For professional investors - Marketing communication - November 2021

PERFORMANCE ATTRIBUTION FOR MULTI-FACTORIAL EQUITY PORTFOLIOS

QUANTITATIVE RESEARCH GROUP
BNP PARIBAS ASSET MANAGEMENT

BNP PARIBAS ASSET MANAGEMENT

The sustainable investor for a changing world

Frédéric Abergel
Senior quantitative analyst - quantitative research group
14 rue Bergère, 75009 Paris, France.
frederic.abergel@bnpparibas.com

Thomas Heckel
co-head of quantitative research group
14 rue Bergère, 75009 Paris, France.
thomas.heckel@bnpparibas.com

ABSTRACT

This paper revisits the cross-sectional approach to the performance analysis of multi-factor investment strategies. Its main contributions are three-fold: First, the use of a cross-sectional projection of asset returns onto the factor portfolio weights to form approximate portfolio returns; second, the introduction of nonlinear, interaction terms between factors that faithfully reproduce the investment portfolio construction; third, a natural and intuitive decomposition of the portfolio performance as the sum of factor contributions. The method we propose has several advantages over other time series-based or general cross-sectional regression models: It faithfully reflects the current state of the investment portfolio; it is parsimonious in the number of explanatory variables; it leads to an approximation of the portfolio returns that has a small residual error; and it provides a straightforward interpretation of the portfolio performances in terms of the factors it is designed from. The method we advocate is first presented and explained in detail. Concrete applications to multi-factor equity strategies are then presented.

KEYWORDS:

Performance attributions
Factor investing
Portfolio construction
Asset management

KEY MESSAGES:

- **Time series-based attribution does not reflect the structure of a multi-factor portfolio**
- **Cross-sectional regression based on a parsimonious set of regressors with interaction terms works well**
- **The role of portfolio construction is key**

1. INTRODUCTION

This article addresses the performance analysis - or, as it is often termed, performance attribution - of multi-factor investment strategies. Multi-factor portfolios are built from style factors and the task at hand is to explain their performance in terms of each factor.

Let us first briefly recall what we mean by "factors".

Factors are characteristics of assets that are important in explaining their risk and performance. For instance, in the stock market, the Capital Asset Pricing Model (CAPM) asserts that the performance of a stock should be determined by a single stock characteristic, the **beta**. However, it is now widely accepted, and vastly documented in the academic literature, see e.g., the seminal papers [6] and [10], that other factors typically classified into styles – such as value, low volatility, quality and momentum – also play a role in explaining stock returns.

Value characteristics such as the price-to-book or the price-to-earnings of a company measure the relative cheapness of a stock. It has been found that on average, over time, cheaper stocks tend to outperform other stocks, in particular expensive stocks. **Volatility** is a measure of a stock's price fluctuations over time. Evidence that less-volatile stocks generate at least comparable returns to riskier stocks renders stocks with lower volatility more attractive for investors: potentially the same returns in the medium to long term but with less uncertainty. It is also known that higher quality stocks, e.g., the most profitable companies, tend to generate higher returns than other stocks, in particular when compared to the least profitable companies. Different measures of profitability can be used, e.g., return-on-equity. Last, stocks with the strongest price trends, i.e., stocks with the strongest outperformance relative to other stocks as measured over the previous 12 months, also tend to continue to outperform. This is known as the **momentum** effect.

The building blocks of factor investing are the **single factor portfolios**. There are many routes one can follow to build these factor portfolios, most of them passing through successive steps of ranking, aggregation, normalization and neutralization. These single factor portfolios are long-short portfolios with weights that reflects the score of an asset according to one particular member of the factor family. From there on, **multi-factor investing** involves building, based on these single factor portfolios, a multi-factor, investable portfolio. This rather involved industrial process can be decomposed into two major steps. The first generates a global, aggregate score for each stock, leading to a theoretical long-short portfolio that is the target of the investment strategy: A portfolio that one would like to invest in, should there be no other goal than generating higher returns than a given benchmark. The second, more technical, step amounts to building an investment portfolio that can fulfill a set of requirements such as weight positivity, diversification, low turnover or controlled volatility.

Let us now turn towards the main topic of this paper, namely, performance attribution.

Performance attribution is the final stage of the investment process. It is essentially an ex post process that aims at providing a quantitative performance analysis of the investment portfolio based on single factors. This is an important task, in that it allows one to understand, and not only benefit - or suffer - from the performance of, an investment strategy. There are potentially many routes to achieving that goal, and the interested reader is referred to [1] for a survey of the field. A highly generic, and somewhat blunt, approach due to Brinson has been formalized in [3]. However, it does not answer the specific questions one may ask in the context of factor investing. For that, there are better adapted, more relevant routes one can follow in order to perform this attribution, the general idea being that one wants to relate the performance of an investment strategy to those of the underlying style factors. One way is to explain via regression the returns of the investment portfolio by the returns of factor portfolios, see e.g. [9] [4]. One thereby relies on a **time series** approach for the attribution. Simple and straightforward as it seems, this method has the dire disadvantage of using information from a distant past to assess the current performance of a portfolio. Therefore, it does not cope well with financial market behavior never being stationary. The path taken in this paper follows the alternative **cross-sectional** route in the line of generalized factor-characteristic models such as described and analyzed e.g., in [7] or [5]. In this approach, one explains at each date the cross-section of stock returns in the investment universe by their characteristics. This cross-sectional approach offers the major advantage of relating the performance of the portfolio to its current tilts towards the single factor portfolios. It is therefore more reactive, more flexible and provides a more direct explanation of the portfolio returns in terms of the current state of the market. This approach in itself is not new, one can already find a general exposition of a cross-sectional performance attribution in e.g. [11]. However, and although the starting points are similar, it will soon become clear that ours is a specific, better adapted and more efficient method in the context of multi-factor investment strategies.

