Factor Investing In South Africa

Emlyn Flint
Peregrine Securities

Anthony Seymour
Peregrine Securities

Florence Chikurunhe
Peregrine Securities

Introduction

Risk factors and the strategies based thereon are fast becoming an integral part of the global asset management landscape.[1] The financial industry has adopted the moniker smart beta to describe such strategies as the term is both highly marketable and sufficiently broad to cover a wide range of investment products. However, in this report we will rather make use of the terms risk factors and/or risk premia when referring to underlying market drivers, and systematic strategies when referring to the dynamic investment strategies followed in order to gain exposure to these underlying risk factors. We do this not only to be more rigorous but also to draw attention to the practical fact that identifying a risk factor and subsequently harvesting returns from such a factor are essentially separate problems and need to be approached as such.

The latest annual smart beta surveys from FTSE Russell, EDHEC and MSCI all show variations of the same two major trends. Firstly, there are already a number of large international institutional investors that have sizeable factor-based portfolios and secondly, that many more investors are either in the process of reviewing such strategies or are looking to do so in the near future. In order to understand why risk factor investing has shown such a remarkable growth in popularity, it is worth briefly considering the greater history of portfolio management and asset pricing.

Nearly 70 years ago, Markowitz (1952) introduced the efficient frontier approach to asset allocation, which is still the most popular framework for constructing portfolios of assets. Under this framework, an optimal portfolio is defined as the combination of assets that maximises the expected return of the portfolio at a given time horizon for a specified level of portfolio risk (Meucci, 2001). In theory then, the portfolio construction problem had been solved. One simply needed to input the expected returns and covariances of the assets into the framework and out would pop an optimal portfolio specific to one's risk preferences. When applied in practice though, the model was found to be incredibly sensitive to small changes in the estimated mean returns and the optimisation procedure would almost certainly output unreasonable allocations. This behaviour led to Michaud (1989) coining the infamous phrase "error maximiser".

As a result, academics and practitioners alike then focussed their efforts into two separate areas in order to address the framework's weaknesses. The first area was based on all things risk-related: risk-based portfolio construction, more efficient risk estimates, and new risk and diversification measures. The result of this work has culminated in a rich risk budgeting and diversification approach. Roncalli (2013) provides an excellent review of generalised risk budgeting and Flint et al. (2015) provides a comprehensive study of diversification in the South African market.

The second area is based on all aspects of creating better expected return estimates. In particular, academics and practitioners went on the hunt for the underlying building blocks of asset classes in a similar manner to the way that physicists have hunted for the increasingly small and elementary particles from which all matter is comprised. The result of this search in the financial industry has given rise to the current factor investing paradigm. Podkaminer (2013) describes risk factors as the "smallest systematic units that influence investment return and risk characteristics" and Cazalet and Roncalli (2014) describe risk factor investing simply as "an attempt to capture systematic risk premia". Homescu (2015) further adds that the aim of factor investing is to construct portfolios in a systematic manner in order to gain exposure to a range of underlying risk factors.

The objective of this report is to construct a comprehensive range of risk factors for the South African equity market, analyse the historical behaviour of these factors and provide an overview of how such factors can be used in risk management and portfolio management. In order to achieve this objective, this research draws heavily on the excellent reviews written by Ang (2014), Cazalet and Roncalli (2014), Amenc et al. (2014), Homescu (2015) and Meucci (2016). We also make reference to Mutswari's (2016) recent work on testing the validity of a number of recent factor models for South African stock returns.

The remainder of this report is set out as follows. Linear Factor Models in Finance reviews the set of linear factor models used in finance and discusses the Fama-French factor models at length. South African Equity Risk Factors discusses the general factor construction process and the Fama-French construction methodology in detail. South African risk factors are introduced and thoroughly analysed. Factor-Based Risk Management then considers the application of these factors in risk management, focussing on risk attribution and returns-based style analysis. Factor-based portfolio management is discussed in Factor-Based Portfolio Management, with emphasis on creating multi-factor portfolios, and then the report concludes.

Linear Factor Models in Finance

Almost all finance studies throughout history have shown that there is a trade-off between risk and return. A natural question for investors then is what level of return can one expect to obtain for exposing oneself to a given level of risk? Traditionally, questions of this nature have been answered by using Linear Factor Models, or LFM's, which posit a linear relationship between an asset's expected return and its covariance with the risk factors incorporated in the model.

Meucci (2016) states that LFM's are used in almost every step of the risk and portfolio management process, including asset pricing, risk attribution and modelling, alpha prediction, portfolio optimisation and asset allocation. LFM's are also the cornerstone of factor investing as they are the main quantitative tool used to create systematic factor strategies. In this section, we briefly review the key LFM's used in the asset pricing literature and discuss at length the commonly used Fama-French-type factor models.

CAPM & APT

The capital asset pricing model (CAPM) was introduced by Sharpe (1964) and serves as the basis for all other factor models of asset returns. Based on the framework defined by Markowitz (1952), Sharpe showed that the risk premium on an asset (or portfolio of assets) was a linear function of a single market risk premium, represented by the market-capitalisation index. Mathematically, the CAPM states that

$$\mathbb{E}\left[R_i\right]-R_f=\beta_i\left(\mathbb{E}\left[R_m\right]-R_f\right), \tag{1}$$

where R_i and R_m are the returns on the i^{th} asset and market portfolio respectively, R_f is the risk-free rate, $\mathbb{E}[\bullet]$ represents the expectation and β_i is the beta – or sensitivity – of the i^{th} asset to the market portfolio, calculated as the ratio of the covariance of the asset and the market portfolio to the variance of the market portfolio:

$$\beta_i=\frac{\mathbb{C}ov\left[R_i,R_m\right]}{\mathbb{V}ar\left[R_m\right]}. \tag{2}$$

Beta thus measures the level of non-diversifiable, systematic risk embedded within any asset. Given that there is only a single market risk factor, the CAPM states that the reward for taking on additional risk is directly proportional to the underlying market risk. Therefore, everyone should hold the market portfolio in equilibrium as it is the only risk that is truly rewarded. While extremely elegant, there have been countless studies since its introduction that have shown that the theoretical CAPM is not validated by empirical evidence.

Ross (1976) proposed an alternative model, known as arbitrage pricing theory (APT) based on the increasing evidence of multiple market risk premia. Ross posited that the return of an asset is driven by a combination of random market factors and that this can be modelled with an LFM:

$$R_i=\alpha_i+\sum_{j=1}^{m}\beta_i^j\mathcal{F}_j+\varepsilon_i, \tag{3}$$

where α_i is a constant, β_i^j is the sensitivity of asset i to factor j, $\mathcal{F}_j$ is the return on factor j, and ε_i is the iid error – or stock-specific risk – term, which is also independent from any of the risk factors. It can be shown from Equation 3 that under APT, the risk premium on an asset is given by

$$\mathbb{E}\left[R_i\right]-R_f=\sum_{j=1}^{m}\beta_i^j\left(\mathbb{E}\left[\mathcal{F}_j\right]-R_f\right) \quad (4)$$

Equations 3 and 4 form the basis of nearly all risk attribution systems and systematic factor strategies. One of the challenges in using the APT though is that it is left to the user to define what the underlying market risk factors really are. In this vein, Cazalet and Roncalli (2014) define three main risk factor categories. The first category comprises factors based purely on statistical asset data – e.g. Principal Components Analysis risk factors. The second category comprises factors based on macroeconomic data – e.g. inflation and GDP growth. The final category comprises factors based on market data. This can be further classified into those factors based on accounting data – e.g. size and value – and those based on price data – e.g. momentum and low volatility. In this work, we focus mostly on the third category of risk factors.

The Fama-French Model and its Extensions

Fama-French Three-Factor Model

Based on the prior empirical studies that analysed numerous potential risk factors, Fama and French (1993) proposed a three-factor model for equity stock returns, which has since become the industry standard. This model linearly combines accounting- and price-based factors in the form

$$R_i-R_f=\alpha_i+\beta_i^m\left(R_m-R_f\right)+\beta_i^{smb}R_{smb}+\beta_i^{hml}R_{hml}+\varepsilon_i \quad (5)$$

R_{smb} is the return on a long/short portfolio of small/big market capitalisation stocks and R_{hml} is the return on a long/short portfolio of high/low book-to-market stocks.[2] These are known as the size factor and value factor respectively. Because market capitalization and value ratio indicators are correlated, Fama and French (1993) use a two-way sorting procedure to strip out any confounding factor effects. The value factor thus captures the value premium that is independent of the effect of size and the size factor captures the size premium that is independent of the effect of value.

There has been much literature aimed at assessing the appropriateness of the Fama-French three-factor model in equity markets worldwide. In the South African context, van Rensburg (2001) and van Rensburg and Robertson (2003) provide some of the earliest comprehensive assessments of Fama-French-based APT models on the Johannesburg Stock Exchange (JSE). Although not testing the exact Fama-French three-factor model, they show convincingly that one needs to incorporate several risk factors in order to accurately model the cross-section of equity returns on the JSE. More recent studies in the same vein include the works of Mutooni and Muller (2007), Basiewicz and Auret (2009, 2010), Strugnell et al. (2011) and Muller and Ward (2013), among others. Although these studies report differences in the magnitudes and significance levels of certain equity risk factors, they all conclude that a broader APT-based factor model is required to model South African equity markets correctly. The difference in study results is also to be expected, given the variations in data period and method across the various studies. As both Amenc et al. (2014) and Cazalet and Roncalli (2014) note, risk factors can be both cyclical and market-specific.

Carhart Four-Factor Model

Motivated by the evidence provided by Jegadeesh and Titman (1993) on the existence of significant medium-term price momentum trends, Carhart (1997) introduced a four-factor model based on Fama and French's work but including a momentum factor. This has since become the standard model used in fund performance and persistence literature. Mathematically, the Carhart four-factor model is given as:

$$R_i-R_f=\alpha_i+\beta_i^m\left(R_m-R_f\right)+\beta_i^{smb}R_{smb}+\beta_i^{hml}R_{hml}+\beta_i^{wml}R_{wml}+\varepsilon_i \quad (6)$$

where R_{wml} represents the return on a long/short portfolio of winner/loser stocks, based on the previous 12-month's price performance. Although initially met with severe scepticism, the momentum factor is now referred to as the "premier market anomaly" (Fama and French, 2008). Studies have confirmed the presence of this anomaly across numerous geographies and asset classes, making it the most prevalent market factor to date (Moskowitz, Ooi and Pedersen (2012), Asness et al. (2013)). Perhaps the reason for this pervasiveness is because the momentum factor is in essence a behavioural artefact, driven by cognitive biases which are unlikely to disappear in the near future (Antonacci, 2013). The same is perhaps not true about the justifications of the size and value factors.

