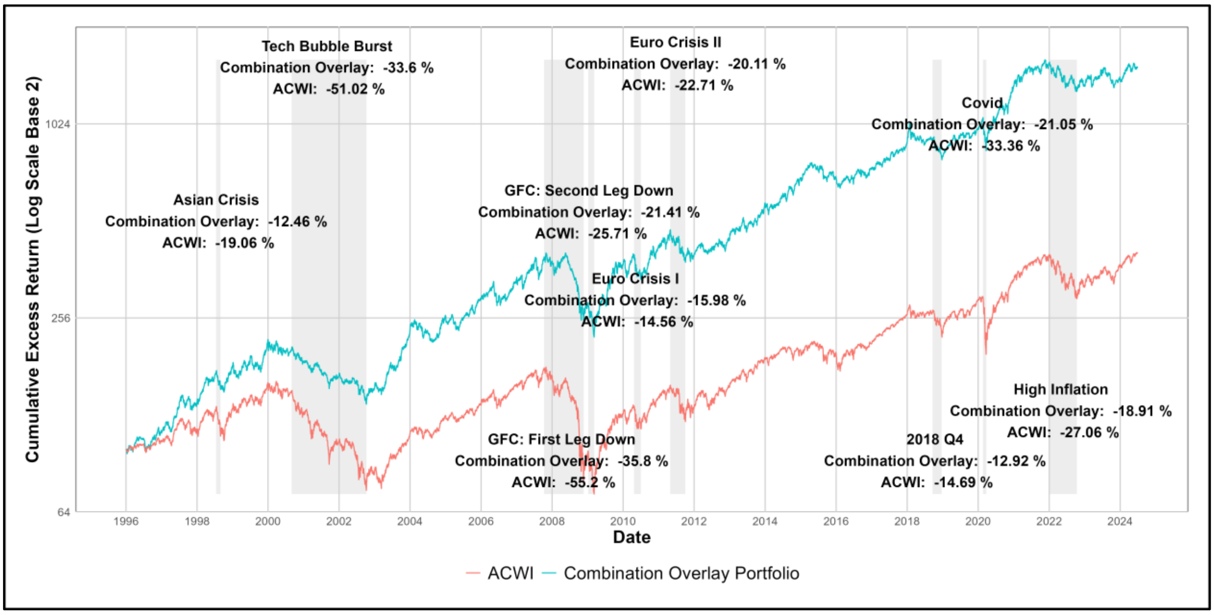

During the Asian Crisis (1998), the Combination Overlay portfolio experienced a drawdown of 12.46%, considerably less than ACWI's 19.06% decline. The drawdown was also shallower than those of both the Tail Overlay and Trend Overlay portfolios, which fell by 15.15% and 16.39%, respectively.

The Tech Bubble Burst (2000–2002), lasting for 548 trading days per Harvey et al. (2021), exemplifies the tendency for trend-following strategies to deliver strong and positive performance during prolonged equity downturns, as documented in the literature. The Combination Overlay portfolio lost 33.60%, markedly less than ACWI's 51.02% drawdown. It also outperformed the Tail Overlay and Trend Overlay portfolios, which declined by 46.68% and 38.69%, respectively. The GFC in its two phases (2007–2008 and 2009) further reveals the impact of the increased convexity present in the Combination Overlay portfolio.

During the first leg of the GFC (lasting 292 days), the Combination Overlay portfolio's drawdown was 35.80%, while ACWI fell by 55.20%. This outcome demonstrates the impact of the embedded tail risk hedging strategy on the Combination Overlay portfolio's returns. Similarly, during the second leg of the GFC, the Combination Overlay portfolio lost 21.41%, outperforming ACWI's decline of 25.71%. It also fared better than the Tail Overlay and Trend

Overlay portfolios, which fell by 24.12% and 23.02%, respectively. The results from the first Euro Crisis (2010) and second Euro Crisis (2011) were mixed.

During the first Euro Crisis, the Combination Overlay portfolio declined by 15.98%, slightly underperforming ACWI's loss of 14.56%, fitting squarely between the Tail Overlay portfolio, which fell the least by 12.16% and the Trend Overlay portfolio, which fell the most by 18.27%. It again straddled the drawdowns suffered by the Tail Overlay and Trend Overlay strategies (drawdowns of 19.82% and 22.93%, respectively) during the second Euro Crisis; however, in this episode, the Combination Overlay portfolio’s 20.11% fall modestly outperformed ACWI's decline of 22.71%.

The shorter-lived market shocks of Q4 2018 also resulted in the Combination Overlay portfolio slightly outperforming ACWI. The portfolio declined by 12.92%, compared to ACWI's loss of 14.69%. It also outperformed the Trend Overlay portfolio, which fell by 14.86%, and performed similarly to the Tail Overlay portfolio, which declined by 12.75%. The Covid Crisis offered further evidence of the embedded tail risk hedging strategy’s influence on the Combination Overlay portfolio.

The strategy fell by only 21.05%, compared to ACWI’s 33.36% loss. This performance was better than that of both the Tail Overlay portfolio, which declined by 23.57%, and the Trend Overlay portfolio, which fell by 30.78%. Finally, the highly inflationary 2022 episode offered further evidence of trend-following's influence on the Combination Overlay portfolio. The portfolio fell by only 18.91%, compared to ACWI’s 27.06% loss. It also outperformed the Tail Overlay portfolio, which declined by 26.34%, and performed slightly better than the Trend Overlay portfolio, which fell by 19.65%.

