Capturing the Black Swan: Scenario-Based Asset Allocation with Fat Tails and Non-Linear Correlations

Vsevolod Gorlach

ERSA working paper 695

August 2017

Economic Research Southern Africa (ERSA) is a research programme funded by the National Treasury of South Africa.

Capturing the Black Swan: Scenario-Based Asset Allocation with Fat Tails and Non-Linear Correlations

Dr Vsevolod I. Gorlach, CFA*†

August 1, 2017

Abstract

This paper highlights the shortfalls of Modern Portfolio Theory (MPT). Amongst other flaws, MPT assumes that returns are normally distributed; that correlations are linear; and that risks are symmetrical. We propose a dynamic and flexible scenario-based approach to portfolio selection that incorporates an investor's economic forecast. Extreme Value Theory (EVT) is used to capture the skewness and kurtosis inherent in asset-class returns; and it also accounts for the volatility clustering and the extreme co-movements across asset classes. The estimation consists of using an asymmetric GJR-GARCH model to extract the filtered residuals for each asset-class return. Subsequently, a marginal cumulative distribution function (CDF) of each asset class is constructed by using a Gaussian-kernel estimation for the interior, together with a generalised Pareto distribution (GPD) for the upper and lower tails. The distribution of exceedance method is applied to find residuals in the tails. A Student's t copula is then fitted to the data; and then used to induce correlation between the simulated residuals of each asset class. A Monte Carlo technique is applied to simulate standardised residuals, which represent a univariate stochastic process when viewed in isolation; but it maintains the correlation induced by the copula. The results are mean-CVaR optimised portfolios, which are derived based on an investor's forward-looking expectation.

Keywords: Portfolio optimisation, Scenario-based, Non-normal distribution, Fat-tails, Non-linear correlation, Extreme Value Theory, Marginal Distribution Modelling, Copula, Simulation, Conditional Value-at-Risk

JEL Classifications: C14, C58, C61, G11 and G17

*Derivative Structuring, Absa Capital. E-mail: vsevolodgorlach@gmail.com
†The author would like to thank Roland Rousseau for his guidance in writing this paper.

1 Background

> *"But he does not wear any clothes"* said the little child in Hans Christian Andersen's "The Emperor's New Clothes"

In 1952 Harry Markowitz published a paper titled: 'Portfolio Selection' which over the next 20 years evolved to become known as Modern Portfolio Theory (MPT) despite Markowitz criticizing that "*there was nothing modern about it*" (Markowitz, 1952).

MPT made investors aware of the benefits of aggregate portfolio effects such as diversification, risk and correlation. Investors realised the need to optimise portfolios; it made sense to be on the 'efficient frontier' – where the maximum returns per unit of risk are achieved; and the Sharpe ratios are maximised.

However, nearly 70 years have elapsed since the concept of MPT first being brought to light; and there have been numerous actual events that MPT cannot explain: The Crash of 1987 (Black Monday), the 1998 EmergingMarket Crises – the Russian Rouble default, the Japanese property bubble in 1991, the Black Wednesday Sterling Crisis (1992), the demise of Barings Bank in 1995, the Asian Financial Crises (1997), the dotcom bubble (2001); and more recently the 2008 Financial Crisis.

So why has MPT not been able to cater for these 'once-in-a-lifetime' extreme events? The answer lies in the fact that the 'efficient frontier', which uses the concept of mean-variance optimisation, is not altogether efficient. The key concepts used to measure risk such as standard deviation, beta, correlations and Sharpe ratios are not just impractical but seriously misleading.

For example, MPT suffers from a number of flaws: it assumes that returns are normally distributed; it requires only three basic inputs: mean, standard deviation and correlation It assumes that correlations are linear; and it views risks as symmetrical; and consequently, it cannot account for fat tails, i.e. extreme observations

MPT is backwardlooking; i.e. it is able to determine what your optimal asset allocation would have been when given the realised returns; however this does not mean that the same allocation will be optimal going forward. The 'optimal' portfolios along the efficient frontier are extremely sensitive to small changes in the input expectations as well as to any correlations amongst the asset classes.

History has shown that MPT not only underestimates the magnitude of such shocks; but it also misjudges the frequency of large unexpected 'blackswan' or 'fattail' events. Consequently the theoretical assumption underlying the symmetrical bell-curve or Gaussian normal distribution to approximate assetclass returns is seriously flawed

Defining risk in terms of standard deviation or 'volatility' instead of 'tail-risk' or 'downside-risk' is also flawed; investors experience 'asymmetrical risk preferences'; where a loss hurts more than the pleasure from an equivalent gain. Finally, correlations are erratic or 'regime-dependent' through time – with risk premiums varying – as investors cycle between risk aversion and risk adoration (Xiong and Idzorek, 2011).

The shortcomings of MPT have been highlighted. These shortcomings should be seen in the light of the fact that investors are shifting to a more flexible approach that is capable of capturing the dynamics in risk and return expectations across an array of asset classes Li and Sullivan (2011) suggest that the traditional strategic approach of fixed-asset allocation is outdated.

The challenge for investors goes beyond merely selecting uncorrelated assets for inclusion in a portfolio that constitutes significant economic exposure; and then specifying a fixed portion of each asset class until optimal economic exposure is achieved (Sullivan, Peterson and Waltenbaugh, 2010)

This paper proposes a dynamic and flexible scenario-based simulation approach to asset allocation This approach incorporates an investor's specific economic forecast as well as the uncertainty surrounding these forecasts. This forward-looking approach relies less on historical optimisation or sampling; and therefore, it can overcome regime shifts or structural breaks inherent in historical data; since these are frequently irrelevant, when going forward.

In addition, through the simulation technique proposed, one is able to capture risks that are statistically more likely to happen than the 'once-off' limited history that is observable. In other words, we are able to define and model events that might never have happened historically; but these are risks that could statistically happen.

History tells us what has happened; scenarios tell us what is likely to happen; and simulations tell us what can happen – even if one has not yet witnessed such an event. Effective asset allocation requires all three of these, in order to succeed.

In summary, this paper proposes a model of portfolio selection with extreme tails and non-linear correlations. The model takes into account skewness and kurtosis, as well as the third and fourth moments of the return distribution – beyond mean and variance. The persistence in volatility, or volatility clustering and correlations of risky asset returns, which tend to increase during times of market turbulence or return dependence is also accounted for (Wang, Sullivan and Ge, 2011).

The non-linear model is thus more dynamic and less restrictive than the traditional static methods that depend on the returns obtained when following a Gaussian process. The assetallocation framework provides a practical application by determining an investor's optimal portfolio in accordance with their explicit view of both the local and the global market environment and conditions.

2 Methodology and Model Specification

The following section highlights the steps involved in the scenario-based assetallocation process. The estimation technique used is consistent with those applied by Wang *et al.* (2011); and it allows the modelling of the market risk of a portfolio containing different asset classes – with a Monte Carlo simulation technique when using a Student's t copula and EVT.

In summary, the estimation methodology consists of the following steps: Firstly the filtered residuals for each asset class return series are extracted by using an asymmetric GJR-GARCH model. Subsequently, the marginal cumulative distribution function (CDF) of each asset class is constructed whereby a Gaussian kernel estimate is used for the interior or 'rump' and a general Pareto distribution (GPD) estimate is used for the upper and lower tails.

A Student's t copula is then fitted to the data; and it is used to induce correlation between the simulated residuals of each asset class. The penultimate step is the scenario-matching piece which identifies and filters out sets of returns which match the investor's unique economic expectations in going forward. Finally, the efficient frontier is constructed by optimising for the highest return per unit of CVaR.

2.1 Extreme Value Theory

In finance, risk can be defined as the variability of returns around the long-run average return. The term 'standard deviation' defines risk or 'volatility' symmetrically – both above and below the mean return. Therefore, it assumes that investors are equally greedy or fearful of the same deviation above or below the long-run average.

However, behavioural finance states that due to the loss aversion-bias exhibited in investors; these investors place more 'value' on losses than on a gain of the same magnitude. Consequently, the asymmetric appetite for risk indicates that investors choose between two portfolios with the same volatility, the portfolio with the least 'downside risk' or the risk of loss.

Consequently, the left-hand part of the asymmetric portfolio return distribution, the 'tail', is what investors are most concerned about. Any portfolio with a long left tail is deemed riskier, regardless of the standard deviation.

Figure 1 illustrates the FTSE/JSE All Share Index (ALSI) monthly returns since 1960. The average return is 1.14% per month and the standard deviation is 569%. If one considers an extreme event as a scenario in which the index experiences a 3-time standard deviation loss this will equate to approximately a -15% loss in a single month. According to the normal distribution, this magnitude loss is extremely rare; and it should only occur 0.13% of the time on average.

With 675 monthly observations since 1960, one would expect to observe a monthly loss greater than three standard deviations – less than once on average. However, Figure 1 shows that the ALSI has experienced 7 occurrences of three standard deviation losses since 196. Therefore, these 'once-in-a-lifetime' events occurred nearly 10 times more frequently than the normal distribution would expect

Figure 2 reiterates the misfit of the normal distribution to the asset class returns; where the monthly All Bond Index (ALBI) returns are fitted with the Gaussian normal distribution. The ALBI returns are much more 'peaked' with more observations falling in the centre than the normal distribution would predict.

On the left, the normal distribution assigns zero probability to the tail-event; and thus, it does not account for extreme observations. This is the actual return when the index lost over 15% in August 1998 during the Emerging Market Crisis when the Russian currency defaulted.