The paper is organized as follows: Section 2 compares portfolio tilts and betas, showing that a time-series based performance analysis does not provide faithful nor intuitive results. Section 3 introduces the mathematical notations used in the paper. Section 4 presents our specific version of the cross-sectional methodology in detail. Section 5 provides some concrete, numerical applications of the cross-sectional method to multi-factor equity strategies. Section 6.3 is a short discussion on an alternative cross-sectional methodology based on portfolio decomposition.

2. PORTFOLIO TILTS OR BETAS?

This introductory section highlights our point of view by giving some concrete examples of the differences between the tilts of an investment strategy towards the single factor portfolios, and the betas of the same strategy's returns.

A classical route to factor-based performance attribution [9] involves using the investment portfolio's betas with respect to the factor portfolios - this approach will be referred to as the time series-regression approach. The portfolio's betas can be computed using sliding windows, full in-sample data, or the risk model. In all cases, the computation requires the use of a long data set to be computed and, as already said in the introduction, they do not faithfully reflect the current state of the investment portfolio. On the other hand, suitably normalized portfolio tilts actually measure the co-linearity between the investment portfolio and the factor portfolios. As such, we think they provide an intuitive and accurate characterization of the portfolio and should therefore appear naturally during the performance attribution process.

We now provide some concrete examples to support our claim. Figure 1 shows the 24-month trailing betas, full in-sample betas and portfolio tilts for a multi-factor investment strategy with respect to the four style factors: Low volatility, momentum, quality and value. The strategy is built according to a Markowitz-type optimization program that will be analyzed in detail in Section 3, see in particular Equation (2) and Equation (4); we do not dwell upon this construction here.

Figure 1: Comparing portfolio tilts and regression betas. The graphs show the 24-months trailing betas and "portfolio" tilts of the investment portfolio towards the single factor long-short "portfolios".

