

February 2020

AUTHORS

Billie Bi

Manager, Equity
Index Research
+44 0 207 866 1955
billie.bi@lseg.com

Andrew Dougan

Director, Research
and Analytics
+44 0 207 866 1975
andrew.dougan@lseg.com

Ross French

Manager, Equity
Index Research
+44 0 207 866 1655
ross.french@lseg.com

Peter Gunthorp

Head of Equity
Index Research
+44 0 207 866 1962
peter.gunthorp@lseg.com

Elizaveta Nebykova

Manager, Equity
Index Research
+44 0 207 866 1940
elizaveta.nebykova@lseg.com

Executive summary

Factor strategies differ significantly in terms of portfolio construction. There is much debate regarding the pros and cons of alternative construction approaches. A vigorous debate has unfurled over the best way to combine factors into a single portfolio, with some practitioners preferring the top-down approach and others favoring a bottom-up "integrated" approach.

In this paper, we:

- explain how the tilt methodology maybe extended to incorporate explicit exposure targets.
- outline use cases and applications of the target exposure methodology that encompass multi-factor portfolio construction, sustainable investment objectives, and reductions in implementation risk
- contemplate how the target exposure framework may be deployed to implement alternative factor allocation regimes and discuss the role of pure factor portfolios, from their creation to use as an investment solution.
- consider absolute return products, then illustrate the role that pure factor portfolios play in constructing genuine Top-down, multi-factor solutions.

Contents

1	Introduction	3
2	Factor portfolio construction	4
3	Applications and Examples	10
4	Conclusions	28
5	Appendix	29
6	References	31

1 Introduction

In recent decades, academic researchers have examined theoretical explanations of, and empirical evidence for the existence of factor risk premia associated with equity characteristics such as Value, Momentum, Size, Low Volatility and Quality. As the popularity of factor investing has risen and the number of factor-based investment products has increased, the debate has shifted from the existence of such factors to assessing alternative strategies designed to capture them [18, 19, 20].

Factor strategies differ significantly in terms of portfolio construction. There is much debate regarding the pros and cons of alternative construction approaches. A vigorous debate has unfurled over the best way to combine factors into a single portfolio, with some practitioners preferring the Top-down approach [11, 15, 16] and others favoring a Bottom-up "integrated" approach [2, 4, 6, 18, 19].

Much of the controversy highlights the outcomes of different methodologies, such as turnover, levels of diversification and tracking error. However, a fundamentally important metric is often disregarded – that of factor exposure. Despite its centrality to factor investing, absolute and relative levels of factor exposure are typically accepted as an artifact of the construction methodology, rather than a conscious and controllable choice.

In order to harvest long-term factor risk premia, factor exposure is a necessary condition. However, more exposure is not necessarily better. For suitably diversified portfolios, both expected return and tracking error are increasing functions of factor exposure. This suggests investors can choose levels of factor exposure consistent with an expected return or tracking error objective and leave the information ratio relatively unaffected. In this paper, we assess the merits of developing portfolios with explicit (factor) exposure targets.

Factor portfolio construction typically focuses on selecting stocks with desirable factor characteristics. In contrast to this, the FTSE Russell framework utilizes the concept of tilting to achieve controlled increases in factor exposure while maintaining appropriate levels of diversification [2, 4, 6]. The tilt framework provides the freedom to adjust the strength of the tilt and therefore the level of exposure, and by increasing (decreasing) the tilt strength, the corresponding factor exposure is elevated (reduced). A fixed tilt strength, in contrast, results in a level of factor exposure that varies through time. This variation is a reflection of changes in the distribution of factor characteristics across stocks and changes in correlation between factors.

The natural alternative to fixed tilts and variable exposures is therefore variable tilts and specific exposure targets. In this paper, we extend the FTSE Russell Tilt framework to encompass explicit exposure targets and examine the implications of this, from both an investment solution and a product development perspective.

This paper is organized as follows; in part 2, we explain how the tilt methodology maybe extended to incorporate explicit exposure targets. Part 3 outlines use cases and applications of the target exposure methodology that encompass multi-factor portfolio construction (3.1), sustainable investment objectives (3.2) and reductions in implementation risk (3.3). In 3.4, we consider how the target exposure framework may be deployed to implement alternative factor allocation regimes. In 3.5, we discuss the role of pure factor portfolios, from their creation to use as an investment solution. In 3.6, we consider absolute return products, while in 3.7, we illustrate the role that pure factor portfolios play in constructing genuine Top-down, multi-factor solutions. Finally, we summarize our findings in part 4.

2 Factor portfolio construction

Most portfolio construction methodologies have some mechanism that permits the levels of factor exposure to be ratcheted up and down; however, factor exposure is rarely the main objective and an alternative decision metric is commonly employed. Examples range from relatively simple selection and weighting techniques (S&W) through to complex, optimized portfolios. The former consists of selecting a set of stocks with desirable factor characteristics and then overlaying some portfolio-weighting scheme. The selection cut-off may be used to control the levels of target exposure, but not off-target exposures. Optimization often does not explicitly target a precise set of factor exposures, but rather a composite factor score, a tracking error or risk budget. Consequently, many factor portfolios display factor exposures that vary through time, reflecting changes in the distribution of factor characteristics and correlations between factors. Many practitioners acknowledge this problem and employ ad hoc factor correction schemes to their S&W portfolios [13] or else add explicit factor constraints in their optimized solutions [21].

In contrast, the FTSE Russell tilt approach provides a transparent mechanism for exercising complete and precise control over single and multiple factor objectives. In [4, 5, 6], we have demonstrated the relative efficiency of the tilt mechanism versus other approaches in terms of the exposure/diversification trade-off and implementation outcomes and emphasized the importance of equalizing the factor exposures before undertaking portfolio comparisons.

In the following sub-sections, we examine the tilt approach in more detail and illustrate how it may be used to construct portfolios with multiple exposure objectives.

2.1 Single Factor

Equation (1) illustrates the basic concept: the application of a factor tilt to an initial set of portfolio weights results in a new portfolio with increased exposure to the factor. Increasing the strength of the tilt increases the level of factor exposure.

$$W_F = W_M \times S_{Value}^p \quad (1)$$

where W_F represents the set of the factor portfolio weights (one for each stock), W_M is an initial set of market capitalization weights and S_{Value}^p is a set of positive numbers representing the factor scores of interest (in this case Value) and p is the strength of the tilt. Equation (1) results in a set of unadjusted stock weights that require normalization or rescaling such that they sum to one. We use this convention throughout this paper, where it is understood that the normalization is implicit. The S in equation (1) arises as a mapping from each stock's factor Z-score to a positive number. The Z-scores themselves result from the cross-sectional normalization of individual stock level characteristics (see Appendix for more details). We return to the precise functional form of this mapping in Section 2.5.

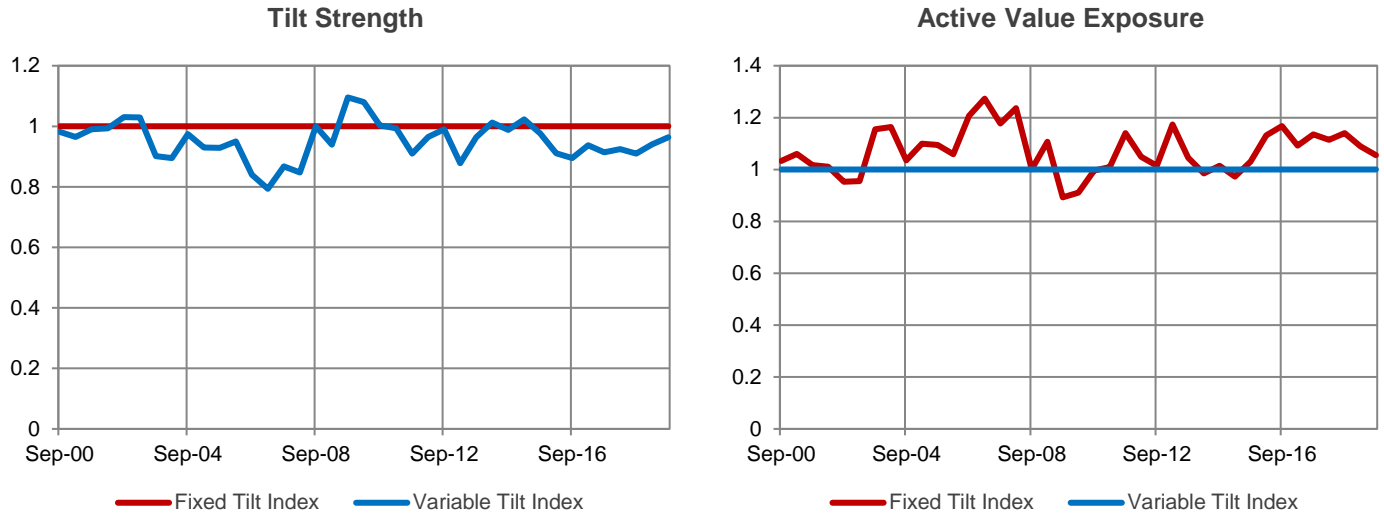
The portfolio level of factor exposure is simply the sum product of the set of portfolio weights W and factor Z-scores Z (see Appendix).

$$\text{Exposure} = W \cdot Z \quad (2)$$

This is a point in time expression for the exposure of any portfolio W and we use it throughout this paper.

In the context of the FTSE Russell Tilt framework [2], a constant or fixed tilt strength results in factor exposures that vary through time. Alternatively, we may fix the level of factor exposure at a constant target level and vary the tilt strength that is required to achieve this target level of exposure. Consider the example in Figure 1 for FTSE Developed: the Fixed Tilt Index applies a constant unit tilt towards Value (shown in red in the left graphic) and as a consequence, the level of active Value exposure varies (shown in red in the right graphic). In contrast, the Variable Tilt Index exhibits a constant level of active Value exposure of 1.0, which is achieved by varying the tilt strength through time (blue lines).

Figure 1: Fixed and variable tilt indexes: Tilt strength and active value exposure



Source: FTSE Russell. Data based on the FTSE Developed Universe and Value factor scores from September 2000 to September 2019.

The tilt strength, p of the Variable Tilt Index is obtained at each rebalance by substituting the expression for the factor weight from equation (1) into the active exposure version of equation (2) and setting the result equal to the target active exposure of 1.0:

$$(W_M \times S_{Value}^p - W_M) \cdot Z_{Value} = 1.0 \quad (3)$$

The solution may be obtained by simple numerical methods. Once the tilt strength is known, the portfolio weights that generate the target level of factor exposure may be obtained from equation (1).