These results emphasise the complementary nature of trend-following and tail risk hedging, offering an attractive blend of risk mitigation in both sudden, sharp drawdowns during crashes and prolonged, deep equity declines during bear markets. The combination of these strategies has empirically introduced convexity with respect to both left tail and right tail global equity returns, serving as a powerful remedy for mitigating the most damaging effects of variance drag, all while preserving competitive long-term absolute returns.

Table 5.6 Cumulative excess returns of the Combination Overlay portfolio and ACWI during specific crises

Period Name	Peak Day	Trough Day	Number of Weekdays	Combination Overlay Portfolio Cumulative Return	ACWI Cumulative Return
Asian Crisis	1998-07-17	1998-08-31	31	-12.46%	-19.06%
Tech Bubble Burst	2000-09-01	2002-10-09	548	-33.60%	-51.02%
GFC: First Leg Down	2007-10-09	2008-11-20	292	-35.80%	-55.20%
GFC: Second Leg Down	2009-01-06	2009-03-09	44	-21.41%	-25.71%
Euro Crisis I	2010-04-23	2010-07-02	50	-15.98%	-14.56%
Euro Crisis II	2011-04-29	2011-10-03	111	-20.11%	-22.71%
2018 Q4	2018-09-20	2018-12-20	65	-12.92%	-14.69%
Covid	2020-02-19	2020-03-23	23	-21.05%	-33.36%
High Inflation	2022-01-03	2022-10-12	202	-18.91%	-27.06%

5.3.5 Relative return comparison

This section evaluates the relative return dynamics of the Combination Overlay portfolio compared to ACWI, with a specific focus on its viability for investors sensitive to tracking error. Based on the data presented in Table 5.7, the Combination Overlay portfolio outperformed ACWI in just under 60% of all quarters over the period under investigation. This demonstrates a consistent degree of outperformance over ACWI. This figure is slightly lower than that of the Trend Overlay portfolio, which outperformed ACWI in about two-thirds of quarters, but much higher than that of the Tail Overlay portfolio, which only generated quarterly outperformance one-quarter of the time.

The longest positive streak of outperformance over ACWI delivered by the Combination Overlay portfolio was 17 consecutive quarters (compared to 13 for the Trend Overlay portfolio and 9 for the Tail Overlay portfolio). The longest negative streak delivered by the Combination Overlay portfolio was 8 consecutive quarters (compared to 7 for the Trend Overlay portfolio and 31 for the Tail Overlay portfolio). Notably, the longest streak of underperformance was much shorter than the longest streak of outperformance.

Table 5.7 Relative return comparison of the Combination Overlay portfolio and ACWI

Statistic	**Value**
Percentage >= 0%	59.36%
Percentage < 0%	40.64%
Maximum Difference	15.72%
Minimum Difference	-8.92%
Longest Positive Streak	17
Longest Negative Streak	8

Figure 5.9 illustrates the difference between the quarterly excess returns of the Combination Overlay portfolio and ACWI over the period. The red dotted line represents the average relative return over the full period. Similar to the Trend Overlay portfolio, this average was positive, which aligns with the findings presented earlier in this chapter. However, the peaks in Figure 5.9 representing the magnitude of outperformance in the Combination Overlay portfolio were greater than those of both the Trend and Tail Overlay portfolios.

The Combination Overlay portfolio's maximum quarterly excess return difference was 15.72%, which was larger than the 10.82% observed in the Trend Overlay and 10.61% in the Tail Overlay. However, the downside risk of underperformance was also greater, with a minimum difference of –8.92% compared to –6.59% for the Trend Overlay and –4.01% for the Tail Overlay. While there were larger swings in relative performance in the Combination Overlay portfolio compared to the individual overlay portfolios, large underperformance mostly followed large outperformance.

Figure 5.9 Quarterly excess return difference between the Combination Overlay portfolio and ACWI

