Defining and measuring portfolio diversification

By E Flint, A Seymour and F Chikurunhe

ABSTRACT

It is often said that diversification is the only 'free lunch' available to investors; meaning that a properly diversified portfolio reduces total risk without necessarily sacrificing expected return. However, achieving true diversification is easier said than done, especially when we do not fully know what we mean when we are talking about diversification. While the qualitative purpose of diversification is well known, a satisfactory quantitative definition of portfolio diversification remains elusive. In this research, we summarise a wide range of diversification measures, focusing our efforts on those most commonly used in practice. We categorise each measure based on which portfolio aspect it focuses on: cardinality, weights, returns, risk or higher moments. We then apply these measures to a range of South African equity indices, thus giving a diagnostic review of historical local equity diversification and, perhaps more importantly, providing a description of the investable opportunity set available to fund managers in this space. Finally, we introduce the idea of diversification profiles. These regime-dependent profiles give a much richer description of portfolio diversification than their single-value counterparts and also allow one to manage diversification proactively based on one's view of future market conditions.

CONTACT DETAILS

Emlyn Flint, Legae Peresec and Department of Actuarial Science, University of Cape Town, Cape Town; Email: emlynf@legaeperesec.co.za
Anthony Seymour, Legae Peresec, Cape Town
Florence Chikurunhe, Legae Peresec, Cape Town

1.   INTRODUCTION

1.1   It's often said that diversification is the only 'free lunch' available to investors; meaning that a properly diversified portfolio reduces total risk without necessarily sacrificing expected return. This is a compelling concept and one that the prudent manager strives towards. However, achieving true diversification is easier said than done, especially when we do not fully know what we mean when we are talking about diversification. Clearly, the qualitative purpose of diversification serves to mitigate the effect of specific sources of risk within any single asset class, and systemic sources of risk across asset classes. But what about a quantitative definition, or more generally, a "broadly accepted, unique, satisfactory methodology to precisely quantify and manage diversification" (Meucci, 2010: p2)?

1.2   Ever since Markowitz (1952) introduced the Modern Portfolio Theory paradigm, diversification has been a cornerstone of well-constructed portfolios. However, in the nearly 70 years that have followed, diversification has usually been defined and targeted in a rather hand-waving fashion. Most research on diversification has generally adopted a heuristic approach and is inevitably littered with time-worn adages, usually involving eggs and baskets.

1.3   Since the advent of the global financial crisis, however, there has been a resurgent interest in the linked ideas of diversification, portfolio construction and risk budgeting. This is most likely due to the fact that portfolios which were thought to be well-diversified did not achieve their expected levels of protection during the financial crisis. As a result, financial practitioners and academics alike have since turned to solving the riddle of diversification. It is their research, as well as our extensions thereof, that we showcase here.

1.4   In this paper, we start by providing a taxonomy of recent advances within the diversification and risk budgeting literature. We then focus our efforts on those measures most commonly used in practice as well as some uncommon but promising metrics. Rather than emphasise mathematical detail, we provide a non-technical overview of each measure and show how they compare under real data by applying them to a range of FTSE/JSE equity indices—specifically, the All-Share (ALSI), Shareholder Weighted All-Share (SWIX), Top40 and Swix40 indices respectively. For the sake of brevity, we only report results for a small selection of these indices in the subsequent sections. Interested readers are welcome to contact the authors for additional results.

1.5   We then end the paper by introducing the idea of diversification profiles, which are built upon a market regime framework. These regime-dependent profiles give a much richer description of portfolio diversification than their single-value counterparts and also allow one to manage diversification proactively based on one's view of future market conditions.

2.   A TAXONOMY OF DIVERSIFICATION MEASURES
2.1   Defining diversification and concentration

2.1.1   Before outlining the various diversification measures, it is worth first exploring the underlying concept of diversification. On an intuitive level, diversification is easily understood and readily followed by professional and non-professional investors alike. In this sense, metaphors and similes are perfect vehicles for describing the idea. However, if one wants to actively manage diversification for a given portfolio of assets in a systematic manner, then one needs something a bit more tangible.

2.1.2   As a starting point, we posit that diversification is a well-defined multivariate function of underlying market variables. The exact functional form would undoubtedly be complex and understanding its output probably even more so. We thus take a reductionist approach and simplify the problem into a collection of several complementary components. Our goal then is to find a set of simpler components which, when taken either separately or in conjunction, provide a sufficiently accurate approximation of the true, more complex definition of diversification.

2.1.3   One possible starting point is to split the full set of market variables into two complementary sets: variables that are intrinsic to the portfolio and those extrinsic to the portfolio. This naturally gives rise to two complementary types of portfolio diversification; namely intrinsic and extrinsic.

2.1.4   Intrinsic portfolio diversification depends on (i) the characteristics of the underlying asset marginal distributions, (ii) the joint dependence structure between these distributions and (iii) the manner in which the assets are combined. In practice, one could thus define intrinsic diversification as a function of the number of assets, nominal weights, expected returns, volatilities and correlations of the assets within the portfolio. One could also consider additional input variables such as tail risk measures—such as Value-at-Risk (VaR) and expected shortfall—or higher-order moments—such as co-skewness and co-kurtosis.

2.1.5   Extrinsic portfolio diversification is defined similarly to intrinsic diversification except that the portfolio assets would be replaced with a finite number of external risk or return factors. For example, one could define extrinsic portfolio diversification as a function relating the proportion of total portfolio risk explained by a set of global risk factors. Another example would be to define extrinsic diversification as a function of the portfolio risk loadings on the seminal Fama–French (1993) equity risk factors.

2.1.6   While it may seem counterintuitive to break a single definition into a potentially large set of subordinate definitions, it does at least provide one with a tractable method of defining and managing distinct diversification components. Furthermore, this framework does not preclude the very likely possibility that some component definitions will be completely dominated—in the stochastic sense—by others. What it does allow though is for more than a single measure to be included when attempting to approximate the true, unknown portfolio diversification function.

2.1.7   Consider the realistic case of a manager that defines diversification purely

in terms of risk contributions from a set of extrinsic return factors.[1] While the equal factor risk contribution portfolio maximises diversification as described by the manager (see Section 6.5), it may only contain a handful of assets. If so, then this portfolio is still highly exposed to asset-specific risk and is thus arguably not in line with the qualitative goal of diversification. A simple solution would then be for the manager to jointly consider a risk-based diversification measure and a cardinality- or weight-based measure.

2.1.8　　　What about the idea of portfolio concentration? In general, practitioners and academics both use the term 'concentration' when referring to weight-based portfolio metrics, while 'diversification' is usually used in a risk setting. The term concentration has also become more prevalent in the fund performance analysis literature due to the recent introduction of metrics such as active share (Cremers & Petajisto, 2009). However, in the framework given above, concentration would simply be equal to the inverse of diversification. Thus, while we follow the trend in recent literature and write only about diversification, note that we are essentially also writing about portfolio concentration.

2.2　　Common diversification measures

2.2.1　　　Similarly to Fragkiskos (2014), we divide a wide range of diversification measures into categories based on each measure's primary portfolio focus—see Table 1. We also further divide the measures into intrinsic and extrinsic sets. As mentioned above, our goal is not to create an exhaustive list but rather to present those measures mostly commonly used in practice, as well as those recently proposed in the literature.

2.2.2　　　Sections 3 to 7 each describe a different category of diversification measure. This includes a description of the respective measures, a brief motivation for why it should be considered a measure of diversification and, finally, a practical application of the measures on historical FTSE/JSE equity index data.

TABLE 1. Categorised intrinsic and extrinsic diversification measures

Intrinsic diversification measures				
Cardinality	Weights	Returns	Risk	Higher moments
No. assets	Stock Herfindahl index	Expected return	Volatility	Portfolio entropy
Effective stock range	Effective no. constituents	Return gaps	Average correlation & intra portfolio correlation	Diversification delta
	Asset portfolio contributions		— Risk contributions — Diversification ratio — Portfolio diversification index — Minimum variance ratio	Asset co-skewness

---

1　These concepts are properly defined and discussed later in the paper.

Extrinsic diversification measures				
Cardinality	Weights	Returns	Risk	Higher moments
Effective factor range	Global factor exposures	Return factor exposures	Effective PCA dimensionality	Diversification delta
	Factor Herfindahl index	Factor model R^2	Effective no. independent bets	Factor co-skewness
	Factor portfolio contributions		— Effective no. minimum torsion bets	
			— Factor model R^2	

2.2.3   The purpose behind applying these measures to historical South African index data is twofold. Firstly, it gives one a diagnostic review of historical South African index diversification and secondly, provides a good description of the investable opportunity set that was available to fund managers in the past. The former should be of general interest to all followers of the markets, while the second should be of specific interest to professional investment practitioners.

3.   CARDINALITY-BASED MEASURES

Cardinality refers to the number of assets within a given portfolio. Markowitz (1952), among others, showed that portfolio variance is a decreasing function of portfolio cardinality unless all stocks are perfectly correlated. Evans & Archer (1968) conducted the first real-world study of this phenomenon by creating random equal-weight portfolios of increasing cardinality from a list of stocks on the S&P 500 Index and plotting the average portfolio volatility as a function of cardinality. They confirmed Markowitz's theoretical findings by showing that the idiosyncratic risk component of total portfolio risk tended towards zero as the number of stocks increased. Because of this inverse relationship, portfolio cardinality can either be regarded directly as a naïve measure of diversification or as another constraint within an existing optimisation framework.

3.1   Number of assets

Since the seminal Evans & Archer paper, this type of study has been replicated across a wide range of countries, asset classes and historical periods. Bradfield & Munro (2011) recently addressed the question of the optimal number of stocks to hold in the South African environment. They showed that, at the time of their study, between 30 and 50 stocks were needed for effective risk reduction in both benchmark and non-benchmark cognisant settings.[2] They also highlighted the role that benchmark concentration and average asset correlation had on their findings.

---

2   Effective risk reduction here is defined as being within 1–2% of the largest *N*-stock portfolio volatility

3.1.1 CONCENTRATION AND AVERAGE CORRELATION EFFECTS

3.1.1.1   We extend Bradfield & Munro's study in two ways. Firstly, we analyse the specific effects of index concentration and average pairwise correlation on portfolio volatility in a more general simulated market environment. Secondly, we apply Bradfield & Munro's effective stock range calculation to the SWIX over the period June 2003 to January 2015 and quantify the effects of average stock correlation and index weight concentration on total index volatility.