Fama-French Five-Factor Model

In the time since Fama and French's (1993) initial work, many authors have shown that the three-factor model and even the four-factor model may well not be sufficient to explain the variation in the cross section of asset returns. To this effect, Fama and French (2014) introduced a novel five-factor model which included factors relating to the profitability and level of investment made by a company. In contrast to their original model, which is based on APT and empirical market research, the justification for the five-factor model stems from the bottom-up dividend discount model. Specifically, Fama and French (2014) suggest that expected stock return, as modelled by the dividend discount model, is based on three variables, namely the book-to-market ratio, expected earnings and expected growth in book equity – what they dub 'investment'. From their investigations, they posit the following five-factor model:

$$R_i-R_f=\alpha_i+\beta_i^m\left(R_m-R_f\right)+\beta_i^{smb}R_{smb}+\beta_i^{hml}R_{hml}+\beta_i^{cma}R_{cma}+\beta_i^{rmw}R_{rmw}+\varepsilon_i \quad (7)$$

where R_{cma} represents the return on a long/short portfolio of conservatively/aggressively invested stocks, and R_{rmw} represents the return on a long/short portfolio of robust/weak profitability stocks. Apart from the dividend discount model, the inclusion of these two factors was also influenced by the work of Novy-Marx (2013) and others, who showed that high profitability (or quality) stocks are rewarded with a significant and consistent premium, even after accounting for the return stemming from the original risk factors. Asness et al. (2013b) have since refined Novy-Marx's proxy of profitability/quality and proposed a new long/short factor of quality/junk stocks, where quality is defined as a composite score based on the dividend discount model and comprising numerous single accounting values. For the remainder of this paper, we will focus only on Fama and French's (2014) version of the profitability (i.e. quality) factor.

Asness et al. Six-Factor Model

Given that the Fama-French five-factor model is motivated by the dividend discount model, which describes the long-term behaviour of expected stock returns, the absence of the shorter-term momentum factor becomes somewhat more understandable. However, its exclusion is still surprising given that these very same authors named momentum as the premier market anomaly. In addition to this observation, Asness et al. (2015) also suggest that value and momentum are complementary risk factors and should be placed together. As a result, they propose a six-factor model extension which includes the momentum factor and makes use of a slightly adjusted value factor:

$$R_i - R_f = \alpha_i + \beta_i^m \left(R_m - R_f\right) + \beta_i^{smb} R_{smb} + \beta_i^{hml} R_{hml}^* + \beta_i^{wml} R_{wml} + \beta_i^{cma} R_{cma} + \beta_i^{rmw} R_{rmw} + \varepsilon_i \quad (8)$$

According to their results, the six-factor model provides a more complete explanation of the variation in historical US stock returns than the five-factor model and the adjusted value factor, which was shown to be nearly redundant by Fama and French (2014) before adjustment, now remains a significant risk factor.

Other Risk Factors

In what has now become one of the classic empirical finance papers, Harvey et al. (2015) surveyed hundreds of asset pricing papers published over the last fifty years and tallied more than 300 factors that are purported to explain the variation in the cross-section of expected returns. This concerted exercise in data mining led to Cochrane (2011) coining the phrase "the factor zoo".

The proliferation of purported factors is also partly a consequence of the popularity of the factor investing paradigm: factors are now everywhere and everything has become a factor. Cazalet and Roncalli (2014) suggest that this is arguably the most pernicious fantasy in the factor investing literature. Instead, they state that there are only a handful of risk factors that represent true risk premia or market anomalies. Ang (2014) suggests four main criteria for determining whether an observed market phenomenon is actually a true risk factor:

1. It should have strong support in academic and practitioner research and strong economic justifications.
2. It should have exhibited significant premiums to date that are expected to persist.
3. It should have history available during both quiet and turbulent market regimes.
4. It should be implementable in liquid, traded instruments.

Although the final criterion is not strictly required if only using the factor model in a risk attribution setting, it is still vitally important for creating tradable systematic factor strategies.

The factors we have discussed so far are all considered to be true risk factors in the sense that they are prevalent across nearly all markets studied to date, have valid economic and/or behavioural justifications and have histories stretching back more than a hundred years in some cases. In addition to these well-established risk factors, there are also a handful of recently discovered factors that are fast becoming accepted as true risk factors.

Two such recent factors attempt to capture the observed empirical phenomena that low volatility stocks outperform high volatility stocks and, similarly, that low beta stocks outperform high beta stocks. Ang et al. (2006) and Blitz and van Vliet (2007) popularised the idea of the low volatility factor and showed significant premium levels attached to this factor across a range of markets. Baker et al. (2014) and Frazzini and Pedersen (2014) among others have since confirmed their results and refined the economic rationale, further justifying the observed risk premia.

The low beta factor can be traced all the way back to Black (1972) and the leverage effect. Despite this lengthy history, the factor has only come back into vogue in the last ten years. Interestingly, van Rensburg and Robertson (2003) showed early on that the low beta anomaly commanded a significant premium in the South African equity market and could be accessed by sorting portfolios into quintiles based on their CAPM betas.

Other common factors not considered in this work are the carry (i.e. dividend yield), liquidity and quality factors. The carry risk factor is perhaps the most easily accepted in South African markets, where both the FTSE/JSE Dividend Plus Index and dividend-based unit trusts have existed for many years already. The liquidity factor is also easily appreciated in South African markets given its extremely high levels of concentration and the constant problem of capacity that many of the larger fund managers are faced with. Even though the strategy is accessed by going long illiquid stocks and shorting liquid stocks, it is unlikely that one could ever easily trade a South African liquidity factor in any decent size. For this reason, we leave this factor for future consideration. Finally, we have the quality factor. As mentioned above, the Fama and French (2014) profitability factor is essentially equivalent to the Novy-Marx (2012) version of quality. Although the more involved definition by Asness et al. (2013b) is arguably a better proxy for the true quality factor, it is also considerably more complicated to manufacture. For the sake of simplicity then, we leave this more advanced quality factor for future consideration.

South African Equity Risk Factors

In Linear Factor Models in Finance, we outlined several of the most popular APT-based factor models used in practice which have become essential risk and portfolio management tools. Although the selection of an optimal model specification remains an open question, it is clear that the underlying risk factors used in these competing models will continue to remain relevant for the foreseeable future. To this end, there are several online, open-source risk factor databases for large international equity markets.[3] However, and despite the South African-based factor studies mentioned earlier, a similar database does not exist – or at least is not publically available – for the South African equity market.

One of the goals of this research is to create a growing database of South African equity risk factors – and underlying stock variables – constructed as per the international asset pricing literature. In particular, we construct seven Fama-French style factors: size, value, momentum, profitability, investment, low volatility and low beta. Our hope in doing so is to make available to market participants an independent factor database that enables one to run a number of important risk and portfolio management factor applications in line with international best practice. This online South African factor data library can be found at https://www.preregrine.co.za/Content/PeregrineSecuritiesResearch.

Generalised Factor and Signal Processing

The factors discussed in this work are based on the Fama-French portfolio sorting methodology, which we will outline shortly. However, it is important to realise this is simply a special case of a more general signal processing framework. Meucci (2016) outlines three steps in the general allocation policy for systematic strategies. Firstly, process the set of current information into one or more factor signals. Secondly, transform these signals into a single set of consistent characteristics (i.e. expected return estimates) on the underlying stocks. Thirdly, construct optimal portfolio weights as a function of the transformed signal characteristics.

The initial step can be broken further into data collection, signal generation and signal processing. Consider a momentum signal for example. After collecting the requisite price data and correcting for any corporate actions and dividend payments, one uses a defined function to create factor scores. This could be as simple as prior 12-month return or something more complicated like a Hull moving average filter. Finally, these scores are filtered over time and/or cross-sectionally in order to create factor signals. Common filtering techniques include smoothing over time, scoring to reduce volatility, ranking cross-sectionally, twisting ranks nonlinearly, and trimming or Winsoring outliers.

The second step is not usually carried out when constructing single factors but is vitally important when considering multiple factors. For example, consider a universe of stocks that have both momentum and value scores. One then needs to define a methodology for creating a single consistent characteristic value for each stock that is consistent with both sets of factor scores. Such methods can vary from basic portfolio sorts to complex nonlinear programming solutions. We revisit this point in the section on Factor Portfolio Mixing and Integrated Factor Scores.

Finally, create an optimal portfolio based the estimated stock characteristics, a given satisfaction index and a set of constraints. This implementation step is ultimately what separates systematic factor strategies from underlying risk factor portfolios. In special cases, one can directly trade the underlying risk factors but usually investors are faced with real-world constraints that make this impossible. For example, long-only investors wanting to gain exposure to the long/short Fama-French value factor need to use optimisation techniques in order to maximise targeted factor exposure while minimising unwanted factor exposures. See the section on Factor Risk Attribution and the Factor Efficiency Ratio for more on this.

Constructing South African Risk Factors

We now consider the Fama-French construction methodology in light of the general factor framework outlined above. The dataset consists of the 383 constituents of the FTSE/JSE All Share Index (ALSI) over the period January 1996 to August 2016. All available total return and fundamental stock data were obtained from Bloomberg and INet for the 20-year period. Due to severe limitations on available fundamental data, the initial starting date had to be moved forward to December 2002, thus yielding a final sample period of just less than 14 years.

The majority of Fama-French risk factors are based on fundamental stock variables, with the remainder based on price information variables. The definitions of each such variable were kept consistent with the relevant international literature. At any particular month in the analysis window, the factor variables are defined as follows:

- **Size** is defined as the market value of the stock as at the end of the previous month. The shares in issue are taken directly from the underlying FTSE/JSE index data and multiplied by the index-recorded share price to obtain the gross market capitalisation.
- **Value** is defined as the ratio of book value to market value (BtM). This ratio is computed by taking the most recent book value six months prior to the current month and dividing it by the market value as at the end of the previous month. This is slightly different to the original definition but is in line with the alteration proposed by Asness and Frazzini (2013).
- **Momentum** is defined as the prior twelve month total stock return, less the prior month's return to account for any short-term reversal effects.
- **Profitability** is defined as the ratio of operating profit (total annual revenue, net of sales and other expenses) to the most recent book value for the previous year.
- **Investment** is defined as the relative growth in total assets six months prior to the current month.
- **Low volatility** is defined as the standard deviation of weekly total stock returns measured over the three years prior to the current month. If three years of weekly return data are not available, a smaller history is used with the minimum period required being one year. This is the factor definition proposed by Blitz and van Vliet (2007).
- **Low beta** is defined as the CAPM beta estimated from weekly excess total stock returns and excess ALSI returns, measured over the three years prior to the current month. If three years of weekly return data are not available, a smaller sample is used with the minimum period required being one year. This is the factor definition proposed by Blitz and van Vliet (2007).

The stock universe available for factor construction at any given month is taken as the historical ALSI constituent basket for that month. In order to isolate the true premia of the underlying factors, Fama and French (1993) employ a basic two-way portfolio sorting methodology. We create long/short factor returns in a consistent manner:

1. First rank all stocks according to their size score. Using the 50th percentile as a break point, create two subsets

	Book-to-Market Value Portfolios		
	Growth	Neutral	Value
Small	SG	SN	SV
Big	BG	BN	BV

	12-1m Momentum Portfolios		
	Losers	Neutral	Winners
Small	SL	SN	SW
Big	BL	BN	BW

Table 1: Depiction of the two-way factor portfolio sorts for the Carhart four-factor model

of stocks, namely **Big** (all the stocks above the break point) and **Small** (stocks below the break point).