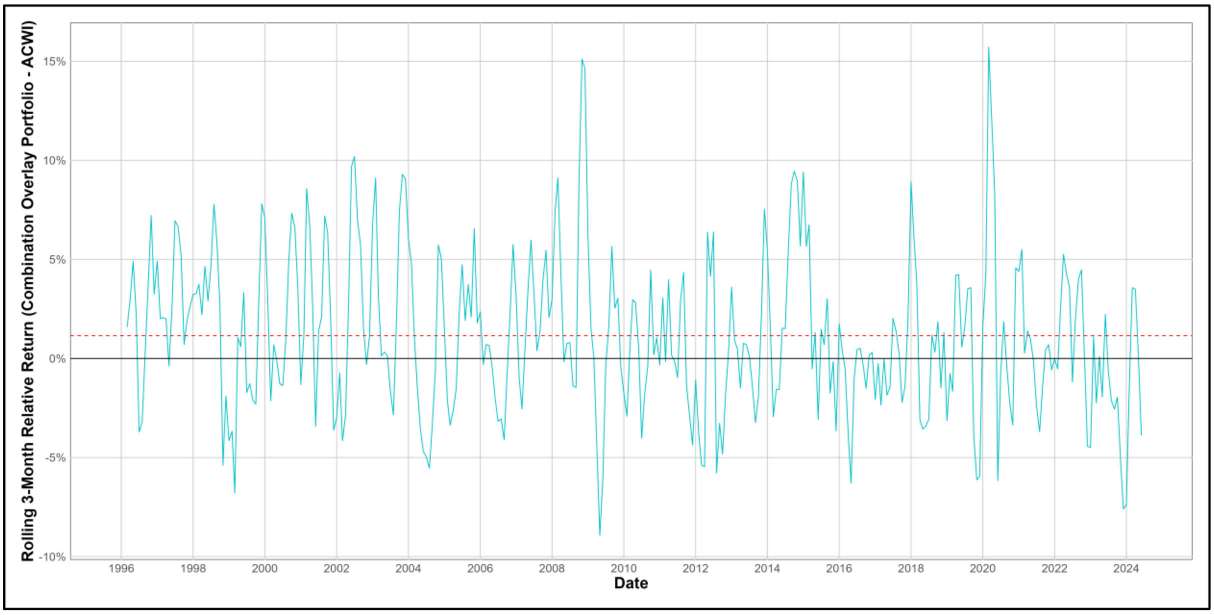

5.3.6 Performance attribution

Table 5.8 below presents a performance attribution of the Combination Overlay portfolio's excess CAGR. The strategy empirically captured the 5.05% excess returns generated by ACWI and 50% of the 11.27% excess CAGR achieved by the trend-following strategy covered in Chapter 3. The tail risk hedging strategy detracted from performance by 1.38% over the period. The combined sum of these returns amounts to 9.30%, while the CAGR of the Combination Overlay portfolio was 10.06%, as illustrated in Table 5.1. The residual, defined as the difference between the 10.06% CAGR of the strategy and the 9.30% accounted for by the sum-of-parts approach, is 0.76% per annum. This residual can be interpreted as the benefit derived from the strategy's reduction in variance drag, which enhanced the compounding profile of the strategy.

The same analysis was conducted for both the Trend Overlay and Tail Overlay strategies, finding residuals of 0.44% each. The greater compounding benefit generated by the Combined Overlay strategy provides further evidence supporting the overarching conclusion: combining trend-following and tail risk hedging strategies and overlaying them on top of a global equity portfolio is more effective than employing each strategy separately.

Table 5.8 Performance attribution of the Combination Overlay portfolio

Source	**Contribution to CAGR**
ACWI	5.05%
Trend-Following	5.64%
Tail Risk Hedging	-1.38%
Residual (Reduction in Variance Drag)	0.76%
CAGR	**10.06%**

5.4 CHAPTER SUMMARY

In this chapter, the impact of overlaying both a trend-following strategy and a tail risk hedging strategy on a global equity portfolio represented by ACWI was examined. The chapter began by constructing the Combination Overlay portfolio, which consists of a 100% long position in ACWI, overlaid with a 50% allocation to the trend-following strategy and a full allocation to the tail risk hedging strategy, each of which was detailed in previous chapters. The trend-following strategy utilises an investment universe of 79 futures contracts across four key asset classes: global equities, government bonds, commodities, and currencies. The tail risk hedging strategy employs long positions in 10-delta, one-year to expiration SPX put options to provide downside protection.

Consistent with Schwalbach and Auret (2023), it was found that the Combination Overlay portfolio exhibited key characteristics that make it a compelling alternative to a traditional equity portfolio. These include empirically higher absolute returns, significantly improved risk-adjusted returns, and a meaningfully more favourable risk profile compared to both ACWI and the individual overlay portfolios (created by overlaying trend-following and tail risk hedging separately onto ACWI). Specifically, the Combination Overlay portfolio delivered a CAGR of 10.06%, outperforming ACWI's 5.05% and the Tail Overlay portfolio's 4.11%. While the Trend Overlay portfolio achieved a higher CAGR of 11.12%, the Combination Overlay portfolio offered a superior Sharpe ratio of +0.72, compared to +0.68 for the Trend Overlay portfolio and +0.32 for ACWI, indicating better risk-adjusted performance.

A key focus of the chapter was the analysis of convexity introduced by the Combination Overlay portfolio with respect to ACWI returns. Hypothesis $H_{0, I}$: Overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently was tested. In addition to meaningfully improved tail risk profile, historical drawdowns and the visual evidence quadratic regressions over multiple return frequencies, the regression analysis revealed a significantly larger positive coefficient on the squared ACWI term for the Combination Overlay portfolio compared to the individual overlay strategies, indicating a greater degree of convexity. These findings led to the rejection of the null hypothesis $H_{0, I}$, demonstrating that the combination of the two strategies provides enhanced convexity benefits over using either strategy alone.

Finally, the primary hypothesis of this thesis was evaluated, H_0: Compounded long-term absolute and risk-adjusted equity returns are not improved by overlaying a combination of a trend-following strategy and a tail risk hedging strategy on top of full exposure to global equity. The empirical evidence showed that the Combination Overlay portfolio not only improved absolute returns but also significantly enhanced risk-adjusted returns and tail risk metrics. The portfolio exhibited lower volatility, shallower drawdowns, reduced negative skewness, and improved VaR and CVaR measures compared to ACWI and the individual overlay strategies. These results led to an unequivocal rejection of the null hypothesis H_0, confirming that overlaying both strategies on a global equity portfolio has empirically enhanced long-term performance.

CHAPTER 6: CONCLUSION

6.1.1 Summary of thesis

Equities play a pivotal role in long-term investment portfolios, serving as the principal engine for returns over extended periods. However, equity returns are particularly vulnerable to the detrimental effects of variance drag, wherein excessively large volatility spikes during market crises reduce the base from which future returns compound, thereby eroding long-term performance. In this thesis, it is proposed that incorporating convexity into a global equity portfolio by overlaying strategies such as trend-following and tail risk hedging can mitigate variance drag and introduce an uncorrelated source of returns, thereby enhancing both the compounding potential and the long-term absolute and risk-adjusted performance of global equity portfolios.