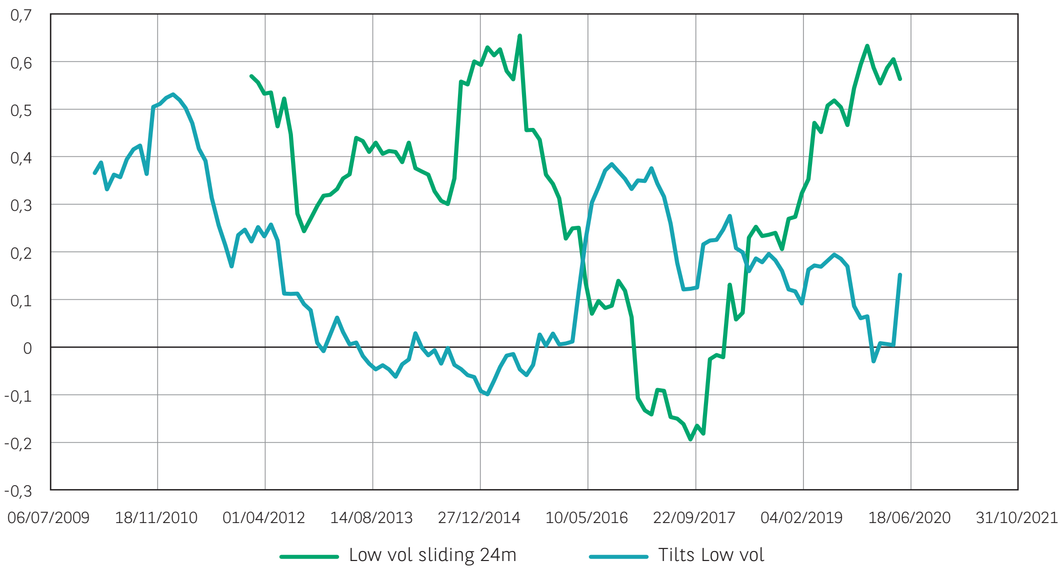

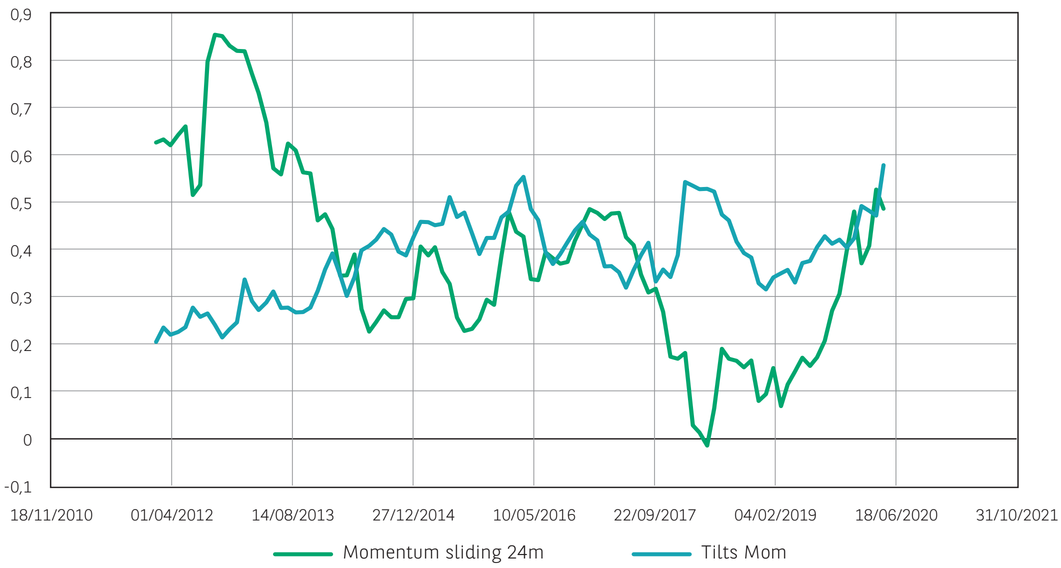

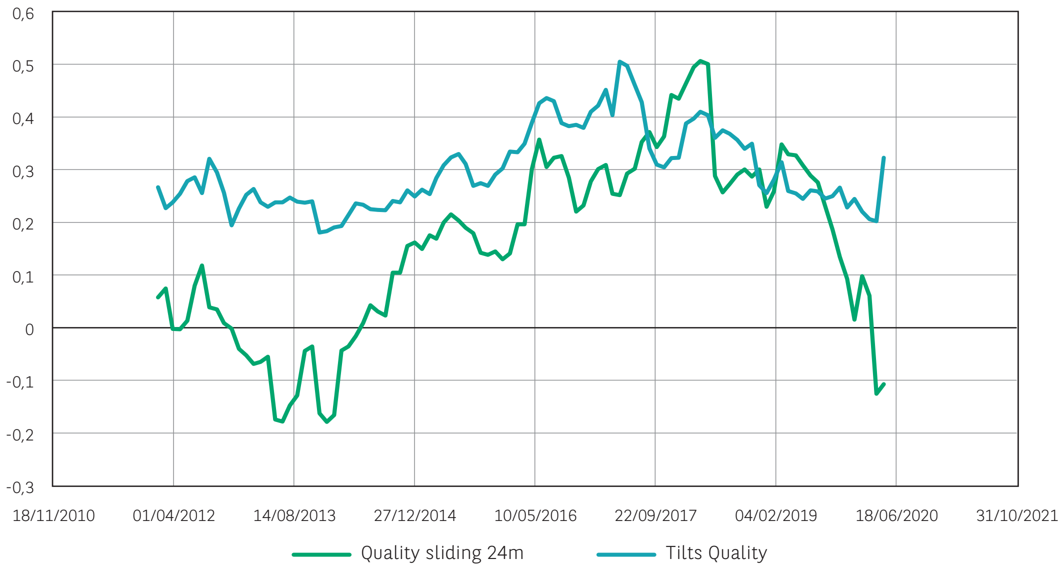

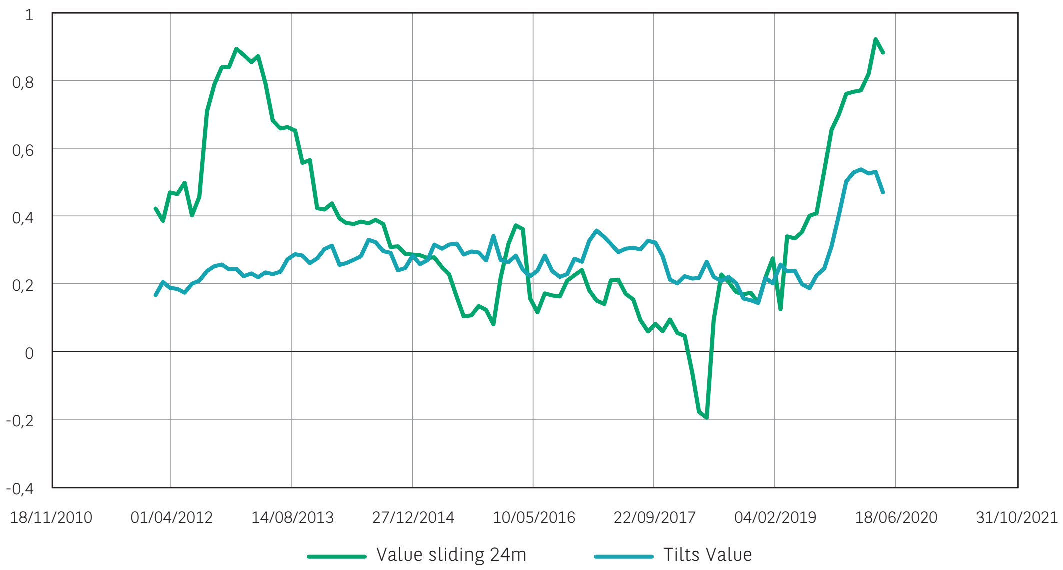

Due to memory effects, the betas do not represent well the portfolio composition: For instance, the low volatility, momentum and quality trailing betas become negative around the time when the corresponding tilts are highest, a fact that is clearly counter-intuitive. Instead of using the trailing betas, one can use the risk model to determine ex ante betas. The corresponding results are shown in Figure 2.

Figure 2: Comparing portfolio tilts and ex ante betas. The graphs show the betas computed using the risk model and the portfolio tilts of the investment portfolio towards the single factor long-short portfolios.