2.2 Multi-factor

The single factor approach may be extended to the case of multiple factors. Additional factor objectives may be incorporated in a multiplicative manner by tilting an initial set of portfolio weights towards multiple factors. This results in a portfolio with increased exposure to all of the target factors.

$$W_F = W_M \times S_{Value}^n \times S_{Quality}^p \times S_{Volatility}^q \times S_{Size}^r \times S_{Momentum}^s \quad (4)$$

In a similar manner to the single factor case, the application of equal tilt strengths to each factor would result in differing levels of exposure to each factor. Furthermore, those exposures would fluctuate through time, reflecting variations in the distribution of factor characteristics and the time varying nature of factor correlations. In order for a portfolio to target fixed, consistent levels of exposure to several factors, multiple time varying factor tilt strengths are required.

Consider the example of Value and Quality, where the target level of active Value exposure is 1.0, as in the earlier single factor example, but an additional active Quality exposure target of 0.5 is specified. The solution now requires us to solve two simultaneous equations; one for the level of active Value exposure, and one for the active Quality objective:

$$(W_M \times S_{Value}^p \times S_{Quality}^q - W_M) \cdot Z_{Value} = 1.0$$


and

$$(W_M \times S_{Value}^p \times S_{Quality}^q - W_M) \cdot Z_{Quality} = 0.5$$

(5)

Figure 2 demonstrates how the resulting variable tilts to Value and Quality achieve both the active factor exposure targets. However, these are not the only factor exposures that exist in the portfolio. Figure 2 illustrates that other non-targeted exposures exist, in particular, negative active Momentum exposure. These additional factor exposures arise as a result of the correlation between factors that are explicitly targeted and those that are not.

Figure 2: Value & Quality: Tilt strengths and active exposures



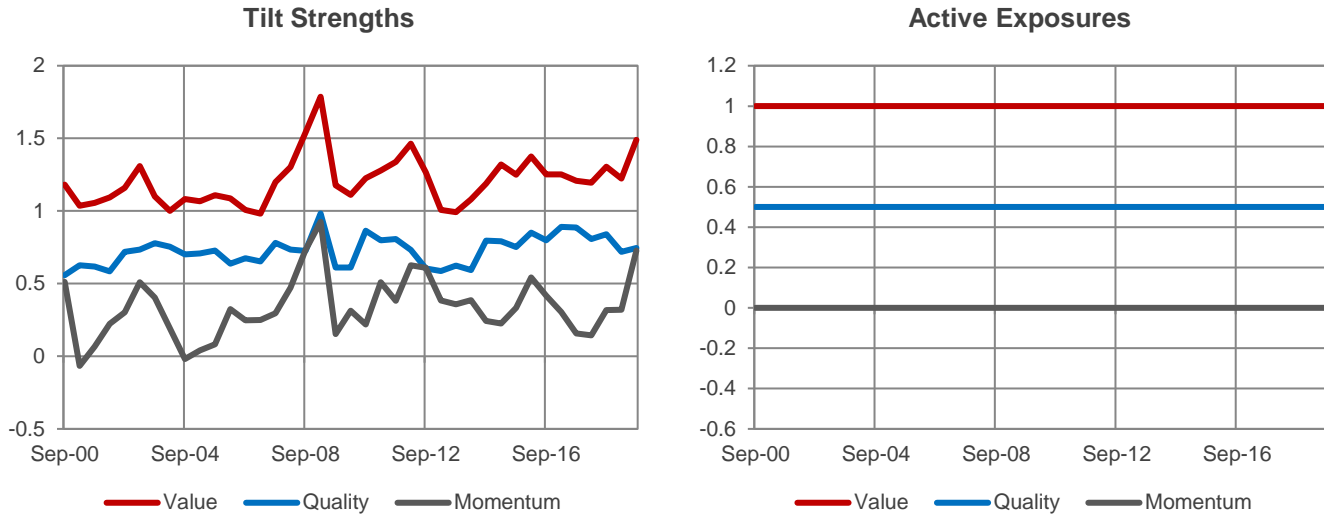
Source: FTSE Russell. Data based on the FTSE Developed Universe and Value and Quality factor scores from September 2000 to September 2019.

2.3 Corrective tilts

We have seen that variable tilts permit specific factor exposure objectives to be achieved. However, they may also result in exposures to factors other than those targeted. These exposures are usually referred to as off-target exposures. Such exposures are common to most methodologies, but are relatively large for S&W approaches (see [5]) where the typical weighting schemes applied introduce uncontrolled off-target factor exposures. Off-target exposures may be controlled through use of constraints in optimized approaches. Within the tilt framework, they may be controlled or entirely eliminated through the use of corrective tilts, frequently by applying low-strength tilts. This is appealing as it retains the same portfolio construction framework and transparency of the tilt approach.

Consider the Value plus Quality portfolio of Section 2.2; Figure 2 shows that, although we achieve the fixed Value and Quality exposure targets, the portfolio also exhibits variable and negative active exposure to Momentum, which may adversely influence performance outcomes. We can correct for this by applying an additional Momentum tilt to neutralize this unwanted (and un-targeted) exposure. We illustrate this in Figure 3, where the tilt power to Momentum varies through time reflecting its changing correlation with Quality and Value. The tilt strengths to Quality and Value also change slightly from those shown in Figure 2 in order to maintain the fixed exposure targets of 1.0 and 0.5 in the face of the additional Momentum objective.

Figure 3: Value & Quality plus neutral momentum: Tilt strengths and active exposures



Source: FTSE Russell. Data based on the FTSE Developed Universe and Value, Quality and Momentum factor scores from September 2000 to September 2019.

2.4 Tilting – A general exposure framework

In the discussion thus far, we have considered tilting towards traditional style or risk factors. The tilt approach may be generalized to target other types of exposure. This may include Sustainable Investment characteristics spanning Carbon and ESG objectives, in addition to country and industry weight objectives. The framework can also incorporate target levels of active market (Beta) exposure.

The general tilt equation is as follows:

$$W_F = W_M \times \underbrace{S_{Val}^n \times \dots \times S_{Mom}^p}_{\text{Style Factors}} \times \underbrace{S_{Carbon}^q \times \dots \times S_{ESG}^r}_{\text{Sustainable Investment}} \times \underbrace{S_{Beta}^s}_{\text{Beta}} \times \underbrace{C \times I}_{\text{Country \& Industry}} \quad (6)$$

Each exposure target merely requires the incorporation of an additional term in the multiplicative formula. All exposure objectives are therefore embedded in *exactly the same way*. Indeed, additional implementation properties such as, levels of investment capacity, maximum stock weights and turnover, may also be expressed and controlled as additional multiplicative tilts in equation (6).

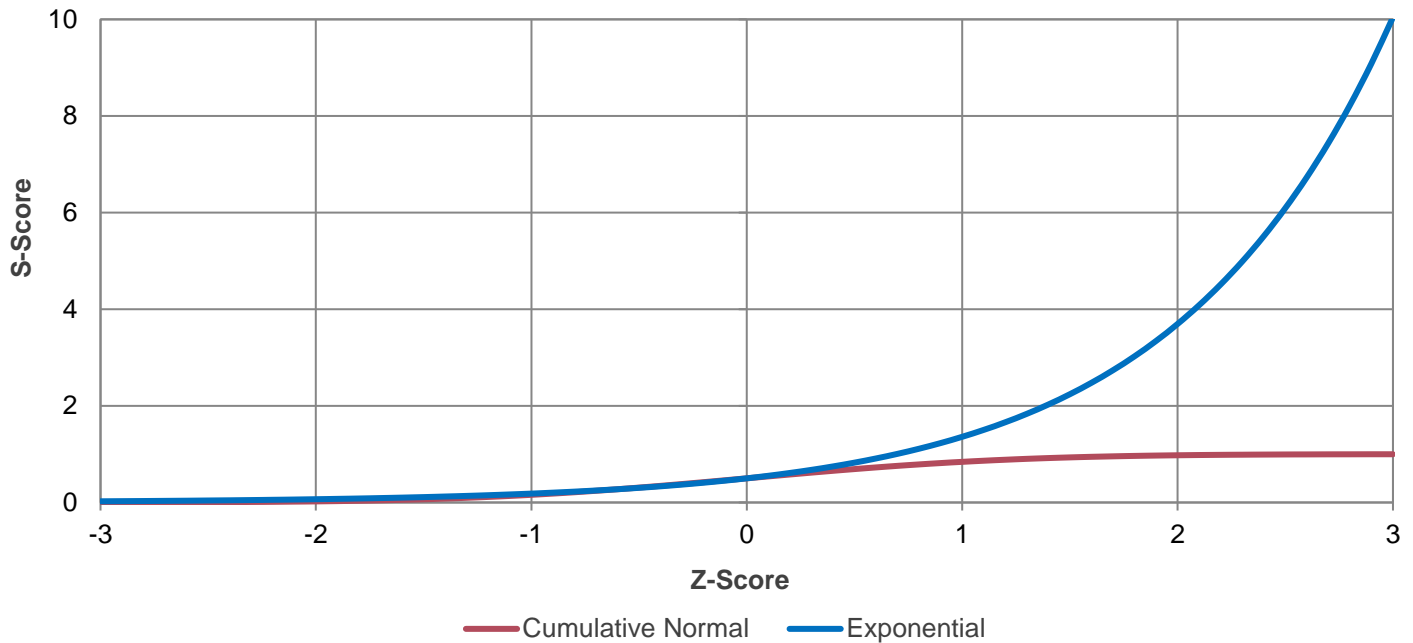
2.5 Z-score mapping

Earlier, we referred to the use of a mapping function that maps Z-scores to positive S-scores. Z-scores themselves result from the cross-sectional normalization of individual stock level characteristics. For example, if our Size metric is minus the natural logarithm of market capitalization, the Z-score normalization converts this raw metric to lie between +/-3 standard deviations.

The mapping function may take several forms, with the only basic requirement that it results in a score that is positive and monotonic. Previously, we have used the Cumulative Normal for fixed tilt indexes [1], but for exposure targeting we will use the exponential function.

It is easy to demonstrate that either mapping function yields more or less equivalent results when there are few exposure objectives and the exposure targets are themselves small. However, in the case of high exposure and/or multiple exposure objectives, the exponential function facilitates the finding of a solution where the cumulative normal finds this difficult or cannot. To see why, one merely needs to examine Figure 4, which shows the profiles of these two tilt functions:

Figure 4: Mapping functions: Cumulative normal and exponential



Source: FTSE Russell.

For high Z-scores, the exponential may take on a much higher range than the cumulative normal, thus allowing it greater freedom to achieve higher factor exposure. However, this freedom comes at a price of possibly greater concentration in individual stocks and therefore the exponential tilt function must always be used in conjunction with an investment capacity constraint.

2.6 How much exposure?

The ability to control the level of factor exposure raises the question of what is an appropriate level of factor exposure. As active factor exposure increases, tracking error also increases. Tracking error and total active factor exposure goals are ultimately determined by the investment objectives. Accordingly, a useful factor framework should provide the flexibility to target different levels of active factor exposure and tracking error. Investment objectives requiring low levels of active risk can be met by targeting lower levels of factor exposures along with tighter industry and country constraints. For investors seeking higher levels of factor exposure, higher levels of associated active risk can be targeted and looser constraints applied.