With such a poor fit, the normal distribution underestimates events around the mean and in the tail which makes risk modelling highly inaccurate. In order to overcome this problem, Extreme Value Theory (EVT) can be applied thereby allowing one to model two additional descriptive parameters (Embrechts, Kluppelberg and Mikosch, 1997).

These parameters measure firstly, the degree of asymmetry i.e. the skewness which accounts for the fat-tails; and secondly, the degree of peakedness in the centre of the distribution i.e. the kurtosis inherent in asset class returns. Simultaneously, EVT accounts for the volatility clustering and extreme co-movements across various asset classes (Sullivan, et al. 2010).

Table 1 provides an overview of the five asset classes included in the analysis and the respective summary statistics. All asset class data was obtained from the JSE and are represented by indices in the following way: equities by the FTSE/JSE All Share Index (ALSI); bonds by the AllBond Index (ALBI); cash by SteFi (Short Term Fixed Interest composite index) (CASH); offshore equities by MSCI World Index in USD which is converted to ZAR in order to represent a South African investors perspective (MSCI) and property according to the JSE SAPI Listed Property Index (PROP)

The descriptive statistics are based on the monthly JSE data from January 1960 to March 2016. Readers can note the negative skewness in all of the asset classes with the exception of cash and property; and the excess kurtosis in both equities and particularly in the bond index.

2.2 Conditional Value at Risk

Having established that downside risk or 'tail-risk' is a more relevant measure of an investor's risk tolerance than the standard deviation for riskmanagement purposes; one needs models to focus exclusively on the left-hand side of the distributions. One such measure is the Value-at-Risk (VaR) approach.

VaR describes the minimum expected loss given a confidence level. For example, -15% for a given confidence level of 95% VaR implies that one can expect a loss of least 15% once every 20 months (i.e. 5% of the time). VaR is therefore a better proxy of risk than standard deviation; since it measure the minimum loss expected for a given 'worst-case' scenario that occurs with a present frequency (i.e. minimum expected loss in the worst 5% of events).

However, VaR can be a highly misleading concept. Although VaR is a marginal improvement on the concept of traditional volatility; since it focuses on the risks that matter to investors (minimum expected loss for a given likelihood), it is not enough to simply know the minimum expected loss. In other words VaR does not provide information on how big the magnitude of loss can be once a tail-event occurs.

Therefore, in order to measure the size of the expected loss, given that i.e. a 5% VaR event occurs, Conditional Value-at-Risk (CVar) is used. Figure 3 illustrates the concept of CVaR. According to Uryasev (2000), CVaR measures the expected loss during a given period at a certain confidence level. As a better alternative to VaR, it incorporates both the possibility and the expected magnitude of such a loss.

Thus, CVaR is known as the mean excess loss for continuous distributions; and it is defined as the weighted average of VaR and losses strictly exceeding VaR for discrete distributions (Rockafellar and Uryasev, 2000).

2.3 Return Filtering

According to Wang et al. (2011), in order to model the tails of a distribution with a GPD, when using EVT; the observations need to be approximately independent and identically distributed (i.i.d.). However, most financial returns exhibit some degree of autocorrelations and more importantly, heteroskedasticity.

Figure 4 plots the sample autocorrelation function of the monthly log returns as well as the squared log returns associated with ALSI. The issue of heteroskedasticity is evident when one considers the squared returns. Therefore, given the degree of persistence in variance; the return series first need to be filtered, in order to produce approximately i.i.d. observations.

To produce a series of i.i.d. observations, we fit a firstorder autoregressive model AR(1) to the conditional mean of the monthly returns for an asset class using equation (1):

$$r_t = c + \theta r_{t-1} \varepsilon_t \tag{1}$$

And an asymmetric GJR-GARCH(1,1) to the conditional variance using equation (2):

$$\sigma_t^2 = k + \alpha_{t-1}^1 + \phi \varepsilon_{t-1}^2 + \psi[\varepsilon_{t-1} < 0]\varepsilon_{t-1}^2 \tag{2}$$

The model addresses the leverage effect where a negative association has been observed to exist between shocks to asset returns and future volatility (Black, 1972). In particular, the last term of equation (2) incorporates asymmetry into the variance through the use of a Boolean indicator that takes the value of 1 if the prior model's residual is negative and 0 otherwise (Glosten, Jagannathan and Runkel, 1993).

Finally, the residuals for each asset class are standardised and modelled by using the standardised Student's t-distribution in order to compensate for the fat tails commonly found in the distribution of most asset class returns using equation (3):

$$z_t = \varepsilon_t / \sigma_t \qquad i.i.d \sim t_v \tag{3}$$

The results from the process described above are illustrated in Figure 5. The figure plots the autocorrelation of both the standardised residuals as well as the squared standardised residuals for the ALSI. It is evident that the filtering process results in approximately i.i.d. observations and thus volatility clustering has been eliminated.

Therefore, the resulting standardised residual returns approximate a zero-mean, unit-variance i.i.d. series which allows for EVT estimation of the tails from the sample CDF (Nystrom and Skoglund, 2002).

2.4 Marginal Distribution Modelling

Given the standardised, i.i.d. residuals, we now estimate the marginal semi-parametric empirical CDF of each asset class. Since EVT allows only for estimation of the tails of the distribution, we combine these tail distributions with a model for the internal part of the distribution.

The non-parametric Gaussian kernel CDF is used to estimate the interior of the distribution, where most of the data are found; and for which it is well suited. However, these perform poorly when applied to the tails of the distribution. Therefore, to better estimate the tails of the distribution, apply EVT to those residuals that fall in each tail (McNeil and Frey, 2000).

In order to find the upper and lower threshold of the residuals in each tail, the peaks-over-threshold; or the distribution of exceedance method is used which is consistent with the findings of McNeil and Saladin, 1997. Specifically, 10% of the residuals is reserved for the left-hand and right-hand tails. The result is a partition of the standardised residuals into three sections: the lower tail; the interior; and the upper tail.

Figure 6 illustrates the results of this procedure for the ALSI. The resulting piecewise distribution object allows interpolation within the interior CDF and extrapolation (function evaluation) in each tail. Extrapolation is beneficial; as this allows for estimation outside the historical data which is critical in risk management application (Hyung and de Vries, 2007)

As previously mentioned, the interior is estimated with a non-parametric Gaussian kernel CDF; and the extreme residuals, or tails beyond the 10% threshold, are fitted using EVT. Therefore, a generalised Pareto distribution (GPD), estimated by maximum likelihood, is used. The CDF of GDP is parameterised by using equation (4):

$$F(y) = 1 - (1 + \zeta y/\beta)^{-1/\zeta} \tag{4}$$
$$\text{Where: } y \text{ exceedances} \geq 0$$
$$\beta \text{ tail index parameter } > 0$$

Figure 7 shows a visual representation of the upper and lower tails of the return distribution of ALSI. The figures illustrate that the GPD fits the historical, or empirical, return distribution much more accurately than the normal CDF. Therefore, the GPD model allows for a much more accurate representation of the reality of fat tails (Hilal, Poon and Tawn, 2011).

2.5 Extreme Dependance Modelling

Having accounted for the fat-tailed conditional distribution of returns; now we address another important element in risk modelling: How asset class returns move together under extreme circumstances. During periods of high market volatility, asset classes exhibit a higher degree of co-dependency; and the probability of joint negative returns increases (Zeevi and Marshal, 2002).

During periods of high market stress, risky asset returns across asset classes abruptly decline in unison; and correlations increase significantly. Therefore, correlations are not static; there are correlations that are relevant under normal market conditions (i.e. close to the mean) and different correlations for the same asset classes when there is an extreme event. Linear correlations are not capable of modelling this type of behaviour.

By way of example, Figure 8 plots the jointmarginal distribution of monthly returns of the ALSI and ALBI. As the figure illustrates, both indices have realised a substantial simultaneous loss event of four standard deviations or more. This indicates that if there is as systematic event or a macro shock in financial markets, correlation suddenly increases; and there is a tendency towards unity.

Empirically, not only do individual asset classes have fatter tails than the normal distribution is able to account for; but the combination of asset classes also exhibits a higher incidence of joint negative returns in times of market stress (Wang *et al.* 2011).

According to Focardi and Fabozzi (2004); in order to account for the incidence of returns suddenly moving in unison, copula theory can be employed This accommodates interrelated and extreme dependencies of returns. Copulas allow for the modelling of fat tails – even when asset class returns present a high degree of co-movement

A t copula is used; since this particular copula better captures the effects of fat-tails in asset class returns; and it does so by allocating non-zero probabilities to those observations which may occur outside the normal range of historical returns (Roncalli, Durrleman and Nikeghbali, 2000). Consequently, by adjusting the copula's degree-of-freedom parameter, it becomes possible to extrapolate the multivariate fat-tailed distribution; so that it is consistent with the observed empirical data (Embrechts, McNeil and Staumann, 1999).

The procedure to calibrate the t copula is as follows: Firstly, the standardised residuals are transformed to uniform variables by the semi-parametric CDF derived in the previous steps. Secondly, the t copula is fitted to the transformed data. The calibration method; the uniform variables being transformed by the empirical CDF of each margin is known as the conicalmaximum likelihood (CML) (Bouye, Durrleman, Nikeghbali, Riboulet and Roncalli, 2000).

Thirdly, having the parameters of the t copula, we simulate jointly dependent asset class returns by first simulating the corresponding dependent standardised residuals. Finally, by extrapolating into the GPD tails and interpolating into the smoothed interior; one can transform the uniform variables to standardise the residuals via an inversion of the semi-parametric marginal CDF of each asset class.