Central to this approach is the utilisation of a portable alpha framework, which maintains full exposure to global equities while overlaying the selected strategies. This framework ensures that the long-term returns of the resulting overlay portfolios remain predominantly driven by global equity beta, as represented by the ACWI. Meanwhile, alpha is earned by improving the compounding profile of the portfolio through the introduction of convexity via the combination of tail risk hedging and trend-following strategies. It is found that this approach has the benefit of introducing significant portfolio diversification without the need for diluting equity exposure.

6.1.2 Trend-following

In Chapter 3, a trend-following strategy was constructed using an investment universe of 79 futures contracts across global equities, bonds, commodities, and currencies. The strategy employed an ensemble of simple long/short trend-following signals, with positions scaled based on ex-ante volatility estimates to ensure balanced risk contributions across assets. The findings provided robust evidence of trending behaviour across markets, leading to the rejection of $H_{0,\ A}$ in favour of the alternative hypothesis that markets do exhibit trending characteristics.

It was statistically validated that the trend-following portfolio used in this research accurately represents broader trend-following strategies. This validation was achieved by regressing the returns of the Trend Portfolio against those of an industry benchmark and subsequently rejecting the null hypothesis $H_{0,B}$ in favour of the alternative hypothesis $H_{1,B}$, which posits that the SG Trend Index returns have a significant relationship with the returns of the Trend Portfolio.

To confirm whether the returns of the trend-following portfolio could be explained by traditional equity market factors, $H_{0,C}$ which posits that trend-following returns are explained by traditional factors was tested. Regression analyses revealed near-zero and statistically insignificant loadings on factors such as the market factor, size (SMB) and value (HML), with a significant positive loading only on the momentum factor (UMD). After adjusting for these factors, the trend-following portfolio still generated a statistically significant monthly alpha of 0.85%, leading to the rejection of $H_{0,C}$. This significant alpha aligns with past research, which has found that trend-following strategies generate strong absolute returns that are uncorrelated with equity returns.

The convexity of trend-following returns with respect to equity returns was also evaluated by testing $H_{0,D}$: Trend-following returns are not convex with respect to global equity returns. The analysis showed that trend-following returns displayed a U-shaped payoff profile, indicating convexity in both the left and right tails of the equity return distribution. This finding was in line with Moskowitz et al. (2012).

A statistically significant positive coefficient on squared ACWI returns further supported the rejection of $H_{0,D}$ in favour of $H_{1,D}$ which states that trend-following returns are convex with respect to global equity returns. Following these conclusions, $H_{0,E}$ Overlaying a trend-following strategy on top of global equity does not enhance risk-adjusted and absolute returns compared to holding global equity alone was evaluated. It was found that overlaying a trend-following strategy on top of a fully collateralised global equity portfolio historically improved absolute and risk-adjusted returns, resulting in the rejection of $H_{0,E}$.

6.1.3 Tail risk hedging

In Chapter 4, tail risk hedging strategies were focused on, specifically using long-dated, far out-the-money SPX put options to hedge a global equity portfolio. First, $H_{0,F}$: SPX options are not suitable to hedge sudden and severe drawdowns in ACWI exposure was addressed. Empirical evidence demonstrated a strong and significant relationship between the returns and volatilities of SPX and ACWI, particularly at extremes, justifying the use of SPX options for hedging ACWI exposure and leading to the rejection of $H_{0,F}$. Hypothesis $H_{0,G}$: Tail risk hedging does not introduce left-tail convexity with respect to global equity returns when overlaid was then evaluated. The tail risk hedging strategy showed clear left-tail convexity with respect to equity returns, significantly mitigating losses during acute equity drawdowns.

This was also confirmed by both visual analyses and regression results showing a positive and significant coefficient on the squared ACWI returns, leading to the rejection of $H_{0,G}$. Further, hypothesis $H_{0,H}$: Overlaying a tail risk hedging strategy on top of global equity does not improve tail risk statistics and leads to a reduction in risk-adjusted returns compared to holding global equity alone was tested. Despite a consistent performance drag during normal market conditions, it was found that the Tail Overlay portfolio provided significant improvements in tail risk metrics without compromising the Sharpe ratio. The strategy exhibited lower volatility, shallower drawdowns, and reduced negative skewness compared to ACWI. Consequently, $H_{0,H}$ was rejected.

6.1.4 Combining trend-following and tail risk hedging

In Chapter 5, the combined effect of overlaying both trend-following and tail risk hedging strategies on a global equity portfolio was explored, forming the Combination Overlay portfolio. Hypothesis $H_{0,I}$: Overlaying a combination of a tail risk hedging strategy and a trend-following strategy on top of global equity does not introduce greater convexity than either strategy independently was tested.

The findings indicate that the Combination Overlay portfolio demonstrated superior performance in terms of absolute returns, risk-adjusted returns, and risk metrics compared to ACWI and the individual overlay strategies. The portfolio exhibited a higher Sharpe ratio, lower volatility, reduced negative skewness, and significantly improved tail risk metrics.