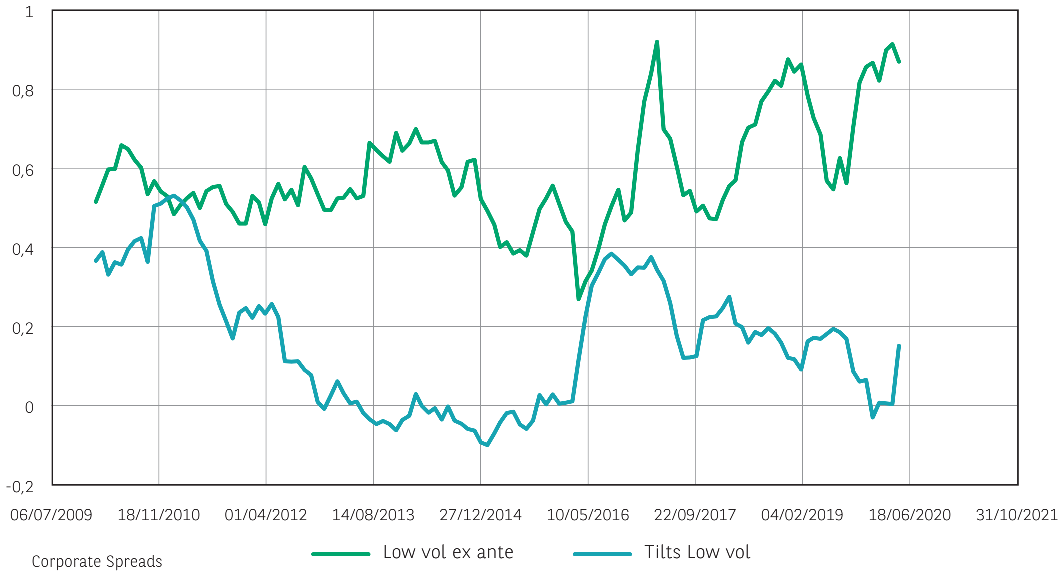

Corporate Spreads

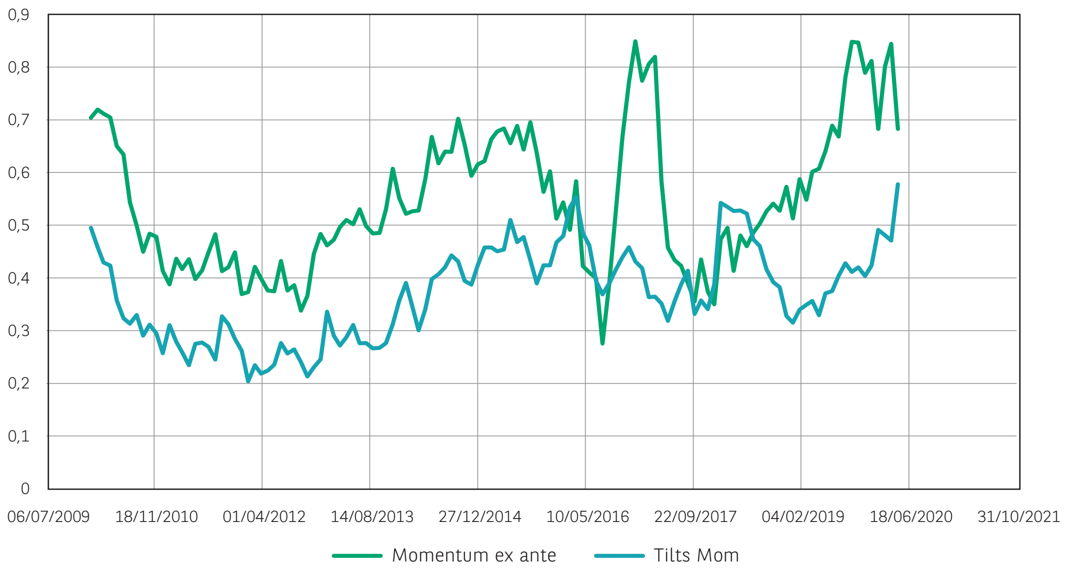

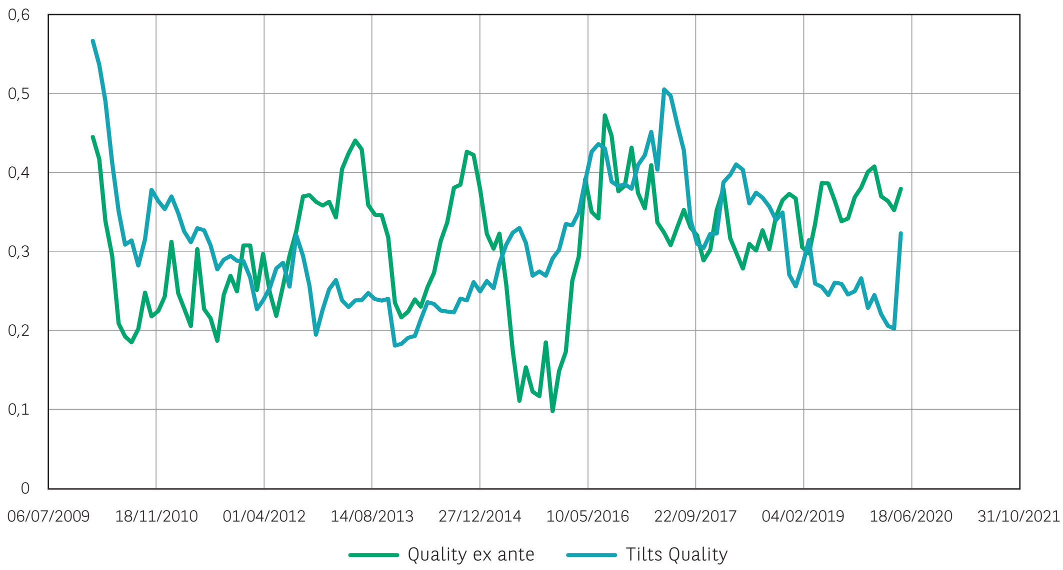

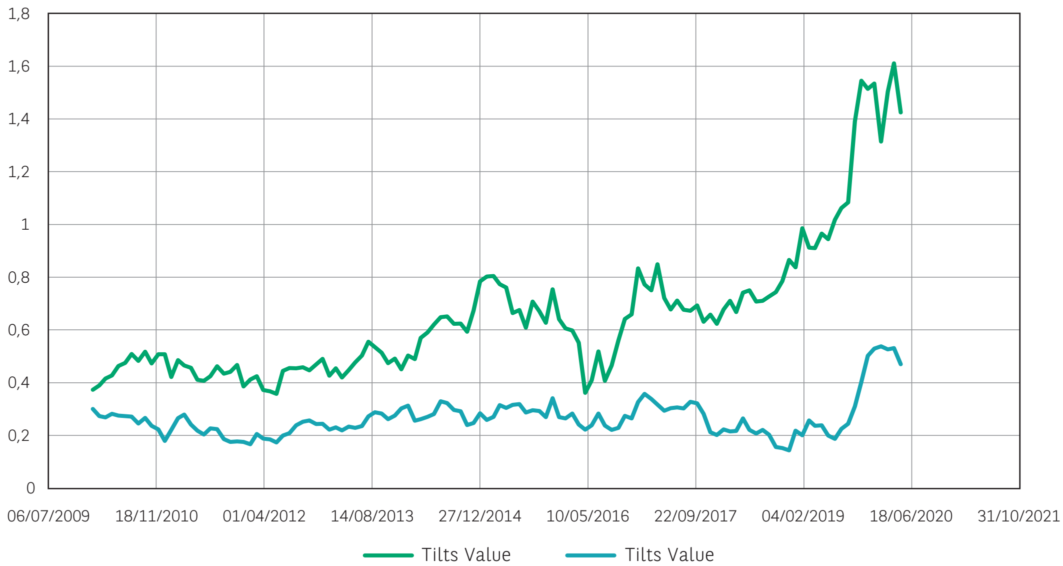

Although more stable over time, the ex ante betas suffer from similar drawbacks as the trailing betas. Similar to the tilts, they depend on the current portfolio composition, but use a metric that mixes the portfolio weight with the history of the stocks in a way that is not related to, and not controlled by, the factor portfolios.

To summarize this short introductory section, it is our opinion that a good performance attribution method should rely on the portfolio tilts rather than its betas. One major question then remains: How can one relate the portfolio tilts to the factor and investment portfolio performances? The cross-sectional approach provides a natural and accurate answer to that question.

3. SOME NOTATIONS

We now introduce some mathematical notations that will be used in the rest of this article. **Bold face** letters stand for vectors or matrices.

- N the total number of assets in the universe
- **R** the vector of stock returns in excess of cash
- $\boldsymbol{\beta}$ the vector of stock return betas with respect to market returns
- **I** the vector of benchmark weights
- RM the benchmark (market) excess return over cash RM = **I.R**
- **e** the normalized unit vector of equal weights $\boldsymbol{e}_i = \frac{1}{N}$
- **W** the investment portfolio weights
- β_W the investment portfolio beta β_w = $\mathbf{W}.\boldsymbol{\beta}$
- **Q** the investment portfolio active weights **Q = W – I**
- $\mathbf{z}_k$, k = 1, ..., 4 the four long-short factor portfolio weights low volatility, momentum, quality, value
- **y** the High Size portfolio weights **y = I – e**
- **A** the theoretical, multi-score long-short portfolio weights, a linear combination of the $\mathbf{z}_k$ with some prescribed weights: $A = \sum_{i=1}^{4} a_k z_k$
- **M** the covariance matrix of the stock return (CAPM) alphas (**R** – $\boldsymbol{\beta}$ RM)
- For any vector **X**, $\mathbf{X}^+$ stands for the positive part of **X** with components $X_i^+ \equiv X_i 1_{\{X_i>0\}}$
- For any vector $\boldsymbol{X}, \widetilde{\boldsymbol{X}}$ stands for the market-neutral version of **X**

4. THE CROSS-SECTIONAL APPROACH TO PERFORMANCE ATTRIBUTION

4.1 The basic approach