Once an appropriate level of exposure has been determined that is consistent with risk budget or tracking error considerations, there are several approaches that may be taken to allocating this risk budget across a set of factors. We distinguish three levels of sophistication and discuss each in more detail in Section 3.7.

Equal exposure – Parsimonious

A simple approach to factor allocation is an analogue of equal weighting or, more accurately, an equal factor exposure allocation. This is consistent with a view of factor exposures as a source of excess return, where the excess return of all factors is proportional to active factor exposure. The only information required is the cross-sectional dispersion in factor scores. The target levels of active exposure can be adjusted for differences in cross-sectional dispersion, ensuring that all are equally attainable.

Inverse risk weighted – Simple risk model

A more complex method is to allocate factor exposure in inverse proportion to factor return volatility – lower volatility factors warrant a greater allocation. This requires a time series of factor returns and is predicated on the persistency of factor return volatility. It can be viewed as analogous to an inverse variance weighting scheme, which itself may be interpreted as a proxy for the minimum variance allocation solution when all correlations are zero.

Equal risk contribution – Sophisticated risk model

A more sophisticated allocation model is to allocate to each factor such that ex ante, the contribution of each factor to tracking error is equal. This requires an estimate of the factor return correlation matrix and stability of both factor correlation and volatility. This is akin to the minimum variance solution, which utilizes the full covariance matrix.

3 Applications and Examples

In this part, we examine applications of the exposure targeting framework described earlier. These use cases span a series of practical issues facing investors. The examples are intended to demonstrate the flexibility of the target exposure approach in incorporating exposure objectives beyond those of traditional compensated Style factors. These include objectives covering Sustainable Investment (SI) criteria and uncompensated risk exposures to countries and industries.

In each case, increasing the level of active exposure increases the level of active risk, but only compensated factor exposures are expected to be rewarded with an associated return premium. Consequently, the ability to constrain uncompensated exposures to limit uncompensated contributions to active risk is desirable. Section 3.1 demonstrates how the target exposure approach may be extended to incorporate additional exposure objectives and highlight the effect of imposing Beta, country and industry constraints on performance and risk outcomes.

Section 3.2 focuses on extending the role of exposure targets to sustainable investment (SI) goals. We illustrate the improvements to the efficiency with which carbon emissions reductions may be achieved by controlling for interactions between SI characteristics and traditional rewarded Style factor exposures.

More generally, in Section 3.3 we introduce the role of factor corrections. We apply a set of factor corrections to a portfolio that is weighted using a common non-market capitalization portfolio weighting scheme. This allows the undesirable factor exposures introduced by the portfolio weighting scheme to be removed and emphasizes the implementation risks associated with off-target factor exposures resulting from the use of non-market capitalization weighting portfolio weighting schemes [11, 12].

In Section 3.4, we illustrate how alternative approaches to the factor allocation decision may be implemented within the target exposure framework. This demonstrates how increasingly sophisticated approaches to the factor allocation decisions result in improved risk control.

In Section 3.5, we develop the idea of pure factor portfolios and demonstrate how the tilt framework may be employed to create portfolios with specific levels of exposure to a target factor but remain neutral with respect to other risk factors. Sections 3.6 and 3.7 deploy the notion of pure factor portfolios as an investment solution.

In Section 3.6, we examine the role of pure factor portfolios in facilitating the creation of absolute return products.

Section 3.7 illustrates the role that pure factor portfolios can play in constructing genuine Top-down multi-factor portfolios. We show that factor purity is essential if the attributional transparency of a Top-down approach is to be maintained. This Section also considers the use of pure factor building blocks for the creation of bulk solutions and illustrates the essential role that pure factor portfolios play in a factor overlay program.

3.1 Multi-factor indexes

The target exposure methodology gives an ability to precisely control factor exposures. In this section, we demonstrate this flexibility and control by examining a five-factor example that targets Value, Quality, Low Volatility, Size and Momentum factor exposures for the FTSE All-World universe of stocks (for factor definitions see [1]). As we discussed briefly in Section 2.6, a simple and parsimonious approach to factor allocation is to allocate to factors in proportion to σ_k , the market capitalization weighted cross-sectional dispersion in factor Z-scores (see Appendix):

$$\text{Active Exposure Target}_K = a \times \sigma_k \tag{7}$$

where a is the raw active factor exposure target. Factors that exhibit greater levels of dispersion will have relatively higher exposure targets. Here we set $a = 0.4$ for all factors. This exposure target level does not favor any one factor over another and is consistent with a typical investor's tracking error targets.


In addition to controlling factor objectives, the target exposure approach also permits the control of additional targets of interest such as Beta, country and industry exposures. In order to highlight this, we examine two additional, but

comparable five factor indexes that incorporate additional exposure targets: the first incorporates Beta neutrality and the second applies Beta, industry and country neutrality.

In this context “neutral” indicates that active Beta, country and industry exposures are zero i.e. the absolute exposures match those of the underlying market capitalization weighted index. Each index is rebalanced in March and September with an annual two-way turnover constraint of 40% imposed at each rebalance.

Figure 5 illustrates the relative performance over the period from September 2000 to September 2019. Despite some variability in the performance of individual factors, each index exhibits consistently a strong performance – factor diversification is beneficial in this example. Interestingly, an increase in the number of objectives or a more constrained solution is consistent with lower performance.

Figure 5: Relative performance: FTSE All-World five factor indexes



Source: FTSE Russell. Data based on the FTSE All World Index Universe from September 2000 to September 2019.

Table 1 examines the performance of each index more deeply. As expected, the additional constraints reduce tracking error and thereby improve information ratios. Conversely, the Sharpe Ratios of the Beta neutral indexes are lower than those of the Factors Only index.

Table 1: Performance Statistics: Factors Only, Beta Neutral and Beta, Country & Industry Neutrality

	Five Factor Index			
	FTSE All World	Factor Only	Beta Neutral	Beta, Country & industry Neutral
Geometric Return (%p.a.)	5.35	8.70	8.29	7.85
Volatility (%p.a.)	15.62	13.80	15.84	15.44
Sharpe Ratio	0.33	0.63	0.52	0.51
Max Draw Down (%)	-57.95	-51.57	-54.90	-54.59
Excess Return (%p.a.)		3.35	2.94	2.49
Tracking Error (%p.a.)		3.53	2.81	2.33
Information Ratio		0.95	1.04	1.07
Beta		0.87	1.00	0.98
Two-Way Turnover (%p.a.)	13.48	82.07	83.50	83.36

Source: FTSE Russell. Data based on the FTSE All-World Index Universe from September 2000 to September 2019.

This arises partly because the Factors Only index exhibits less absolute risk than the FTSE All-World benchmark due to its sub-market Beta of 0.87, whereas the Beta neutral indexes display approximately the same levels of absolute risk by design. The incremental reduction in the performance of the Beta, Country and Industry Neutral Index is the result of not taking country and industry bets that actually paid off for the other indexes during the period between September 2000 and September 2007. Active country and industry exposure is typically associated with unrewarded active risk, rather than as a source of long-term risk premia.

Finally, we note that the annual two-way turnover for each factor index is slightly greater than 80% due to intra-rebalance deletions and corporate actions.

3.2 Sustainable investment objectives

The target exposure approach is not limited to compensated factor exposure targets but is equally applicable to other investment goals such as Sustainable Investment objectives, including Carbon Emissions and ESG targets. In this section, we discuss how target exposure portfolios may be formed to meet Sustainable Investment objectives.


One approach to achieving a Carbon Emissions goal is to form portfolios that increase (reduce) the weighting of companies exhibiting lower (higher) levels of carbon intensity (see [8] for definitions). For example, equation (1) may be applied with a tilt strength p that is consistent with a given reduction in Carbon Emissions relative to the benchmark index.

This approach may be extended to target an equivalent level of Carbon Emissions reduction, while simultaneously ensuring that all other active exposures are zero. Hence, we can utilize equation (6) to simultaneously target Carbon Emissions and factor, industry and country exposures, thereby solving any difficulties arising from the interaction of Sustainable Investment and traditional factor objectives.

We refer to these two portfolios as the Carbon Emissions and Pure Carbon Emissions portfolios respectively. Figure 6 illustrates the trade-off between the tracking error for various levels of Carbon Emissions reductions. Both approaches are

able to achieve the same Carbon Emission reduction goals. However, the tracking error of the Pure Carbon Emissions portfolio is always lower for a given emissions reduction target as a result of setting all active factor exposures to zero.


Figure 6: Tracking error versus reduction in emissions intensity



Source: FTSE Russell, from March 2012 to September 2019, based on the FTSE Developed Index Universe. Trucost data, September 1, 2019.

Industry, country and style factors are considered risk factors. Consequently, an active allocation to these factors contributes to an increase in both portfolio volatility and tracking error. Figure 7 reports the average industry weights and style exposures. The Carbon Emissions portfolio exhibits active industry positions because it simply underweights carbon intensive industries such as oil & gas, basic materials and utilities and is overweight financials, technology and healthcare. It also exhibits a large cap bias, since banking and healthcare constituents are generally larger than the average stock. In contrast, the Pure Carbon Emissions portfolio holds the benchmark industry and country weights and sets all active style factor exposures to zero. The tracking error of the Pure Carbon Portfolio is lower as a result, because it is exposed to lower levels of systematic risk.

Figure 7: Time averaged Industry and Style exposures



Source: FTSE Russell, from March 2012 to September 2019, based on the FTSE Developed Index Universe. Trucost data, September 1, 2019.

The flexibility of the portfolio construction approach outlined, allows us to achieve more balanced outcomes than would result from the simple elimination of those sections of the economy that are intrinsically carbon intensive. Carbon intensive companies are concentrated in a relatively small number of industries and naive construction approaches tend to simply re-allocate capital from carbon intensive industries to carbon light industries.