Regression analyses also showed a larger positive coefficient on the squared ACWI term, indicating greater convexity than either strategy alone. The additional benefit derived from reducing variance drag was quantified as 0.76% per annum. These findings led to the rejection of $H_{0,\ I.}$

The primary hypothesis and research question of this thesis, H_0: Compounded long-term absolute and risk-adjusted equity returns are not improved by overlaying a combination of a trend-following strategy and a tail risk hedging strategy on top of full exposure to global equity was evaluated. The empirical evidence from the analyses demonstrated that the Combination Overlay portfolio not only improved absolute returns but also significantly enhanced risk-adjusted returns and tail risk metrics. The portfolio successfully mitigated the adverse effects of variance drag by reducing the impact of extreme volatility and large drawdowns, thereby enhancing long-term compounded returns by 0.76% per annum. In addition, the Combination Overlay portfolio generated a large, positive, and statistically significant alpha of 0.38% per month after controlling for traditional equity factors, global government bond, and commodity excess returns. The primary hypothesis H_0 is rejected in favour of the alternative hypothesis that overlaying both strategies improves long-term absolute and risk-adjusted returns.

This thesis provides empirical support for the effectiveness of overlaying trend-following and tail risk hedging strategies onto a global equity portfolio to mitigate variance drag and enhance long-term performance. It was found that the combination of these strategies introduces convexity with respect to both left-tail and right-tail equity returns, offering protection during extreme market conditions without sacrificing upside potential over the long-term. By utilising a portable alpha framework, investors can efficiently implement these overlays without diluting existing equity exposure, maintaining full participation in the equity market's growth and enhancing compounding potential through a meaningfully improved risk profile.

6.1.5 Implications for equity investors

These findings hold significant implications for equity investors. It is demonstrated that investors aiming to maximise long-term wealth can benefit from incorporating convex strategies that overcome the limitations of traditional diversification methods. The combination of trend-following and tail risk hedging offers a robust approach to mitigating the erosive effects of variance drag on compounded returns. By applying a portable alpha framework and overlaying trend-following and tail risk hedging strategies on top of global equity exposure, it is shown that investors can enhance expected returns. This approach provides equity investors with an alternative to security selection as the primary method for attempting to outperform passive equity benchmarks.

6.1.6 Future research

Future research could further develop the portable alpha framework utilised in this study by incorporating additional hedge fund strategies that offer desirable diversification characteristics relative to global equity portfolios. Exploring alternative specifications of the strategies discussed in this thesis would provide a deeper understanding of their potential applications and effectiveness. Investigating the behavioural factors that influence investor adoption of such overlay strategies could yield valuable insights into portfolio construction and risk management practices.

Future studies could examine the impact of varying allocation sizes to the trend-following and tail risk hedging overlays, to determine the optimal balance that maximises returns while minimising risk. Allocations could also be made dynamically, based on a set of criteria. Research could also explore the effectiveness of these overlay strategies in different market conditions or across different asset classes, such as emerging markets or fixed income portfolios. Another avenue for future research is the assessment of transaction costs and implementation challenges associated with these strategies, to evaluate their practicality.

Future research could explore the applicability of the combined tail risk and trend-following strategy in less liquid markets, such as South Africa. Specifically, the application of a tail risk hedging overlay using liquid options on local broad equity indices, such as the Top 40 Index warrants investigation and comparison with the use of SPX options.

Alternative approaches to constructing trend-following portfolios warrant exploration. One potential avenue involves testing for trend effects by examining whether groups of autocorrelation coefficients are jointly significantly different from zero. Given that different strategy specifications can yield varying results, this introduces the potential for dispersion in performance. Future research could focus on evaluating strategies designed to replicate trend-following industry benchmarks, utilising robust regression techniques to minimise the risk of deviations from benchmark performance.

Future research could investigate the use of market timing signals to dynamically adjust allocations between the trend-following and tail risk hedging strategies. In particular, changes in implied volatility surfaces could be explored as potential indicators for regime shifts, guiding when to emphasise one strategy over the other.

LIST OF REFERENCES

Alquist, R., Frazzini, A., Ilmanen, A., & Pedersen, L. H. (2020). Fact and fiction about low-risk investing. *Journal of Portfolio Management, 46*(6), 72-92.

Anand, A., Irvine, P., Puckett, A., & Venkataraman, K. (2013). Institutional trading and stock resiliency: Evidence from the 2007–2009 financial crisis. *Journal of Financial Economics, 108*(3), 773-797.

AQR Capital Management. (2024, July 30). *Portable alpha: Still a great solution for improving return outcomes*. Retrieved from AQR: https://www.aqr.com/Insights/Research/White-Papers/Portable-Alpha-Still-A-Great-Solution-For-Improving-Return-Outcomes

Asness, C. S. (2004). An alternative future. *The Journal of Portfolio Management, 30*(5), 94-103.

Asness, C. S., Frazzini, A., & Pedersen, L. H. (2012). Leverage aversion and risk parity. *Financial Analysts Journal, 68*(1), 47-59.

Asvanunt, A., Nielsen, L. N., & Villalon, D. (2015). Working Your Tail Off: Active Strategies. *The Journal of Investing*, *24*(2), 134-145.

Babu, A., Hoffman, B., Levine, A., Ooi, Y. H., Schroeder, S., & Stamelos, E. (2020). You Can't Always Trend When You Want. *The Journal of Portfolio Management*, *46*(4), 52-68.