3.1.1.2   To study the effects of concentration and average correlation on portfolio volatility, we simulate a 170-asset market. Table 2 details the simulated market parameters and their possible ranges.

TABLE 2. List of simulated market variables and defined ranges

Market variable	Range
Asset volatility	$\sigma \in [20\%, 50\%]$
Average asset correlation	$\bar{\rho} = 0, 0.1, \ldots, 0.8$
Market concentration factor	$C = 0, 2 \ldots, 18$
	0 = equal weights – 20 ≈ extreme concentration
Total number of stocks	$N = 170$

3.1.1.3   Asset volatility is sampled uniformly from a given range of 20% to 50%. A random positive definite 170-asset correlation matrix is constructed so that the average correlation value, $\bar{\rho}$, is equal to the predefined value between 0 and 0.8. The portfolio weighting structure is defined by a concentration factor, C, and individual weights, w_i, are calculated by means of an exponential decay function:

$$w_i = \frac{\left[(N+1) - i \right]^C}{\sum w_i}. \tag{1}$$

3.1.1.4   As an empirical reference, the 167-stock ALSI C factor has ranged from 13 to 20 historically, the SWIX C factor from 9 to 15 and the Top40 C factor from 4 to 7.

3.1.1.5   For each generated covariance matrix, we then simulate 1000 random portfolios of different sizes ranging from 1 to 170 stocks for each pairwise combination of $\bar{\rho}$ of C. There are two aspects of randomness embedded within these portfolios. The first aspect relates to the specific assets selected from the total 170-asset universe in each run. The second aspect relates to the weights attached to those selected assets.

3.1.1.6   Figure 1 graphs average portfolio volatility as a function of number of stocks for a selection of concentration-correlation pairs. Specifically, we look at $\bar{\rho} = \{0, 0.2, 0.4\}$ and $C = \{0, 6, 12, 18\}$. Average correlation scenarios are identified by colour and concentration scenarios by shading, with equal weight given as the darkest line and extreme concentration as the lightest line.
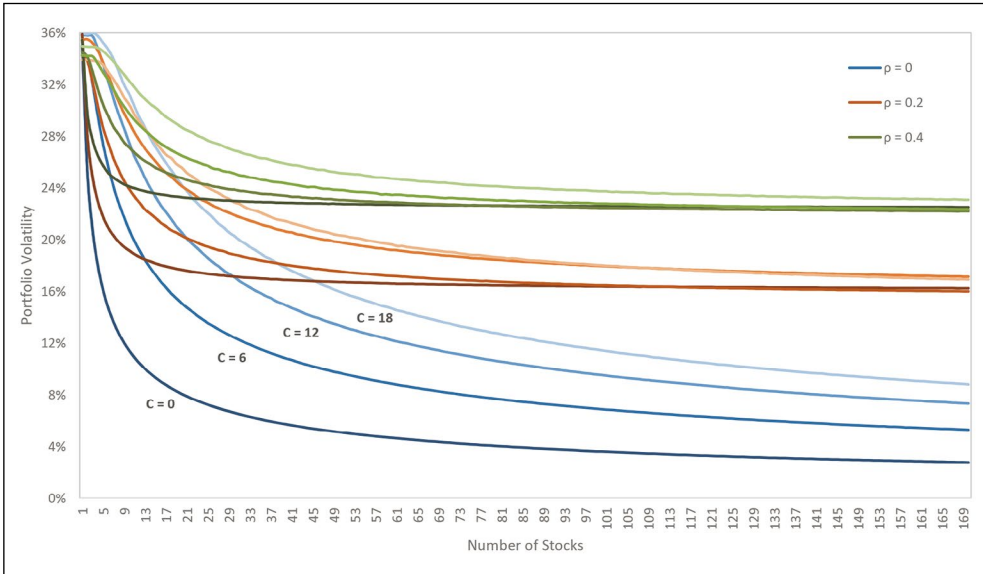
FIGURE 1. Portfolio volatility versus number of stocks for a selection of
concentration-correlation pairs

3.1.1.7   In all cases, average volatility is highest for the 1-stock portfolio at around 35%—the mid-point of the specified volatility range—and reaches a minimum for the 170-stock portfolios, allowing for some sampling error. Consider first the uncorrelated asset scenarios. The lowest line represents the equal-weight and uncorrelated portfolio. In this case, portfolio volatility is simply equal to the average asset volatility divided by the number of stocks. If we then increase portfolio concentration, we see that the average volatility curve shifts upwards and also flattens out. One would therefore need a larger number of assets in order for portfolio volatility to be within a given distance of the limiting 170-asset portfolio volatility.

3.1.1.8   If we now repeat the experiment but increase average asset correlation, then we see a significant increase in minimum portfolio volatility. Because of this, the volatility shift effect noted above is essentially removed and only the concentration gradient effect remains observable.

3.1.1.9   Following Bradfield & Munro (2011), we define the effective stock range (ESR) as the number of stocks required to reduce portfolio volatility to within 1–2% of total 170-asset portfolio volatility. Figure 2 displays ESR versus portfolio concentration for three average correlation scenarios.

3.1.1.10   We observe that ESR is an increasing function of portfolio concentration and a decreasing function of average correlation. Furthermore, we also observe that the ESR band narrows for higher correlations and lower concentration (bottom left of the graph). This is because higher correlations decrease the difference between minimum and maximum portfolio volatilities, while lower concentrations increase the curvature of the
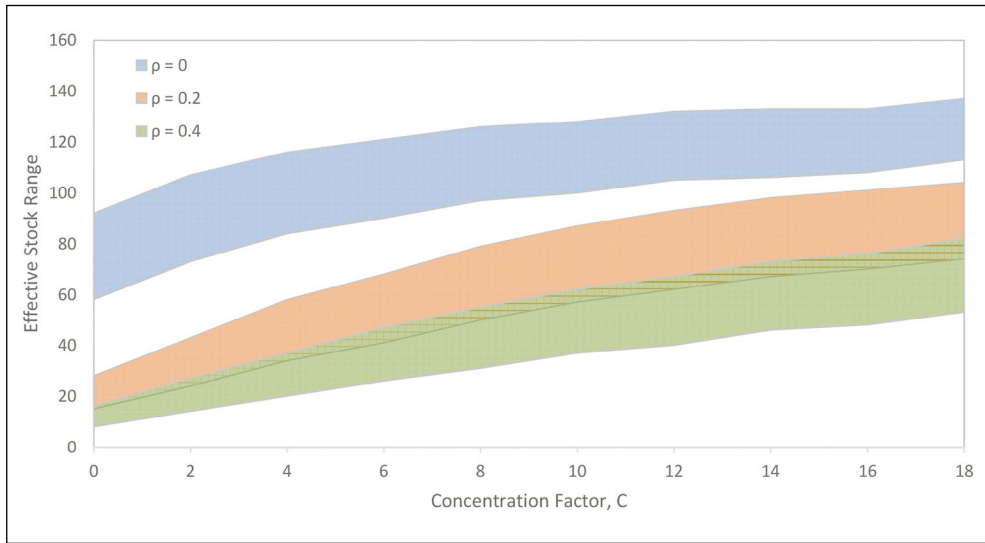
FIGURE 2. Effective stock range versus portfolio concentration for different average market correlations

volatility function—see the darkest line in the most correlated, upper group of Figure 1. For markets with average correlation around 0.2, Bradfield & Munro's 30–50 ESR occurs at a concentration factor of just below 4. In contrast, if average correlation is closer to 0.4, then we would only find an ESR of 30–50 for portfolios with a concentration factor around 8. Figure 2 therefore clearly highlights that index ESR should vary quite considerably over history in line with changes in average correlation and index concentration.

3.1.2   HISTORICAL EFFECTIVE STOCK RANGE FOR THE ALSI AND SWIX

3.1.2.1   To calculate a history of monthly index ESR, we generate 500 random portfolios of different sizes ranging from 1 to N_t stocks, where N_t represents the number of stocks in the index at a given time t. For each month from June 2003 to January 2015, we calculate a covariance matrix from the previous three years of weekly stock total return data—roughly 156 observations—and record the month-end index weights. Using the exact index constituents at each point in time avoids any survivorship bias or look-ahead bias. However, this does mean that certain stock histories may be shorter than the requisite three years. In these cases, we calculate pairwise correlations and volatilities based on the available return history. If the available data are not statistically sufficient, then we fill in the missing values with the cross-sectional average volatility and pairwise correlation values. This happens very rarely in our sample. In order to ensure that the final covariance matrix is positive semi-definite, we filter any eigenvalues below a small positive threshold. We are thus using a form of Bayesian shrinkage estimator for the covariance matrix (*sensu* Ledoit & Wolf, 2004).[3]

---

3   This same covariance estimation procedure is used throughout the paper.

3.1.2.2   The left-hand panel of Figure 3 displays historical SWIX volatility assuming (i) equal weighting and zero correlations (EWZC), (ii) equal weighting and market correlations (EW) and, finally, (iii) market weightings and correlations (SWIX). EWZC volatility thus represents the contribution of average asset volatility to total index volatility. The average correlation contribution is then represented by the difference between the EW and EWZC volatility curves, and the final concentration contribution is given by the difference between the SWIX and EW volatility curves. The right-hand panel of Figure 3 displays these proportions over time.

3.1.2.3   Given the large number of stocks in the SWIX, the average volatility contribution remains consistently small and comprises 19.5% of index volatility at the end of our sample. The correlation contribution is much larger, at an average of 54%, but does display slightly more variation over time. The concentration contribution is currently 34%, which is somewhat higher than its historical 30% average, having grown by nearly 4.8% in the last 18 months of the sample as the index has become increasingly concentrated in terms of stock weightings—see Section 4. Interestingly, even though the absolute index volatility has ranged between 10% and 25% over the last 13 years, the individual volatility contributions have remained relatively consistent. This is perhaps because average correlation and average volatility both increase significantly during market down-turns (Chua et al., 2009), thus maintaining the relative balance across the contributions. We analyse this further in Figure 4.