2. Independently rank all the stocks according to their value score. Taking the 30th and 70th percentiles as break points, construct three value subsets; namely, **High** value above the 70th percentile, **Neutral** value between the 30th and 70th, and **Low** value (i.e. growth) stocks below the 30th percentile.
3. Repeat the previous step to construct stock subsets on the basis of momentum, profitability, investment, low volatility and low beta scores respectively. Note that in the case of investment, low volatility and low beta, the portfolio below the 30th percentile is the one which is expected to render the positive return.
4. Use the two-way size/factor sort in order to create equally-weighted factor portfolios, as depicted in Table 1. For example, the size/value sorting procedure gives one six portfolios: namely, Small Value, Small Neutral and Small Growth, and Big Value, Big Neutral and Big Growth.
5. Construct long/short factor returns by averaging the returns on the Small High and Big High factor portfolios and subtracting the average of the returns on the Small Low and Big Low factor portfolios. Repeat this for each set of sorting tables to create the six size-agnostic factor portfolios.
6. Construct long/short size factor returns for each of the independent two-way sorting tables by averaging the returns on the Small High, Small Neutral and Small Low factor portfolios and subtracting the average of the returns on the Big High, Big Neutral and Big Low factor portfolios. The final long/short size factor return is then calculated as the average of the various size factor returns across all factors included in the model.

Following Step 5 above, the long/short value factor return is calculated as (9)

$$R_{hml} = \frac{1}{2}\big(R(SV)+R(BV)\big) - \frac{1}{2}\big(R(SG)+R(BG)\big)$$
$$= R_{hml}^{+} - R_{hml}^{-} \qquad (10)$$
$$= \frac{1}{2}\big(R(SV)-R(SG)\big) + \frac{1}{2}\big(R(BV)-R(BG)\big)$$
$$= \frac{1}{2}\big(R_{hml}^{s} + R_{hml}^{b}\big)$$

Equations 9 and 10 show how to decompose the long/short factor return into separate long and short components as well as into separate size components. These decompositions also represent perhaps the two most common constraints faced by investors in the risk factor space: namely, long-only and capacity constraints. We will revisit this in the section on Factor Analysis.

Following Step 6, the size factor return from the size/value portfolios is calculated as (11)

$$R_{smb}^{val} = \frac{1}{3}\big(R(SV)+R(SN)+R(SG)\big) - \frac{1}{3}\big(R(BV)+R(BN)+R(BG)\big).$$

A similar calculation is done for the size/momentum portfolios and the final size factor return is thus given as (12)

$$R_{smb} = \frac{1}{2}\big(R_{smb}^{val} + R_{smb}^{mom}\big).$$

One departure from the methodology of Fama and French is the continued use of two-way rather than n-way sorts for the larger factor models. We do this because of the discrepancy between the size of the SA stock universe, which ranges from 150 – 171 stocks over the 14 year period, and the size of the US stock universe, which numbers in the thousands. Even if one were to use only two portfolios per factor, a four-way sort would cause the average portfolio size to drop to only ten stocks. This is clearly not large enough to ensure a well-diversified portfolio free from stock-specific risk.

Rebalancing of the value, profitability and investment factors occurs annually at each December-end. The low volatility and low beta factors are rebalanced quarterly, beginning from December-end, and the momentum factor is rebalanced monthly. As noted in Step 4, the standard methodology is to create equally-weighted factor portfolios, although one can also consider value-weighted portfolios. If any constituents of the factor portfolios delist during the holding period, an appropriate portfolio rebalance is done as at the close on the day prior to delisting as per standard indexing rules.

In summary, the process outlined above ensures that we create realistic and tradable daily risk factor returns over the complete sample period. Finally, we use the ALSI total return less the three-month NCD rate as a proxy for the excess market factor.

Factor Analysis

Figure 1 displays the cumulative log-performance of the eight South African long/short risk factors over the full 14-year sample period. Equal-weighted factors are represented by the solid lines and cap-weighted factors by the dashed lines. The most striking observation is that the scale of the momentum factor is significantly larger than any of the other factors, including the (excess) market factor. Apart from the international evidence that suggests that momentum generally does command the largest risk premium (Antonacci, 2013), the strong performance is likely also due to the underlying equity market's strong performance over the

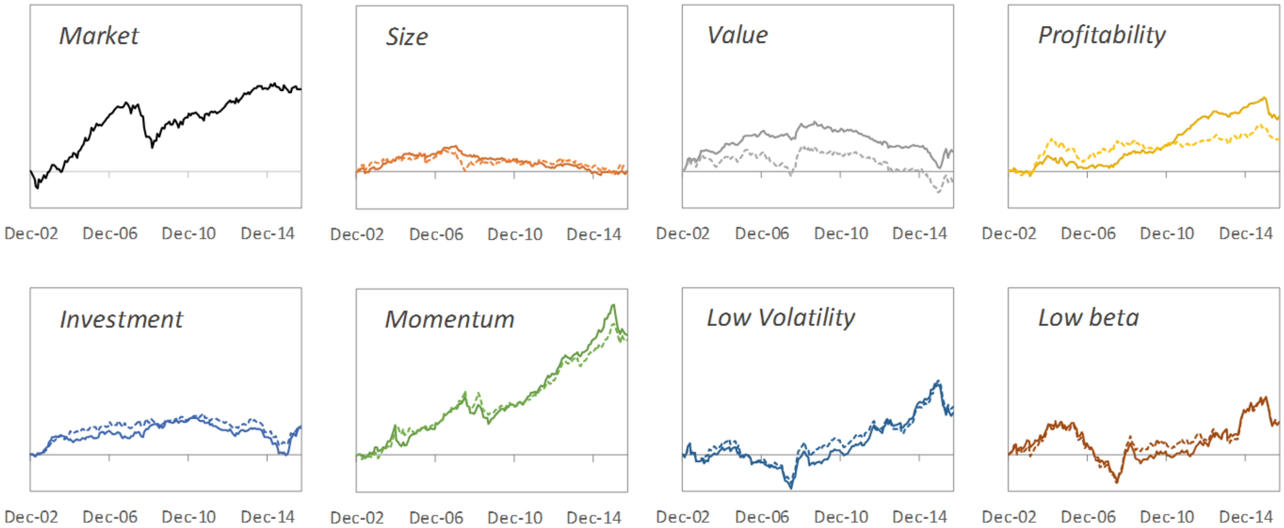

Figure 1: Cumulative log-performance of equal-weight (solid) and cap-weight (dashed) South African risk factors, Dec 2002 to Aug 2016

sample period, coupled with the extreme level of concentration. On average, the ten largest stocks in the ALSI have historically accounted for nearly 60% of the total index value (Flint et al., 2013). Therefore, any strong underlying equity market trend – positive or negative – is almost certainly driven by this handful of large counters. Such a feature is exactly what the momentum factor attempts to capture. Lastly, one must also remember that the momentum portfolio rebalances monthly and thus a large proportion of this return could be lost in practice due to high turnover costs.

Figure 1 also shows that the weighting scheme used in the Fama-French sorting procedure can impact the performance of the risk factor, although the magnitude of the effect is very factor-dependent. The discrepancy in equal- and cap-weighted factors is most obvious for the momentum factor but also affects value and profitability factors to some extent. Interestingly, we note almost no difference in either the trend or return magnitude for the low volatility and low beta factors.

Over the complete period, the size premium has remained consistently small and has in fact been slightly negative since the 2008 financial crisis; in line with the findings of Strugnell et al. (2011). As Table 2 shows, the expected return on the size factor is only 0.1%, a stark contrast to the 12.4% return on the momentum factor. The value factor, arguably the most well-known and accepted risk premium, has also struggled since the financial crisis, thus giving only a 2% annual return over the full period. This perhaps explains the poor performance of many South African value funds over the last decade.

We also note that the investment factor has not been particularly well rewarded over the last five years, showing a similar contraction as in the value premium. This is perhaps somewhat understandable as the level of annual asset growth and the book value of a company are surely somewhat connected on a fundamental level. This hypothesis is also supported by the fact that investment is the only factor to show a positive correlation 0.31 to value, even if small in absolute terms.

In contrast to the size, value and investment factors, profitability has shown strong performance over the last decade, particularly over the financial crisis and recovery period. This makes intuitive sense though as this factor essentially proxies the quality of a company's earning streams and one would expect high quality earnings streams to have been the least affected by the crisis and also to have participated strongly in the subsequent recovery rally. It also supports the recent industry trend in international markets of focussing on quality-sorted versions of the other factors (Gray (2014), Vogel and Gray (2015)).

	Market	*Size*	*Value*	*Profitability*	*Investment*	*Momentum*	*Low Volatility*	*Low Beta*
Exp. Return (CAGR)	8.66%	0.11%	1.98%	5.70%	2.68%	12.43%	4.25%	3.38%
Volatility	15.95%	7.21%	10.45%	9.21%	10.18%	14.86%	15.90%	18.04%
Kurtosis	0.55	0.36	1.66	3.14	5.66	2.11	1.67	0.22
Skewness	-0.12	-0.14	0.23	-0.93	1.01	-0.60	-0.22	-0.13
Return Range	27.31%	12.24%	20.54%	19.43%	24.35%	29.58%	33.23%	30.65%
Min Return	-14.25%	-7.63%	-11.08%	-12.90%	-7.68%	-17.63%	-17.11%	-13.71%
Max Return	13.05%	4.61%	9.46%	6.53%	16.68%	11.96%	16.13%	16.94%
Sharpe Ratio	0.54	-1.01	-0.52	-0.18	-0.46	0.34	-0.20	-0.22
Max Drawdown	-47.4%	-32.8%	-47.6%	-26.6%	-40.6%	-28.8%	-45.4%	-56.2%

Table 2: Equal-weight long/short factor summary statistics, Dec 2002 to Aug 2016

Equal-Weight	Market	Size	Value	Profitability	Investment	Momentum	Low Volatility	Low Beta
Mkt	1.00							
Size	-0.37	1.00						
Value	-0.20	0.09	1.00					
Profitability	-0.14	0.08	-0.26	1.00				
Investment	0.01	-0.04	0.31	-0.51	1.00			
Momentum	-0.01	-0.04	-0.45	0.42	-0.75	1.00		
Low Vol	-0.49	0.20	-0.03	0.62	-0.46	0.33	1.00	
Low Beta	-0.48	0.20	0.02	0.55	-0.33	0.35	0.78	1.00

EW vs CW	Market	Size	Value	Profitability	Investment	Momentum	Low Volatility	Low Beta
	0.94	0.73	0.84	0.68	0.70	0.74	0.95	0.92

Table 3: Equal-weight factor correlation matrix and correlations between equal-weight and cap-weight factor returns, Dec 2002 to Aug 2016

Tables 2 and 3 also highlight some interesting points about the low volatility and low beta factors. In contrast to what one might expect, Table 2 shows that these two factors have the second highest and highest return volatility respectively. However, this phenomenon actually confirms the rationale motivating these factors; namely that there is an inverse relationship between volatility or beta and the actual risk premium awarded to the stock. Whatever the economic reasoning though, we note that both factors have performed strongly since the financial crisis. The strong positive correlation of 0.78 between the returns of these two factors suggests that they are capturing overlapping parts of the same underlying factor, which one would expect. However, we do note a higher kurtosis, lower volatility and lower maximum drawdown attached to the low volatility factor. One final point of interest with these factors is their strong positive correlations of 0.62 and 0.55 respectively to the profitability factor. We leave this observation for future research.