3.3 Implementation risk and corrective tilts

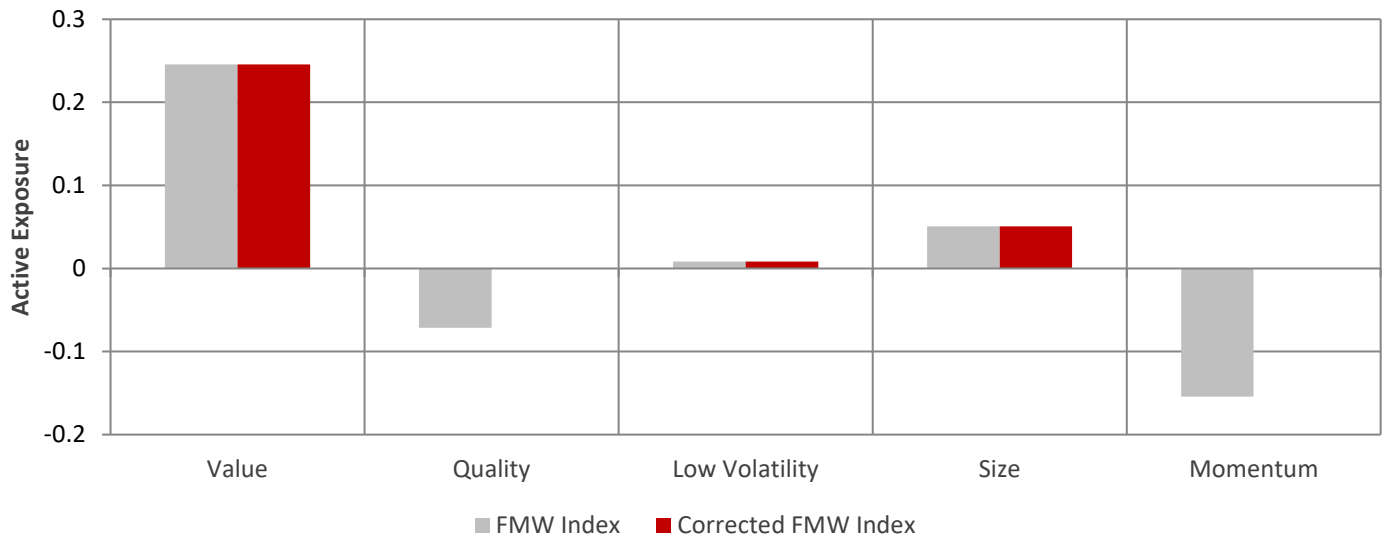
It is well known that non-market capitalization and diversified portfolio weighting schemes result in exposure to off-target and unwanted factors [3]. This is not surprising as such schemes are often designed with different objectives in mind e.g. improved levels of diversification or risk reduction rather than to harvest factor risk premia. Nevertheless, such unwanted factor exposures frequently result in implementation risk with a negative effect on portfolio performance, particularly if such exposures are systematically in the wrong direction.

As an example, consider the FTSE Developed Financial Metrics Weighted (FMW) Index, which weights companies according to the size of their economic footprint via a set of financial metrics, including net income, adjusted cash flow,

dividend plus buybacks and book value. The index is an equally weighted composite of four sub-indexes, each weighted in proportion to one of the above measures of economic size. For more details, please refer to the FTSE Global Financial Metrics Weighted Index Series Ground Rules [7].

The grey bars in Figure 8 show the average active factor exposures of the FMW index. The dominant exposure is to Value, which is positive as one would expect given the weighting scheme. However, there are also small positive exposures to Size and Low Volatility and more importantly, significant negative exposures to both Quality and Momentum. Again, this is unsurprising as both Quality and Momentum tend to be negatively correlated with Value. Negative exposure to factors that have historically exhibited positive factor premiums will detract from performance. In the remainder of this section, we propose one means of overcoming this problem.

Figure 8: Average active exposure: FMW and corrected FMW Indexes



Source: FTSE Russell. Data based on the FTSE Developed Index Universe at rebalance from September 2001 to September 2019.

The target exposure methodology allows us to neutralize any or all of these unwanted factor exposures. In this example, we correct the unwanted (negative) active Quality and Momentum exposures, while maintaining the remaining active factor exposures at their original levels. This is done by applying a set of corrective tilts to the initial FMW portfolio weights at each periodic rebalance. Mathematically, this is equivalent to solving the following system of equations:

$$(W_{FMW} \times S_{Val}^n \times S_{Qual}^p \times S_{Vol}^q \times S_{Size}^r \times S_{Mom}^s - W_M) \cdot Z_{Fac} = \begin{cases} (W_{FMW} - W_M) \cdot Z_{Fac} & \text{where } Fac = Val, Size, Vol \\ 0 & \text{where } Fac = Qual, Mom \end{cases} \quad (7)$$

where W_{FMW} is the initial set of FMW portfolio weights and Z_{Fac} are factor Z-scores. The average active exposures of the Corrected FMW Index are shown in red in Figure 8; the active exposures to Quality and Momentum are eliminated (set to zero) and all other factor exposures are equal to those of the original FMW portfolio.

Table 2 shows the diversification and implementation metrics of the corrected FMW index are more or less the same as those of the original FMW index. It also illustrates a slight uplift in performance and lower tracking error culminating in a near doubling of the information ratio.

Table 2: Performance statistics: FTSE Developed, FMW and Corrected FMW Indexes

	FTSE Developed	FMW Index	Corrected FMW Index
Geo. Return (%p.a.)	7.79	7.99	8.12
Volatility (%p.a.)	15.43	15.92	15.87
Sharp Ratio	0.50	0.50	0.51
Max. Drawdown (%)	-57.37	-58.80	-58.10
Excess Return (%p.a.)		0.19	0.31
Tracking Error (%p.a.)		1.79	1.57
Information Ratio		0.11	0.20
Beta		1.03	1.02
Effective N	355	320	316
Capacity (%)	100.00	69.37	68.84
Active Share (%)	0.00	21.38	21.72
2-Way Turnover (%p.a.)	12.41	25.04	27.22

Source: FTSE Russell. Data based on the FTSE Developed Index Universe from September 2001 to September 2019.

The performance differential is more apparent in Figure 9, which shows the performance of the Corrected FMW Index relative to the original FMW Index. The former outperforms the original FMW index over the whole period and particularly during the volatile periods where Quality exposure would have provided valuable downside protection and up-trending markets, where Momentum exposure was rewarded.

In this example, we have corrected for unwanted factor exposures but in line with the consideration of Section 3.2 we could have equally applied sustainable investment corrections to the FMW portfolio.

Figure 9: Relative performance: FMW versus Corrected FMW Index

Source: FTSE Russell. Data based on the FTSE Developed Index Universe from September 2001 to September 2019.

3.4 Factor allocation schemes

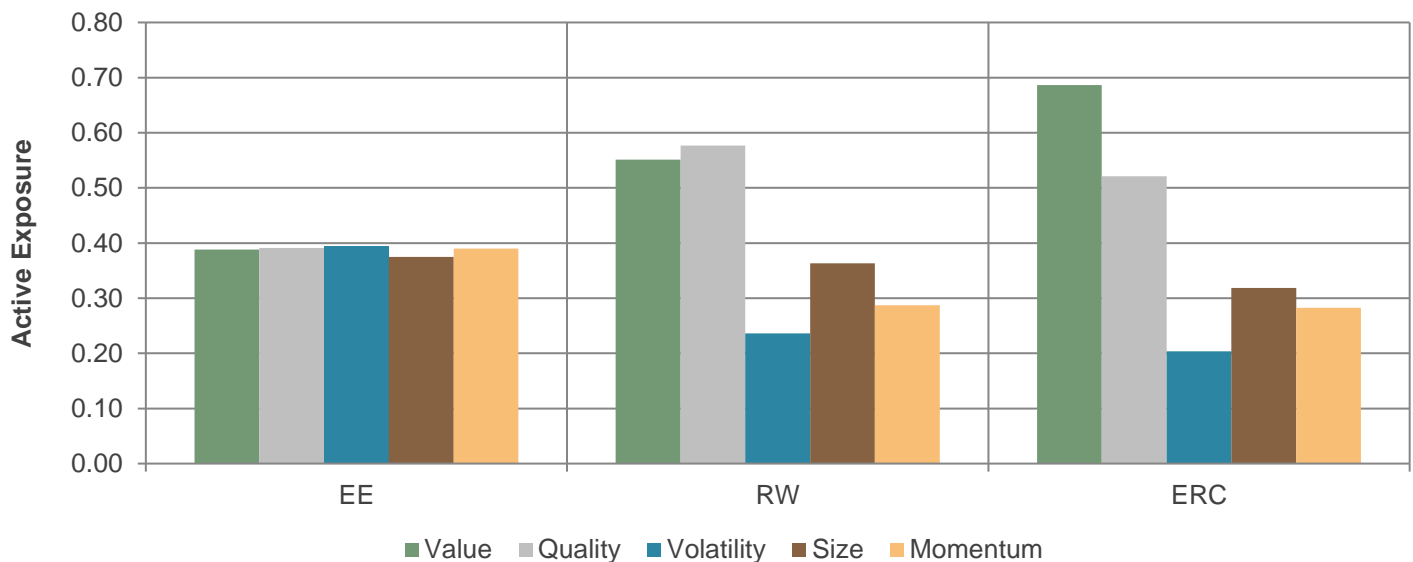
The target exposure methodology is particularly useful in the context of multi-factor strategies as the ability to precisely control factor exposures permits alternative factor allocation mechanisms to be incorporated within the same Bottom-up framework. Three possible allocation schemes are equal factor exposure (EE), risk-weighted factor exposure (RW) and factor exposures designed to deliver equal risk contributions (ERC). We examine and compare these three approaches for five factor combinations of Value, Quality, Volatility, Size and Momentum for the FTSE All-World Universe.

The total active factor exposure (the sum of the active exposures to the five factors) of each portfolio is set to two at each bi-annual rebalance and country, industry and Beta neutrality conditions are imposed. For the EE portfolio, the exposure targets are simply 1/5th of the total active factor exposure (this is the Factors Only portfolio from Section 3.1). The factor allocation targets of the RW portfolio are determined by allocating the total active factor exposure in inverse proportion to the realized (two-year daily) volatility of individual factor returns. Factor allocation in the ERC scheme is derived such that each factor, ex ante, contributes equally to the overall tracking error of the portfolio. The factor covariance matrix is constructed using two years of daily factor returns.

The factor returns themselves are derived from daily cross-sectional regressions of stock returns on prior month-end factor exposures and industry and country dummy variables. We will present a possible alternative to this technique using pure factor indexes in Section 3.5.

Figure 10 shows the average active factor exposures under the three factor allocation schemes. Active factor exposures under the EE regime are, by design all approximately equal. Conversely, the RW and ERC active factor exposures exhibit very unequal and different outcomes. In particular, the active Quality and Value exposures are substantially larger than for other factors, implying that they are less volatile and relatively uncorrelated to other factors.

Figure 10: Average active exposure: equal, risk-weighted and equal risk contribution exposure



Source: FTSE Russell. Data based on the FTSE All World Index Universe from September 2000 to September 2019.

Figure 11 displays the average percentage ex-ante contribution to tracking error resulting from the active factor exposures shown in Figure 10. The pattern is reversed with more variable contribution by factor for the EE portfolio. Low Volatility dominates and Value contributes almost nothing to the tracking error of the EE portfolio. The factor contributions to tracking error for the RW portfolio are far more balanced, and the ERC allocation scheme, results in broadly equal contributions to risk as intended. All three approaches exhibit negligible contributions to risk from country and industry exposures and approximately 30% from idiosyncratic sources.

Figure 11: Contribution to tracking error: Equal, risk-weighted and equal risk contribution exposure



Source: FTSE Russell. Data based on the FTSE All World Index Universe from September 2000 to September 2019.

Ignoring differences in the sources of active risk, there is little to differentiate the three portfolios in performance terms. Table 3 shows that the EE portfolio has lower turnover, whereas RW and ERC portfolios have lower tracking errors and therefore higher information ratios.