3.1.2.4   The left-hand panel of Figure 4 graphs the evolution of average portfolio volatility curves through time, plotted against the number of stocks included in the portfolio and normalised by the actual SWIX volatility at each given date. This normalisation is done to aid comparison over time and has the effect that the right-most point of each curve will always tend to 1, representing 100% of actual index volatility. These curves are thus a realised version of the theoretical average portfolio volatility curves given in Figure 1, and take into account the actual concentration and correlation within the universe of SWIX stocks over time.

3.1.2.5   The left-hand panel of Figure 4 lends further credence to the idea that correlations increase in tandem with volatilities during market crises. This can be seen by
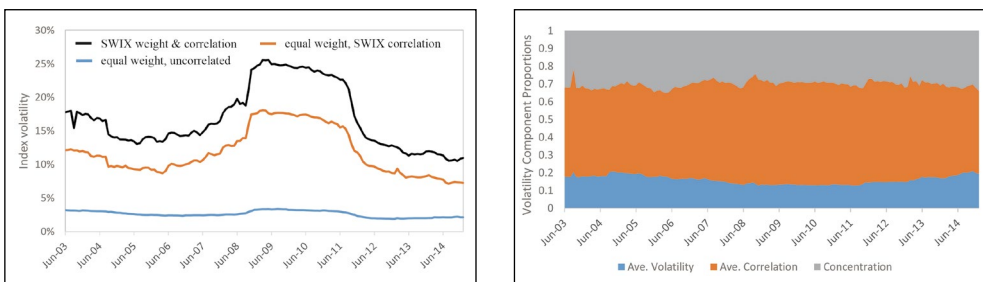

FIGURE 3. SWIX, equal-weight SWIX and equal-weight uncorrelated SWIX volatility (left), and contributions from each index volatility component (right), June 2003–January 2015

the much lower average normalised volatility values from mid-2008 to mid-2011.[4] The subsequent increase in the average values then coincides with the decline in total index volatility shown in Figure 3. The current high peaks are thus indicative of the low-volatility regime we are experiencing.

       3.1.2.6   The right-hand panel of Figure 4 displays the historical ESR for the SWIX index. As expected, we see a high degree of variability in both the bounds and size of this range. There is a steady decline in ESR until July 2012, where it reaches a minimum of 12–26 stocks. Thereafter, ESR increases to a range of 23–51 stocks by the end of the sample. The evolution of ESR roughly mirrors the evolution of the average one-stock normalised volatility, shown in the left-hand panel of Figure 4. This because both values are ultimately volatility-based. However, the reflection is not exact as ESR also includes the effects of changes in correlations, concentrations, expected returns (indirectly via the standard deviation calculation) and the number of index constituents. Therefore, ESR can actually be considered a fairly inclusive measure of total portfolio diversification.

4.   WEIGHT-BASED MEASURES

4.1   Despite recent advances, weight-based metrics are probably still the most commonly used diversification measures in practice. Evans & Archer (1968) were the first to suggest using the Herfindahl–Hirschman Index (HHI) from information theory as a means of summarising the dispersion within the portfolio weight distribution. The HHI is defined as the sum of the squared portfolio weights. The equal-weight portfolio thus gives the maximal HHI value of N while a single-asset portfolio gives the minimum value of 1. HHI-inspired metrics are now commonplace in the concentration literature.
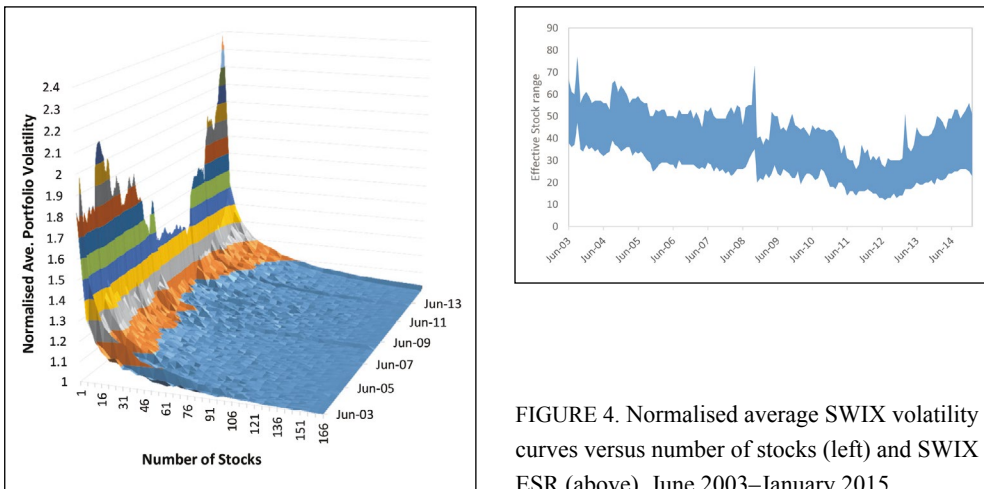

FIGURE 4. Normalised average SWIX volatility curves versus number of stocks (left) and SWIX ESR (above), June 2003–January 2015

---

4   Given that we use three years of return data to estimate the covariance matrix, we would expect to see amplified volatilities for the three years following the start of the financial crisis.

4.2   Another measure borrowed from information theory is Shannon entropy, which measures the total uncertainty within a given distribution.[5] Deguest et al. (2013) define the effective number of constituents (ENC) as the exponential of the weight entropy,

$$ENC = \exp\left(-\sum w_i \ln w_i\right). \tag{2}$$

4.3   Taking the exponential ensures that the measure is bounded between 1 and *N*. Higher ENC values indicate higher weight diversification and, as with the HHI, the measure is also maximised for an equal-weight portfolio.

4.4   Below, we analyse the historical stock and sector weights of the ALSI and SWIX. We also compare the inverse HHI and ENC measures for both indices. Figures 5 and 6 graph the ALSI and SWIX industry weights back to January 1996 and June 2003 respectively.[6] The colours represent the three 'super sectors': Resources at the bottom, Industrials in the middle and Financials at the top. Industries within each sector are then denoted by the different colour shadings.

4.5   Figure 5 paints a very interesting picture about the evolution of the index over a 20-year period. Evidently the aphorism that the ALSI was historically a bet on Resources only really holds from 2001 up to the end of 2010. Before the millennium, the Industrial sector held the highest weighting, followed generally by Financials. In January 1998, Industrials accounted for 46% of the ALSI, Financials another 29% and Resources only made up the remaining 25%. Fast forward to June 2008—and to the end of a 7-year commodities bull-run—and Resources now accounted for 57% of the ALSI market cap. Industrials were a clear second at 28% and Financials made up the remaining 15%.

4.6   Since the advent of the global financial crisis however, Resources have been under considerable pressure and large multinational industrial counters such as Naspers, Richemont, SAB and MTN have taken over as the main return generators in the South African equity market (see Section 5.2 for more on this). The market is now well and truly dominated by Industrials, with an ALSI weighting of 54% as of January 2015. Financials are once again the second largest sector at 32% and Resources account for the remaining 22%. A similar sector ratio for the SWIX gives Industrials at 55%, Financials at 29% and Resources at 15%. Therefore, from a sector concentration standpoint, the SWIX is effectively just as concentrated as the ALSI at present. This is worth bearing in mind when selecting an appropriate fund benchmark or hedging basket.

---

5   See Zhou et al. (2013) for a detailed discussion on entropy and its uses in finance.

6   The FTSE/JSE indices were only launched in June 2002. Prior to this, the index series was known as the JSE Actuaries IndicES; IT Included all JSE listed companies and there was no free float for constituents. Because of this, any data before 2002 can only be considered an approximation.

4.7   While Figures 5 and 6 allow one to examine sector and industry concentrations, it is difficult to surmise how the underlying stock concentration has changed over the period. We now consider the weight-based diversification measures outlined above. Figure 7 plots the normalised inverse HHI and normalised ENC measures for the ALSI and SWIX. Normalised values can range between 0 and 1 with 1 being an equal-weight portfolio and 0 being a single-asset portfolio.
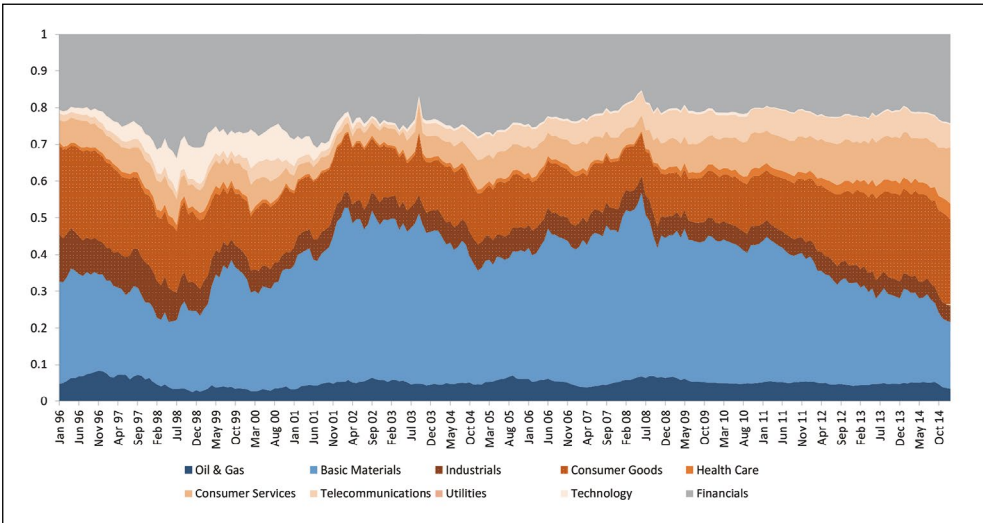

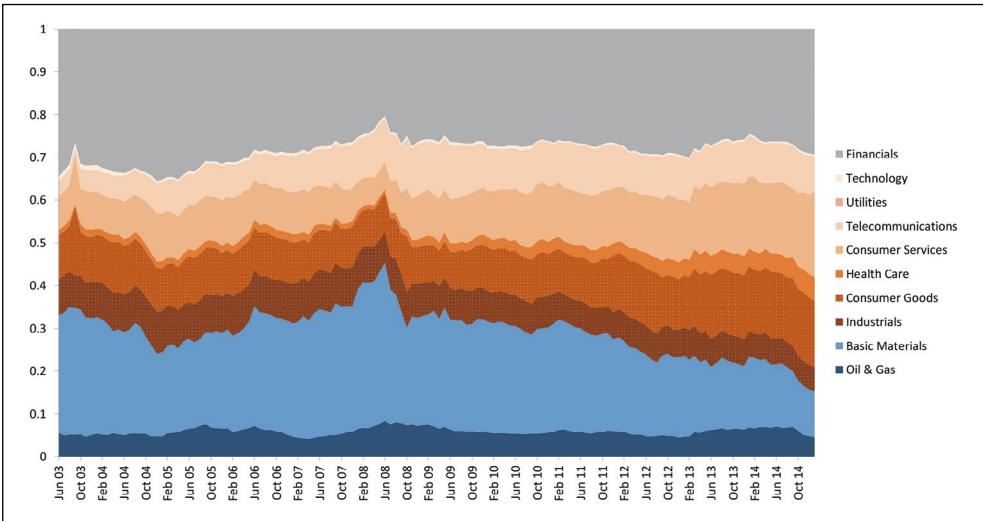FIGURE 5. ALSI weights per industry, January 1996–January 2015

FIGURE 6. SWIX weights per industry, June 2003–January 2015

FIGURE 7. ENC and HHI weight measures for the ALSI and SWIX,
January 1996–January 2015

4.8   The HHI measures (dark blue and dark brown) are considerably lower than the ENC measures (light blue and orange) for both indices. However, both measures display very similar patterns over time and thus encode similar information from a relative perspective. As expected, the SWIX is more weight-diversified then the ALSI although the spread between the two indices has declined considerably in the last two years of our sample, particularly as measured by the HHI.