As with all asset classes, risk factors also display varying degrees of cyclical behaviour. Although this is graphically evident in Figure 1, we provide more tangible evidence of this feature in Table 4, which presents factor statistics for three contiguous sub-periods of 4 1/2 years. In particular, we consider the bull market from December 2002 to June 2007, the crisis and recovery rally from June 2007 to December 2011, and the positive but slowing market from December 2011 to August 2016.

There are clear and meaningful differences in nearly all factors and statistics across the sub-periods. In particular, notice that the largest drawdown for most of the factors has actually occurred in the most recent sub-period and specifically over the last two years. Two of the main reasons for this – although certainly not the only ones – are that the proportion of SA-specific risk to global risk in the local market has been consistently increasingly since 2012 (Flint et al., 2015), and that some of the largest ALSI constituents have experience significant company-specific events in the recent past. This observation highlights the general need to ensure that one is effectively diversified against those risks which do not carry any discernible risk premia as well as being diversified across the risk factors that do carry a positive premium over the long-term. It is this last reason that has driven the rise of multi-factor portfolios, discussed further in in the section on Factor-Based Portfolio Management.

Statistic	Period	Market	Size	Value	Profitability	Investment	Momentum	Low Volatility	Low Beta
Exp. Return (CAGR)	1	21.64%	7.59%	11.40%	1.88%	6.86%	13.41%	-3.50%	-0.91%
	2	-2.52%	-4.47%	0.78%	11.57%	3.29%	7.25%	7.45%	2.27%
	3	8.26%	-2.29%	-5.24%	3.93%	-1.76%	18.38%	9.06%	8.81%
Volatility	1	16.13%	7.89%	9.50%	9.75%	7.70%	15.52%	13.49%	17.44%
	2	19.61%	6.68%	8.98%	7.18%	8.56%	12.30%	15.69%	16.64%
	3	10.67%	6.67%	12.16%	10.31%	13.29%	16.63%	18.13%	20.00%
Sharpe Ratio	1	1.34	0.02	0.42	-0.57	-0.07	0.39	-0.81	-0.48
	2	-0.13	-1.78	-0.74	0.58	-0.48	-0.01	0.00	-0.31
	3	0.77	-1.45	-1.04	-0.34	-0.69	0.66	0.09	0.07
Max. Drawdown	1	-21.3%	-8.0%	-8.0%	-16.4%	-9.8%	-24.9%	-26.8%	-39.3%
	2	-47.4%	-15.2%	-15.2%	-15.5%	-9.5%	-31.1%	-45.4%	-56.2%
	3	-15.4%	-32.8%	-47.6%	-26.6%	-40.6%	-34.2%	-35.5%	-38.1%

Table 4: Long/Short Factor performance across three sub-periods: Dec 2002 – Jun 2007, Jun 2007 – Dec 2011, Dec 2011 – Aug 2016

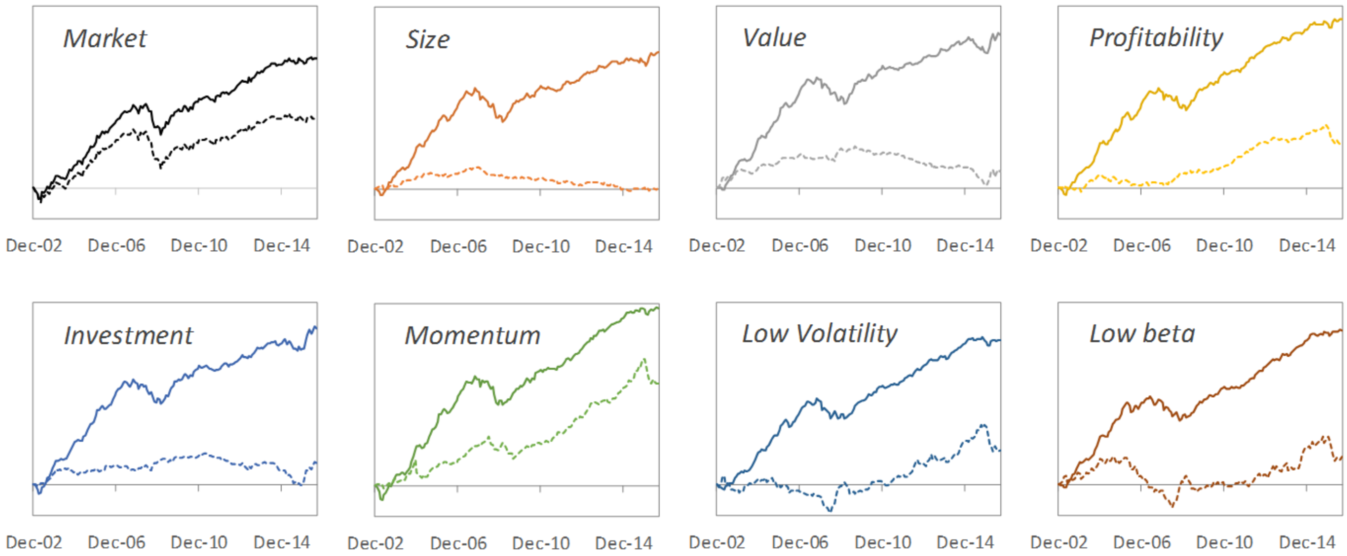

Figure 2: Cumulative log-performance of long-only (solid) and long/short (dashed) South African risk factors, Dec 2002 to Aug 2016

Table 5: Equal-weight long-only factor summary statistics, Dec 2002 to Aug 2016

	Market	*Size*	*Value*	*Profitability*	*Investment*	*Momentum*	*Low Volatility*	*Low Beta*
Exp. Return (CAGR)	16.92%	17.85%	20.22%	22.49%	20.66%	23.74%	18.97%	20.40%
Volatility	15.32%	13.62%	15.19%	13.51%	15.51%	15.32%	11.98%	12.74%
Kurtosis	0.34	1.31	0.60	0.53	0.14	1.41	1.47	1.35
Skewness	-0.18	-0.73	-0.20	-0.42	0.01	-0.58	-0.70	-0.59
Return Range	26.09%	22.33%	25.13%	21.94%	26.90%	26.74%	21.78%	21.86%
Min Return	-13.10%	-14.07%	-13.14%	-11.87%	-10.01%	-14.27%	-13.07%	-11.58%
Max Return	12.99%	8.26%	11.99%	10.07%	16.89%	12.47%	8.71%	10.28%
Sharpe Ratio	1.10	0.77	0.84	1.12	0.86	1.07	0.97	1.02
Max Drawdown	-40.4%	-42.2%	-35.1%	-30.8%	-33.3%	-37.5%	-27.5%	-32.5%

Factor Robustness

As with any empirical financial study, one needs to address the question of robustness. In particular, one should always be cognisant of the fact that the constructed factor portfolios will always only be noisy proxies of the true underlying risk factors. To this end, we consider the robustness of such factors to the choices made during the construction process. We have already highlighted one such choice in Figure 1 by showing the effect that weighting scheme can have. In this section we scrutinise a number of other important construction choices.

Long-only versus Long/Short Factors

One of the most pertinent constraints for many investors is the inability to short sell assets either at all or to the extent that they would wish. This raises the issue of whether long-only factor proxies are able to provide similar risk factor exposure in comparison to their long/short counterparts. A fundamental challenge in factor investing is the investability of the underlying factor portfolios. It is all well and good to create theoretically appealing long/short factor portfolios and use these for risk attribution – see the section on Factor Risk Attribution and the Factor Efficiency Ratio – but this may all for nought if one cannot effectively allocate capital to such portfolios. Hence the proposal of long-only factor portfolios. Although such portfolios will contain residual market risk by construction, we believe that their interpretation as risk factors still remains valid. Furthermore, given that all the factors will on average have similar levels of market risk exposure, this residual risk should largely cancel out in any risk attribution exercises.

Figure 2 compares the performance of the long-only component of each factor (solid lines) against the complete long/short portfolios (dashed lines), and Table 5 gives the long-only factor summary statistics. In the case of the market factor, we are comparing the absolute market return with its excess-to-cash counterpart. There is a stark contrast in performance between all the long-only and long/short portfolios. It is also clear that the long-only risk factors – barring size – comfortably outperform the absolute market return.

Table 6 gives the correlation matrix of the long-only factors as well as the correlations between the long-only and long/short versions of each factor. The supposition of contaminating latent market exposure is proven by the strong positive correlations with the market factor. Furthermore, the correlations between each risk factor are now also very high as a result. Considering the correlations between long-only and long/short factor versions, it is interesting to note that despite the similarity in trend between the two momentum factors, the correlation between these two factors is only mildly positive at 0.29. This serves as a

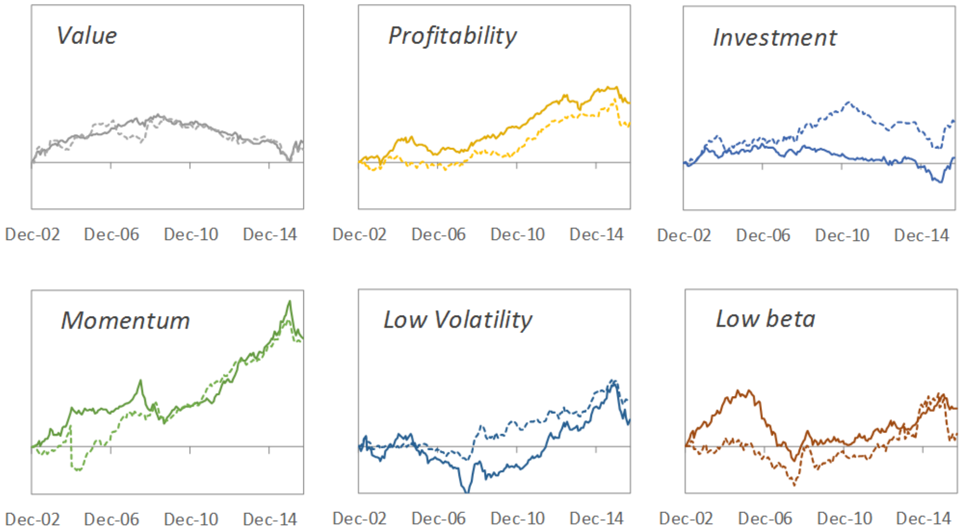

Figure 3: Cumulative log-performance of big (solid) and small (dashed) South African risk factors, Dec 2002 to Aug 2016

Long-Only	*Market*	*Size*	*Value*	*Profitability*	*Investment*	*Momentum*	*Low Volatility*	*Low Beta*
Market	1.00							
Size	0.71	1.00						
Value	0.66	0.90	1.00					
Profitability	0.76	0.91	0.84	1.00				
Investment	0.75	0.88	0.90	0.81	1.00			
Momentum	0.80	0.87	0.76	0.91	0.80	1.00		
Low Volatility	0.62	0.86	0.80	0.89	0.69	0.81	1.00	
Low Beta	0.55	0.81	0.74	0.83	0.66	0.78	0.85	1.00

L-O vs L/S	*Market*	*Size*	*Value*	*Profitability*	*Investment*	*Momentum*	*Low Volatility*	*Low Beta*
		0.18	0.41	0.17	0.42	0.29	0.07	0.16

Table 6: Long-only factor correlation matrix and correlations between long-only and long/short factor returns, Dec 2002 to Aug 2016

poignant reminder about the pitfalls of conflating price trend and correlations. What Figure 2 does suggest though is that the short component of the momentum factor provides only limited benefit across the period.