Table 3: Performance: Equal exposure, risk-weighted and equal risk contribution

	FTSE All World	Five Factor Index		
		EE	RW	ERC
Geometric Return (%p.a.)	5.35	7.85	7.51	7.44
Volatility (%p.a.)	15.62	15.44	15.50	15.52
Sharpe Ratio	0.33	0.51	0.48	0.48
Max Draw Down (%)	-57.95	-54.59	-54.71	-54.81
Excess Return (%p.a.)		2.49	2.16	2.09
Tracking Error (%p.a.)		2.33	1.75	1.75
Information Ratio		1.07	1.23	1.19
Beta		0.98	0.99	0.99
Two-Way Turnover (%p.a.)	13.48	83.36	101.69	101.84

Source: FTSE Russell. Data based on the FTSE All World Index Universe from September 2000 to September 2019.

The EE portfolio marginally outperforms the other two portfolios, primarily due to its relatively high exposure to Low Volatility and relatively low exposure to Value (see Figure 10).

3.5 Pure single factor portfolios

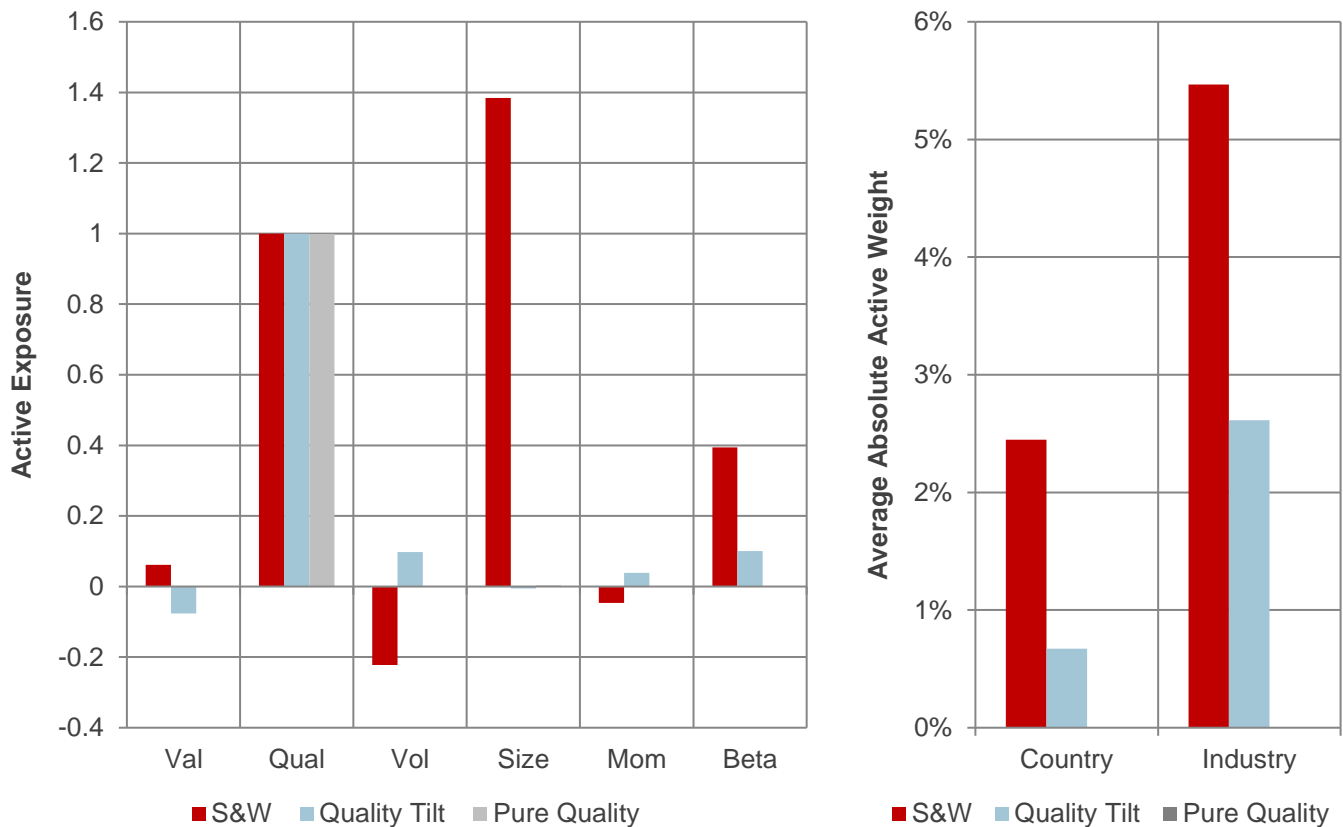
The precise control over factor exposures afforded by the target exposure framework permits the creation of long only, pure single factor portfolios. Purity requires that the factor portfolio exhibits a specific level of active factor exposure to a target factor and zero active exposure to all other factors. In addition, market Beta, industry and country neutrality conditions may be imposed so that the portfolio is pure with respect to other sources of risk and return. In this section, we introduce the notion of pure factor portfolios and examine some of their key properties.

For an investor seeking exposure to a particular factor and wishing to limit implementation risks, the pure factor approach offers certain advantages. By removing off-target factor, country and industry exposures, the excess return of a pure factor outcome will more closely reflect the behavior of the underlying factor premium. In contrast, factor portfolios with off-target factor exposures exhibit significant implementation risks as unwanted and uncontrolled outcomes are introduced.

To highlight the effects of off-target factor and country and industry exposures, we construct three semi-annually rebalanced Quality portfolios using the FTSE Developed Universe. First, we construct a Selection and Weighting (S&W) portfolio by selecting the highest ranked stocks by Quality and equal weighting the portfolio. The proportion of the universe chosen is such that the resulting portfolio has unit active Quality exposure.

The second, which we refer to as a Quality Tilt portfolio, is the result of a tilt towards Quality, where the tilt strength is chosen so that the level of active Quality exposure is again one. We contrast these two portfolios with a Pure Quality portfolio that is again constructed using the tilt mechanism to achieve unit active exposure to Quality, but also zero active exposures to all other factors (including market Beta), countries and industries. Figure 12 below displays the average active factor exposures and average (absolute) active country and industry weights of each portfolio.

Figure 12: Average active factor, country and industry exposure



Source: FTSE Russell. Data based on the FTSE Developed Index Universe at rebalance from September 2000 to September 2019.

The S&W portfolio displays the largest off-target style exposures; most remarkably it has a higher exposure to Size than to Quality, a consequence of its equal weighting methodology [5]. The Quality Tilt portfolio exhibits smaller off-target exposures (the largest being its exposure to Beta), while the off-target exposures are negligible for the Pure Quality portfolio. This highlights a desirable property of the tilt methodology; one obtains target exposures by actively tilting towards them; off-target attributes that are not actively tilted towards tend to be relatively small. The Pure Quality portfolio can therefore be interpreted as a Quality Tilt portfolio to which a series of (small) corrective tilts have been applied in order to remove any unwanted exposures.

The primary distinction between the Quality Tilt and Pure Quality portfolios can be seen in terms of their active country and industry weights; they are relatively large for the former whereas, by design, they are exactly zero for the latter. The S&W portfolio assumes the largest country and industry bets. This becomes apparent when we examine the sources of active risk associated with each portfolio in Table 4.

Table 4: Percentage Contribution to Active Risk

	Contribution to Tracking Error					Total
	Tracking Error (%p.a.)	Quality Exposure	Off-Target Exposures	Country & Industry	Residual	
S & W Portfolio	6.97%	11.2%	23.0%	62.0%	3.8%	100.0%
Quality Tilt Portfolio	2.97%	44.4%	11.4%	26.1%	18.0%	100.0%
Pure Quality Portfolio	2.38%	62.0%	0.0%	0.0%	38.0%	100.0%

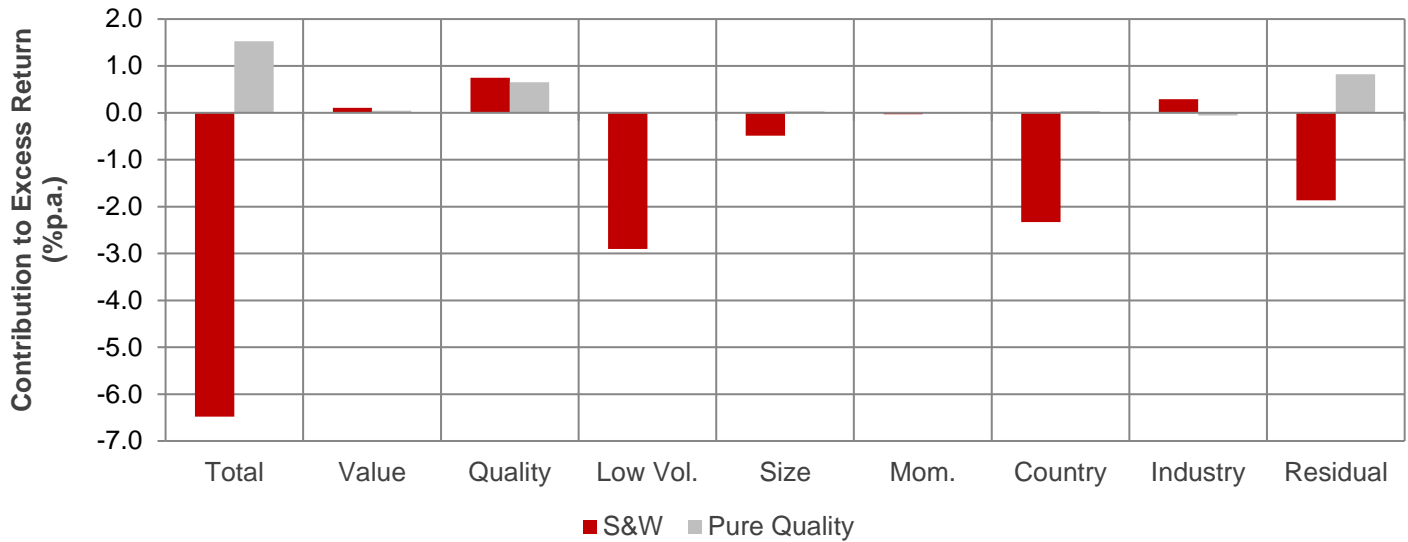
Source: FTSE Russell. Data based on the FTSE Developed Index Universe from September 2000 to September 2019.

The major source of active risk for the S&W portfolio is a result of unintentional exposures; most notably those that arise from country and industry bets. On-target Quality exposure is a minor source of active risk, suggesting such a methodology is unsuitable for the creation of single factor products. Quality exposure is the main source of active risk for the Quality Tilt portfolio, which also displays smaller, but significant contributions from country and industry sources. The Pure Quality portfolio effectively removes all sources of active risk with the exceptions of the targeted factor and stock specific (residual) risk.

Pure factor indexes will display lower levels of tracking error all else being equal by design, as the contribution to tracking error from unintentional exposures is removed. The performance of pure factor indexes is therefore more closely aligned with stated investment objectives, thereby reducing implementation risks.