5.   RETURN-BASED MEASURES

Diversification measures based predominantly on returns are quite rare. However, they do capture an aspect of diversification that many other measures often miss. Consider a market made up of two perfectly correlated assets. The majority of diversification metrics would find no diversification benefit in combining the two assets because the second asset is simply a scaled version of the first asset. However, this view misses two very important aspects of portfolio mathematics. Firstly, two assets can move in opposition and still be perfectly correlated so long as their expected returns are sufficiently different and their volatilities are sufficiently small. Secondly, even if the assets do move together, the magnitude of the returns may still differ significantly. By combining these assets, one eliminates the potential for extreme losses by trading off the potential for extreme gains. Thus, if the purpose of diversification is to mitigate extreme portfolio shocks—positive or negative—then portfolio diversification must surely be greater than the individual assets' diversification. One can thus conclude that return-based measures view diversification from a distinctly wealth-based perspective, where risk is defined as the probability of not being able to reach one's financial goals.

5.1 Number of assets

5.1.1 Statman & Scheid (2008) proposed the return gap (RG) measure as a means of capturing this aspect of diversification. RG is defined as the difference between the returns of a pair of asset over a given time period.[7] They show that RG is equal to twice the cross-sectional volatility in the two-asset case. Motivated by this link, Statman & Scheid propose that the expected RG can be written as a simplified function of average asset volatility and average asset correlation:

$$\mathbb{E}\left[RG\right] = 2\bar{\sigma}\sqrt{\frac{1-\bar{\rho}}{2}}. \tag{3}$$

5.1.2 The empirical, descriptive analogue to this would be the realised RG over a given period. In practice, even if two assets do exhibit higher correlations, there will remain the potential for return diversification as long as the realised RG is sufficiently large.

5.2 Realised index contributions

5.2.1 Another return-based diversification measure—again linked to cross-sectional volatility—is the index point contributions per constituent. As with the RG measure, we are interested in the cross-sectional difference in stock movements. However, instead of considering direct changes in asset value—i.e. returns—we consider the changes in the index point value of each asset. Doing so ensures that we include the effects of index concentration as well as the underlying asset return and risk characteristics—both of which are responsible for the realised asset returns.

5.2.2 Index contributions are also very useful in a performance evaluation framework as it gives one a very good indication of the investable opportunity set that was available to investors for a given basket of assets over a particular period. Knowing this helps one to evaluate investment decisions by identifying those assets which added—or removed—the most value to the portfolio over the period. The amount of value-add is also directly quantified.

5.2.3 Figure 8 displays the evolution of the top 10 and bottom 10 ALSI point contributors from February 2009—the starting point of the financial recovery—to December 2014. Industrial and Financial stocks account for the majority of the top 10 contributors, while the bottom five contributors are all Resource stocks. Diversification is then measured as a function of the mean and dispersion of this term-specific point contribution distribution, either in the form of point contribution gaps or cross-sectional contribution volatility.

5.3 Return factor diversification and adjusted-R^2

5.3.1 Another measure which can arguably be included in the return-based category is the adjusted-R^2 from a factor model of the underlying portfolio returns. The R^2

---

7 Note that this is known as alpha in the context of performance evaluation and is also commonly used when trying to find a suitable hedge instrument for a given portfolio of assets.
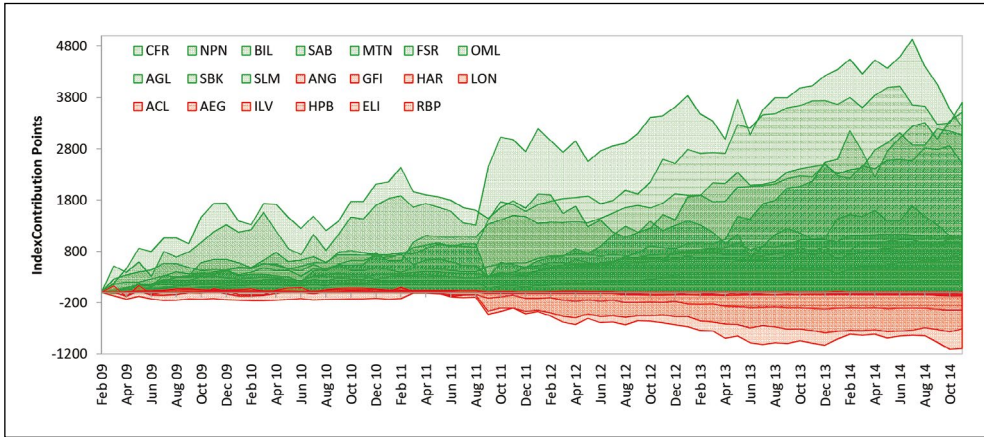
FIGURE 8. Distribution of top and bottom 10 ALSI point contributors,
February 2009–December 2014

value tells one exactly how much of the variation in portfolio returns over a given period was explained by a specified set of extrinsic return factors. The unexplained variability is then interpreted as the idiosyncratic portion of portfolio returns, given the specified set of return factors. This last point emphasises the fact that the choice of return factors is essentially one of the main variables in this diversification measure and also determines which aspect of portfolio diversification is being considered.

5.3.2   As an example, consider the simple yet robust three-factor global model given in Polakow & Flint (2015), which was found to explain a high percentage of South African large-cap equity index movements. A best-subsets regression (BSR) model was fitted to local index returns using three global factors—global equity momentum, global commodity pressure, and the USDZAR exchange rate—and allowed for some interaction. Given this choice of factors, the model approximates the global return component embedded in local index returns and the unexplained variability approximates the local idiosyncratic return component. In this model, one is thus estimating the degree of global integration of the South African equity market.

5.3.3   Figure 9 displays the monthly rolling from the global three-factor BSR model fitted to Top40 weekly returns. Since 2003, the global factors have accounted for an average of 70% of the variation in Top40 returns, leaving only 30% attributable to local effects.[8] Since mid-2010 though, the South African-specific component has grown to the point where, by the end of our sample, Top40 return movements were equally driven by local and global factors.

---

8   Results from a four-factor model including global bonds and a five-factor model including global risk sentiment are very similar to those from the three-factor model.
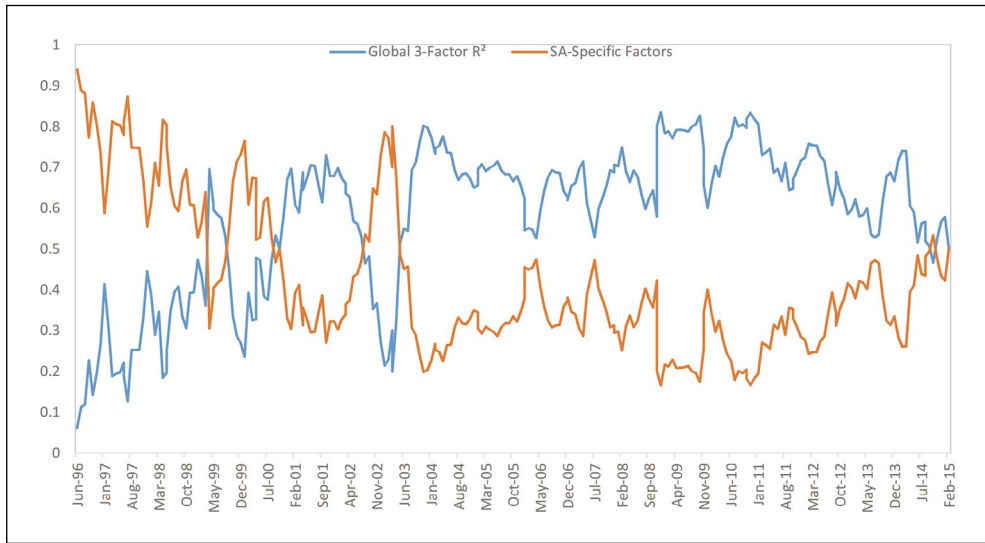
FIGURE 9. Rolling adjusted-R^2 and SA-specific risk component for a three-factor global model fitted to the Top40, June 1996–February 2015

6.   RISK-BASED MEASURES

6.1   The second most common method for defining diversification is from a portfolio risk perspective. The premise for risk-based diversification measures comes directly from Modern Portfolio Theory, namely, that portfolio volatility decreases as portfolio diversification increases. Therefore, portfolios that provide lower risk metrics must imply higher diversification.

6.2   In general, diversification risk measures are almost exclusively based on the second moment of the portfolio distribution and/or the underlying asset distributions. Fragkiskos (2014) neatly partitions risk measures into two categories: risk contribution measures and risk ratio measures. The former follows the risk budgeting framework set out by Maillard et al. (2010) and the latter follows the framework introduced by Tasche (2006). We firstly describe these measures in a manner similar to Fragkiskos (2014) and then apply them to historical South African index data.