Factor Size Effects

Another constraint faced by many investors is that of capacity. Even if one has the ability to short, it may be that the majority of a factor's return stems from the Small sub-portfolios of the factor. Such a size bias would imply limited investment capacity owing to the small market capitalisation of the underlying stocks and potential illiquidity issues. Several authors have suggested that such factor size biases exist in many markets (Homescu, 2015). If present in the highly concentrated SA equity market, this bias would have serious ramifications on the prospect of large-scale SA factor investing. Figure 3 breaks down each factor return into its Big (solid line) and Small (dashed line) sub-portfolio as per Equation 10. Note that these sub-portfolios are still long/short combinations and hence are of similar magnitudes to the complete factor returns shown in Figure 1.

Momentum and value don't display any significant size bias. Of the remaining four, profitability displays a small, persistent bias towards large stocks, while investment displays a persistent bias towards small stocks. Low volatility and low beta display discrepancies between big and small long/short portfolios that vary over the sample period.

Rebalancing Frequency & Date

Value, profitability and investment portfolios are rebalanced annually at the beginning of each year. Low volatility and low beta portfolios are rebalanced quarterly with the first rebalance occurring at the beginning of the year, and momentum portfolios are rebalanced at the end of each month. The choice of rebalance frequency for each factor is driven by the time frame over which the factor signal decays. There is also the more practical issue that any benefit gained from more frequent rebalancing may be offset by the additional transaction costs. For the majority of our factors, the time frame of the risk premia is well established. However, given the relatively new 'discovery' of the low volatility and low beta factors, the effect of rebalance frequency is less well documented. To this end, we compared the returns from the low volatility and low beta factors when rebalancing monthly, quarterly, biannually and annually and found only minor differences.

Another rebalancing issue to consider for those factors with longer holding periods is the choice of month in which to enact the rebalance. As above, we test how much of an impact moving rebalance dates has by considering the returns from twelve value factors each rebalanced in different months of the year and again find no significant return differences. Although it may seem odd to include such a non-result in our research, it is an incredibly important one from a practical implementation perspective.

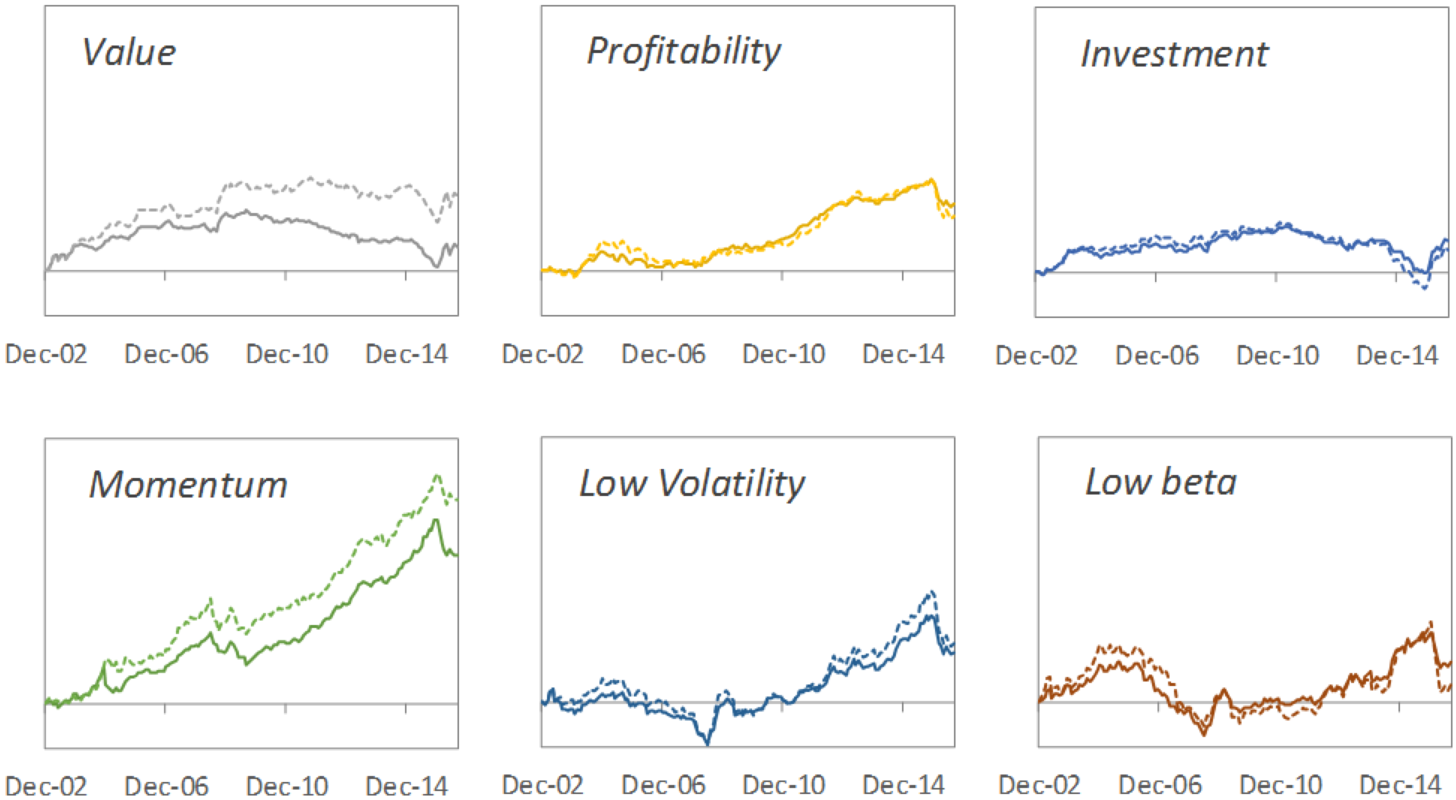

Figure 4: Cumulative log-performance of standard (solid) and extreme (dashed) South African risk factors, Dec 2002 to Aug 2016

Furthermore, it showcases the fact that the factor construction methodology outlined in in the section on Constructing South African Risk Factors is generally robust to rebalancing choices.

Portfolio Extremity

The standard Fama-French two-way sorting procedure uses the 50th percentile of the size score and the 30th and 70th percentiles of the factor scores as the relevant sorting break points. A natural question then is whether using more extreme percentile break points results in larger factor risk premia. The trade-off here is that one essentially creates 'purer' factor portfolios but at the cost of increasing the portfolio's idiosyncratic risk. This is particularly pressing in the South African equity market, which only contains around 160 counters.

To test the robustness of the factors to the sorting methodology, we create extreme factor portfolios using the 20th and 80th percentiles of the relevant factor scores as sorting break points. Figure 4 gives the comparison between the standard (solid line) and extreme (dashed line) factors. Somewhat surprisingly, only the extreme value and momentum factors show any significant difference to their standard counterparts. In both cases, the divergence of the extreme factor performance is most evident in the last ten years and seems to be linked to outperformance during periods of financial stress. We leave further investigation of this phenomenon for future research.

Alternative Factor Definitions

Although varying the choice of sorting percentile can in some respects be considered as using an alternative factor definition, the more obvious alternative is to use a different fundamental stock characteristic as a proxy for the underlying factor score. As an example, we have already discussed the multiple definitions of the quality factor in the section on The Fama-French Model and its Extensions. In a similar vein, a number of authors have considered alternative measures for value and for low volatility. Popular alternative value score candidates include earnings-to-price, cash flow-to-price and a composite scored based on these two metrics as well as the original book-to-market ratio (Amenc et al., 2014). In the low volatility literature, the alternatives are not different risk measures but rather different calculation methods for volatility; the main variables being the length of the historical estimation window and the frequency of return data.[4] Blitz and van Vliet (2007) suggest using three years of weekly data, Baker et al. (2014) suggests using either 60 monthly or 60 weekly return observations, local research considers three years of monthly data, while Frazzini and Pedersen (2014) suggest one year of daily return data.

Figure 5 gives the cumulative log-performance of long/short factors based on these alternative value and low volatility scores. The variant return range for both factors is fairly substantial and particularly so for the value factor. Furthermore, the behaviour of the variant value factors differs significantly throughout the period, which suggests that the selected stock characteristics capture different aspects of the true value risk factor. The relative outperformance of the composite value score supports this suggestion and also highlights the importance of reducing signal noise; in this case achieved by averaging out the characteristic-specific noise.

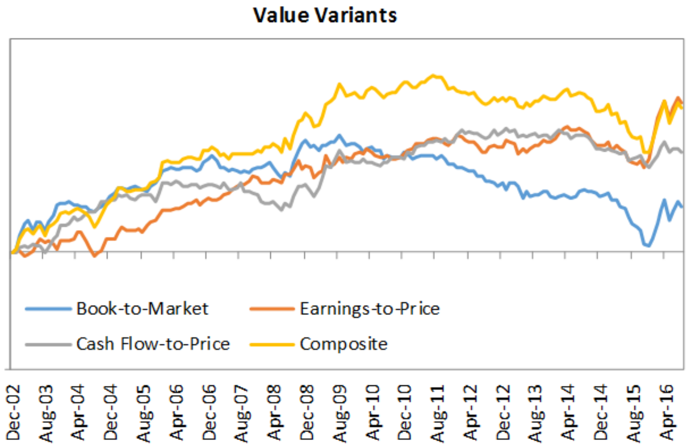

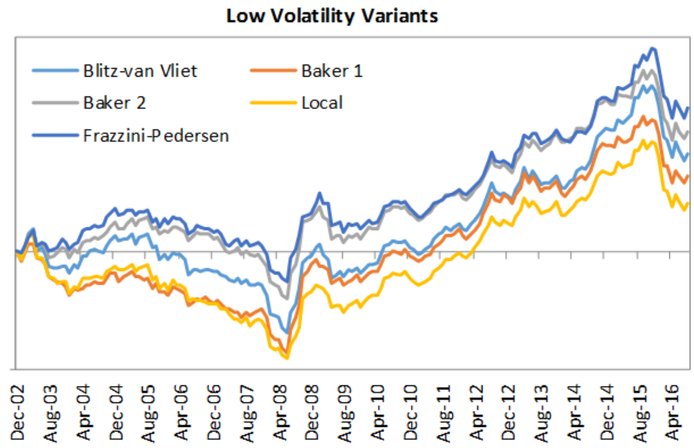

Figure 5: Cumulative log-performance of value factor variants (left) and low volatility factor variants (right), Dec 2002 to Aug 2016

For the low volatility factor, performance of the factors all show the same pattern, indicative of the fact that only the calculation method is changing, rather than the measure itself. Interestingly, both of the top performing variants are those that use the smallest estimation window – 1 year and 60 weeks respectively – as well as higher frequency data – daily and weekly respectively.