Let us now describe the cross-sectional approach to performance attribution: At any given time, each asset in the investment universe possesses a set of K = 4 characteristic scores corresponding to the K factors. Once market-neutralized, these scores transform into the weight vectors zk of the single factor long-short portfolios. A general factor and characteristic-based model can be decomposed into a time series-based component, with factor returns and betas, and a cross-sectional component with factor loadings and implied factor returns, see e.g., [5]. With the specific application to performance attribution in mind, we recast this general approach into a pure cross-sectional framework. This consists of explaining the cross-section of the stock excess returns **R** by the vector $\boldsymbol{\beta}$ of betas against the market returns, and the single factor, long-short portfolio weights. Such a model can be written in its simplest form as

$$\boldsymbol{R} = \widehat{\boldsymbol{R}} + \in = v\boldsymbol{\beta} + \sum_{k=1}^{k} \lambda_k \mathbf{z}_k + \hat{\epsilon} \tag{1}$$

where v and the λ's are determined by OLS regression. Actually, since the factor portfolios are market-neutral: $z_k.\beta = 0$, the λ_k's are precisely determined by the system of linear equations

$$\forall k \in 1, \dots, K, z_k.R = k \sum_{k=1}^{k} \lambda_p z_p.z_k.$$

The λ_p have a natural interpretation as the cross-sectional returns of the single factor portfolios. In fact, should these portfolios be cross-sectionally uncorrelated:

$$\forall p \neq q, z_p . z_q = 0$$

then the OLS matrix would be diagonal and the λ_p would be equal, up to some normalization, to the realized portfolio returns

$$\lambda_k^{Diagonal} = \frac{z_k . R}{z_k . z_k}$$

Although desirable, this cross-sectional orthogonality will generally not hold true, and the cross-sectional returns are not equal to the realized returns. It is however important in factor design that the single factor portfolios be as cross-sectionally independent (as orthogonal) as possible, so as to avoid the artificial oscillations and compensations that arise in the regression when strong co-linearities are present.

Equation (1) is at the core of the cross-sectional approach to performance analysis. From an econometric perspective, such a relation is overly simplistic; additional explanatory variables need to be included to capture the actual returns distribution, not to mention its dynamics. However, simple as it is, this model possesses the fundamental property that the **realized returns** of each single factor portfolio are matched exactly:

$$z_k . R = z_k . \widehat{R}.$$

This property is the cornerstone of cross-sectional performance attribution for factor investing: Any portfolio with weights that are a linear combination of the single factor portfolio weights will receive a perfect explanation, and one hopes that if the investment portfolio stays "close enough" to such a combination, it, too, will be well explained using this regression.

Consider now an investment strategy with weight vector **W** and active weights portfolio **Q** ≡ **(W – I)**, and assume for simplicity that $\beta_W = 1$.