6.3   A correlation interlude

6.3.1   Before introducing the full gamut of risk measures though, it is worth considering the most well-known—and misused—of such metrics: correlation. This is undoubtedly the market standard for evaluating dependence between two assets, with an increase in correlation signalling a decrease in diversification potential. Correlation studies are widespread, particularly since the advent of the global financial crisis. The time- and regime-dependent nature of correlations is generally regarded as a stylised market fact

nowadays and the idea of 'correlations tending to one' has become synonymous with periods of financial stress.[9]

6.3.2   The simplest measure of risk diversification is thus given by the average pairwise correlation of the assets comprising the portfolio. Figure 10 graphs the historical equal-weighted average correlation for the ALSI and Top40 asset universes respectively. It then further graphs the cap-weighted average correlations for the ALSI and SWIX, and the Top40 and Swix40 respectively. Cap-weighted average correlation has been dubbed intra-portfolio correlation (IPC) by Livingston (2013). The inclusion of index weights in IPC should lead to a better measure of the 'investable' dependence structure within each index.

6.3.3   Several observations are readily apparent from Figure 10. Firstly, from the left panel we note that the effect of the assumed weighting structure is significant, particularly for the much larger ALSI universe. Secondly, when comparing panels we find little difference between the ALSI (SWIX) and Top40 (Swix40) weighted correlations. This is not surprising though, given that the largest 40 stocks in the ALSI (SWIX) comprise over 80% of the total index. Finally, it is apparent from both panels that there has been a significant decline in all average correlation measures since mid-2011, with the ALSI value ending the sample at a low of 0.21 and the Top40 similarly ending at 0.24.

6.4   Risk ratios

6.4.1   Tasche (2006) formalised the idea of risk ratios as coherent measures of diversification. Specifically, for a generic risk measure β, Tasche defined the diversification factor as the ratio of the portfolio risk value, β_p, to the sum of the individual asset risk values, β_i:

$$DF = \frac{\beta_p}{\sum \beta_i}. \tag{4}$$

6.4.2   The idea of using risk ratios stems from the sub-additive property of coherent risk measures. That is, the sum of the individual asset risk values will always be greater than or equal to the total portfolio risk value. The difference between the two values is
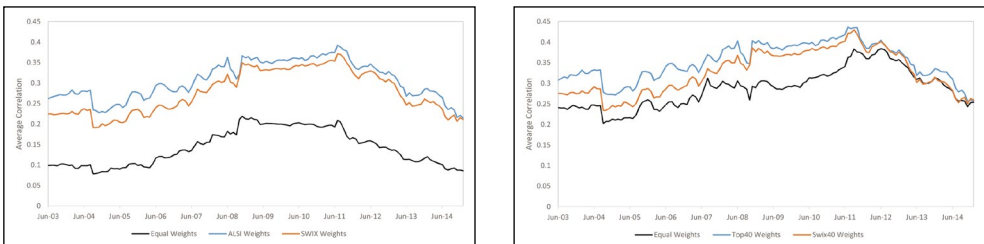

FIGURE 10. Cap-weighted and equal-weighted average correlation for the ALSI and SWIX (left) and the Top40 and Swix40 (right), June 2003–January 2015

---

9   See Ang & Chen (2002) and Chua et al. (2009) amongst others.

driven by how the assets have been combined and is thus taken as a measure of the portfolio's diversification.

6.4.3   The most well-known of the risk ratio measures is Choueifaty & Coignard's (2008) diversification ratio (DR), defined as the (inverse) ratio of weighted asset volatility to portfolio volatility,

$$DR = \frac{\sum w_i \sigma_i}{\sigma_p}. \tag{5}$$

6.4.4   Note that DR is actually very similar in concept to the analysis conducted in Section 3.1.2, which quantified the proportions of total portfolio volatility attributable to the average volatility, average correlation and concentration components respectively. DR has been well received in practice as it fits easily into the existing mean-variance paradigm, either as the new objective function to be maximised—giving rise to the most diversified portfolio (MDP)—or as another portfolio constraint.

6.4.5   Choueifaty & Coignard (2013) identified a number of important properties of DR and the MDP. Firstly, DR^2 can be reframed as a linear combination of the volatility-and-cap-weighted average correlation and the volatility-weighted concentration ratio. Secondly, DR^2 can be taken as a measure of the effective degrees of freedom within a given asset universe. Choueifaty & Coignard motivated this interpretation by showing that for a universe of N independent risk factors, the portfolio that weighted each factor by its inverse volatility would have $DR^2 = N$. We revisit this interpretation in later sections. Finally, they motivated for the MDP to be considered the 'undiversifiable' portfolio based on the fact that all assets not in the MDP are more correlated to the MDP than all assets within the MDP portfolio, which in fact all have the same correlation to the MDP.

6.4.6   Other popular risk ratios include Frahm & Wiechers' (2013) ratio of the minimum variance portfolio to actual portfolio variance, and Pérignon & Smith's (2010) ratio of the sum of individual VaRs to portfolio VaR. The idea of using tail risk measures rather than the symmetric volatility measure is compelling, particularly in the case where investors are looking at directly managing tail risk exposure. However, it must be noted that VaR is not sub-additive, which makes interpreting the ratio quite difficult. We instead suggest that the ratio of portfolio expected shortfall (or conditional VaR) to the sum of the component expected shortfalls would be a better tail risk diversification measure, as expected shortfall is sub-additive and arguably also gives one a better indication of actual tail risk.

6.5   Risk contributions

6.5.1   While the concept of portfolio risk budgeting is definitely not new, it has had a very strong resurgence in recent years largely due to the landmark studies of DeMiguel et al. (2007), Maillard et al. (2009) and Lee (2011). This has not gone unnoticed by practitioners. A significant proportion of global funds is now housed in portfolios which follow a risk budgeting mandate. In such a framework, diversification is defined in terms of the dispersion of the underlying asset (or factor) risk contributions. In this sense, the diversification

measures used here are very similar to the older weight-based measures, the only difference being that one is now considering weights in a risk space rather than a value space.

6.5.2   In the following subsections, we briefly outline the evolution of risk budgeting and provide some mathematical underpinnings. We then consider a particular subset of risk contribution measures based on the independent risk factors extracted via principal components analysis (PCA) and more advanced multivariate transformation methods.

6.5.3   A BRIEF INTRODUCTION TO RISK BUDGETING

6.5.3.1   Portfolio volatility, σ_p, is a function of underlying constituent volatilities, σ_i, weights, w_i, and correlations, ρ_{ij}, written as

$$\sigma_p = \sqrt{\sum_{i=1}^{N} w_i^2 \sigma_i^2 + \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j \sigma_i \sigma_i \rho_{ij}}. \tag{6}$$

6.5.3.2   The marginal contribution to risk per asset, MCR_i, quantifies the amount by which portfolio volatility will increase for a small increase in a given asset weight and is defined as

$$MCR_i = \frac{1}{\sigma_p}\left(w_i \sigma_i^2 + \sigma_i \sum_{j \neq i}^{N} w_j \sigma_j \rho_{ij} \right). \tag{7}$$

6.5.3.3   The total contribution to risk per asset is given by $TCR_i = w_i \times MCR_i$. It therefore follows that portfolio volatility can be reformulated as the sum of total risk contributions:

$$\sigma_p = \sum_{i=1}^{N} TCR_i. \tag{8}$$

6.5.3.4   The risk weights of a portfolio can thus be taken as the TCRs divided by the total portfolio volatility. Note that this framework works equally well for the case where portfolio volatility is described in terms of extrinsic risk factors.

6.5.3.5   The minimum variance approach was the first risk budgeting portfolio introduced in the literature as it is defined as the portfolio which gives equal MCRs. The next to be introduced was the naïve risk parity approach, where portfolio asset weights are inversely proportional to their respective volatilities. This can be thought of as the poor man's approximation to the equal-weight TCR portfolio but under the simplifying assumption of independent assets. In contrast, the more general equal risk contribution approach (ERC) introduced later ensures that all assets contribute equally to portfolio volatility under actual market conditions.

6.5.3.6   In a similar fashion to measuring portfolio weight concentration, one can characterise a portfolio's underlying risk concentration with the HHI or ENC measures but calculated directly from the portfolio's risk weights. In the case of the entropy-based measure,

we refer to this as the risk weight entropy (RWE). Under such a measure, the ERC portfolio would be considered most risk-diversified because it has equal risk weights by construction.

6.5.4 PCA- AND FACTOR-BASED RISK CONTRIBUTIONS

6.5.4.1   In uncorrelated markets, portfolio variance is simply equal to the weighted sum of underlying asset variances. In this case, the optimally diversified portfolio is the one in which asset weights are inversely proportional to asset variances (note, not volatilities). However, as noted in Section 6.3, assets in actual markets are not uncorrelated. A number of statistical methods exist to transform the original correlated asset space into an independent factor space. The most popular of these is PCA. We give a brief non-technical overview of PCA below and refer the interested reader to Joliffe (2002) for more information.

6.5.4.2   PCA is a multivariate statistical technique that allows one to factorise an N-asset covariance or correlation matrix into a set of N orthogonal (i.e. independent) eigenvectors and eigenvalues. The eigenvectors are typically referred to as 'component loading vectors' and describe a series of uncorrelated long/short asset portfolios, dubbed principal portfolios by Partovi & Caputo (2004). The variance of each principal portfolio is given by the corresponding eigenvalue. In practice, the principal portfolios are usually sorted in descending order, meaning that the first principal portfolio describes the highest proportion of the variation in the underlying covariance/correlation matrix. Ordinarily, only a handful of principal portfolios are needed to explain most of the underlying assets' variation. Thus, one can use PCA to significantly reduce the dimensionality of the original universe without any material loss of information. In the case of PCA on the correlation matrix, deciding on the number of principal portfolios to retain is often based on the Keiser–Guttman stopping criterion, which states that one should only keep those principal portfolios with eigenvalues (i.e. variances) greater than or equal to one.[10]

6.5.4.3   Polakow & Gebbie (2008) define effective dimensionality (ED) as the integer number of retained principal portfolios from a PCA on the asset correlation matrix. They found that only one dimension was necessary to explain the majority of movements and co-movements within the top five capitalised stocks on the JSE, and that there were only eight dimensions within the Top40 asset universe. Adding South African cash and bonds as well as a range of 10 broad international indices to the Top40 universe increased the dimensionality to 13. ED is thus qualitatively similar to the cardinality measures given in Section 3, but applied instead in an uncorrelated risk factor universe (as opposed to the correlated asset universe).