Factor-Based Risk Management

At its core, portfolio management is about making decisions: when to buy or sell any given asset and in what quantity. These decisions are made in order to add value to a passive benchmark, be it a nominated index or a cash-based rate.[5] In this setting, 'adding value' is usually defined in two ways. The first is by achieving a positive return, or alpha, over and above the nominated benchmark at an acceptable level of risk. The second is by achieving a specified target return at a lower level of risk than that of comparable passive market products.

In both cases, the strength of any portfolio decision should be measured by how much value it generates for the fund, conditional on the market and fund constraints faced by the manager over the performance period. In prior Peregrine Securities research, we showed how one could use the fundamental law of active management (FLOAM) framework of Clarke et al. (2002) in order to decompose a fund's relative return and risk into contributions from each of the underlying fund constituents (Flint et al., 2015).

We build on this work here but consider instead the idea of risk attribution rather than risk decomposition. In particular, we consider how to attribute a fund's risk – absolute or relative – to a given set of external risk factors. Such an attribution lets one identify what kinds of factor risk a fund is exposed to and furthermore calculate how large these factor bets are. Knowing this allows one to make informed and efficient investment decisions.

Factor Risk Attribution and the Factor Efficiency Ratio

Given a series of fund returns – absolute or relative – we can use one of the LFM's in The Fama-French Model and its Extensions to attribute risk to the underlying risk factors constructed in Constructing South African Risk Factors. Although more difficult than attributing risk to the fund's constituents, Meucci (2007, 2016) describes how one can still attribute fund risk to a set of external risk factors in an additive fashion. Furthermore, if one does have sight of the fund's holdings, it is possible to attribute risk similarly for each of the underlying constituents so that the fund's factor risk contributions can be written as a linear combination of the constituents' factor risk contributions (see also Roncalli and Weisang, 2012). This is perhaps the most important factor application in the risk management space. Consider the pedagogical example below.

We select the Carhart four-factor risk model and make use of long-only risk factors. Let us assume that there are four funds that are currently under investigation. We simulate monthly returns for these funds using the factor exposures given in Table 7. A small random alpha term (centred at 0.25%) and a larger random noise term (centred at zero) are added to each fund's monthly return.

	Fund1	Fund2	Fund3	Fund4	Benchmark
Market	0.5	0.1	0.1	0.2	0.25
Size	0.2	0.5	0.1	0.2	0.25
Value	0.2	0.2	0.5	0.1	0.25
Momentum	0.1	0.2	0.2	0.5	0.25

Table 7: Simulated fund risk factor exposures

Betas	Fund1	Fund2	Fund3	Fund4	Benchmark
Alpha	0.38%	0.10%	0.35%	0.18%	0.00%
Market	0.52	0.05	0.14	0.23	0.23
Size	0.29	0.48	0.11	0.24	0.26
Value	0.15	0.24	0.51	0.06	0.24
Momentum	0.05	0.22	0.15	0.49	0.25
R^2	94.2%	95.5%	94.6%	95.4%	95.9%
Risk (Volatility)	14.27%	13.88%	13.01%	14.68%	13.86%
Tracking Error	5.37%	4.71%	4.89%	4.48%	n.a.
	Risk Contributions				
Market	52.9%	4.0%	13.8%	21.3%	22.8%
Size	23.8%	44.9%	10.2%	20.2%	24.0%
Value	12.9%	24.3%	55.5%	5.4%	23.4%
Momentum	4.6%	22.3%	15.1%	48.4%	25.7%
Residual	5.8%	4.5%	5.4%	4.6%	4.1%
	Tracking Error Contributions				
Market	30.9%	26.9%	14.3%	-0.4%	n.a.
Size	-0.8%	9.5%	12.7%	-0.7%	n.a.
Value	3.9%	0.1%	-1.6%	4.7%	n.a.
Momentum	4.6%	0.8%	15.8%	29.5%	n.a.
Residual	61.3%	62.6%	58.8%	66.8%	n.a.

Table 8: Carhart risk factor attribution

Table 8 gives a comprehensive factor risk attribution for both the absolute and relative risk of each fund based on the Carhart four-factor model. By construction, the estimated betas are very similar to the input fund exposures and the R^2 of the risk model is very high. Table 8 also shows the risk contributions of each factor as well as the catch-all residual term. These values are also closely related to the estimated beta levels owing to the high correlation between the risk factors as well as their similar volatility levels. Finally, contributions to tracking error are also calculated across the funds for each risk factor. Because of the good fit of the risk model, most of the tracking error stems from the fund-specific noise term.

In the context of factor investing, where investors are actively seeking exposure to the underlying risk factors, risk and tracking error contributions become incredibly important as they provide a means of quantifying and thus evaluating such exposure. To this end, Hunstad and Dekhayser (2015) introduce the Factor Efficiency Ratio (FER) as a means of gauging the amount of intended versus unintended factor risk exposure in a given fund (or asset). Letting $\mathcal{F}_d$ represent the set of k desired factors, we can write

$$FER(\mathcal{F}_d) = \frac{\sum_{i=1}^{k} RC_i}{1 - \sum_{i=1}^{k} RC_i}, \qquad (13)$$

where RC_i is the generic risk contribution stemming from the i^{th} desired risk factor. Hunstad and Dekhayser originally consider the contributions to active risk (i.e. tracking error) but one can just as easily use any convex risk measure to calculate risk contributions.[6] This FER is interpreted as follows: for every *X*% of risk stemming from the desired factor set, the fund takes on an additional 1% of risk from undesired factors. Therefore, the higher the FER, the more efficient the fund is at gaining desired factor exposure.

Consider the four fund example and further assume that all of these funds are marketed as composite value/momentum indices. Using this as our desired factor set, we calculate FER's of 0.21, 0.87, 2.40 and 1.17 for each of the funds respectively. Based on these scores, it is clear that Fund 3 provides one with the most efficient exposure to the desired value and momentum factors.

Return-Based Style Analysis & Fund Replication

Sharpe (1992) introduced the concept of returns-based style analysis (RBSA) used extensively in the fund management literature. In essence, RBSA is a form of constrained regression

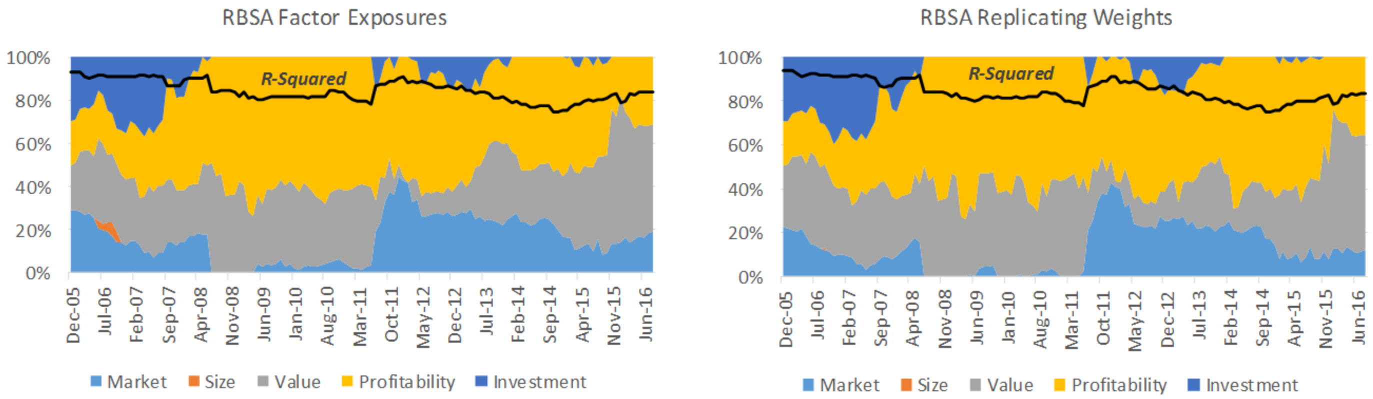

Figure 6: RBSA betas (left) and end-of-period weights (right) for the FTSE/JSE Dividend Plus Index and the long-only Fama-French five-factor model, Dec 2005 to Aug 2016

that allows one to draw inference on funds for which only historical return data is available. Sharpe suggested using factors based on asset classes and interpreted the model output as being indicative of a manager's style mix. Ultimately, given a set of historical fund returns, RBSA estimates the static mix of tradable market indices or factors that most closely replicates the fund's returns, R_{pt}. Letting β represent the vector of factor exposures, we can formulate the RBSA estimation problem as follows:

$$\begin{aligned} &\underset{\beta}{\operatorname{argmin}} \sum_{t=1}^{T} \left(R_{pt} - \sum_{j=1}^{m} \beta_j \mathcal{F}_{jt} \right)^2 \\ &s.t. \qquad \beta_j \geq \\ &\sum \beta_j = 1. \end{aligned} \tag{14}$$

In a sense, the RBSA betas represent the long-only weights of the replicating style portfolio. However, this is not strictly true because the betas remain fixed across the estimation window whereas portfolio weights would change in line with the performance of the underlying factors. Several improvements to the initial RBSA methodology have been suggested to address this (and other) issues. These include the use of the Kalman filter, corrections for heteroscedasticity and the inclusion of structural break detection mechanisms. Another point which is common to all regression but generally not considered in RBSA is that of confidence intervals around the estimated betas.[7] For example, a style weight of 30% with a confidence interval of +/- 2% should be viewed very differently to a weight of 30% with a confidence interval of +/- 20%.

A variation of RBSA that is particularly relevant in the index tracking space is to solve for the initial number of 'shares' (rather than betas) of each factor that minimises the tracking error (rather than sum of squared errors) of the estimated style portfolio to the given fund returns. Therefore, one can not only use the RBSA framework to measure a given fund manager's style mix but also – after some adjustment – to create tradable replicating portfolios for a fund. This alternative usage has been explored at length in connection with hedge fund replication.

As in Factor Risk Attribution and the Factor Efficiency Ratio, we illustrate RBSA with an illustrative example. We attempt to uncover the style mix of the FTSE/JSE Dividend Plus Index by making use of the long-only Fama-French five-factor model. Figure 6 displays the RBSA factor exposures (left panel) and the adjusted-RBSA replicating weights (right panel) from December 2005 onwards. We fit both models using rolling 36-month windows and record the static betas and end-of-period weights respectively.