Then, the excess return **W.R** – *RM* over the benchmark of the portfolio, being exactly equal to **Q.R**, is approximated using the cross-sectional regression by $Q \cdot \hat{R} \equiv \sum_{K=1}^{k} \lambda_k (Q.z_k)$, and a natural candidate for the performance attribution follows from this linear decomposition. Each term λ_k $(Q.z_k)$ represents the individual contribution of the k_{th} factor, decomposed as the product of its cross-sectional factor return λ_k multiplied by the kth active portfolio tilt $(Q.z_k)$. This yields a natural explanation for the performance of a portfolio based on its tilts towards each individual factor portfolio. Now, in our view, a sensible portfolio construction is one that tilts the investment portfolio towards the single factor portfolios. Hence, one should expect positive portfolio tilts - we want to be long the stocks with higher scores - so that the contribution of a given factor to the portfolio performance is higher as the tilt gets higher and the cross-sectional factor return gets higher.

Now, were the regression exact, the performance attribution would provide a perfect explanation with no error term. Since it is not, there remains a gap between the realized and estimated portfolio returns: It is that part of the portfolio performance that cannot be explained by the factors. Because of the various constraints impacting the design of an investment portfolio, it is natural to expect an imperfect attribution to the factors but, to the extent that the investment portfolio stays close enough to the single factor portfolios, the residual should remain under control.

To summarize, we will consider the performance analysis satisfactory provided two things occur: First, that the overall approximation error is small, unbiased and has a small variance; second, that each factor contribution is consistent with the realized returns of the corresponding factor portfolio. The latter criterion is mostly linked to factor design and investment portfolio construction. As for the former, error control is an important task and we now proceed to show how it can be improved by enlarging the regression set.

4.2 Enhancing the regression: The Size factor

The first extra variable we consider is the High Size portfolio. Although the multi-factor strategies we analyze in the paper do not explicitly "play" such a factor (nor its opposite), it may come in handy for the performance attribution of a well-diversified portfolio in a very concentrated market such as the US equity market. In fact, there are some hard constraints on the maximum weight one can assign to a given stock, and this may create some exposure to the factor.

Using the neutralized High Size weight vector y as an extra explanatory variable, the regression model becomes

$$R = \hat{R} + \in$$
$$\hat{R} \ = k\beta + \ \sum_{i=1}^{4} \lambda_k z_k + \lambda_s \tilde{y}$$

4.3 Enhancing the regression: The interaction terms

As already mentioned, the closer the investment portfolio remains to some linear combination of the single factor portfolios, the more accurate the attribution. However, an actual investment portfolio has to satisfy a set of constraints – positivity, diversification, beta, etc. – that move it away from a theoretical long-short portfolio. Moreover, not only does the investment portfolio satisfy such constraints, but it is often built from a rather involved Markowitz-type optimization program. This is indeed the case for the investment portfolio we analyze in this article.

This section highlights two additional explanatory variables that enhance the cross-sectional model in the context of performance attribution. By carefully scrutinizing the portfolio construction, we propose using these additional variables to better "explain" the cross-section of asset returns. These new regressors are portfolio-specific, but the method is general and can be adapted to many different portfolio constructions, as we shall see in Section 6.

Leaving aside for the sake of clarity the turnover and transaction cost constraints, we focus on the rather general case where the optimal portfolio **W** and its active weight vector **Q** can be written as the solution to a Markowitz portfolio optimization with the alpha variance-covariance matrix **M** and a set of expected returns **μ**:

$$\begin{aligned}
&sup\ \boldsymbol{q}.\boldsymbol{\mu} \\
&q = \boldsymbol{w}\ - \boldsymbol{I} \\
&0\ \leq w_i \\
&q_i \leq\ q_{max} \\
&\boldsymbol{q}.\boldsymbol{e}\ = 0 \\
&\boldsymbol{Mq}.\boldsymbol{q}\ \leq TE^2.
\end{aligned} \qquad (2)$$

In Problem (2), the objective function is naturally expressed in terms of the **active weight q**, as are most of the constraints. In fact, only the non-negativity constraint involves the original weight vector **w**.

1 by "arbitrary", we mean one that has no relation whatsoever to the factor portfolios

The performance attribution of a portfolio solution to Equation (2) in the case of an arbitrary set of expected return[1] is beyond the scope of our methodology, as will easily be shown in the trivial counter-example in Section 6: In order for the cross-sectional performance attribution to work, the solution portfolio should stay somewhat close to some function of the factor portfolio weights. There are however many possibilities that come to mind when one wishes to introduce a set of expected returns derived from the factor portfolios, and we shall focus on a special case – and one that is very important in practice – where the expected returns $\boldsymbol{\mu}$ come from an explicit transformation of the theoretical long-short portfolio.

Reverse optimization and portfolio construction. The portfolio construction generating the investment strategies presented in this article rests on the useful and intuitive concept of reverse optimization as introduced by [8], see e.g. [2] for an in-depth analysis. When specific to the case of multi-factor investing, the key underlying assumption is that the theoretical long-short portfolio is, up to a multiplicative constant, the active weight vector solution to the Markowitz optimization problem with the sole risk constraint:

$$\begin{aligned} &sup\ \boldsymbol{q}.\boldsymbol{\mu} \\ &\boldsymbol{Mq}.\boldsymbol{q}\ \leq TE^2 \end{aligned} \qquad (3)$$

the solution of which is

$$A = C\boldsymbol{M}^{-1}\boldsymbol{\mu}$$

for some TE-related constant C.

This process is termed the reverse optimization because it builds a set of expected returns from a candidate portfolio rather than, more usually, the other way around.

Since Problem (2) is invariant by multiplicative scaling of the expected returns, one simply imposes the relation

$$\boldsymbol{\mu} = \boldsymbol{MA}$$

and Problem (2) becomes

$$\begin{aligned} &sup\ \boldsymbol{q}.\boldsymbol{MA} \\ &\boldsymbol{q} = \boldsymbol{w}\ - \boldsymbol{I} \\ &0\ \leq w_i \\ &q_i \leq\ q_{max} \\ &\boldsymbol{q}.\boldsymbol{e}\ = 0 \\ &\boldsymbol{Mq}.\boldsymbol{q}\ \leq TE^2. \end{aligned} \qquad (4)$$

We now focus on Problem (4) and derive the two non-linear regressors. In Section 4.3, we will touch upon the extension to a different set of factor-based expected returns.