This produces simulated standardised residuals consistent with those obtained from the AR(1) + GJR-GARCH(1,1) filtering process. Finally, by using the simulated standardised residuals as the i.i.d. input noise process, we reintroduce the autocorrelation and heteroskedasticity observed in the original asset class returns.

It is very important to note at this point that from the above procedure, the residuals that are derived are independent in time; but they are dependent at any point in time. In other words, each array of simulated standardised residuals represents an i.i.d. univariate stochastic process when viewed in isolation However, each array maintains the rank correlation induced by the copula.

3 Scenario-Based Asset Allocation

The primary objective of this paper is to overcome the limitations of MPT; and to demonstrate a practical and robust procedure whereby investors have the ability to introduce a specific economic forecast or macroeconomic view based on their unique outlook. The result is a forward-looking optimal asset allocation that accounts for fat-tails and non-linear correlation amongst the various asset classes.

This asset allocation is dynamic in the sense that by adjusting the investor's forecasted economic view, the optimal portfolio asset allocation would change in accordance with the changing variance and covariance across time. This forward-looking scenario approach requires a number of steps; and a specific procedure These steps are outlined below.

The first step is to capture the investor's forecasted economic scenario. An input scenario is a range of subjective views or forecasts of certain input variables that have a statistically meaningful relationship – correlation – with the asset classes being considered as the output variables.

One of the advantages of using a forward-looking scenario-based approach in asset allocation is that it is unrestricted and completely dependent on the investor's specific view. For example, even if a key input variable has only been marginally significant recently; but an investor deems it to be highly relevant when going forward, this variable can then be included.

Table 2 highlights the input variables used in the analysis along with a brief motivation for why each variable was selected. The input variable data was obtained from the JSE and is of a monthly frequency. All the variables have been tested to, on average, have a statistically meaningful relationship to the asset classes historically; and they are believed to be relevant when going forward.

Note that the input variables selected to express a view are financial variables and not economic variables There are a number of reasons for this; but the salient point being the frequency of reporting of the data as well as the variability of the data themselves.

For example, to express an interest rate view, the 10year government bond yield is used instead of the prime interest rate. The 10year yield is the most liquid point on the yield curve; and its value changes daily; whereas the prime interest rate resets relatively infrequently. Therefore, a meaningful correlation between asset classes returns, which are generally volatile and a 'static' series such as the prime interest rate, is difficult to establish.

Having selected the significant input variables that will influence the dependent variables, the asset classes, the next step involves using Monte Carlo simulation to simulate jointly the dependent and independent variables.

Given the parameters of the t copula, we simulate 12 monthly (which are then aggregated to an annual return) dependent uniform variables of both the dependent and the independent variables jointly 100,000 times.

Then via the inversion of the semi-parametric marginal CDF of each variable, transform the uniform variables to standardise residuals, in order to be consistent with those obtained from the AR(1) + GJR-GARCH(1,1) filtering process. These residuals are independent in time; but they are dependent at any given point in time. Next, reintroduce the autocorrelation and volatility clustering that is observed in the historic returns of the dependent and independent variables.

Given that there are five asset classes (Table 1) and seven independent variables (Table 2), the resulting matrix is of 100,000 x 12 dimension of annualised returns which include possible fat-tail events. A crucial observation to take note of is that the said matrix is not random.

As previously mentioned, the 'rows' of annualised returns are independent from one another, i.e. row 1 is independent of row 2, of row 3, ... of row 100,000. However, the observation within a particular row, the 'columns', share the rank correlation induced by the t copula; and they are thus not random. Consequently, the non-random nature of the information maintained within a row of observations is what allows for the forward-looking process to hold.

Having jointly simulated the dependent and the independent variable, the next step in the framework involves introducing the investor's view or forecast over a particular horizon. In this paper a 12month ahead forecast period is selected. Therefore, the investors take a view on the direction they believe a particular variable would probably move over the coming year.

However, since monthly observations are simulated, the horizon period is completely flexible; and investors can choose, i.e. 3, 6, or 24month horizons; and the monthly returns are annualised accordingly We have assumed a strategic rather tactical asset allocation framework and have thus used monthly observations. However, the model can cater to more frequent data such as daily and weekly observations.

Table 3 illustrates a scenario for a local investor that is bearish on the outlook for South Africa over the next 12 months. The input variables are consistent with those highlighted in Table 2. Table 3 demonstrates the use of range inputs as forecasts – rather than using the traditional pointforecast approach.

For example a statement such as*, "We forecast the rand to be R15.50 to the dollar this time next year"* has very little practical use. Firstly, what is the probability of the rate actually realising at that level? Secondly, what likelihood is attached to the forecast? Finally, what are the best- and the worst-case scenarios if this is the anticipated case?

This paper demonstrates that one does not need to be precise with one's forecast in order to get a precise assetallocation recommendation. In fact, the analysis finds that range forecast reduce variability in the output; and it produces more reliable and stable asset allocations compared with an overly specific and unreliable point forecast.

In other words, getting the direction of an input variable over the forecast horizon is sufficient. A value higher/lower than the current spot; or being more precise and specifying a range of by how much higher/lower the variable will be over the forecast horizon is adequate. A view that the independent variable will increase by between 5% and 15% from the current spot over the next 12 months is preferred rather than a point forecast as a forward-looking hypothesis

Ultimately, if USDZAR is forecast to weaken from the current spot of i.e. R14.50 to either R15.20 or R15.80 to the dollar in 12months' time; this would have very little impact on one's optimal asset allocation between bonds and equities. However, the direction of the rand in one year's time, weaker or stronger, will indeed have a meaningful influence on the relative allocation

The final step of the scenario-based assetallocation process involves isolating from the matrix of simulated returns, the returns that match the investor's view. For illustrative purposes, Table 4 provides a snapshot of this step. The top part of Table 4 recaps the investor's expectations; while the lower part of the table is a sample of the jointly simulated returns. The rows that are shaded are the observations that exhaustively match the investor's forecast views.

To reiterate, by using the parameters of the t copula, with a Monte Carlo simulation approach, the input and output variables have been jointly simulated. Therefore, the rows are independent of one another; however, there is dependency across the rows at a point in time – the observations in a particular row are correlated with one another

Subsequently, from all possible simulated returns, only the rows which wholly match the investor's view (shaded in Table 4) are isolated; and only the corresponding asset class returns that match that row are considered for the portfoliooptimisation process.

Given that there is dependency between the variables across a particular row, i.e. the correlation structure has been maintained during simulation, the asset class returns that correspond with a particular view can be seen as possible and statistically likely returns for those asset classes that could occur when all the conditions of the view are met.

For example, row 2 in Table 4 is highlighted as a set of returns that completely match the investor's bearish view (i.e. all the returns of the input variables in that row fall within the lower and upper forecast range imposed by the investor). The asset class returns for that particular row are as follows: ALSI 12%, ALBI -9%, CASH 4%, MSCI 15% and PROP -13%. However, this is only one of the sets of possible annual returns that match this particular view.

Consequently, the simulation process has created over 270 times more observations than had only historical data been considered. In fact, the historical data are of no use; since the probability of a specific set of returns occurring jointly or repeating is highly unlikely.

For example, due to the limited history available, there may never have been a period in the past where the following scenario happened simultaneously – and thus matching the investor's expectations; the rand depreciated by 15%; yields rose by 15%; the gold price fell by 10%; and the base metals index fell by 15%. However, there is no reason why this joint scenario cannot happen, even if it has not yet been observed.

Similarly, the annual assetclass returns highlighted in row 2 (ALSI 12%, ALBI -9%, CASH 4%, MSCI 15% and PROP -13%) may never have been realised jointly in the limited history available However, they are statistically probable; or it is possible for them to occur. The constraint of historical data has been overcome via simulation.

Therefore, these asset class returns that filter through as having matched the forward-looking view are called 'conditional returns'; as they are conditional on an investor's specific scenario. It is from these returns that the mean-CVaR efficient portfolios are constructed.

Figure 9 illustrates the difference between the conditional and the historical returns of the ALSI. Given the specific co-variation imposed, it is expected that the conditional returns would differ from the historic returns. Therefore, the portfolio optimisation is based on the correlation between the asset classes from the investor's specific scenario; but not on the historic correlations.

4 Mean-CVaR Portfolio Optimisation

The final stage of the assetallocation framework is to construct efficient portfolios that are mean-CVaR optimal. Through an optimisation framework, an efficient frontier is constructed that generates the optimal asset allocation of 10 portfolios ranked from least risky to most risky. Thus all the portfolios on the efficient frontier represent the optimal asset allocation for a particular risk-return trade-off (Rachev, Stoyanov and Fabozzi, 2008)

Figure 10 depicts the efficient frontier and the asset class weights for the 10 portfolios on the frontier that correspond with the bearish view imposed. Table 5 shows the corresponding expected return and CVaR of each portfolio along the efficient frontier with the respective asset class weights attached to each portfolio.

The above assetallocation optimisation imposes a constraint on the offshoreequity component to allow only a maximum of 25% allocation to this asset class. In order to illustrate a realistic asset allocation, the asset allocation complies with Regulation 28 of the Pension Fund Act which states a maximum of 25% of an institutional investor's portfolio is permitted to have any offshore exposure

For example, the forecast scenario of a weaker rand would generally result in the optimisation – by allocating a larger portion to offshore equities – due to the benefit of a weakening rand on one's offshore holdings. However, with the constraint imposed at 25%, the optimisation maxes out the allocation to offshore equities in portfolio 7; and it optimises all the remaining portfolios to the right of portfolio 7 according to the constraint.