To illustrate this, imagine that an investor holds the conviction that high Quality stocks will outperform over the six-month period between March 2018 and September 2018. Figure 13 compares the results of employing S&W Quality and Pure Quality portfolios as vehicles for this strategy.

Figure 13: Performance attribution: S&W Quality and Pure Quality portfolios



Source: FTSE Russell. Data based on the FTSE Developed Index Universe from March 2018 to September 2018.

The performance attribution illustrates that the Pure Quality portfolio outperforms the benchmark (“Total” bar) and that the investor was indeed correct in his prediction. However, had the S&W portfolio been used as an implementation vehicle, the result would have been to underperform the benchmark, primarily as a result of the negative performance contributions arising from Low Volatility and country exposures.

3.6 Pure Single Factor Portfolios: Absolute Return Applications

We have seen that the Target Exposure approach allows us to build pure single factor indexes with Beta equal to one. However, such indexes are “long only” meaning that while they capture the intended factor return, this is dominated by the market return. Investors may therefore be interested in portfolios that remove this “market beta” in the expectation of obtaining returns that more clearly represent the factor risk premia and which have low correlation to the return of other long-only portfolios.

Academics have traditionally formulated such factor risk premia as the difference in performance between a basket of stocks with favorable factor scores and a basket of stocks with poor factor scores. Typically, however, no attempt is made to ensure that the market Beta of the long and short sides of the portfolio is matched or off-target exposures are controlled. The Fama-French double sort construction recognises the confounding influence of off-target exposures and attempts to neutralise for Size, for example, when measuring the performance of Value [9]. However, the exposure matching is relatively crude and restricted to a limited category of exposures in a long short context.

The Factor Mimicking Portfolio (FMP) represents a precise means of measuring factor performance [10]. The FMP is a long / short portfolio with unit exposure to the factor of interest, no off-target exposures and minimizes (idiosyncratic) risk. The performance of the FMP therefore allows us to remove the influence of overall market movements and draw clear inferences regarding the performance outcomes of factors.

Despite this, the FMP is impractical from an implementation perspective, since although it is an excellent tool to assess factor performance, it is not explicitly designed for investment in factors. In general, it will contain small illiquid stocks and will require monthly or even daily rebalance. More importantly, the shorting of individual stocks required to implement the short side of the FMP is expensive and may even be proscribed by restrictions of a particular investment mandate.

However, the notion of pure long only factor indices introduced earlier raises the possibility of mirroring of such long short outcomes in a way that can be accessed by investors. Recall that, in the target exposure framework, we can neutralise a comprehensive set of off-target exposures, while achieving a specific level of exposure to a target factor in a long only context. This is the pure factor portfolio described in Section 3.5 and is analogous to the FMP but is long only with a market Beta of one.

Thus, we can consider three possible proxies for the FMP:


- Long a pure factor index with positive factor exposure and short a pure factor index with negative factor exposure.
- Long a pure factor index with positive factor exposure and short the underlying market index.
- Long the underlying market index and short a pure factor index with negative factor exposure.

The returns to such long-short portfolios could therefore be used as a basis for a Fama-French type analysis of the factor betas of other portfolios which is much closer in outcomes to the Factor Mimicking Portfolio. Moreover, the returns to such portfolios can be used as proxies for the factor returns used in Section 3.4. Hence pure factor portfolios have a further possible use case in analyzing factor risk as well as factor return.

Implementing the short side of options one and three is likely to be relatively expensive compared to shorting a futures contract (or ETF) on a benchmark index such as, the Russell 1000 under option two. However, each approach should (before costs) result in the same outcome, if each exhibits the same total active factor exposure and the factor premium on the long and short sides is symmetric. A key consideration is therefore the extent to which any factor premium is primarily a short side phenomenon.

Figure 14 shows how excess return and tracking error vary with with levels of active exposure for a pure Quality portfolio. Empirically, excess return is more or less a linear function of active exposure, implying that, at least for Quality, the payoff is symmetrical. The “V” shape of the tracking error versus active Quality exposure chart demonstrates that active risk is also an approximately linear function of active exposure.

Figure 14: Excess return and tracking error as a function of active quality exposure



Source: FTSE Russell. Data based on the FTSE Developed Index Universe at rebalance from September 2000 to September 2019.

In Table 5, we compare the performance outcomes of a long position in a Pure Quality portfolio with positive active exposure and a short position in a Pure Quality portfolio with negative active exposure for the FTSE Developed universe. We examine various combinations such that the net active Quality exposure is always one. Thus, for example, “Long +1.0 Short Bmk” in Table 5 corresponds to option two discussed above. The long and short sides of each portfolio are rebalanced and reset to equal size on a bi-annual basis.

Table 5: Performance: Long-short combinations of Pure Quality: FTSE Developed

	Benchmark	Long +1.0 Short Bmk	Long +0.75 Short -0.25	Long +0.5 Short -0.5	Long +0.25 Short -0.75	Long Bmk Short -1.0
Geo. Return (%p.a.)	5.09	0.88	0.83	0.86	0.89	0.85
Volatility (%p.a.)	15.53	1.72	1.55	1.53	1.73	2.18
Sharpe Ratio	0.33	0.51	0.54	0.56	0.52	0.39
Beta	1.00	0.00	0.00	-0.01	-0.01	-0.01
Correlation with Bmk	1.00	-0.03	-0.04	-0.05	-0.06	-0.05

Source: FTSE Russell. Data based on the FTSE Developed Index Universe from September 2000 to September 2019.

Note that the long-term geometric mean return is positive and approximately the same for all combinations. This is consistent with the previously observed symmetry to payoffs to good and poor Quality exposure. We also observe the major benefits of such an absolute return portfolio; the volatilities (in Table 5) of all the long-short combinations are significantly smaller than that of the FTSE Developed benchmark, resulting in improved Sharpe Ratios and both Beta and the correlation to the FTSE Developed benchmark index are close to zero.

3.7 Pure single factor portfolios: Multi-factor applications

A simple approach to the creation of multi-factor portfolios is to average the weights of separate single factor portfolio sleeves in some preferred proportion. Such an approach is often referred to as Top-down because it uses pre-existing portfolios as the building blocks to create multi-factor solutions. The alternative approach is to create multi-factor portfolios from the Bottom-up, where each stock is weighted in consideration of all its factor characteristics simultaneously.

Research indicates that the Top-down approach is a less efficient way of creating multi-factor portfolios than the Bottom-up approach [4, 6]. The primary reason for this is that the averaging of sleeve weights results in the averaging of exposures. This dilution effect can result in relatively low levels of factor exposure [4, 6].

The efficiency of Top-down construction is further compromised when multi-factor portfolios are composed of single factor sleeves built using construction methods, which introduce significant off-target factor exposures and unwanted industry and country biases. The Top-down portfolio will inherit any biases present in the single factor sleeves. As an example of this, the Selection & Weighting (S&W) approach introduced in Section 3.5 results in significant uncontrolled factor exposures. A Top-down multi-factor portfolio constructed using individual portfolio sleeves that are themselves created using a S&W approach is unlikely to achieve the desired levels of exposure. This arises as a result of interactions between the on and off-target exposures in the component portfolio sleeves. Similarly, industry and country biases may become material sources of unwanted and unrewarded active risk.

These problems may be alleviated by using high exposure pure factor indexes to construct the component single factor portfolio sleeves. This mitigates both the dilution effect and ensures there is no factor overlap between the individual portfolio sleeves. Investors may therefore build multi-factor portfolios with their own desired factor allocations, simply through mixing a set of pure single factor indexes in the desired proportions.

If the objective is exposure to many factors, dilution effects may still be apparent. However, if relatively low levels of exposure are sought, this is unlikely to be a problem. Indeed, it is possible to show that for low exposure targets, the difference in efficiency between this Top-down approach and a Bottom-up approach is small [4,6].

The irony of the above considerations is of course that the pure factor indexes necessary for the creation of a genuine Top-down solution themselves require a Bottom-up construction approach such as the one outlined in Section 3.5. In the following sub-sections, we discuss why appropriately constructed Top-down approaches may be convenient for investors.

We characterize each sub-section under: Bulk Solutions, Attributional Transparency and Factor Overlays. To illustrate, we create five pure factor exposure portfolios that are market Beta, industry and country neutral for Value, Quality, Momentum and Size and Low Volatility respectively. Each pure factor portfolio is rebalanced bi-annually and targets an exposure of 1.0 with the exception of Low Volatility, which targets 0.5. The reduced Low Volatility exposure target is a consequence of the tension between maintaining both Low Volatility exposure and Beta neutrality without relying on overly concentrated outcomes.

Bulk solutions

Given a comprehensive set of pure single factor building blocks, a Top-down construction approach allows an investor to create any multi-factor portfolio that is required, subject to the limitations imposed by exposure dilution effects. This is made possible by altering the proportions in which the single factor building blocks are combined.

A set of pure single factor portfolios, therefore, allows the creation of an infinite number of Top-down multi-factor combinations reflecting the varying preferences of individual investors—hence the term “bulk solutions.” Given this freedom, it would be philosophically incoherent for an index provider to promote the benefits of a Top-down approach to multi-factor index construction, while offering single factor sleeves and multi-factor products that are not linear combinations of and inconsistent with the individual single factor sleeves [11,12,13].

Examples of the type of investment objectives that a bulk solution could solve:

- 50/50 Blend of Pure Quality and Pure Low Volatility – two factor portfolio for a cautious investor requiring downside protection.
- Equal Weight Blend of Pure Value, Pure Quality and Pure Momentum – three factor portfolio for an investor with core small cap and low volatility holdings, requiring exposure to the other risk premia.
- Equal Weight Blend of Pure Value, Pure Quality, Pure Momentum, Pure Low Volatility and Pure Size – five factor portfolio for an investor requiring a fully factor diversified portfolio.

Table 4 shows the performance outcomes of the component building block pure factor portfolios and the two, three and five factor examples for the period September 2000 to September 2019. Each portfolio is rebalanced on a bi-annual basis in line with the pure single factor portfolios. It is noteworthy that the composite solutions inherit the Beta, country and industry neutral properties of the component portfolios.

Table 4: Performance: Pure factor portfolios and composite bulk solutions

	FTSE Dev.	Pure Component Portfolios					Bulk Solutions		
		Value	Quality	Mom.	Low Vol.	Size	2 Factor	3 Factor	5 Factor
Geo. Return (%p.a.)	5.09	5.80	6.16	6.05	5.33	6.38	5.75	6.04	5.97
Volatility (%p.a.)	15.53	15.87	15.46	15.57	15.43	15.57	15.37	15.50	15.44
Sharpe Ratio	0.33	0.37	0.40	0.39	0.35	0.41	0.37	0.39	0.39
Max. Drawdown (%)	-57.37	-58.22	-54.59	-55.61	-55.02	-56.42	-54.80	-56.14	-55.97
Rel. Return (%p.a.)		0.67	1.02	0.92	0.23	1.23	0.63	0.91	0.84
Tracking Error (%p.a.)		2.22	1.75	3.17	2.38	2.02	1.47	1.40	1.06
Information Ratio		0.30	0.59	0.29	0.10	0.61	0.43	0.65	0.80
Beta		1.01	0.99	0.98	0.98	0.99	0.99	0.99	0.99
2-way T/O (%p.a.)	12.98	109.41	82.06	211.36	75.80	52.76	67.03	112.53	80.82
Effective N	350.29	147.56	144.01	110.14	150.82	1030.86	202.74	214.39	326.70
Capacity (%)	100.00	18.55	21.85	13.66	15.34	29.25	28.92	34.77	45.13
Active Share (%)	0.00	45.41	40.28	50.46	37.86	44.54	30.69	34.39	25.85

Source: FTSE Russell. Data based on the FTSE Developed Index Universe from September 2000 to September 2019.