Babu, A., Levine, A., Ooi, Y. H., Pedersen, L. H., & Stamelos, E. (2019, May 15). *Trends everywhere.* Retrieved from AQR: https://www.aqr.com/Insights/Research/Working-Paper/Trends-Everywhere

Baker, H. K., & Filbeck, G. (2013). *Portfolio Theory and Management.* Oxford: Oxford University Press.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *The Journal of finance, 52*(5), 2003-2049.

Baltas, N., & Kosowski, R. (2013). Momentum Strategies in Futures Markets and Trend Following Funds. *December 2012 Finance Meeting EUROFIDAI-AFFI Paper*. Paris.

Bates, D. S. (1991). The crash of '87: was it expected? The evidence from options markets. *The Journal of Finance, 46*(3), 1009-1044.

Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of econometrics, 94*(1-2), 181-238.

Beck, N., Hsu, J., Kalesnik, V., & Kostka, H. (2016). Will your factor deliver? An examination of factor robustness and implementation costs. *Financial Analysts Journal, 72*(5), 55-82.

Berkowitz, J. P., & DeLisle, R. J. (2018). Volatility as an asset class: Holding VIX in a portfolio. *The Journal of Alternative Investments, 21*(2), 52-64.

Bhabra, G. S., Gonzalez, L., Kim, M. S., & Powell, J. G. (2001). Volatility prediction during prolonged crises: evidence from Korean index options. *Pacific-Basin Finance Journal, 9*(2), 147-164.

Bhansali, V. (2008). Tail Risk Management. *The Journal of Portfolio Management, 34*(4), 68-75.

Bhansali, V. (2013). The Four Modes of Practical Risk Management. *Journal of Portfolio Management, 39*(2), 3-4.

Bhansali, V. (2014). *Tail Risk Hedging: Creating Robust Portfolios for Volatile Markets.* McGraw Hill Professional.

Bhansali, V. (2015a). Tail-Risk Management for Retirement Investments. *The Journal of Retirement, 2*(3), 78-86.

Bhansali, V. (2015b). Behavioral Perspectives on Tail-Risk Hedging. *The Journal of Investing, 24*(2), 122-133.

Bhansali, V. (2018). Right Tail Hedging: Managing Risk When Markets Melt Up. *The Journal of Portfolio Management, 44*(7), 55-62.

Bhansali, V., & Davis, J. (2010a, March 24). *Offensive Risk Management: Can Tail Risk Hedging Be Profitable?* Retrieved from SSRN: https://ssrn.com/abstract=1573760

Bhansali, V., & Davis, J. (2010b). Offensive Risk Management II: The Case for Active Tail Hedging. *The Journal of Portfolio Management, 37*(1), 78-91.

Bhansali, V., & Harris, L. (2018). Everybody's Doing It: Short Volatility Strategies and Shadow Financial Insurers. *Financial Analysts Journal, 74*(2), 12-22.

Birru, J., & Figlewski, S. (2012). Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008. *Journal of Financial Markets, 15*(2), 151-180.

Black, F. (1972). Capital market equilibrium with restricted borrowing. *The Journal of Business, 45*(3), 444-455.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of political economy, 81*(3), 637-654.

Black, F., Jensen, M. C., & Scholes, M. S. (1972). *The Capital Asset Pricing Model: Some Empirical Tests.* New York: Praeger.

Blitz, D., van der Grient, B., & van Vliet, P. (2010). Fundamental indexation: Rebalancing assumptions and performance. *The Journal of Index Investing, 1*(2), 82-88.

Bouchey, P., Nemtchinov, V., Paulsen, A., & Stein, D. M. (2012). Volatility harvesting: Why does diversifying and rebalancing create portfolio growth? *The Journal of Wealth Management, 15*(2), 26-35.

Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 621-651.

Brière, M., Burgues, A., & Signori, O. (2010). Volatility exposure for strategic asset allocation. *The Journal of Portfolio Management, 36*(3), 105-116.

Brooks, J. (2017, August 10). *A Half Century of Macro Momentum.* Retrieved from AQR: https://www.aqr.com/Insights/Research/White-Papers/A-Half-Century-of-Macro-Momentum

Brush, J. S. (1997). Comparisons and combinations of long and long/short strategies. *Financial Analysts Journal*, *53*(3), 81-89.

Butler, A., & Butler, A. (2023). *Peering around Corners: How to Replicate Trend-Following Managed Futures.* Retrieved from Resolve Asset Management: https://investresolve.com/

Carhart, M. (1997). On persistence in Mutual Fund Performance. *Journal of Financial Economics*, *52*(1), 57-82.

Cavaglia, S., Fan, J. H., & Wang, Z. (2022). Portable beta and total portfolio management. *Financial Analysts Journal, 78*(3), 49-69.

Chekhlov, A., Uryasev, S., & Zabarankin, M. (2005). Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance, 8*(01), 13-58.

Clare, A., Seaton, J., Smith, P. N., & Thomas, S. (2016). The trend is our friend: Risk parity, momentum and trend following in global asset allocation. *Journal of Behavioral and Experimental Finance, 9*, 63-80.

Clewlow, L., & Hodges, S. (1997). Optimal delta-hedging under transactions costs. *Journal of Economic Dynamics and Control, 21*(8-9), 1353-1376.

Daigler, R. T., & Rossi, L. (2006). A portfolio of stocks and volatility. *The Journal of Investing, 15*(2), 99-106.

Daniel, K. D., & Moskowitz, T. J. (2016). Momentum Crashes. *Journal of Financial Economics, 122*(2), 221-247.