The remaining four asset classes are unconstrained; but constraints can be applied in conjunction with a particular investor's mandate, i.e. a balanced manager may not want to hold more than 75% of equities in their portfolio. In this case additional constraints would be applied to ensure the equity cap as well as the offshore allowance is not breached while deriving an optimal asset allocation that is mean-CVaR optimal.

The output achieved, Figure 1 along with Table 5, is beneficial to investors for a number of reasons: Firstly, investors are now able to visualise the efficient frontier and asset class weights which are based on their unique view or on the scenario they have imposed. For example, portfolio 1 has 100% allocated to cash; hence the expected return is the cash return without any tail risk.

However, as one moves upwards and rightwards along the efficient frontier, the tail-risk increases significantly but not at a linear rate. Therefore, the investor can choose the 'sweet spot' of how much expected return for a given level of tail risk they are willing to trade off.

For example, the right-most portfolio is allocated 100% to the propertyasset class; and it has the highest expected return as well as the highest CVaR. The CVaR of portfolio 10 indicates that there is a 1 in 20 (CVaR 95%) probability that the investor can lose, on average, 35% of his portfolio given that a 5% event occurs.

Secondly, risk-return profiles are easily quantifiable; and investors can see the asset weights proposed which match their benchmark along with the tail-risk associated with the particular expected return.

For example, a CPI + 3% benchmark is a common benchmark for South African fund managers. Therefore, given this target benchmark, investors can now see what their optimal asset allocation should be – given their forecasts. In Table 5, portfolio 6 has an expected return of 9% which meets the investor's benchmark. Therefore, if the investor is confident of his view, then his optimal asset allocation should be in line with that of portfolio 6

Thirdly, the assetallocation process is completely dynamic; in the sense that by changing the percentage range or the direction of one or more of the input variables, the efficient frontier as well as the optimal asset allocation would change dynamically. This process holds provided no variables are removed or replaced; and they are only adjusted.

However, removing or replacing an independent variable would require the process to be re-simulated; as the correlation structure between the new set of input variables and the asset classes would have changed.

For example, if an investor is less concerned with the absolute change in the US 10year yields; but they prefers to take a view on the slope of the curve instead; then, the US 10year input variable is replaced by a variable which captures the slope of the curve, i.e. 2s10s [1].

The co-variation structure between the independent and the dependent variables will now change; since the relations of both the asset classes as well as the remaining input variables to a 2s10s variable are different from an outright yield variable. The implications of a steepening treasury curve and an increase

[1] An increase in 2s10s implies the difference between the 10 year yield and the 2 year yield is widening and thus the yield curve is steepening. A decrease in 2s10s implies the curve is flattening

in the 10year yield would have different implications for ALBI returns; and thus the co-variation structure needs to be re-estimated.

Figure 11 and Table 6 are a further demonstration of the scenario-based allocation framework. Here the asset classes have been replaced by FTSE/JSE equity sectors as the output variables. The independent variables as well as the forward-looking bearish view have been left unchanged from Table 3.

An equity investor may apply this scenario-based framework to impose their macro view; and may conduct an initial filter of the equity sectors based on the risk-return before proceeding to the stock picking of individual shares within the sectors. Unfortunately, due to the large unsystematic risk inherent in each stock, the framework is not able to provide a statistically sound asset allocation to any individual equities.

5 Empirical Evidence

Given that the asset allocation framework is dependent on the selection of independent variables as well as the forecast of these variables, backtesting the model is challenging. An unbiased and objective manner that removes as much subjectivity as possible without forgoing statistical accuracy is needed in order to test the model results. In other words: How well does the model perform given that the investor got their input assumption correct?

In order to overcome the issue of accounting for investors' forecasts and conducting the analysis in an objective manner, we begin by splitting the distribution of all the input variables into quartiles. To reduce any further complications by selecting the significant input variables for each backtesting period, we have left them consistent with those in Table 2; as these have proved, over the long-run to be correlated with the asset classes.

Subsequently, we assume that for each input variable the investor has the skill to predict the return of the variable over a oneyear horizon with a degree of accuracy – within the correct quartile. For example, in 2008 the rand depreciated by 27% against the dollar. Therefore, the investor, displaying a moderate amount of skill would select the third quartile of the historic USDZAR returns as his input scenario. Similarly, the oil price fell by 50% over the same period. The investor in this instance would select the first quartile.

A further challenge in backtesting is in selecting the expected return-CVaR target asset allocation that an investor is hoping to achieve over every period. To overcome this we assume that at each period the investor is targeting a CPI+3% return; and we thus select the assetallocation weights each period that correspond with the expected return.

In order to test the asset allocation proposed; we compare it with a benchmarkasset allocation that most closely resembles a balanced fund: equities (45%), bonds (30%), cash (5%), offshore equities (15%) and property (5%).

As an example, in an extreme year such as 2008, a balanced fund with the above allocation would have returned -17%. However, under the asset allocation weights proposed by the model, a return of -5% would be realised. Although a loss was realised, the model was able to avoid a tail-event loss.

How would the model perform under non-extreme market conditions, i.e. a range-bound market and a bull market? In 2011 the ALSI grew from 31,400 to 32,000 representing a flat market year-on-year. Over this period a balanced fundasset allocation yielded 5%, underperforming a CPI+3% benchmark by 3%. Assuming an investor had reasonable input scenarios, the asset allocation corresponding to a CPI+3% would have returned 7.5%, 50 basis points below the expectations.

The bull market of 2012 saw the ALSI rallying 16%. Over this period an average balanced fund allocation would have outperformed a CPI+3% benchmark significantly: 15% vs 9%. The asset allocation proposed under the model would also outperform the benchmark returning 12.5%; however the difference is a reduction in the tail risk. The model was able to significantly reduce the tail risk of the portfolio, without compromising the returns, by proposing a reduced allocation to equities.

Figure 12 illustrates the returns of a portfolio allocation in line with an average balanced fund vs the asset allocation corresponding closest to a CPI+3% allocation as proposed by the model. From the figure it is evident that a balanced portfolio has tails that are much fatter than those of the model allocation. This is an indication that the model was able to consistently avoid extreme losses, assuming the investor's forecast was positioned in the correct quartile of the independent variables' distribution.

Figure 13 measures how closely the two asset allocations, the balanced and the model, came within their targeted benchmark. A negative return indicates that the asset allocation outperformed the CPI benchmark. The figure demonstrates that the balancedfund allocation had significantly more observations that outperformed the benchmark but also more observations that underperformed the benchmark.

The model assetallocation returns are a lot more peaked and centred around zero, indicating that the proposed allocation under the model was able to much more closely achieve its target level. Table 7 highlights the descriptive statistics of the portfolios described above.

6 Conclusion – The Emperor's New Clothes

Akin to the Emperor continuing his procession, since turning back would be to admit one's own blindness, so also have investors relied for too long on MPT. Better by far to continue on in the pretence that they are the only ones with the wisdom to see the clothes than to admit that MPT is flawed.

This paper set out to firstly, highlight the flaws in MPT and then to propose a technique to overcome these shortcomings; and secondly, to recommend a procedure that does not rely solely on historic observations for optimisation – and thus allowing investors to apply a forward-looking approach to asset allocation.

The paper has shown that by applying EVT, the skewness and kurtosis inherent in asset class returns as well as the volatility clustering and extreme co-movement across asset classes can be accounted for. By partitioning the distribution of each class, using the Gaussiankernel estimation for the interior and a GPD for the upper and lower tails, a better approximation can be achieved.

Using a Student's t copula to fit the data, the correlation between the simulated residuals of each asset class is induced. A Monte Carlo simulation technique is used to simulate the standardised residuals which represent a univariate stochastic process when viewed in isolation; but maintains the correlation induced by the copula.

Given the dependency between the simulated returns, the investor's forward-looking scenario is isolated from the simulated observations; and the corresponding assetclass returns are used to construct mean-CVaR optimised portfolios.

Backtesting revealed that the model was able to consistently reduce extreme losses, indicating that the tail risk was mitigated under the framework. In addition, the backtesting results showed that the model returns, at a statistically significant level of confidence, were extensively in line with those of the target benchmark return of CPI + 3%

References

[1] Black, F. (1972). "Capital Market Equilibrium with Restricted Borrowing." *The Journal of Business,* 7(2): 86-99.

[2] Bouye, E., Durrleman, V., Nikeghbali, A., Riboulet, G. and Roncalli, T. (2000). "Copulas for Finance: A Reading Guide and Some Applications." Financial Econometrics Research Centre. City University Business School, London.

[3] Embrechts, P., Kluppelberg, C. and Mikosch, T. (1997). *Modelling Extreme Events for Insurance and Finance.* Berlin/Heidelberg: Springer-Verlag.

[4] Embrechts, P., McNeil, A. and Straumann, D. (1999). "Correlation and Dependence in Risk Management: Properties and Pitfalls." Risk Management: Value-at-Risk and Beyond, University of Cambridge.

[5] Focardi, S.M. and Fabozzi, F.J. (2004). *The Mathematics of Financial Modelling and Investment Management.* John Wiley & Sons. Hoboken, New Jersey.

[6] Glosten, L.R., Jagannathan, R. and Runkel, D.E. (1993). On the Relation between Expected Value and the Volatility of the Nominal Excess Return on Stocks." *The Journal of Finance,* 48(5): 2374-2387.