6.5.4.4   Another measure based on eigenvalues (i.e. principal portfolio variances) is Rudin & Morgan's (2006) portfolio diversification index (PDI),

$$PDI = 2\sum_{i=1}^{N}\left(i \bullet \tilde{v}_i\right) - 1, \tag{9}$$

---

10 Other more sophisticated methods include Velicer's MAP criteria and Horn's parallel analysis, both of which are applicable to PCA on either the covariance or correlation matrices.

where $\tilde{v}_i$ is the eigenvalue of the i^{th} principal portfolio normalised so that the sum of all eigenvalues equals one. A correlation matrix with equal eigenvalues would maximise the PDI at a value of N, while a correlation matrix with only a single principal portfolio will minimise PDI at a value of 1.[11] PDI is thus the risk-analogue of weight-based diversification, giving one a measure of how concentrated the distribution of principal portfolio variances are.

6.5.4.5   The problem with PDI is that it doesn't account for the underlying portfolio weights, which as we have already shown can have a large impact on the actual diversification level of a portfolio. Meucci (2010) addressed this issue by combining the above PCA approach with the ERC approach. In essence, Meucci followed the risk budgeting approach but applied directly to the independent principal portfolios rather than the correlated assets.[12] He then defined the effective number of independent bets (ENB), as the exponential entropy of the principal portfolio risk weights p_i :

$$ENB = \exp\left(-\sum_{i=1}^{N} p_i \ln p_i \right). \tag{10}$$

6.5.4.6   As with the original ENC, the ENB measure is bounded between 1 and N, with higher values representing higher diversification.

6.5.4.7   Meucci's ENB measure is arguably one of the most comprehensive diversification measures to date as it accounts for weights, volatilities, correlations and the underlying risk factors. However, it is not without its problems. A number of authors have since conducted backtests of the optimally diversified ENB portfolio and have reported several incongruities including: severe portfolio instability, non-unique weight solutions and counter-intuitive diversification results (see, for example, Lorhe et al., 2012). Our own unpublished testing has also confirmed these points as valid criticisms. Furthermore, principal portfolios are statistical constructs and are generally not easily interpretable in a financial sense.

6.5.4.8   Meucci et al. (2014) addressed these issues by suggesting a new method for transforming the original portfolio factors into the 'closest' set of uncorrelated factors. Closeness here is defined as the minimum tracking error between the original and transformed factors. Such a transformation solves the original PCA issues and also ensures that the new uncorrelated factors by and large maintain their financial interpretation. This latest diversification measure is called the effective number of minimum-torsion bets. While this method is indeed promising, it does add the further complication of selecting a suitable initial set of portfolio risk factors. Furthermore, the diversification measure is only comparable across portfolios as long as one uses the same set of underlying risk factors. That being said, it remains a highly promising (factor-conditional) measure of total portfolio diversification.

---

11 Note that one can also apply a stopping criteria before measuring PDI.
12 Because of the links to the ERC risk budgeting approach in an independent factor setting, Lorhe et al. (2012) dubbed this approach diversified risk parity.

6.5.5   HISTORICAL SOUTH AFRICAN RISK-BASED DIVERSIFICATION

6.5.5.1   We apply a range of risk-based diversification measures to the ALSI and SWIX. In particular, we calculate DR^2, ED, PDI, RWE and ENB. Before displaying these results though, we graph the industry risk weights from June 2003 onwards for the ALSI in Figure 11 and for the SWIX in Figure 12.

6.5.5.2   It is interesting to compare the industry risk weights to the industry cap-weights given in Figures 5 and 6. As one would expect, the risk weights show a much higher
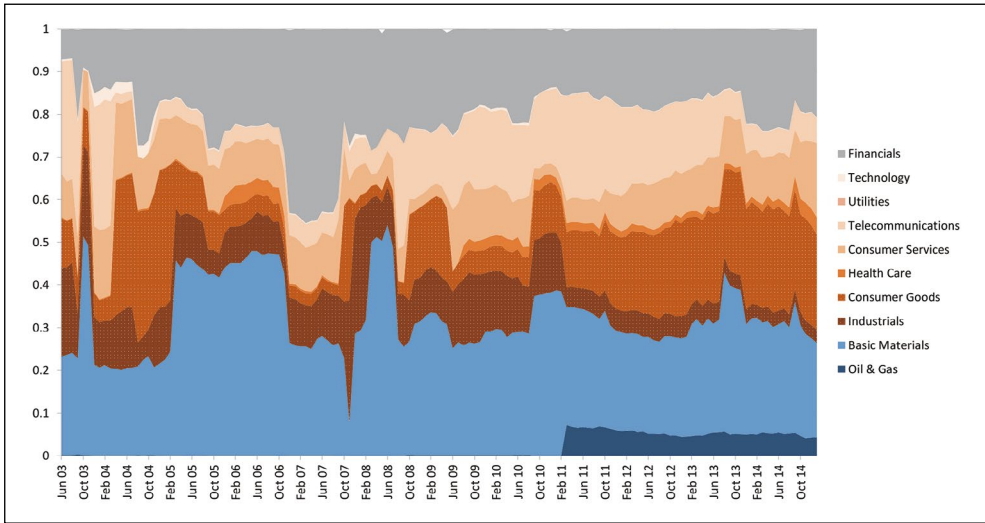

FIGURE 11. ALSI risk weights per industry, June 2003–January 2015
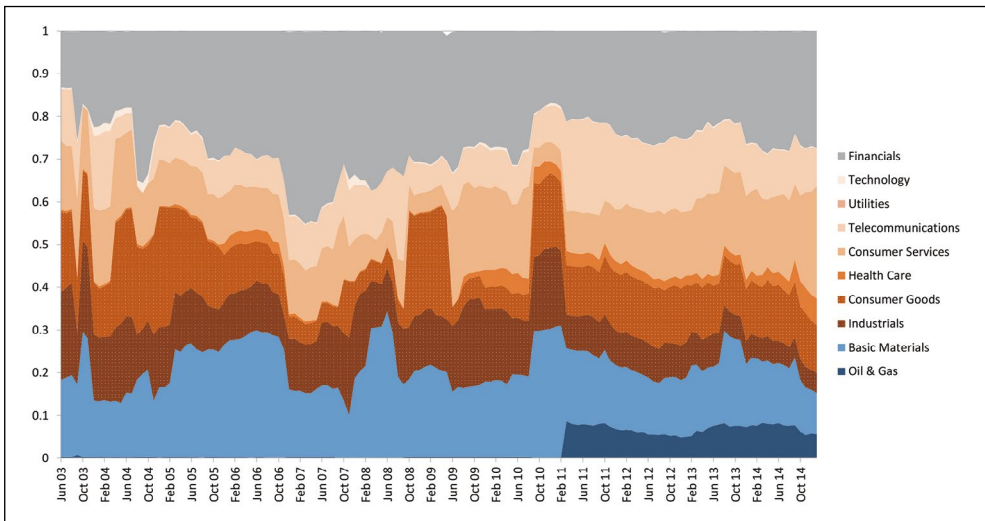

FIGURE 12. SWIX risk weights per industry, June 2003–January 2015

degree of variability over time than the index weights. The total Resources risk weight for the ALSI doesn't really display the same high values seen in its cap weight counterpart between 2001 and 2010. This is indicative of the fact that the mid-2000's bull run was fairly smooth and consistent. Industrial risk weights have comprised the greatest share of index risk since mid-2008, in line with a similar Industrials cap-weight trend. The underlying industry risk weights are still quite different relative to their cap-weights though. The Financials risk weight has generally always been smaller than the other sectors, apart from 2007. At the end of our sample, ALSI index volatility was split into 26% from Resources, 53% from Industrials and 21% from Financials.

6.5.2.3   In general, the SWIX sector and industry risk weights display slightly less variation than their ALSI counterparts. Furthermore, the industry risk weights also seem to be more in line with their industry cap-weights counterparts. As one would expect, Resources are clearly a tertiary concern for the SWIX with the majority of risk attributable to Industrial counters. The end-of-sample sector risk ratio for the SWIX shows Industrials at 58%, Financials at 27% and Resources at 15%; very similar to the cap-weight equivalent.

6.5.2.4   Figure 13 displays the time series of historical risk-based measures applied to the ALSI and SWIX. All the measures are bounded between one and the number of stocks within the index at the given date, ranging between 155 and 167. In general, all the measures agree on the high-level findings that (i) diversification decreased from the start of the financial crisis to slightly beyond its end, (ii) ALSI diversification has increased since 2013, and (iii) the gap between SWIX and ALSI diversification is closing.

6.5.2.5   The extreme range of values displayed above emphasises the fact that diversification has many different aspects and its quantification depends significantly on the specified definition. Consider first the DR^2 value. This is by far the lowest of all the measurements, with a range of 2.2–4.8 for the ALSI and 2.5–5.6 for the SWIX. While this may seem strange, Choueifaty et al. (2013) report that the MSCI World Index only had a DR^2 value of around 2.9 at the end of 2010. The higher range for the SWIX confirms our intuition
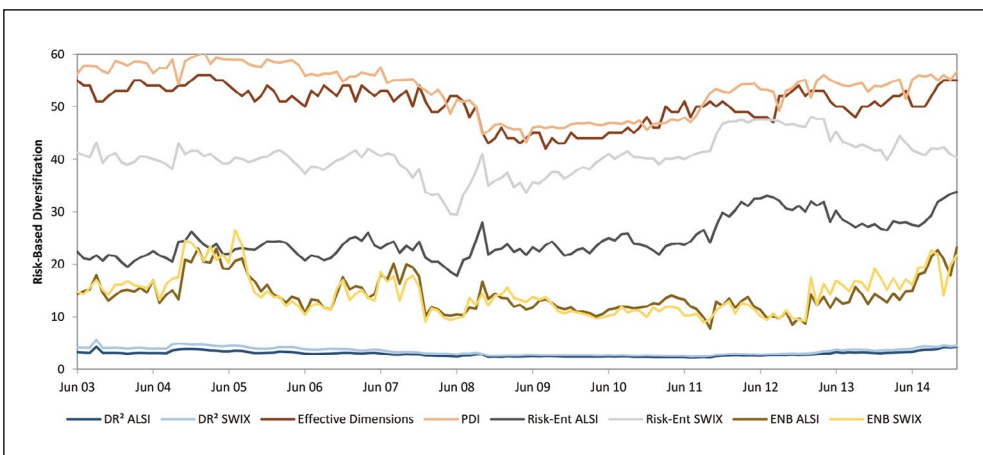

FIGURE 13. Risk-based diversification measures for ALSI and SWIX, June 2003–January 2015

that the SWIX has historically been more diversified than the ALSI. That being said, the recent decrease in the spread between ALSI and SWIX measures suggests that this may not always be so in the future.