The two-way turnover of the composite bulk solutions is slightly less than the average of the individual pure component portfolios from which they are constructed. This results from the ability to “cross” turnover arising in the different single factor component sleeves. We also observe an increase in the information ratio as the number of factors increases, highlighting the merits of factor diversification.

It is interesting to compare the appropriately constructed Top-down, multi-factor portfolios created here with ones created using the Bottom-up approach described in Section 3.1. For example, the Bottom-up five factor portfolio with identical factor exposures and with country, industry and Beta neutrality has a two-way turnover of 71.15%, Effective N of 388.65, investment capacity of 66.00% and an active share of 22.92%. Such increased efficiency is typical of Bottom-up constructions when one compares them to Top-down ones on a like-for-like basis [5, 6].

Finally, we remark that multi-factor portfolios created on demand using this approach may be useful from the perspective of those who wish to engage in factor timing. It is essential that such portfolios are constructed using either pure single factor portfolios (or an appropriate Bottom-up portfolio) to ensure that factor portfolios consist only of the exposures of interest in order to limit implementation risk. We discussed earlier a single factor example of implementation risk at the end of Section 3.5.

Attributional transparency

In a Top-down multi-factor context, the performance contribution of an individual factor is often associated with the performance of a particular factor portfolio sleeve. However, this is correct only if the factor portfolio sleeves are (relatively) pure. Clearly, performance attribution to a sleeve is always possible in the Top-down approach but this is only consistent with performance attribution to a factor if each sleeve has no other factor exposures apart from the one it is intended to represent.

The attributional simplicity afforded by constructing Top-down, multi-factor portfolios from pure single factor portfolios is illustrated in Table 5. For a two factor 50% Quality and 50% Low Volatility portfolio, the performance contribution of

Quality is 50% of the excess return of the Pure Quality portfolio ($0.50 \times 1.02 = 0.51\%$ p.a.). Similarly, the contribution from Low Volatility exposure is 50% of the excess return of Pure Low Volatility portfolio ($0.50 \times 0.23 = 0.115\%$ p.a.). There is no contribution from any other factors as the multi-factor construct arises from pure Quality and pure Low Volatility portfolio sleeves. The total of the Quality and Low Volatility contributions is therefore 0.51% p.a. + 0.115% p.a. = 0.625% p.a., which is close to the actual excess return of 0.63% p.a. The small differences result from the treatment of corporate events (such as deletions) for the separate sleeves compared to the multi-factor composite.

Table 5: Factor Attribution: Pure and composite factor portfolios

	Contribution to Excess Return (% p.a.)								
	FTSE Dev	Pure Single Factor					Top-down Multi-factor		
		Value	Quality	Mom.	Low Vol.	Size	2 Factor	3 Factor	5 Factor
Value	-	0.67	-	-	-	-	-	0.22	0.13
Quality	-	-	1.02	-	-	-	0.51	0.34	0.20
Mom.	-	-	-	0.92	-	-	-	0.31	0.18
Low Vol.	-	-	-	-	0.23	-	0.12	-	0.05
Size	-	-	-	-	-	1.23	-	-	0.25
Total							0.63	0.87	0.81
Actual	0.00	0.67	1.02	0.92	0.23	1.23	0.63	0.91	0.84

Source: FTSE Russell. Data based on the FTSE Developed Index Universe from September 2000 to September 2019.

We can contrast this with attempts to attribute the performance of Quality to a sleeve constructed using a Selection and Weighting approach with its multiple off-target outcomes. We have seen in Section 3.5 that such a sleeve has multiple factor exposures and consequently a formal performance attribution as in Figure 13 reveals contributions from off-target exposures that dwarf the on-target contribution from Quality meaning any claim to attributional transparency is forgone.

Factor overlays

Additional factor objectives may be overlaid on an existing portfolio using an appropriate set of pure single factor portfolios. This includes factor completion portfolios designed to ensure that an aggregate portfolio or overall strategic factor objective is achieved. For example, consider the FMW portfolio described in Section 3.3 and the additional objectives of correcting the negative active exposures to Quality and Momentum. We can achieve this objective, while holding all other exposures unchanged, without modifying the weights of the FMW portfolio by overlaying pure single factor portfolios in the appropriate proportions.

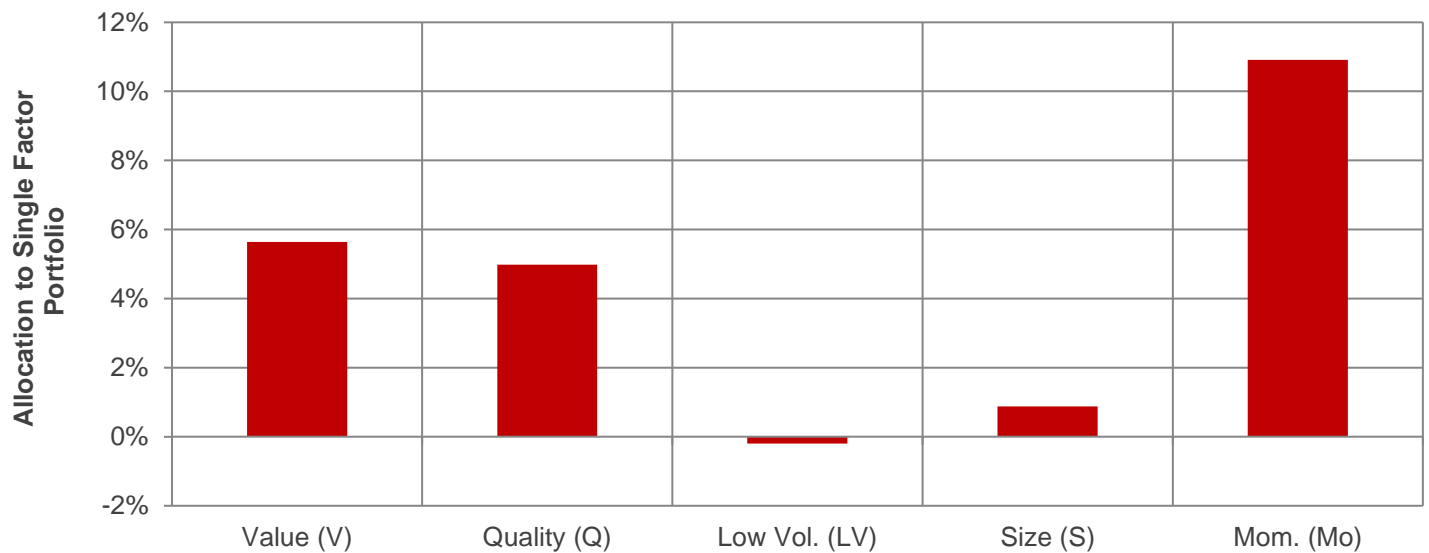
The proportions to be allocated to each pure factor portfolio and held alongside the original FMW portfolio result in the following aggregate portfolio:

$$W = Q * W_Q + Mo * W_{Mo} + V * W_V + S * W_S + LV * W_{LV} + (1 - Q - M - V - S - LV) * W_{FMW} \quad (8)$$

where Q , Mo , V , S and LV are the proportions allocated to each of the pure factor portfolios, W_Q , W_{Mo} etc represent the pure factor portfolios and W_{FMW} is the FMW portfolio. Substituting this expression in place of that for the tilt portfolio in equation (7) gives the system of equations from which these proportions can be found. Figure 15 shows the time-averaged proportions of the pure single factor portfolios required.

The key advantage of this approach is one of convenience as it may be used as a solution to any number of different portfolios with different factor exposure objectives, using only a fixed set of five pure factor portfolios. The downside is that such a solution is not as efficient an implementation as holding the single modified Bottom-up portfolio with factor corrections described in Section 3.3. The positive position in the Pure Value portfolio is to compensate for the reduced contribution of the Value exposure in the FMW portfolio caused by the re-allocation to the pure factor portfolios. Furthermore, the small negative position in Low Volatility (Figure 15) demonstrates the possibility that short positions in the single factor portfolios may be required. This contrasts the Bottom-up method of Section 3.3, which is an entirely long-only construction.

Figure 15: FMW corrective overlays: Time-averaged allocation to pure single factor portfolios



Source: FTSE Russell. Data based on the FTSE Developed Index Universe at rebalance from September 2001 to September 2019.

Finally, we remark that where a portfolio overlay requires many or large factor exposure corrections, the Bottom-up approach described in Section 3.3 is likely to be successful, whereas the Top-down overlay approach may not be feasible.

4 Conclusions

This paper introduces a refinement to the FTSE Russell Tilt framework. The original tilt approach uses fixed tilt strengths and results in varying levels of factor exposure. A simple inversion of this, where the tilt strength is varied, allows us to obtain fixed levels factor exposure. More generally, varying the tilt strengths to obtain fixed exposures leads to the notion that we may target multiple exposures simultaneously within the same framework.

The benefits of such precise control are clear if factor exposure is a primary determinant of the risk and return characteristics of diversified portfolios. The approach can be generalized to control and target other types of exposure such as sustainable investment objectives, along with country, industry and market exposures. The ability to exercise explicit control over the allocation to factors and other characteristics ensures that portfolio construction is more closely aligned with investment objectives and outcomes and reduces implementation risks.

We have presented a number of different applications and investment solutions that are possible using this portfolio construction technique including:

- Multi-factor exposure targets with Beta, country and industry neutrality
- Portfolios embedding explicit carbon intensity improvements
- Factor corrections to remove unwanted factor exposures
- Alternative factor allocation regimes
- Pure single factor indexes
- Absolute return applications of pure single factors
- Composites of pure single factor portfolios for bulk solutions, attributional transparency and factor overlays

These represent a handful of the many possible uses of the target exposure methodology.

5 Appendix

This Appendix contains the definitions of the quantities used in this document.