Introducing the non-linear terms. We analyze two limiting cases that highlight the role played by the constraints and introduce two extra regressors that we shall use in the cross-sectional regression model. In a nutshell, these two new regressors are (some explicit functions of) $\mathbf{A}^+$ and $\boldsymbol{\mu}^+$, respectively the positive part of the theoretical long-short portfolio and of the expected returns. These ad hoc variables can be seen as interaction terms that mimic the way individual factors are combined together during the portfolio construction phase.

Large TE bound. When there is no tracking error constraint – or when it is large compared to the maximum weight constraint – Problem(4) approaches the limit problem

$$
\begin{aligned}
&sup\ \boldsymbol{q}.\boldsymbol{MA} \\
&\boldsymbol{q} = \boldsymbol{w}\ - \boldsymbol{I} \\
&0\ \leq w_i \\
&q_i \leq\ q_{max} \\
&\boldsymbol{q}.\boldsymbol{e}\ = 0.
\end{aligned}
\tag{5}
$$

the solution $W_{Large\ TE}$ of which is an explicit function of **μ**+ = (**MA**)+. More precisely, one obtains the solution to (5) by reordering the values in **μ**+ in decreasing order and assigning the largest possible weight $w_k = q_{max} + I_k$ for all k until the constraint **w.e** = 1 is saturated. Although the solution to (5) is clearly not equal to **μ**+, we will continue using **μ**+ as a shorthand notation for this first non-linear regressor.

Large maximum weight constraint. When no pointwise upper bound constraint is enforced – or when it is large compared to the TE bound – Problem (4) resembles

$$
\begin{aligned}
&sup\ \boldsymbol{q}.\boldsymbol{MA} \\
&\boldsymbol{q} = \boldsymbol{w}\ - \boldsymbol{I} \\
&0\ \leq w_i \\
&\boldsymbol{q}.\boldsymbol{e}\ = 0. \\
&\boldsymbol{Mq}.\boldsymbol{q}\ \leq TE^2.
\end{aligned}
\tag{6}
$$

Due to non-locality (in general, M is not diagonal) there is no explicit solution to (6), but the problem is equivalent to the variational inequality:

$$
\begin{aligned}
&(\mathbf{Mq})_i - c_1(\mathbf{MA})_i - c_2 \quad = \quad 0,\ w_i > 0 \\
&(\mathbf{Mq})_i - c_1(\mathbf{MA})_i - c_2 \quad \quad\ \ 0,\ w_i = 0
\end{aligned}
$$

for some Lagrange multipliers c_1, c_2. A rather good approximate solution of (6) can be obtained by assuming that the set of positive portfolio weights coincides with those in A, in which case the corresponding weight vector $W_{No\ Max\ Weight}$ is well approximated by **A**+, the positive part of the theoretical long-short weight vector.

A conclusion from this analysis is that the solution **W** to (4) behaves like a mixture of the single factor portfolios and the two long-only portfolios **μ**+ and **A**+. These two portfolios represent the interaction terms referred to earlier, formed by nonlinear combinations of the single factor portfolios. These interaction terms greatly reduce the unexplained part of the investment portfolio returns and – since they are functions of the single factor portfolios – their contributions to the performance can easily be re-allocated to the original factors.

The graphs in Figure 3 should help one convince oneself of the relevance of these nonlinear terms, instead of: on top of these purely visual impressions, on can compute the time-series correlation between these model portfolio performances and the long-only portfolio return alphas. They are quite high indeed: We observe a 65% realized correlation with the **A**+ portfolio, rising to 78% with the **μ**+ portfolio.

Exhibit 3: Investment portfolio and interaction terms. The graphs show the cumulative alpha of the investment portfolio and of the two additional regressors.