6.5.2.6   The next lowest measure is ENB, which reaches a minimum of 7.8 in October 2011 and maximum of 26.5 in July 2005. We generally see very little difference between the ENB profiles for the ALSI and SWIX, both of which remained depressed from the start of 2008 to the start of 2013. Interestingly, it seems that the cap-weight and risk-weight differences between the two indices have little effect when measuring diversification and working in the independent PCA factor space.

6.5.2.7   RWE values for the ALSI reached a minimum of 17.8 just prior to the advent of the financial crisis and have been increasing—albeit rather bumpily—since then. ALSI RWE ended the sample at a maximum of 33.8. The SWIX RWE has historically ranged between 29.4 and 48.1, significantly higher than the ALSI. However, the SWIX and ALSI profiles have decoupled since the beginning of 2013, with SWIX RWE continuing to decline. The difference between ALSI and SWIX metrics also ended the sample at its lowest point, indicative of the fact that as Industrials have continued to gain traction at the expense of Resource counters, the SWIX will likely only become increasingly weight- and risk-concentrated.

6.5.2.8   The ED and PDI measures are based solely on the eigenvalues of the correlation matrix and are thus equivalent for both indices. Given their calculation methods, it is not surprising then that the ED and PDI profiles are essentially the inverse of the equal-weight average correlation profile displayed earlier in Figure 10; particularly so for ED. The PDI profile—ranging between 43.2 and 60.5—is generally slightly higher than the ED profile—between 42 and 56—although both measures are clearly capturing similar underlying information. The increasing trend from mid-2008 is again readily apparent and in line with the decrease in average asset correlations over the same period.

7.   HIGHER-MOMENT MEASURES

7.1   Samuelson (1967) was one of the first to observe that measurement of diversification using only the first two moments of a distribution may be too restrictive and crude. Several new higher-moment measures have been proposed in the literature, which we mention here for completeness sake. One such example, already discussed in Section 6.4, is to consider ratios of tail risk measures. Such measures implicitly account for the skewness and kurtosis of the portfolio and its underlying assets.

7.2   Another candidate is the negative entropy of the terminal portfolio distribution, introduced by Kirchner & Zunckel (2011). Diversification is thus defined as the amount of uncertainty within the portfolio. The more peaked the terminal distribution, the lower the uncertainty and the greater the negative entropy. In this instance, a full cash portfolio would be considered maximally diversified because it has the least uncertainty attached to its future value.

7.3　　Another higher-moment measure which combines entropy and the method of risk ratios is the diversification delta (DD), proposed by Vermorken et al. (2012). DD is defined as the ratio of the weighted average entropy of the individual assets to the entropy of the full portfolio. DD thus associates the diversification benefit with the difference between component entropies and portfolio entropy. Vermorken et al. (2012) suggested that the measure was bounded between 0 and 1, with diversification being greater for higher DD numbers. However, Salazar et al. (2014) recently showed that the DD measure has several flaws, the most worrying of these being that the measure is not sub-additive and can become negative under very common market conditions. Because of this, DD loses its straightforward interpretation as a bounded measure of diversification. Salazar et al. (2014) thus introduced a revised measure, DD*, which fixed these issues. However, in an application based on US stocks and bonds it turned out that the DD* measure was highly correlated with the Sharpe ratios calculated across the stock/bond weight spectrum and actually gave exactly the same portfolio weights when used to optimise the portfolio. This finding highlights the fact that if the assets under question are close enough to normality, there is no real benefit in considering higher order-moments. The true power of such a measure may be more apparent for portfolios including nonlinear derivatives, the return distributions of which are usually far from normal and generally bimodal. We leave this idea for further research though.

7.4　　Lastly, one could also use the risk ratio method on risk measures that directly focus on higher moments. One such metric is co-skewness, which quantifies the symmetry of an asset distribution relative to a benchmark asset's distribution symmetry (Harvey & Siddique, 2000). The risk ratio of this metric would therefore define diversification in terms of the effect that combining assets has on the overall symmetry of the portfolio distribution in relation to the average of the individual asset symmetries.

8.　　INTRODUCING DIVERSIFICATION PROFILES

8.1　　Up to now, we have applied a wide range of diversification measures in an *ex-post*, descriptive sense. That is, we have examined the historical diversification potential of a number of South African equity indices. However, most of these diversification metrics can also be applied in an *ex-ante*, inferential manner. A number of recent papers study the real-world performance of portfolios that maximise diversification as their main objective. Results to date have been mixed.

8.2　　DeMiguel et al. (2009), Demey et al. (2010), Lee (2011) and Homescu (2014) have shown that portfolios based solely on risk or diversification management produce superior risk and return statistics relative to cap-weighted or mean-variance alternatives. On the other hand, Scherer (2011) suggests that risk minimisation on its own is a meaningless objective. Furthermore, the reason behind the apparent outperformance is not because of an inherent diversification premium, but rather because the portfolios simply load onto existing return factor premia (Scherer, 2011; Lorhe et al., 2012). Poddig & Unger (2012) also questioned the

supposed robustness of such strategies by showing that empirical out-of-sample performance can be significantly worse than the traditional mean-variance optimal strategies, and that this is usually seen during crisis periods, exactly when the benefits of diversification are most wanted.

8.3   As with all financial optimisation though, the problem boils down to estimation error. There is no *a priori* reason to expect that diversification measures would be any less susceptible to estimation error than other common risk and return measures. However, instead of trying to find the single true estimate of future market conditions—a very difficult task—one can instead recast the problem into a more robust market-regime framework in an attempt to alleviate this issue.

8.4   If one assumes that the actual future market distribution is an unknown blend of known market regime distributions, then one can create a comprehensive range of future market possibilities over which to analyse any target portfolio. We showcase the potential of such a regime-dependent framework by introducing a regime-enhanced diversification measure within a simple balanced portfolio setting.

8.5   Flint et al. (2014) showed that the South African equity market was well partitioned by Kritzman & Li's (2010) turbulence index. This measure is essentially the multivariate equivalent of a univariate z-score, with the index, d_t, being defined as

$$d_t = (R_t - \mu_t) \sum\nolimits_t^{-1} (R_t - \mu_t)',$$

(11)

where R_t is the vector of asset returns, μ_t is the vector of expected returns and Σ_t is the asset return covariance matrix.

8.6   Financial turbulence allows one to classify each period in history as either turbulent or quiet by measuring how 'unusual' a given period is. Unusualness here is defined by the size of return deviations from their historical long-term norms, atypical correlations between assets relative to the historical long-term norms, or a combination of both features. Flint et al. (2014) showed that the ALSI returns during turbulent periods were significantly more volatile than usual and had a negative expected return. The differences between the identified regimes were thus significant enough to warrant considering regime-specific optimal portfolios.

8.7   In a similar manner to Flint et al. (2014), we construct a monthly turbulence index starting from May 1983 based on five major asset classes usually present in a generic South African balanced fund—namely, South African bonds and equities, global bonds and equities, and commodities. Months are identified as turbulent or quiet based on whether their turbulence index scores are above or below a fixed threshold value.[13] Of the total 420

---

13 The threshold value is calculated as the 75th percentile from a chi-squared distribution with five degrees of freedom. See Kritzman et al. (2012) for the technical justification.

months within the period, 91 are identified as turbulent. Table 3 gives the expected returns, volatilities and correlations for the turbulent subsample and Table 4 gives the same for the quiet subsample.

TABLE 3. Statistics for turbulent market regime, May 1983–December 2014

Turbulent	SA equity	SA bonds	Global equity	Global bonds	Commodities
Expected return	−0.64%	13.52%	−6.94%	−6.16%	−5.56%
Volatility	30.00%	13.46%	23.07%	9.11%	19.02%
Correlations					
SA equity	1				
SA bonds	0.41	1			
Global equity	0.56	0.19	1		
Global bonds	0.11	0.24	0.16	1	
Commodities	0.4	0	0.36	0.11	1

8.8 Except for South African bonds, each asset class displays significantly reduced expected return and increased volatility during times of turbulence. South African bond performance remains robust over both regimes, although quiet periods are characterised by much lower volatility. Correlations are also generally higher over the turbulent subsample, particularly so for the South African equity correlation pairs. The average correlation increases from 0.18 during quiet periods to 0.25 during turbulent periods.

TABLE 4. Statistics for quiet market regime, May 1983–December 2014

Quiet	SA equity	SA bonds	Global equity	Global bonds	Commodities
Expected return	31.09%	13.91%	17.73%	6.80%	6.12%
Volatility	17.56%	5.22%	11.70%	5.39%	9.64%
Correlations					
SA equity	1				
SA bonds	0.17	1			
Global equity	0.45	0.14	1		
Global bonds	0.03	0.21	0.2	1	
Commodities	0.29	0.01	0.14	0.13	1

8.9 A realistic range of possible future market conditions can then be calculated by blending the two identified regimes. Chow et al. (1999) suggest an elegant procedure for producing a blended estimate that accounts for one's view on the likelihood of each regime occurring in future, as well as one's relative aversion towards each regime. Focussing only on the risk aspect, the future market covariance matrix, Σ^*, may be calculated as

$$\Sigma^* = \lambda_T p \Sigma_T + \lambda_Q (1 - p) \Sigma_Q, \tag{12}$$

where Σ_T and Σ_Q are the estimated turbulent and quiet covariance matrices respectively, p is the probability of a turbulent regime occurring, and λ_T and λ_Q are the normalised risk aversion parameters for each regime. The future market expected return input can be calculated in a similar fashion to above.

8.10   We test this framework by calculating RWE (the exponential entropy of the risk weights) for a generic balanced portfolio across the different regimes. Table 5 reports the portfolio weights, risk weights and RWE values based on the full-sample, turbulent and quiet covariance matrices. The balanced portfolio highlights the well-known fact that the equity risk weight is generally much larger than the equity portfolio weight, resulting in much lower RWE numbers relative to portfolio ENC. The portfolio shows a significantly higher RWE value during turbulent regimes. This is due to the fact that South African bond volatility increases relatively more during turbulent regimes than other asset volatilities, which leads to a more equitable dispersion of risk weights.