5.1 Z-scores

We define a set of factor characteristics f_i for each stock labeled by $i = 1, \dots, N$ where N is the number of stocks in our universe. Let ω_i be any set of positive weights which sum to one, then the weighted Z-score is defined by:

$$Z_{\omega_i} = (f_i - \mu_{\omega}) / \sigma_{\omega} \quad (9)$$

where the weighted mean and variance are:

$$\mu_{\omega} = \sum_{i=1}^N \omega_i * f_i \quad \text{and} \quad \sigma_{\omega}^2 = \sum_{i=1}^N \omega_i * (f_i - \mu_{\omega})^2 \quad (10)$$

The set of weights ω_i may take any form, but in this paper we employ either a set of equal weights defined by $E_i = 1/N$ or market capitalization weights:

$$W_{M,i} = \frac{M_i}{\sum_{i=1}^N M_i} \quad (11)$$

where M_i is the free-float adjusted market capitalization of the i^{th} stock. When equal weights are employed to calculate the weighted mean, standard deviation and Z-Score, they reduce to the usual un-weighted expressions, so we write $\mu_E = \mu$, $\sigma_E = \sigma$ and $Z_{E_i} = Z_i$.

5.2 Exposure

To assess how much of a characteristic is embedded in a portfolio at a given point in time, we define exposure as:

$$\text{Exposure} = W.Z = \sum_{i=1}^N W_i * Z_i \quad (12)$$

where W is a set of portfolio weights and Z is as set of Z-scores. The active exposure, relative to set of market capitalization weights W_M is defined by:

$$\text{Active exposure} = W.Z - W_M.Z \quad (13)$$

5.3 Diversification

To assess the degree of diversification in portfolio, we define Effective N of a portfolio as the inverse of the Herfindahl measure of concentration [17]:

$$\text{Effective N} = 1/(W \cdot W) = 1 / \sum_{i=1}^N W_i^2 \quad (14)$$

Effective N attains its maximum under an equal weighting scheme when it is equal to the actual number of stocks. Hence, Effective N can be seen as a measure of “how far” a given portfolio is from this maximally diversified portfolio.

5.4 Active share

The active share is defined as half the sum of the absolute weight differences of two portfolios:

$$\text{Active Share} = \frac{1}{2} \sum_{i=1}^N |W_i - \widehat{W}_i| \quad (15)$$

where W and \widehat{W} are two sets of portfolio weights.

5.5 Capacity

Portfolio capacity is defined as the reciprocal of the weighted sum of stock capacity ratios:

$$\text{Capacity} = 1 / \left[\sum_{i=1}^N W_i * \frac{W_i}{W_{M,i}} \right] \quad (16)$$

where $W_{M,i}$ are the market capitalization weights. This yields a number between 0% and 100% and reflects the ease of investment relative to a market capitalization weighting (100%) scheme.

6 References

- [1] FTSE Global Factor Index Series Ground Rules, FTSE Russell, May 2018.
- [2] “Factor Exposure Indexes - Index Construction Methodology,” FTSE Russell, Aug 2014.
- [3] “Factor Exposures of Smart Beta Indexes,” FTSE Russell, MMM 2015.
- [4] “Factor Exposure and Portfolio Concentration,” FTSE Russell, Apr 2017.
- [5] “Factor Indexes and Factor Exposure Matching: Like-for-Like Comparisons,” FTSE Russell, October 2018.
- [6] “Alternative Approaches to Multi-Factor Index Construction: Like-for-Like Comparisons,” FTSE Russell, Nov 2018.
- [7] FTSE Global Financial Metrics Weighted Index Series Ground Rules, FTSE Russell, May 2018.
- [8] FTSE Global Climate Index Series Ground Rules, FTSE Russell, Jul 2019.
- [9] Fama, E. F. and French, K. R., “Common Risk Factors in the Returns of Stocks and Bonds,” *Journal of Financial Economics*, Vol. 33, p3-56, Feb. 1993.
- [10] Fama, E. F. and MacBeth, J. D., “Risk, Return and Equilibrium Empirical Tests,” *Journal of Political Economy*, Vol. 81, No. 3, p607-636, 1973.
- [11] Amenc, N., Goltz, F., Lodh, A. and L. Martellini, L., “Towards Smart Equity Factor Indices: Harvesting Risk Premia without Taking Unrewarded Risks,” *Journal of Portfolio Management*, Vol. 40, No. 4, p106-p122, 2014.
- [12] Amenc, N., Ducoulombier, F., Goltz, F., Lodh, A. and Sivasubramanian, S., “Diversified or Concentrated Factor Tilts,” *The Journal of Portfolio Management*, Vol. 42, No. 2, p64-p76, 2016.
- [13] Amenc, N., Ducoulombier, F., Esakia, M., Goltz, F. and Sivasubramanian, S., “Accounting for Cross-Factor Interactions in Multifactor Portfolios without Sacrificing Diversification and Risk Control,” *Journal of Portfolio Management*, Special Issue, Vol. 43, No. 5, p99-p114, 2017.
- [14] Amenc, N., Goltz, F. and Sivasubramanian, S., “Multifactor Index Construction: A Skeptical Appraisal of Bottom-Up Approaches,” *The Journal of Index Investing*, 9 (1) p6-p17, Summer 2018.
- [15] Chow, T., Li, F. and Shim, Y., “Smart Beta Multifactor Construction Methodology: Mixing versus Integrating,” *The Journal of Index Investing*, 8 (4) p47-p60, Spring 2018.
- [16] “Goldman Sachs ActiveBeta® Equity Indexes Methodology,” Goldman Sachs Asset Management, November 2016.
- [17] Woerheide, W. and Persson, D. “An Index of Portfolio Diversification”, *Financial Services Review*, Vol. 2, p73-p85, 1993.
- [18] Fitzgibbons, S., Friedman, J., Pomorski, L. and Serban L., “Long-Only Style Investing: Don’t Just Mix, Integrate,” *The Journal of Investing*, Vol. 26, No. 4, p153-p164, Winter 2017.
- [19] Bender, J. and Wang, T., “Can the Whole Be More Than the Sum of the Parts? Bottom-Up versus Top-Down Multifactor Portfolio Construction,” *The Journal of Portfolio Management Special QES Issue*, Vol. 42, No. 5, p39-p50, 2016.
- [20] Ghayur, K., Heaney, R. and Platt, S., “Constructing Long-Only Multifactor Strategies: Portfolio Blending vs. Signal Blending,” *Financial Analysts Journal*, Vol. 74, No. 3, p70-p85, 2018.
- [21] Doole, S., Chia, C., Kulkarni, P. and Melas, D. “The MSCI Diversified Multi-Factor Indexes,” MSCI White Paper, May 2015.

About FTSE Russell

FTSE Russell is a leading global provider of benchmarks, analytics and data solutions with multi-asset capabilities, offering a precise view of the markets relevant to any investment process. For over 30 years, leading asset owners, asset managers, ETF providers and investment banks have chosen FTSE Russell indexes to benchmark their investment performance and create investment funds, ETFs, structured products and index-based derivatives. FTSE Russell indexes also provide clients with tools for performance benchmarking, asset allocation, investment strategy analysis and risk management.

To learn more, visit ftserussell.com; email info@ftserussell.com; or call your regional Client Service Team office

EMEA

+44 (0) 20 7866 1810

North America

+1 877 503 6437

Asia-Pacific

Hong Kong +852 2164 3333

Tokyo +81 3 4563 6346

Sydney +61 (0) 2 8823 3521

© 2020 London Stock Exchange Group plc and its applicable group undertakings (the "LSE Group"). The LSE Group includes (1) FTSE International Limited ("FTSE"), (2) Frank Russell Company ("Russell"), (3) FTSE Global Debt Capital Markets Inc. and FTSE Global Debt Capital Markets Limited (together, "FTSE Canada"), (4) MTSNext Limited ("MTSNext"), (5) Mergent, Inc. ("Mergent"), (6) FTSE Fixed Income LLC ("FTSE FI"), (7) The Yield Book Inc ("YB") and (8) Beyond Ratings S.A.S. ("BR"). All rights reserved.

FTSE Russell® is a trading name of FTSE, Russell, FTSE Canada, MTSNext, Mergent, FTSE FI, YB and BR. "FTSE®", "Russell®", "FTSE Russell®", "MTS®", "FTSE4Good®", "ICB®", "Mergent®", "The Yield Book®", "Beyond Ratings®" and all other trademarks and service marks used herein (whether registered or unregistered) are trademarks and/or service marks owned or licensed by the applicable member of the LSE Group or their respective licensors and are owned, or used under licence, by FTSE, Russell, MTSNext, FTSE Canada, Mergent, FTSE FI, YB or BR. FTSE International Limited is authorised and regulated by the Financial Conduct Authority as a benchmark administrator.

All information is provided for information purposes only. All information and data contained in this publication is obtained by the LSE Group, from sources believed by it to be accurate and reliable. Because of the possibility of human and mechanical error as well as other factors, however, such information and data is provided "as is" without warranty of any kind. No member of the LSE Group nor their respective directors, officers, employees, partners or licensors make any claim, prediction, warranty or representation whatsoever, expressly or impliedly, either as to the accuracy, timeliness, completeness, merchantability of any information or of results to be obtained from the use of the FTSE Russell products, including but not limited to indexes, data and analytics or the fitness or suitability of the FTSE Russell products for any particular purpose to which they might be put. Any representation of historical data accessible through FTSE Russell products is provided for information purposes only and is not a reliable indicator of future performance.

No responsibility or liability can be accepted by any member of the LSE Group nor their respective directors, officers, employees, partners or licensors for (a) any loss or damage in whole or in part caused by, resulting from, or relating to any error (negligent or otherwise) or other circumstance involved in procuring, collecting, compiling, interpreting, analysing, editing, transcribing, transmitting, communicating or delivering any such information or data or from use of this document or links to this document or (b) any direct, indirect, special, consequential or incidental damages whatsoever, even if any member of the LSE Group is advised in advance of the possibility of such damages, resulting from the use of, or inability to use, such information.

No member of the LSE Group nor their respective directors, officers, employees, partners or licensors provide investment advice and nothing contained herein or accessible through FTSE Russell products, including statistical data and industry reports, should be taken as constituting financial or investment advice or a financial promotion.

Past performance is no guarantee of future results. Charts and graphs are provided for illustrative purposes only. Index returns shown may not represent the results of the actual trading of investable assets. Certain returns shown may reflect back-tested performance. All performance presented prior to the index inception date is back-tested performance. Back-tested performance is not actual performance, but is hypothetical. The back-test calculations are based on the same methodology that was in effect when the index was officially launched. However, back-tested data may reflect the application of the index methodology with the benefit of hindsight, and the historic calculations of an index may change from month to month based on revisions to the underlying economic data used in the calculation of the index.

This document may contain forward-looking assessments. These are based upon a number of assumptions concerning future conditions that ultimately may prove to be inaccurate. Such forward-looking assessments are subject to risks and uncertainties and may be affected by various factors that may cause actual results to differ materially. No member of the LSE Group nor their licensors assume any duty to and do not undertake to update forward-looking assessments.

No part of this information may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the applicable member of the LSE Group. Use and distribution of the LSE Group data requires a licence from FTSE, Russell, FTSE Canada, MTSNext, Mergent, FTSE FI, YB, BR and/or their respective licensors.