[7] Hilal, S., Poon, S. and Tawn, J. (2011). "Hedging the Black Swan: Conditional Heteroskedasticity and Tail Dependence in the S&P500 and VIX." *Journal of Banking and Finance,* 35(9): 383-400.

[8] Hyung, N. and de Vries, Casper. (2007). "Portfolio Selection with Heavy Tails." *Journal of Empirical Finance,* 14(3): 383-400.

[9] Li, X. And Sullivan, R. (2011). "A Dynamic Future for Active Quant Investing." *Journal of Portfolio Management,* 37(3): 29-36.

[10] Markowitz, H. (1952). "Portfolio Selection." *The Journal of Finance,* 7(1): 77-91.

[11] McNeil, A.J. and Frey, R. (2000). "Estimation of Tail Related Risk Measurement for Heteroskedastic Financial Time Series: An Extreme Value Approach.' *Journal of Empirical Finance,* (7): 271-300.

[12] McNeil, A.J. and Saladin, T. (1997). "The PeaksoverThreshold Method for Estimating High Quantiles of Loss Distributions." Proceedings of 28^{th} International ASTIN Colloquium.

[13] Nystrom, K. and Skoglund, J. (2000). "Univariate Extreme Value Theory, GARCH and Measures of Risk." Preprint, Swedbank.

[14] Rachev, S.T., Stoyanov, S.V. and Fabozzi, F.J. (2008). *Advanced Stochastic Models, Risk Assessment, and Portfolio Optimisation: The Ideal Risk, Uncertainty, and Performance Measure.* John Wiley & Sons. Hoboken, New Jersey.

[15] Rockafellar, T.R. and Uryasev, S. (2002). "Conditional Value-at-Risk for General Loss Distribution." *Journal of Banking and Finance,* 26(7): 1443-1471.

[16] Roncalli, T., Durrleman, A. and Nikeghbali, A. (2000). "Which Copula is the Right One?" Université d'Évry - Centre D'Etudes des Politiques Economiques et de L'Emploi (EPEE); Lyxor Asset Management.

[17] Sullivan, R.N., Peterson, S.P. and Waltenbaugh. (2010). "Measuring Global Systematic Risk: What are Markets saying about Risk?" *Journal of Portfolio Management,* 37(1): 67-77.

[18] Uryasev, S. (2000). "Conditional Value-at-Risk: Optimisation Algorithms and Applications." *Financial Engineering News,* (14): 1-5.

[19] Wang, P., Sullivan, R.N. and Ge, Y. (2011). "Risk-Based Dynamic Asset Allocation with Extreme Tails and Correlations. *Journal of Portfolio Management,* 38(4): 26-42.

[20] Xiong, J.X. and Idzorek, T.M. (2011). "The Impact of Skewness and Fat Tails on Asset Allocation Decisions." *Financial Analysts Journal,* 67(2): 23-35.

[21] Zeevi, A. and Marshal, R. (2002). "Beyond Correlation: Extreme Comovement between Financial Assets." Preprint, Columbia University.

Appendix

Table 1: Asset Class Summary Statistics of Monthly Returns (January 1986 – March 2016)

	FTSE/JSE All Share Index (ALSI)	South African All Bond Index (ALBI)	Short Term Fixed Interest Index (CASH)	MSCI World Equity Index (MSCI)	SA Listed Property Index (PROP)
Mean	1.19%	1.12%	0.92%	1.13%	0.82%
Std Deviation	5.58%	2.42%	0.34%	4.85%	4.54%
Median	1.42%	1.17%	0.87%	1.10%	1.06%
Min	-29.45%	-14.43%	0.42%	-19.71%	-14.37%
Max	17.51%	11.32%	1.74%	18.69%	18.19%
5th Percentile	-7.54%	-2.75%	0.45%	-6.15%	-6.52%
95th Percentile	9.71%	4.35%	1.53%	8.73%	7.66%
Skewness	-0.70	-0.49	0.37	-0.64	-0.10
Kurtosis	3.13	6.03	-0.85	2.07	1.06

Table 2: Input Variables and Description

	Variable	Description
1	USDZAR exchange rate	An exchange rate component to capture the relationship between the currency and the fixed income market. The exchange rate is also a meaningful driver of the local equity market. Due to the large weight of the dual listed stocks in the ALSI (having costs in rand and revenues in hard currency), a movement in the currency has a direct impact on the performance of these shares. Both the USDZAR and EURUSD rates are included in order to capture information on currency movements attributed to either domestic or foreign factors, i.e. distinguish rand weakness which could be attributed to South Africa or Emerging Market specific reasons from general dollar strength.
2	EURUSD exchange rate	
3	SA 10 year government bond yield	A domestic interest rate variable which is proxied by the 10 year government bond yield. Interest rates show significant evidence of a long-term relationship with most asset classes included in the analysis.
4	US 10 year Treasury yield	It is important for local investors to account for the movement of foreign yields. An increase or an anticipation of rising US interest rates, negatively affect the South African bond market. Foreign investor's sell-out of local SA government bonds when the carry-trade becomes relatively less attractive, and prefer to earn returns in hard currency which are far less volatile.
5	Brent Crude Oil spot price ($/bbl)	The change in oil price impacts the expected rate of inflation which directly translates into a significant impact on bond returns
6	Gold spot price ($/oz)	Gold, considered a safe haven asset as well as a proxy for global inflation expectations.
7	The Economics' Metals Index	The price of base metals is a good proxy for broad Emerging Market performance. For example, investors are able to take a view on the Chinese economy through commodity prices. Secondly, a number of large South African mining and industrial company's profitability is linked to base metal prices.

Table 3 Scenario Forecast

Variable	USDZAR	EURUSD	SA 10 yr	US 10 yr	Oil	Gold	Metals
Upper	15%	10%	15%	10%	10%	0%	-5%
Lower	0%	0%	5%	5%	-10%	-10%	-15%

Table 4 Scenario-Based Asset Allocation Framework

Output Variables (Asset Classes) Returns					Input Variable Returns with Forecasts						
ALSI	ALBI	CASH	MSCI	PROP	USDZAR	EURUSD	SA 10 yr	US 10 yr	Oil	Gold	Metals
?	?	?	?	?	15% 0%	10% 0%	15% 5%	10% 5%	10% -10%	0% -10%	-5% -15%
Jointly simulated returns											
0.12	-0.04	0.04	-0.15	-0.19	-0.03	-0.09	0.03	0.17	-0.09	-0.11	-0.18
0.12	-0.09	0.04	0.23	-0.13	0.03	0.01	0.06	-0.01	-0.06	0.08	-0.01
-0.03	0.05	0.04	0.17	0.15	0.19	-0.01	0.03	0.00	-0.12	0.05	-0.32
0.12	-0.09	0.04	0.15	0.16	0.06	0.03	0.13	0.04	0.07	-0.06	-0.13
-0.14	-0.17	0.04	0.18	0.17	0.16	-0.12	-0.06	-0.08	-0.02	0.12	-0.23
-0.03	-0.07	0.09	0.12	-0.29	0.06	-0.02	-0.15	-0.08	0.02	0.16	0.16
0.27	-0.24	0.05	0.12	-0.04	-0.03	-0.02	-0.21	0.09	0.23	0.07	-0.16
0.08	-0.06	0.05	-0.14	-0.15	-0.01	0.10	-0.19	-0.06	0.29	0.01	-0.06
-0.08	-0.12	0.07	0.01	0.34	0.38	-0.02	0.11	-0.12	0.04	-0.09	-0.01
-0.29	0.05	0.07	0.13	0.08	0.02	0.08	-0.09	0.03	-0.14	-0.09	0.07
-0.03	-0.05	0.06	0.11	-0.25	0.05	0.02	0.07	0.06	0.08	-0.07	-0.08
-0.13	0.05	0.04	-0.03	0.06	0.07	-0.02	-0.02	-0.08	-0.26	0.17	0.26
0.15	-0.04	0.06	0.22	-0.23	0.01	-0.18	-0.14	0.10	0.31	-0.12	0.02
0.09	-0.13	0.05	0.05	0.19	0.12	0.04	0.14	0.08	0.05	-0.04	-0.12
0.02	-0.06	0.07	0.08	0.22	0.43	-0.19	-0.05	0.02	0.32	-0.03	0.03
0.27	-0.10	0.04	0.14	0.11	-0.04	0.00	-0.05	-0.04	0.30	0.11	0.08
-0.01	-0.18	0.08	-0.08	-0.01	0.11	-0.08	0.23	0.02	0.05	-0.11	-0.04
-0.04	-0.10	0.07	0.10	0.18	0.08	0.08	0.09	0.08	0.08	-0.01	-0.06
-0.05	-0.14	0.03	-0.01	-0.03	0.04	-0.20	-0.11	-0.04	-0.15	0.15	0.05
0.00	0.16	0.06	0.11	-0.07	-0.03	0.06	0.15	-0.09	0.27	-0.06	-0.02
0.03	-0.20	0.04	0.13	-0.4	0.03	0.12	-0.12	-0.05	-0.16	0.01	0.03
-0.12	-0.13	0.06	0.02	-0.14	0.11	0.10	0.06	0.09	-0.05	-0.02	-0.14
0.10	-0.16	0.07	0.16	0.36	0.17	-0.02	-0.04	-0.09	0.05	-0.29	-0.30