Although the exposures are similar to the replicating weights, one can still easily see the discrepancies in Figure 6. The R^2 of both models is consistently high, meaning that the majority of variation in the index is well-captured by the five-factor model. The style mix of the index varies considerably over the period, which suggests that the dividend yield measure is actually a composite signal for a number of underlying risk factors. The largest exposure over the period has been to the profitability factor – in line with the yield-driven nature of the index – with the remainder mostly split between the value and market factors. Investment exposure is sporadic and has been absent over the last three years. Size is irrelevant for the Dividend Plus index, which is to be expected given that the index is limited to large- and mid-cap stocks.

Table 9 gives the RBSA betas and end-of-period weights for the 36-month period ending at 31 August 2016. Although similar in nature, there is still an absolute difference of 13.7% across the factors. This difference is driven by the varying performance of the underlying factors and is directly related the level of factor dispersion over the period.

	Market	*Size*	*Value*	*Profitability*	*Investment*
RBSA Betas	19.0%	0.0%	50.0%	31.0%	0.0%
95% Conf. Interval	7.8% – 30.3%	-11.9% – 11.9%	40.3% – 59.7%	21.5% – 40.5%	-10.7% – 10.7%
RBSA Weights	12.2%	0.0%	52.6%	35.3%	0.0%
Weight-Beta Spread	-6.9%	0.0%	2.6%	4.3%	0.0%

Table 9: RBSA betas and replicating weights for the Dividend Plus Index as at 31 Aug 2016

Factor-Based Portfolio Management

In addition to the risk management applications given above, risk factors are also used extensively in portfolio management. And while the concept of factor investing is definitely not new, the rise of the smart beta phenomenon has attracted significant attention to this area.

In the last several years, the focus has started to move away from identifying additional risk factors and towards constructing optimal multi-factor portfolios. While some authors have said that there is no formal framework in place for combining systematic factor strategies (De Franco et al., 2016), the fact of the matter is that the majority of the existing optimisation frameworks – risk/return or risk-only – are fully capable of incorporating both factor portfolios and factor-based risk/return views. Furthermore, the allocation policy for systematic strategies outlined by Meucci (2016) provides one with a fully general framework for creating optimal multi-factor portfolios in the presence of transaction costs and fund constraints.

In this section we discuss several ideas on how to create such multi-factor portfolios, ranging from the very simple to the fairly complex. Note that most of these are based on concepts that we have already introduced and analysed in preceding sections.

Factor Portfolio Mixing and Integrated Factor Scores

According to Fitzgibbon et al. (2016), two of the most common approaches for creating multi-factor portfolios are the 'portfolio mix' and 'integrated score' methods. Portfolio mixing is simply the linear combination of factor portfolios constructed from single-variable sorting procedures. For example, consider a value portfolio based solely on the top quintile of book-to-market stocks and a momentum portfolio based solely on the top quintile of twelve month return stocks. These portfolios would then be taken as existing building blocks and the only challenge facing the investor would be to set an appropriate weight for each portfolio. Viewed in this light, portfolio mixing can be thought of in a similar manner to the decisions made in strategic asset allocation.

The integrated score approach goes one step further by mixing the underlying factor scores ex ante rather than mixing given factor portfolios ex post. The Fama-French two-way sorting methodology – whereby stocks are selected based on their respective factor score ranks relative to a set of constant percentile break points – is perhaps the simplest example of the integrated score approach. In general, the integrated score approach combines individual factor scores in some manner to create a single, unified score. Figure 7 displays this concept graphically and confirms that the field of (non)linear programming provides investors with a natural set of tools for creating optimal integrated multi-factor scores, and thus optimal multi-factor portfolios.

Lastly and very importantly, Hoffstein (2016) points out that one needs to consider the speed of factor decay when creating these integrated signals. This is particularly relevant when combining the fast-decaying momentum signal with slower signals like value or profitability, for example.

Constrained Risk Factor Optimisation

A more technically rigorous approach than those given above is to view the construction of an efficient multi-factor portfolio as a constrained optimisation problem. Although more complex, this approach allows an investor to construct a multi-factor portfolio that is as consistent with their return objectives and risk preferences as their constraint set will permit. There are a number of optimisation frameworks available to investors, including classical mean-variance and risk-based investing (Richard and Roncalli, 2015), among others.[8] Below we sketch out two candidate optimisation approaches that could be used to create constrained optimal multi-factor portfolios.

The first approach makes use of the risk attribution framework introduced in in the section Factor Risk Attribution and the Factor Efficiency Ratio. Assuming that one is given a risk factor model, the problem then becomes finding the underlying stock weights that provide the requisite exposure to the targeted risk factors, whilst minimising undesired factor exposures. If exposure is defined in terms of beta, then one needs to solve for the portfolio of assets that minimises the total distance between estimated and targeted betas, where the target levels for the undesired factors are set to zero. Alternatively, if exposure is defined in terms of risk contributions, then there two options available. The first option is similar to the beta optimisation but where one instead specifies target risk contribution levels. The second option is to solve for the portfolio of assets that maximises the FER for the set of desired factors. FER optimisation is arguably more intuitive and will likely provide more robust results due to the fact that it simultaneously accounts for the desired and undesired factor exposures in a single monotonic metric. Of the two approaches, we therefore favour FER maximisation.

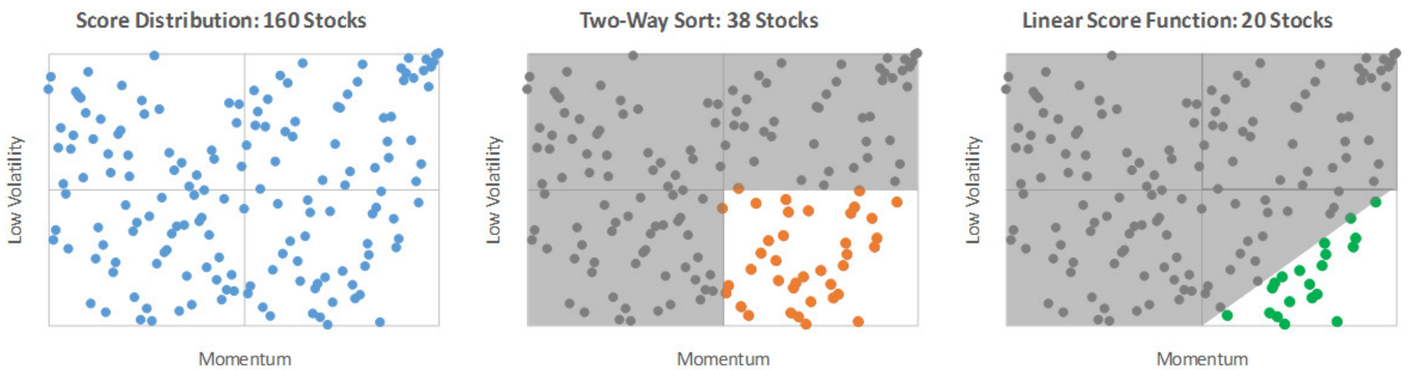

Figure 7: Integrated scoring examples for momentum and low volatility

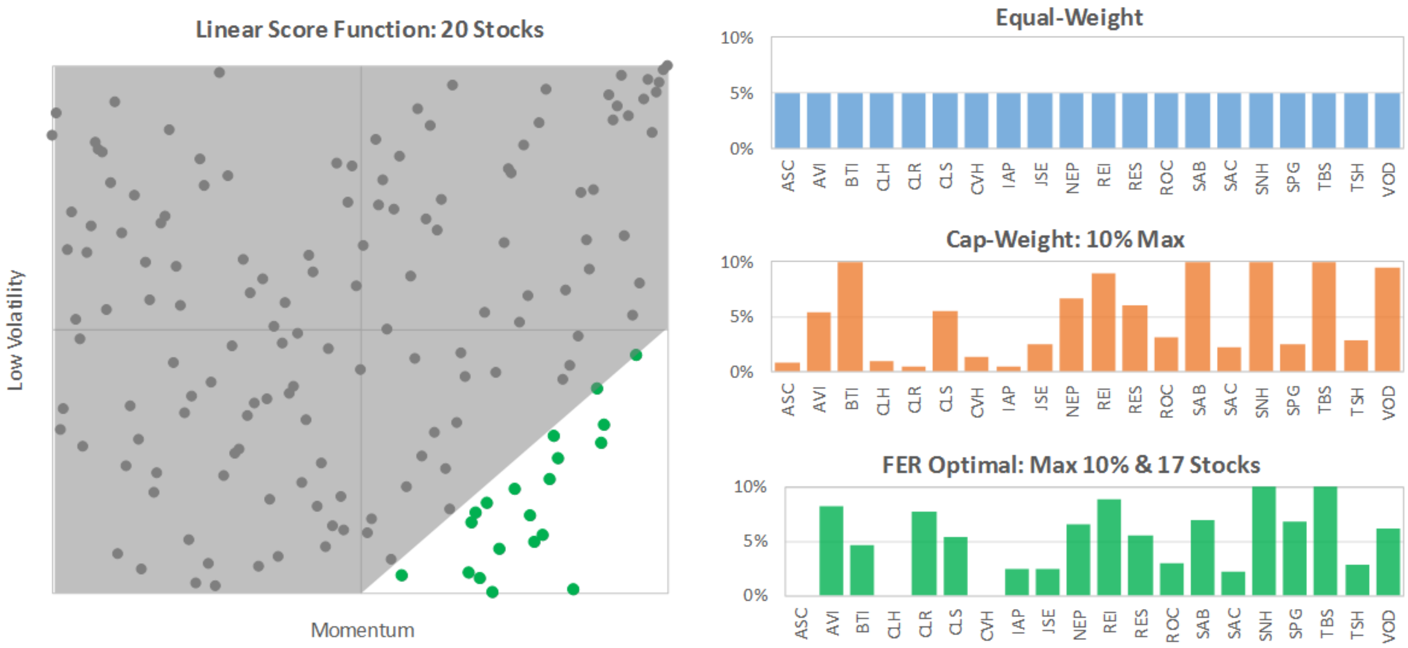

Figure 8: Creating a multi-factor portfolio by combining an Integrated Scoring screen with an MILP optimisation of the portfolio's Factor Efficiency Ratio

The second optimisation approach makes use of mixed integer programming (MIP). A mixed integer program is one in which some variables are continuous and some are integers. Such a setting is ideal for problems in which one has to first select a subset of assets from the available universe – the integer variables – and subsequently search for the set of weights – the continuous variables – that minimises an objective function under a set of constraints. In general, mixed integer programs can be quite hard to solve unless one can formulate the problem in a very particular way. Thankfully, we are able to set up both linear (MILP) and quadratic (MIQP) mixed integer programs for most portfolio construction problems which can be solved fairly easily – albeit slowly – with freely available optimisation toolboxes and heuristic solvers. In prior Peregrine Securities research, we have successfully used the MIQP approach to replicate the Top40 index with only a small number of stocks and also construct optimal hedging baskets for active funds (Flint et al., 2015).