Darius, D., Ilhan, A., Mulvey, J., Simsek, K. D., & Sircar, R. (2002). Trend-following hedge funds and multi-period asset allocation. *Quantitative Finance, 2*(5), 354-361.

Dash, S., & Moran, M. T. (2005). VIX as a companion for hedge fund portfolios. *The Journal of Alternative Investments, 8*(3), 75-80.

Davis, J., Aliaga-Diaz, R. A., Westaway, P., Wang, Q., Patterson, A. J., Ahluwalia, H., . . . Lemco, J. (2018, December 6). *Vanguard economic and market outlook for 2019: Down but not out.* Retrieved from Vanguard: https://pressroom.vanguard.com/nonindexed/Research-Vanguard-Economic-and-Market-Outlook-2019-120618.pdf

Demeterfi, K., Derman, E., Kamal, M., & Zou, J. (1999). A guide to volatility and variance swaps. *The Journal of Derivatives, 6*(4), 9-32.

Dennis, P., & Mayhew, S. (2002). Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis, 37*(3), 471-493.

Doran, J. S., Peterson, D. R., & Tarrant, B. C. (2007). Is there information in the volatility skew? *Journal of Futures Markets: Futures, Options, and Other Derivative Products, 27*(10), 921-959.

Eifert, B. (2020, September 30). *Common Hedging Discussions.* Retrieved from QVR Advisors: https://www.qvradvisors.com/research

Faber, M. (2015). *Global Asset Allocation: a survey of the world's top asset allocation strategies.* El Segundo: The Idea Farm LP. Retrieved from Meb Faber.

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance, 47*(2), 427-465.

Fernandes , M., Medeiros, M. C., & Scharth, M. (2014). Modeling and predicting the CBOE market volatility index. *Journal of Banking & Finance, 40*, 1-10.

Figlewski, S., Chidambaran, N., & Kaplan, S. (1993). Evaluating the Performance of the Protective Put Strategy. *Financial Analysts Journal, 49*(1), 46-56.

French, K. R. (2024, June 30). *Current Research Returns*. Retrieved from Dartmouth College: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

Fung, W., & Hsieh, D. A. (2001). The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers. *The Review of Financial Studies, 14*(2), 313-341.

Goldwhite, P. (2009). Diversification and risk management: What volatility tells us. *The Journal of Investing, 18*(3), 40-48.

Hafner, R., & Wallmeier, M. (2007). Volatility as an asset class: European evidence. *The European Journal of Finance, 13*(7), 621-644.

Hafner, R., & Wallmeier, M. (2008). Optimal investments in volatility. *Financial Markets and Portfolio Management, 22*(2), 147-167.

Harvey, C. R., Rattray, S., & Van Hemert, O. (2021). *Strategic Risk Management: Designing Portfolios and Managing Risk.* John Wiley & Sons.

Herbst-Bayliss, S., & Delevingne, L. (2020, March 18). *Up 3,000%: The tail risk funds that mastered coronavirus market mayhem*. Retrieved from Reuters: https://www.reuters.com/article/us-health-coronavirus-ppestocks/did-hedge-funds-score-on-masks-and-sanitizer-not-so-much-idUSKBN22W0IU/?il=0

Hobson, D., & Klimmek, M. (2012). Model-independent hedging strategies for variance swaps. *Finance and Stochastics*, *16*(4), 611-649.

Hoffstein, C., Sibears, D., & Faber, N. (2019). Rebalance timing luck: The difference between hired and fired. *Journal of Beta Investment Strategies*, *10*(1), 27-36.

Hong, J., & Satchell, S. (2015). Time series momentum trading strategy and autocorrelation amplification. *Quantitative Finance, 15*(9), 1471-1487.

Hsu, S. D., & Murray, B. M. (2007). On the volatility of volatility. *Physica A: Statistical Mechanics and its Applications*, 366-376.

Hull, J. C. (2003). *Options, futures and other derivatives.* Pearson Education India.

Hull, J., & White, A. (2017). Optimal delta hedging for options. *Journal of Banking & Finance, 82*, 180-190.

Hurst, B., Ooi, Y. H., & Pedersen, L. H. (2017). A century of evidence on trend-following investing. *The Journal of Portfolio Management, 44*(1), 15-29.

Ilmanen, A. (2012). Do financial markets reward buying or selling insurance and lottery tickets? *Financial Analysts Journal, 68*(5), 26-36.

Ilmanen, A., Thaper, A., Tummala, H., & Villalon, D. (2021). Tail risk hedging: Contrasting put and trend strategies. *Journal of Systematic Investing, 1*(1), 111-124.

Israelov, R. (2018). Pathetic Protection: The Elusive Benefits of Protective Puts. *The Journal of Alternative Investments, 21*(3), 6-33.

Israelov, R., & Nielsen, L. N. (2015). Still not cheap: Portfolio protection in calm markets. *The Journal of Portfolio Management, 41*(4), 108-120.

Israelov, R., Nielsen, L. N., & Villalon, D. (2016). Embracing Downside Risk. *The Journal of Alternative Investments*, *19*(3), 59-67.

Jarrow, R. A., & Turnbull, S. M. (1994). Delta, gamma and bucket hedging of interest rate derivatives. *Applied Mathematical Finance, 1*(1), 21-48.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance, 48*(1), 65-91.

Kapadia, N., & Szado, E. (2012). Fifteen Years of the Russell 2000 Buy-Write. *The Journal of Investing, 21*(4), 59-80.