Table 5 Asset Allocation Weights, Risk and CVaR

	Portfolio 1	Portfolio 2	Portfolio 3	Portfolio 4	Portfolio 5	Portfolio 6	Portfolio 7	Portfolio 8	Portfolio 9	Portfolio 10
E(Return)	5.5%	6.2%	6.8%	7.5%	8.2%	8.9%	9.5%	10.2%	10.9%	11.6%
CVaR	-4.6%	-2.1%	0.9%	3.9%	7.0%	10.0%	13.0%	17.2%	24.0%	35.4%
ALSI	0.0%	3.7%	6.9%	10.2%	13.5%	16.7%	19.8%	28.6%	28.6%	0.0%
ALBI	0.0%	4.4%	9.0%	13.7%	18.3%	22.9%	27.6%	21.2%	2.9%	0.0%
CASH	100%	82.1%	65.7%	49.3%	32.9%	16.5%	0.1%	0.0%	0.0%	0.0%
MSCI	0.0%	5.2%	9.1%	13.0%	16.8%	20.6%	24.5%	6.0%	0.0%	0.0%
PROP	0.0%	4.7%	9.4%	14.0%	18.6%	23.3%	28.0%	44.2%	68.5%	100.0%

Table 6 Equity Sector Allocation Weights, Risk and CVaR

	Portfolio 1	Portfolio 2	Portfolio 3	Portfolio 4	Portfolio 5	Portfolio 6	Portfolio 7	Portfolio 8	Portfolio 9	Portfolio 10
E(Return)	8.7%	9.5%	10.3%	11.1%	11.9%	12.7%	13.5%	14.3%	15.1%	15.9%
CVaR	23.6%	23.7%	24.0%	24.6%	25.3%	26.3%	27.4%	29.0%	34.2%	54.3%
Resources	14.0%	10.0%	6.2%	2.6%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Banks	3.6%	6.0%	8.4%	11.0%	13.9%	18.4%	23.1%	31.2%	56.5%	100.0%
Insurance	37.4%	32.1%	26.6%	21.2%	14.8%	6.5%	0.0%	0.0%	0.0%	0.0%
Con.Services	2.9%	5.4%	8.0%	10.1%	12.9%	15.4%	17.1%	9.3%	0.0%	0.0%
Con.Goods	17.1%	22.8%	28.5%	34.4%	40.2%	45.8%	52.0%	59.5%	43.5%	0.0%
Industrials	21.7%	20.7%	19.6%	18.2%	16.2%	12.1%	6.4%	0.0%	0.0%	0.0%
Telecoms	3.4%	3.2%	2.7%	2.5%	2.1%	1.8%	1.3%	0.0%	0.0%	0.0%

Table 7 Portfolio Return Comparison

	Balanced Fund Portfolio Allocation	Model Asset Allocation	Relative Return (CPI+3% to Balanced Fund)	Relative Return (CPI+3% to Model Allocation)
Mean	11.30%	10.10%	-3.01%	-1.80%
Std Deviation	10.17%	5.04%	11.63%	6.49%
Median	12.90%	9.50%	-4.74%	-2.18%
Min	-17.86%	-3.80%	-31.38%	-15.03%
Max	37.23%	19.70%	28.54%	16.69%
5th Percentile	-9.34%	-0.11%	-20.15%	-10.83%
95th Percentile	26.14%	17.50%	21.10%	11.70%
Skewness	-0.64	-0.58061	0.73	0.76
Kurtosis	0.50	-0.01146	0.48	0.25

Figure 1 Monthly ALSI Returns (January 1960 – March 2016)

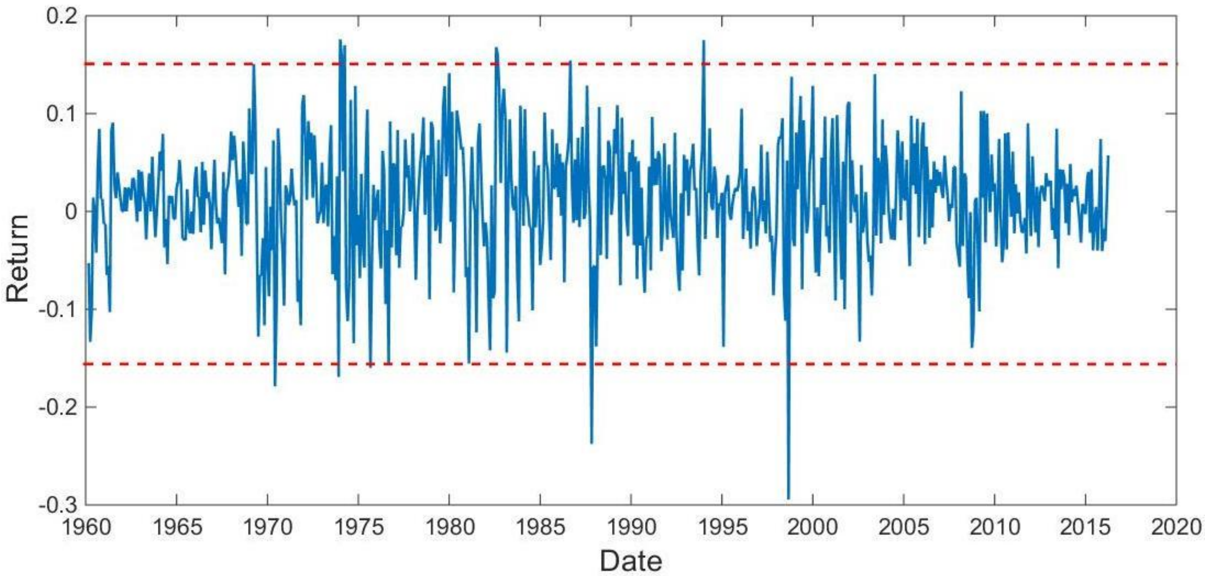

Figure 2 Monthly ALBI returns fitted with a Gaussian Normal Distribution

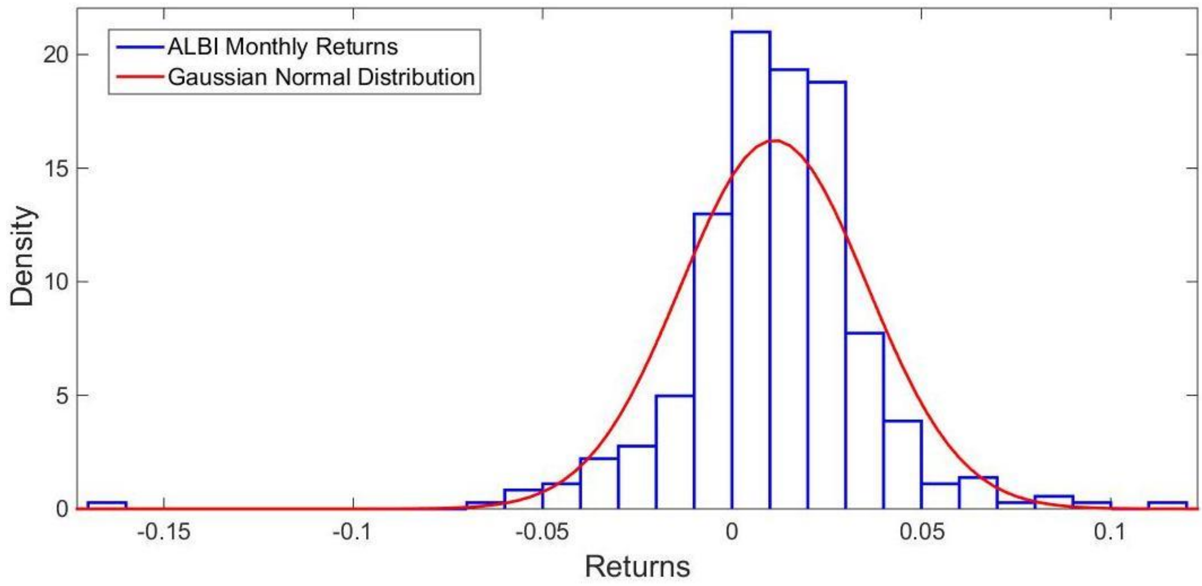

Figure 3 Conditional Value-at-Risk of a Hypothetical Portfolio Returns

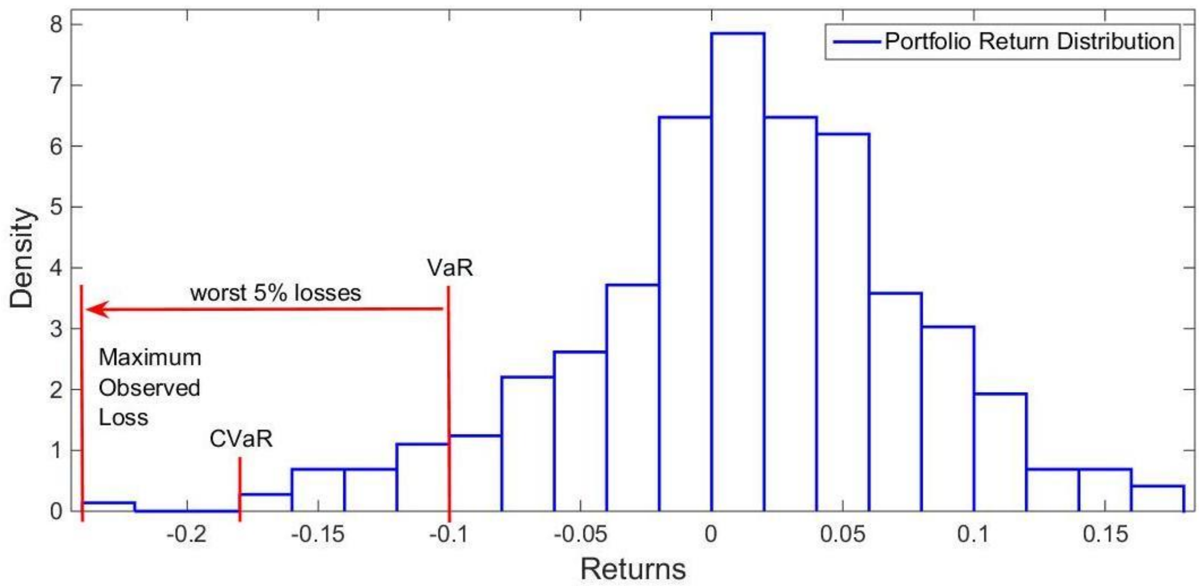

Figure 4 Autocorrelation of Monthly Returns and Squared Monthly Returns (ALSI)