One of the main issues with multi-factor investing is smoothly transitioning between risk and return preferences in the factor space to risk and return preferences in the asset space. This is not a trivial exercise. One way of linking the factor and asset spaces in a manner which does not add additional estimation error would be to combine the integrated score approach with the risk attribution optimisation by means of an MILP. Figure 8 presents an example of this combined approach for a low volatility and momentum multi-factor portfolio using scoring data as at August 2016.

One uses the integrated score as a screening tool to find the subset of assets that display the fundamental factor characteristics most in line with the desired factor set. Taking this subset of factor-screened assets as an input, one then solves the MILP problem for the maximum FER portfolio under the given constraints, where the choice of assets included in the portfolio and the subsequent weights attached to the chosen assets are both variables in the optimisation. Introducing the integrated score screen and subsequently maximising the portfolio's FER obviates the need to explicitly assign factor-consistent expected return estimates to each asset – a difficult task – and thus also reduces the potential for estimation error in the optimisation.

Conclusion

Risk factors and systematic factor strategies are fast becoming an integral part of the global asset management landscape. In this report, we have attempted to provide an introduction to, and critique of, the factor investing paradigm in a South African setting.

We created a range of long/short and long-only risk factors for the South African equity market according to the standard Fama-French factor construction methodology: size, value, momentum, profitability, investment, low volatility and low beta. Historical risk and return characteristics varied significantly across the factors as well as across market regimes. Momentum has been the most rewarded factor historically. Low volatility, profitability and low beta have also shown positive risk premia, while the size factor seems to be non-existent in South Africa. We then tested factor robustness at length and showed the effect that each of the major decisions taken in the factor construction process can have. The largest such effect stems naturally from the choice of long-only or long/short factors. Interestingly, we found that, barring size, all long-only factors handily outperformed the market.

In addition to constructing this factor database, we also showcased several risk factor applications. In the risk management space, we considered risk attribution to factors and introduced the Factor Efficiency Ratio as a measure of how efficiently a fund gained exposure to a set of desired risk factors. We also considered returns-based style analysis with long-only risk factors and showed how this could be used to estimate a manager's style mix or to create a replicating factor portfolio for an index.

In the portfolio management space, we considered the issue of creating multi-factor portfolios. We discussed simple approaches such as portfolio mixing and integrated scoring, and more complex approaches based on solving for target risk contributions or optimising the factor efficiency ratio for the desired factors. Finally, we introduced the mixed integer programming framework as a means of combining the integrated scoring approach with the risk attribution optimisation approach in a robust manner, thus allowing one to smoothly transition between preferences and constraints in the non-tradable factor space and the tradable asset space.

Endnotes

1. The factors and strategies are known by many names. Some of these include: risk factors, risk premia, smart beta, alternative beta, systematic strategies, quantitative strategies and rule-based strategies.

2. Factor construction is discussed at length in Section 3.

3. For example, see the comprehensive risk factor databases maintained by Kenneth French (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html) and Andrea Frazzini (http://www.econ.yale.edu/~af227/data_library.htm).

4.Similar calculation method alternatives apply to the beta factor score.

5. All portfolio management should be considered benchmark-relative, even if the selected benchmark is a constant value of zero.

6. Note that one has to treat negative risk contributions with caution when calculating the FER as they can materially change its interpretation. The simplest solution is to take absolute values of all risk contributions and replace the '1' in the denominator with the sum of the absolute risk contributions.

7. See Lobosco and diBartolomeo (1997) for an approximation formula for constructing confidence intervals around the constrained betas.

8. Please see Homescu (2014) for a comprehensive review of the available portfolio construction frameworks.

References

1. Amenc, N., Deguest, R., Goltz, F., Lodh, A., Martellini, L., & Shirbini, E. (2014). Risk allocation, factor investing and smart beta: Reconciling innovations in equity portfolio construction. Working paper, EDHEC-Risk Institute, July 2014.

2. Ang, A. (2014). Asset Management: A Systematic Approach to Factor Investing. Oxford University Press.

3. Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. The Journal of Finance, 61(1), 259-299.

4. Antonacci, G. (2013). Absolute momentum: a simple rule-based strategy and universal trend-following overlay. Available at SSRN 2244633.

5. Asness, C. S., Frazzini, A., Israel, R., & Moskowitz, T. J. (2015). Fact, fiction, and value investing. Forthcoming, Journal of Portfolio Management, Fall.

6. Asness, C. S., Moskowitz, T. J., & Pedersen, L. H. (2013). Value and momentum everywhere. The Journal of Finance, 68(3), 929-985.

7. Asness, C., & Frazzini, A. (2013c). The devil in HML's details. Journal of Portfolio Management, 39(4), 49.

8. Asness, C., Frazzini, A., & Pedersen., L. (2013b) Quality minus junk. Available at SSRN 2312432.

9. Baker, M., Bradley, B., & Taliaferro, R. (2014). The low-risk anomaly: A decomposition into micro and macro effects. Financial Analysts Journal,70(2), 43-58.

10. Basiewicz, P. G., & Auret, C. J. (2009). Another look at the cross-section of average returns on the JSE. Investment Analysts Journal, 38(69), 23-38.

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

12. Blitz, D., & Van Vliet, P. (2007). The volatility effect: Lower risk without lower return. Journal of Portfolio Management, 102-113.

13. Carhart, M. M. (1997). On persistence in mutual fund performance. The Journal of finance, 52(1), 57-82.

14. Cazalet, Z., & Roncalli, T. (2014). Facts and fantasies about factor investing. Available at SSRN 2524547.

15. Clarke, R., De Silva, H., & Thorley, S. (2002). Portfolio constraints and the fundamental law of active management. Financial Analysts Journal, 58(5), 48-66.

16. Cochrane, J. H. (2011). Presidential address: Discount rates. The Journal of Finance, 66(4), 1047-1108.

17. De Franco, C., Monnier, B., Nicolle, J., & Rulik, K. (2016). How Different Are Alternative Beta Strategies?. The Journal of Index Investing, 7(2), 57-77.

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

19. Fama, E. F., & French, K. R. (2008). Dissecting anomalies. The Journal of Finance, 63(4), 1653-1678.

20. Fama, E. F., & French, K. R. (2014). A five-factor asset pricing model, Fama-Miller Working Paper.

21. Fitzgibbons, S., Friedman, J., Pomorski, L., & Serban, L. (2016). Long-Only Style Investing: Don't Just Mix, Integrate. Integrate (June 29, 2016).

22. Flint, E., Chikurunhe, C., & Seymour, A. (2015). The Cost of a Free Lunch: Dabbling in Diversification. Peregrine Securities Research Report.

23. Flint, E., Seymour, A., & Chikurunhe, C. (2013). Trends in the South African Hedging Space Part I. Peregrine Securities Research Report.

24. Flint, E., Seymour, A., & Chikurunhe, C. (2015). In Search of the Perfect Hedge. Peregrine Securities Research Report.

25. Frazzini, A., & Pedersen, L. H. (2014). Betting against beta. Journal of Financial Economics, 111(1), 1-25.

26. Harvey, C. R., Liu, Y., & Zhu, H. (2015). ... And the cross-section of expected returns. Review of Financial Studies, hhv059.

27. Hoffstein, C. (2016). Multi-Factor: Mix or Integrate? Newfound Research Report. July.

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

29. Homescu, C. (2015). Better investing through factors, regimes and sensitivity analysis. Available at SSRN 2557236.

30. Hunstad, M., & Dekhayser, J. (2015). Evaluating the Efficiency of "Smart Beta" Indexes. The Journal of Index Investing, 6(1), 111-121.

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

32. Lobosco, A., & DiBartolomeo, D. (1997). Approximating the confidence intervals for Sharpe style weights. Financial Analysts Journal, 53(4), 80-85.

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

34. Meucci, A. (2001). Common pitfalls in mean-variance asset allocation. Wilmott Magazine.

35. Meucci, A. (2007). Risk contributions from generic user-defined factors. The Risk Magazine, 84-88.

36. Meucci, A. (2016). Advanced Risk and Portfolio Management Online Lab. Available at http://www.arpm.co/lab/overview/.

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

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

39. Muller, C., & Ward, M. (2013). Style-based effects on the Johannesburg Stock Exchange: A graphical time-series approach. Investment Analysts Journal, 42(77), 1-16.

40. Mutooni, R., & Muller, C. (2007). Equity style timing. Investment Analysts Journal, 36(65), 15-24.

41. Mutswari, P. (2016). On the hunt for a factor model of South African stock returns. University of Cape Town, MPhil Dissertation.

42. Novy-Marx, R. (2013). The other side of value: The gross profitability premium. Journal of Financial Economics, 108(1), 1-28.

43. Podkaminer, E. L. (2013). Risk factors as building blocks for portfolio diversification: The chemistry of asset allocation. Investment Risk and Performance Newsletter, 2013(1).

44. Rensburg, P. V., & Robertson, M. (2003). Size, price-to-earnings and beta on the JSE Securities Exchange. Investment Analysts Journal, 32(58), 7-16.

45. Richard, J. C., & Roncalli, T. (2015). Smart Beta: Managing Diversification of Minimum Variance Portfolios. Available at SSRN 2595051.

46. Roncalli, T., & Weisang, G. (2012). Risk Parity Portfolios with Risk Factors, Available at SSRN 2155159.

47. Ross, S. A. (1976). The arbitrage theory of capital asset pricing. Journal of economic theory, 13(3), 341-360.

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

49. Sharpe, W. F. (1992). Asset allocation: Management style and performance measurement. The Journal of Portfolio Management, 18(2), 7-19.

50. Strugnell, D., Gilbert, E., & Kruger, R. (2011). Beta, size and value effects on the JSE, 1994–2007. Investment Analysts Journal, 40(74), 1-17.

51. Van Rensburg, P. (2001). A decomposition of style-based risk on the JSE. Investment Analysts Journal, 30(54), 45-60.

Authors' Bios

Emlyn Flint
Peregrine Securities

Emlyn Flint joined Peregrine Securities as a derivative, risk and portfolio management analyst in 2012 and has won a number of South African financial industry awards in this role. He is a regular speaker at industry conferences and has authored a number of academic papers in the quantitative finance field. Prior to this, Emlyn he lectured in the Finance and Tax Department at the University of Cape Town, during which time he completed an MCom in actuarial science. He is currently studying towards a PhD in Applied Mathematics at the University of Pretoria.

Anthony Seymour
Peregrine Securities

Anthony Seymour joined Peregrine Securities in 2006 after working as a quantitative analyst at Cadiz Specialised Asset Management. He holds master's degrees in Chemistry and Mathematics of Finance, both from the University of Cape Town.

Florence Chikurunhe
Peregrine Securities

Florence Chikurunhe completed her BSc (Hons) in Advanced Mathematics of finance at the University of the Witwatersrand. She joined Peregrine Securities in 2009 as a quantitative and derivatives analyst.