8.11   The regime-specific values in Table 5 provide one with very useful information about how a given portfolio may perform in future market scenarios. However, this only identifies two particular future market scenarios. By using the blending procedure outlined above, we can calculate the balanced portfolio's diversification across a continuum of possible market conditions.

TABLE 5. Portfolio and risk weights for full sample, turbulent regime and quiet regime

	w_i	Full sample rw_i	Turbulent rw_i	Quiet rw_i
SA equity	45.0%	77.4%	73.1%	81.6%
SA bonds	25.0%	7.4%	10.3%	4.5%
Global equity	15.0%	12.2%	13.2%	11.3%
Global bonds	10.0%	1.0%	1.1%	0.9%
Commodities	5.0%	2.0%	2.3%	1.7%
ENC/RWE	3.94	2.17	2.38	1.94

8.12   Figure 14 displays the risk weights and RWE profile across this continuum, represented by the turbulent regime probability. Thus, we move from a perfectly quiet market on the left extreme to a perfectly turbulent market on the right extreme.

8.13   These profiles extend the existing diversification measures in two ways. Firstly, diversification profiles give a much richer quantitative description of portfolio diversification at any given time by illustrating how it changes across potential market regimes. In our example, we see that the balanced portfolio diversification varies the most during quieter
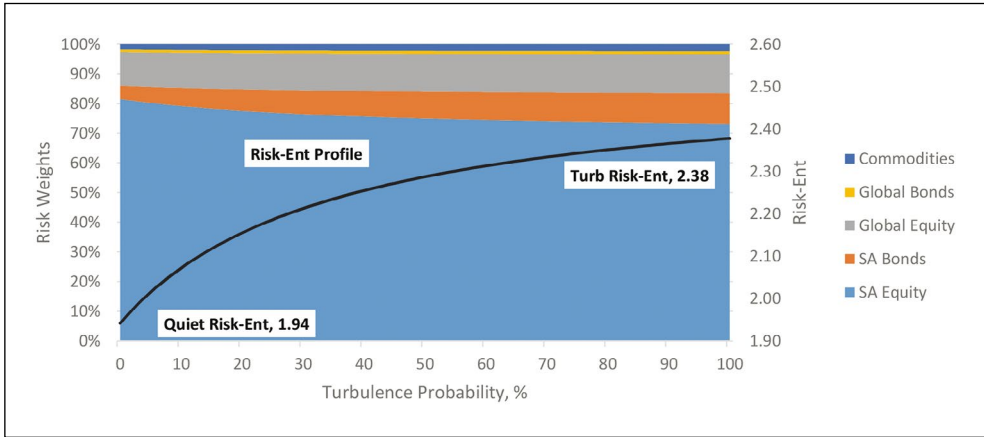
FIGURE 14. Portfolio RWE diversification profile versus turbulent regime probability

markets and flattens out at a higher level during more turbulent markets. Secondly, diversification profiles allow one to select the optimally diversified portfolio conditional on one's future view of the market expressed in terms of the likelihood of given market regimes. For example, if one assumes that there is an 80% likelihood that the future market will be turbulent, one can easily solve for the portfolio that maximises RWE for that particular market scenario. The resulting profile would also allow one to examine the robustness of the optimal portfolio in the event that the actual future scenario diverges from one's expectations.

9.   CONCLUSION

9.1   Diversification is a mercurial concept. Clearly, the qualitative purpose of diversification serves to mitigate the effect of specific sources of risk within any single asset class, and systemic sources of risk across asset classes. But a universally accepted quantitative measure of diversification has yet to be found. We have taken a reductionist approach to the problem of defining diversification. By focusing separately on different aspects of diversification measures, we showed that all existing diversifications are valid given the context in which they are applied and that each measure has its own advantages and disadvantages. A prudent course of action would thus be to consider several of these measures in combination when trying to analyse and understand portfolio diversification in a complete manner.

9.2   In this research, we summarised, motivated and applied a wide range of commonly used diversification measures to South African equity indices. Each measure was categorised based on which portfolio aspect it focused on: cardinality, weights, returns, risk or higher moments. The application of these measures to a range of South African equity indices served two purposes. Firstly, it acts as a diagnostic review of historical South African index diversification and secondly, provides a necessary description of the investable opportunity set available to fund managers in the South African equity space.

9.3   Finally, we introduced a methodology for extending the current suite of diversification measures to include regime-dependence. These regime-dependent diversification profiles give a much richer description of portfolio diversification than their single-value counterparts and also allow one to manage diversification proactively based on one's view of future market conditions.

AUTHOR DECLARATION
We declare no conflicts of interest in publishing this article and welcome any and all comments. The views expressed here are only our own and any errors are our responsibility.

REFERENCES
Ang, A & Chen, J (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, **63**(3), 443–94

Bradfield, D & Munro, B (2011). Number of stocks … the South African story. SAFA Conference proceedings, 2011

Choueifaty, Y & Coignard, Y (2008). Toward maximum diversification. *Journal of Portfolio Management*, **35(**1), 40–51

Choueifaty, Y, Froidure, T & Reynier, J (2013). Properties of the most diversified portfolio. *Journal of Investment Strategies*, **2**(2), 49–70

Chow, G, Jacquier, E, Kritzman, M & Lowry, K (1999). Optimal portfolios in good times and bad. *Financial Analysts Journal*, **55**(3), 65–73

Chua, DB, Kritzman, M & Page, S (2009). The myth of diversification. *Journal of Portfolio Management*, **36**(1), 26–35

Cremers, KM & Petajisto, A (2009). How active is your fund manager? A new measure that predicts performance. *Review of Financial Studies*, **22**(9), 3329–65

Deguest, R, Martellini, L & Meucci, A (2013). Risk parity and beyond-from asset allocation to risk allocation decisions. Available at SSRN 2355778

Demey, P, Maillard, S & Roncalli, T (2010). Risk-based indexation. Available at SSRN 1582998

DeMiguel, V, Garlappi, L & Uppal, R (2009). Optimal versus naive diversification: how inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, **22**(5), 1915–53

Evans, JL & Archer, SH (1968). Diversification and the reduction of dispersion: an empirical analysis. *Journal of Finance*, **23**(5), 761–7

Fama, EF & French, KR (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Econometrics* **33**(1), 3–56

Flint, E, Chikurunhe, F & Seymour, A (2014). Adapting to market regimes with timed hedging. Peregrine Securities, October 2014

Fragkiskos, A (2014). What is portfolio diversification? Available at SSRN 2331475

Frahm, G & Wiechers, C (2013). A diversification measure for portfolios of risky assets, in Batten, J, Mackay, P & Wagner, N (eds). *Advances in Financial Risk Management: Corporates, Intermediaries and Portfolios*. Palgrave Macmillan, 312

Harvey, CR & Siddique, A (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, **55**(3), 1263–95

Homescu, C (2014). Many risks, one (optimal) portfolio. Available at SSRN 2473776

Jolliffe, I (2002). *Principal component analysis*. John Wiley & Sons, Ltd

Kirchner, U & Zunckel, C (2011). Measuring portfolio diversification. arXiv preprint arXiv:1102.4722

Kritzman, M & Li, Y (2010). Skulls, financial turbulence, and risk management. *Financial Analysts Journal*, **66**(5), 30–41

Kritzman, M, Page, S & Turkington, D (2012). Regime shifts: implications for dynamic strategies (corrected). *Financial Analysts Journal*, **68**(3), 22–39

Ledoit, O & Wolf, M (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, **88**(2), 365–411

Lee, W (2011). Risk-based asset allocation: a new answer to an old question? *The Journal of Portfolio Management*, **37**(4), 11–2

Livingston, L (2013). Intraportfolio correlation: an application for investments students. *Business Education & Accreditation*, **5**(1), 91–105

Lohre, H, Neugebauer, U & Zimmer, C (2012). Diversified risk parity strategies for equity portfolio selection. *The Journal of Investing*, **21**(3), 111–28

Maillard, S, Roncalli, T & Teiletche, J (2010). The properties of equally weighted risk contribution portfolios. *Journal of Portfolio Management*, **36**(4), 60–70

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

Meucci, A (2010). Managing diversification. www.symmys.org

Meucci, A, Santangelo, A & Deguest, R (2014). Measuring portfolio diversification based on optimized uncorrelated factors. Available at SSRN 2276632

Partovi, MH & Caputo, M (2004). Principal portfolios: recasting the efficient frontier. *Economics Bulletin*, **7**(3), 1–10

Pérignon, C & Smith, DR (2010). Diversification and value-at-risk. *Journal of Banking & Finance*, **34**(1), 55–66

Poddig, T & Unger, A (2012). On the robustness of risk-based asset allocations. *Financial Markets and Portfolio Management*, **26**(3), 369–401

Polakow, DA & Flint, EJ (2015). Global risk factors and South African equity indices. *South African Journal of Economics*, **83**(4), 598–616

Polakow, D & Gebbie, T (2008). How many independent bets are there? *Journal of Asset Management*, **9**(4), 278–88

Rudin, AM & Morgan, JS (2006). A portfolio diversification index. *Journal of Portfolio Management*, **32**(2), 81–89

Salazar, Y, Bianchi, RJ, Drew, ME & Trueck, S (2014). The diversification delta: a different perspective. CIFR Paper, (047)

Samuelson, PA (1967). General proof that diversification pays. *Journal of Financial and Quantitative Analysis*, **2**(1), 1–13

Scherer, B (2011). A note on the returns from minimum variance investing. *Journal of Empirical Finance*, **18**(4), 652–60

Statman, M & Scheid, J (2008). Correlation, Return gaps, and the benefits of diversification. *Journal of Portfolio Management*, **34**(3), 132–9

Tasche, D (2006). Measuring sectoral diversification in an asymptotic multi-factor framework. Quantitative Finance Papers, http://arXiv.org/

Vermorken, MA, Medda, FR & Schröder, T (2012). The diversification delta: a higher-moment measure for portfolio diversification. *Journal of Portfolio Management*, **39**(1), 67–74

Zhou, R, Cai, R & Tong, G (2013). Applications of entropy in finance: A review. *Entropy*, **15**(11), 4909–31