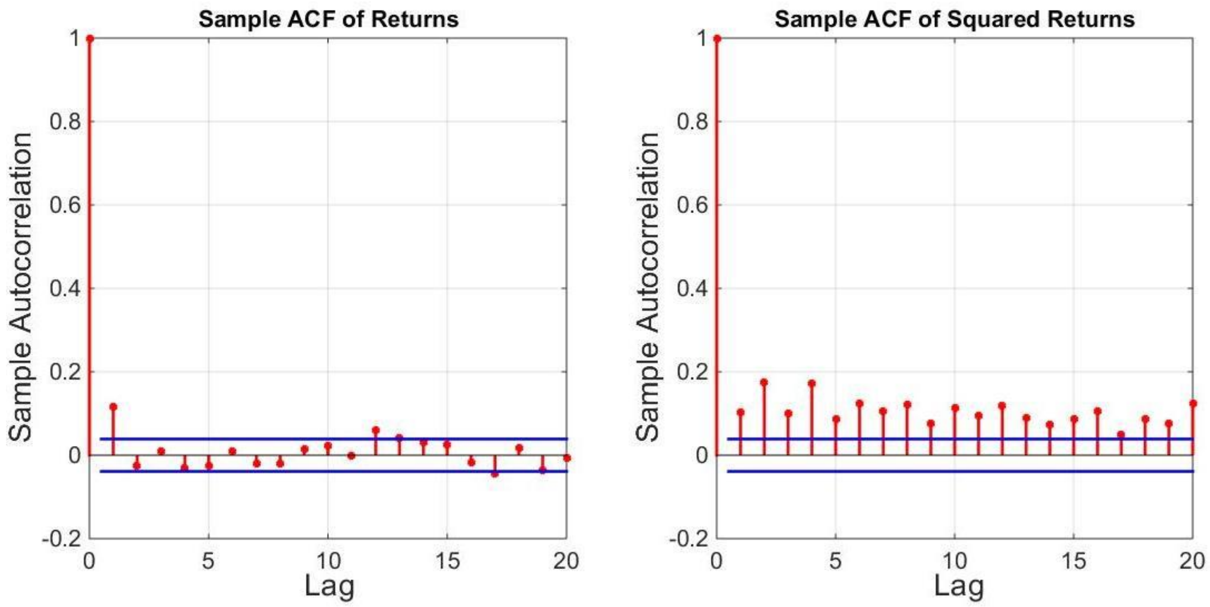

Figure 5 Autocorrelation of Standardised Monthly Returns and Standardise Squared Monthly Returns (ALSI)