Konstantinidi, E., & Skiadopoulos, G. (2016). How does the market variance risk premium vary over time? Evidence from S&P 500 variance swap investment returns. *Journal of Banking & Finance, 62*, 62-75.

Krishnan, H. P. (2017). *The Second Leg Down: Strategies for Profiting After a Market Sell-off.* John Wiley & Sons.

Kulp, A., Djupsjöbacka, D., & Estlander, M. (2005). Managed Futures and Long Volatility. *AIMA Journal*, 27-28.

Lintner, J. (1965). Security prices, risk, and maximal gains from diversification. *The Journal of Finance, 20*(4), 587-615.

Maeso, J. M., & Martellini, L. (2020). Measuring portfolio rebalancing benefits in equity markets. *Journal of Portfolio Management*, *46*(4), 94-109.

Malek, M. H., & Dobrovolsky, S. (2009). Volatility Exposure of CTA Programs and Other Hedge Fund Strategies. *The Journal of Alternative Investments*, *11*(4), 68-89.

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

Marshall, B. R., Nguyen, N. H., & Visaltanachoti, N. (2017). Time series momentum and moving average trading rules. *Quantitative Finance, 17*(3), 405-421.

Mastinšek, M. (2006). Discrete–time delta hedging and the Black–Scholes model with transaction costs. *Mathematical Methods of Operations Research, 64*(2), 227-236.

McQuinn, N., Thapar, A., & Villalon, D. (2021). Portfolio protection? It’s a long (term) story... *The Journal of Portfolio, 47*(1), 35-50.

Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum. *Journal of financial economics, 104*(2), 228-250.

Neely, C., Rapach, D., Tu, J., & Zhou, G. (2014). Forecasting the equity risk premium: The role of technical indicators. *Management Science, 60*, 1772-1791.

Niederhoffer, R., & Weddepoel, C. (2014, April). *CTAs and Rising Interest Rates: Is the Party Over?* Retrieved from Daily Speculations: https://dailyspeculations.com/CTAsAndRisingInterestRates_040714.pdf

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of financial economics, 63*(1), 3-50.

Petersen, M. A. (2008). Estimating standard errors in finance panel data sets: Comparing approaches. *The Review of financial studies, 22*(1), 435-480.

Prondzinski, D., & Miller, M. (2018). Active Versus Passive Investing: Evidence From The 2009-2017 Market. *Journal of Accounting and Finance, 18*(8), 119-143.

Qian, E. (2011). Risk parity and diversification. *Journal of Investing, 20*, 119-127.

Raju, S. (2012). Delta gamma hedging and the black-scholes partial differential equation (pde). *Journal of Economics and Finance Education, 11*(2), 51-62.

Schwalbach, B., & Auret, C. (2023). Leveraging defence into offence: enhancing absolute and risk-adjusted equity returns with tail risk management overlays. *Investment Analysts Journal, 52*(3), 246-258.

Schwalbach, J. M., & McClelland, D. (2018). An analysis of short put strategies and their role in asset allocation. *Investment Analysts Journal, 47*(3), 272-283.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance, 19*(3), 425-442.

Shepherd, S., Ko, A., & Kunz, B. (2018, September). *Alternative risk premia: valuable benefits for traditional portfolios.* Retrieved from Research Affiliates: https://www.researchaffiliates.com/en_us/publications/articles/683-alternative-risk-premia-valuable-benefits-for-traditional-portfolios.html

Stein, D. M., Nemtchinov, V., & Pittman, S. (2009). Diversifying and rebalancing emerging market countries. *The Journal of Wealth Management, 11*(4), 79-88.

Szado, E. (2009). VIX futures and options: A case study of portfolio diversification during the 2008 financial crisis. *The Journal of Alternative Investments, 12*(2), 68-85.

Szado, E., & Kazemi, H. B. (2009). Collaring the Cube: Protection Options for a QQQ ETF Portfolio. *The Journal of Alternative Investments, 11*(4), 24-42.

Szado, E., & Schneeweis, T. (2010). Loosening your collar: alternative implementations of QQQ collars. *The Journal of Trading, 5*(2), 35-56.

Thapar, A., Nielsen, L., & Villalon, D. (2019, November 19). *Chasing Your Own Tail (Risk), Revisited.* Retrieved from AQR: https://www.aqr.com/Insights/Research/White-

Papers/Chasing-Your-Own-Tail-Risk-Revisited#:~:text=One%20of%20the%20most%20important,Risk)%E2%80%9D%20was%20our%20response.

Tian, Y. S. (2011). Extracting risk-neutral density and its moments from American option prices. *The Journal of Derivatives, 18*(3), 17-34.

Till, H. (2016). What Are the Sources of Return for CTAs and Commodity Indexes? A Brief Survey of Relevant Research. *The Journal of Wealth Management, 18*(4), 108-123.

Vähämaa, S. (2004). Delta hedging with the smile. *Financial Markets and Portfolio Management, 18*(3), 241-255.

Warren, G. J. (2012). Can investing in volatility help meet your portfolio objectives? *The Journal of Portfolio Management, 38*(2), 82-98.

Whaley, R. E. (2013). Trading volatility: At what cost? *The Journal of Portfolio Management, 40*(1), 95-108.

Zakamulin, V., & Giner, J. (2020). Trend following with momentum versus moving averages: A tale of differences. *Quantitative Finance, 20*(6), 985-1007.