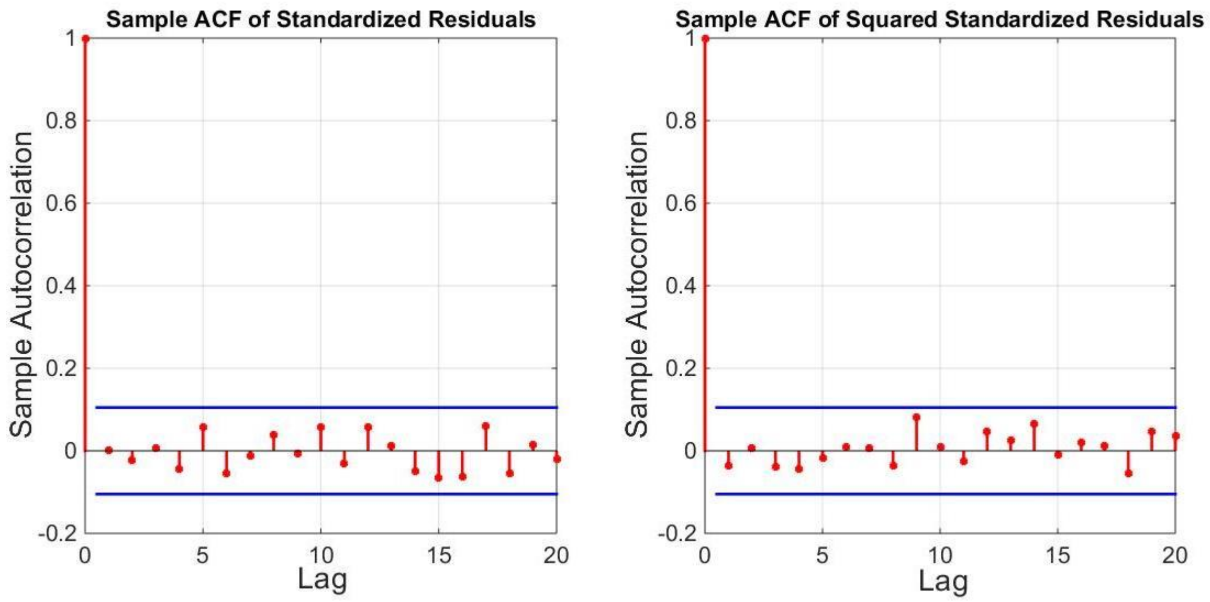

Figure 6 Pareto Lower and Upper Tail Fit and Gaussian Kernel Interior Fit

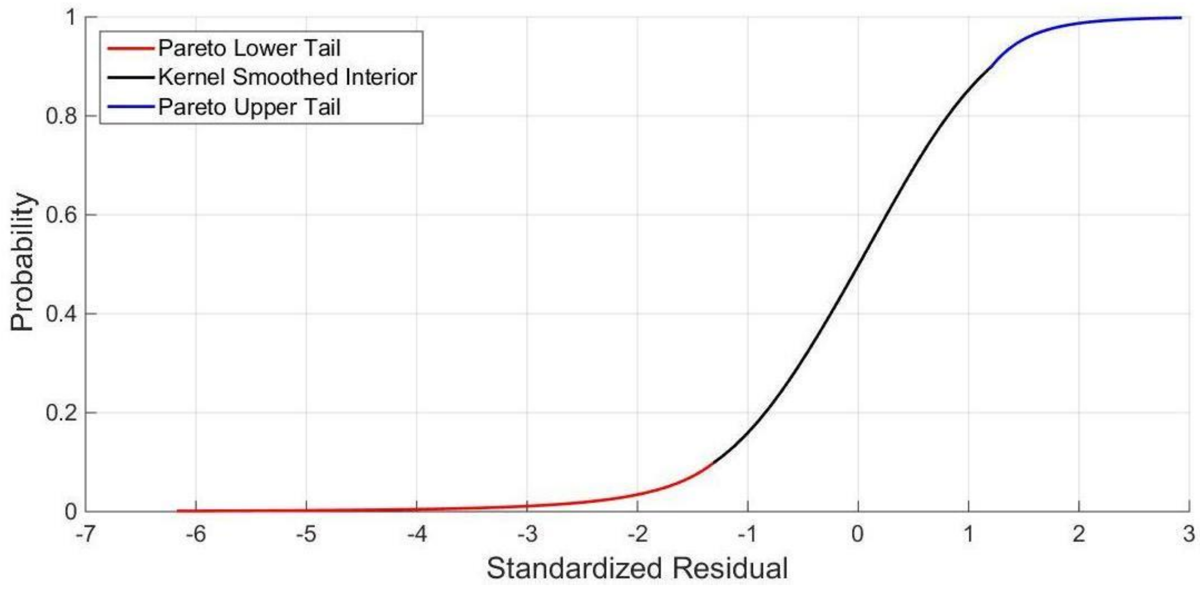

Figure 7 Assessing GPD: Lower and Upper Tail Fit

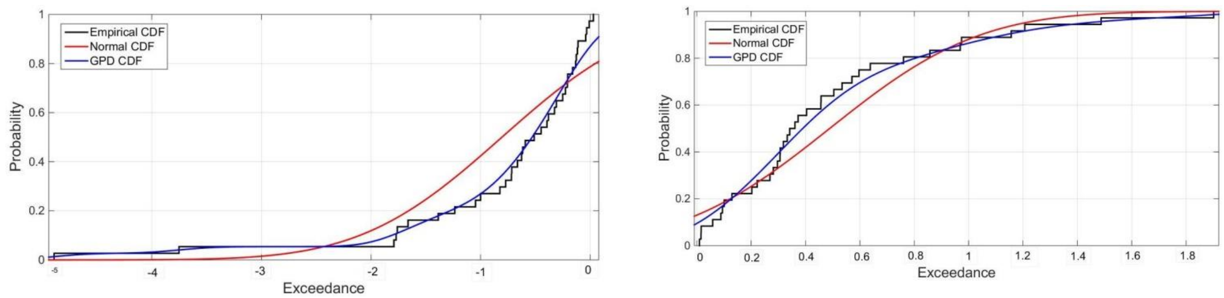

Figure 8 Scatter Plot of ALSI vs. ALBI Monthly Returns

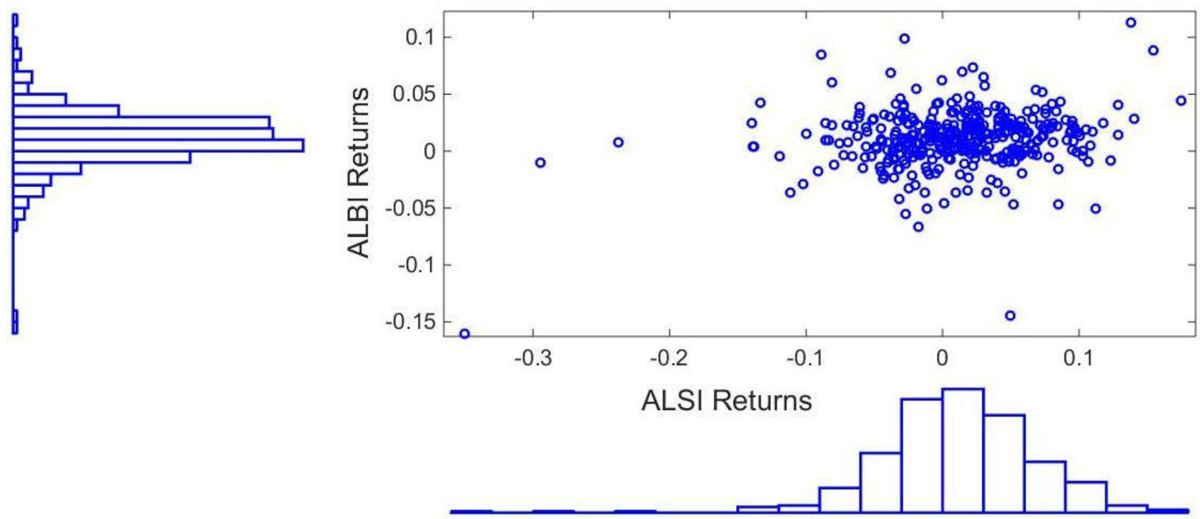

Figure 9 ALSI Historic Monthly Returns vs. Conditional Monthly Returns

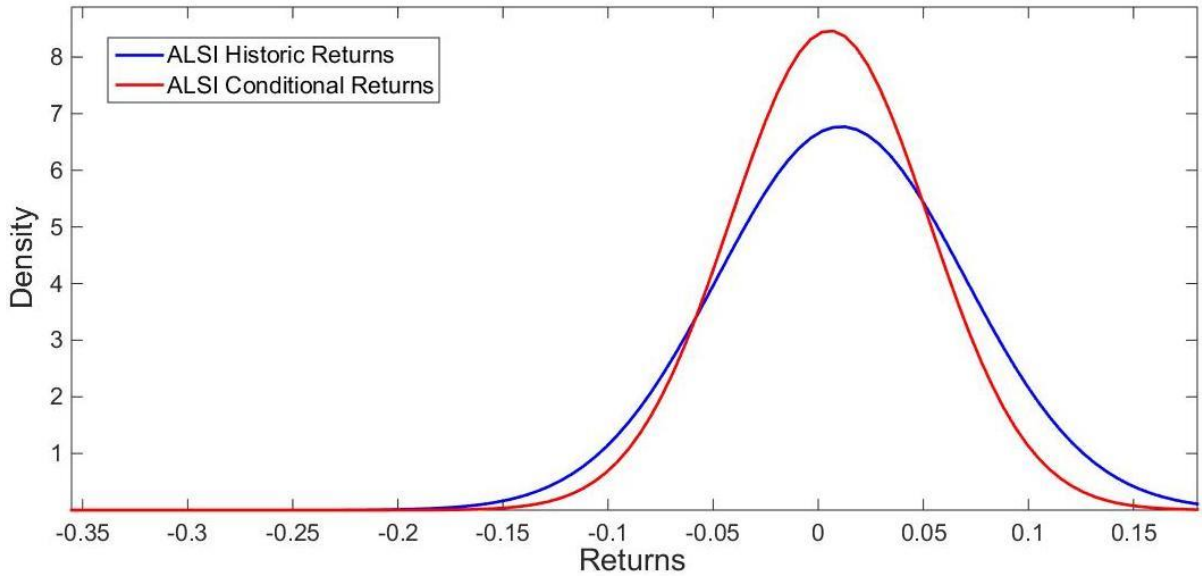

Figure 10 Efficient Frontier and Allocation Weights – Asset Classes

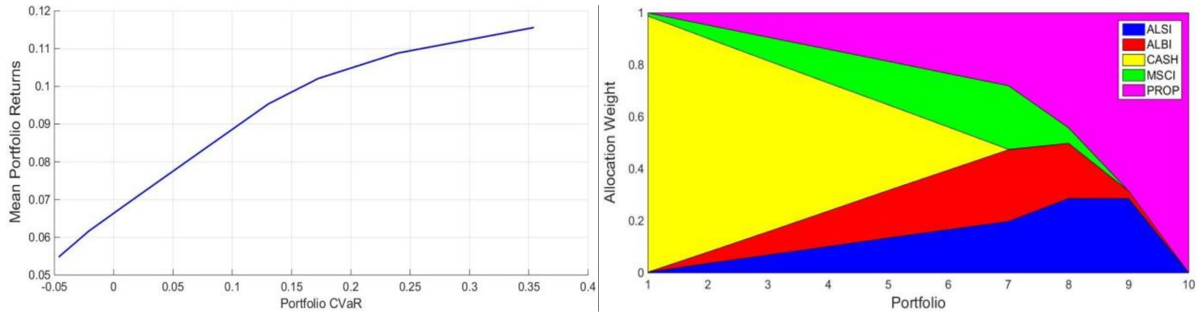

Figure 11 Efficient Frontier and Allocation Weights – Equity Sectors

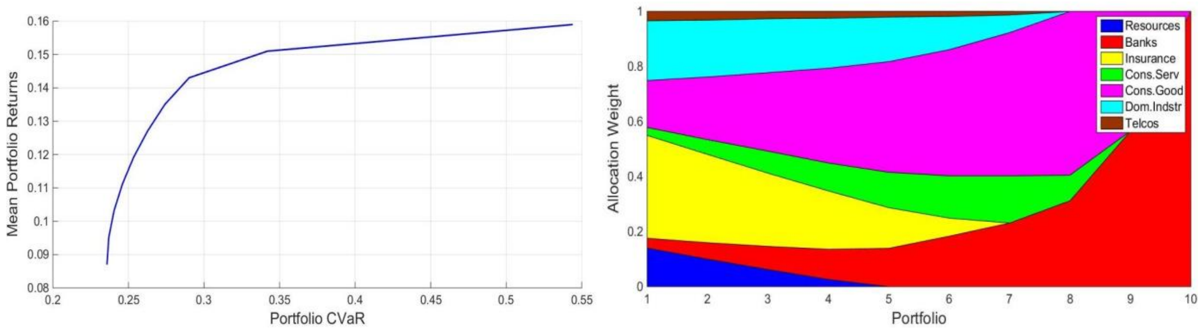

Figure 12 Portfolio Return Comparison

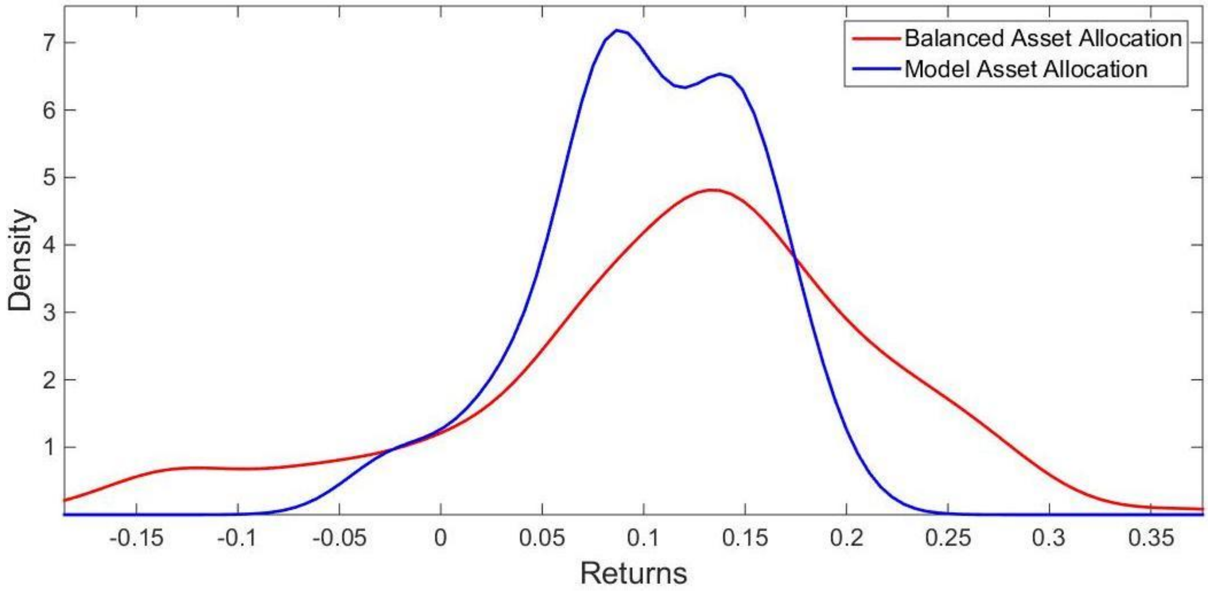

Figure 13 Portfolio Returns Relative to CPI + 3% Benchmark