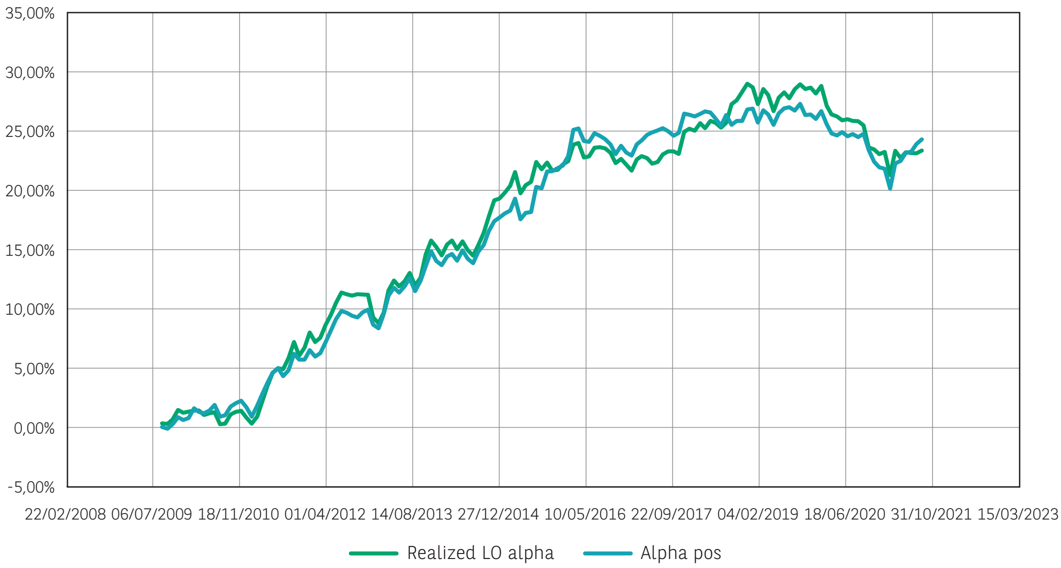

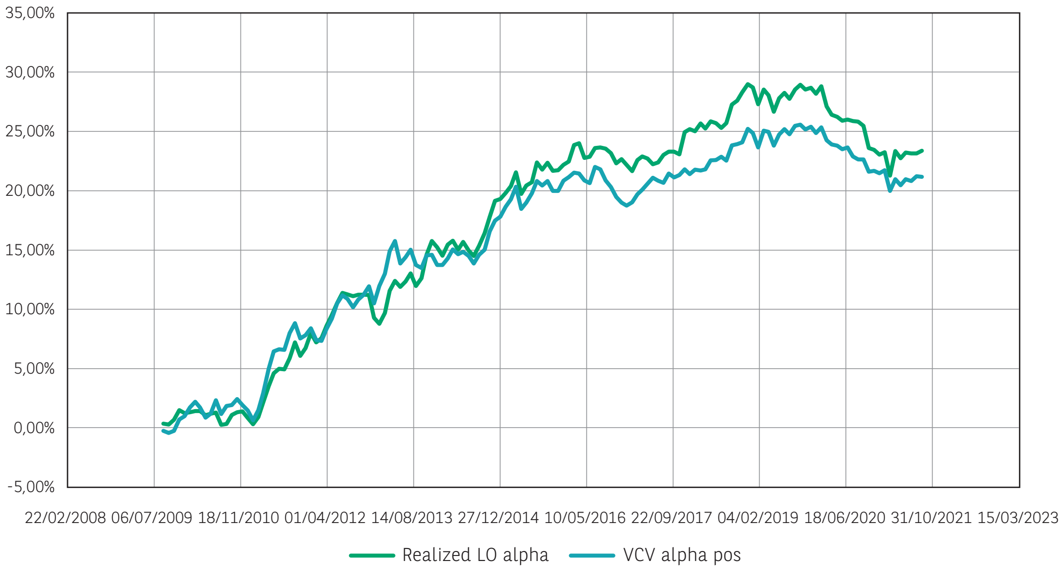

The case of more general expected returns. The approach just introduced to building interaction terms is quite general, and not specific to the particular choice of reverse optimization. What is important, however, is that there should exist some more or less simple but explicit relationship between the factor portfolio weights and the set of expected returns. Bar this, the cross-sectional performance attribution is likely to fail. In Section 6, we will apply the same cross-sectional performance attribution methodology to a multi-factor investment strategy based on a different set of expected returns, showing that the cross-sectional method offers some degree of generality as long as the expected returns remain analytically close to the factor portfolios.

4.4 The full cross-sectional model

As a conclusion to the previous analyses, we now write the full cross-sectional model used in this paper for performance attribution:

$$R = \hat{R} + \epsilon \qquad (7)$$
$$\hat{R} = v\beta + \sum_{i=1}^{4} \lambda_k z_k + \lambda_s \tilde{y} + \gamma \widetilde{\alpha^+} + v\widetilde{\mu^+}$$

where we recall that "^" stands for the market neutralization operator.

As already explained in Remark 4.3, Model (7) may seem specific to a particular portfolio construction, but the methodology is easily generalized to different constructions, the main point being that knowing how the portfolio is built is very helpful in performance attribution.

4.5 Computing the factor contributions

The final stage of performance attribution is to make explicit the contribution of each factor to the overall performance of the investment portfolio. In the time series regression approach, each factor contributes to the extent of **(beta w.r. to the factor)** * **(realized factor portfolio return)**, whereas in the cross-sectional view, the factors' contributions are equal to **(portfolio tilt w.r. to the factor)** * **(cross-sectional factor portfolio return)**.

Starting from (7) and scalar-multiplying by the active weight vector **Q** as explained in Section 4.1, one first obtains a "split" performance attribution

$$Q.\hat{R} = \sum_{i=1}^{4} \lambda_k \left(Q.z_k\right) + \lambda_s (Q.\tilde{y}) + \gamma\left(Q.\widetilde{\alpha^+}\right) + v\left(Q.\widetilde{\mu^+}\right)$$

At this stage, some reworking of the interaction terms is necessary in order to explain all the performance via the factors. Consider for instance the term $\alpha+$: Since, by construction

$$\alpha^+ = \sum_{i=1}^{4} \alpha_k z_k 1_{A>0}$$

$$\widetilde{\alpha^+} = \sum_{i=1}^{4} a_k (\widetilde{z_k 1_{A>0}})$$

it is reasonable to define a "projection" of $\left(Q.\widetilde{\alpha^+}\right)$ onto the kth factor by

$$\alpha_k\left(Q.(\widetilde{z_k 1_{A>0}})\right)$$

Of course, due to the indicatrix being a non-linear function, this projection is not a linear operation, and it cannot be uniquely defined. The solution we propose is natural, self-explanatory and does not require any complicated reorganization of the terms involved in the split attribution.

A similar decomposition is performed on the terms coming from $\left(Q.\widetilde{\mu^+}\right)$. Note once again that a more general set of expected returns might be considered and the same attribution methodology will work as long as there is an explicit mapping from the expected returns to the factor portfolios.

5. PERFORMANCE ANALYSIS OF MULTI-FACTOR INVESTMENT STRATEGIES

Here, we show some concrete applications of the cross-sectional method to analysing the performance of multi-factor investment strategies.

5.1 The global approximation

We first analyze the global approximation property. Simple statistics on the monthly approximation error summarize the clearly better performance of the cross-sectional method when compared to the approaches based on betas: The mean monthly approximation error drops from 0.12% (time series regression) or 0.15% (ex ante betas) to 0.02%, whereas its standard deviation decreases from 0.56% (time series regression) or 0.81% (ex ante betas) to 0.51%.

As a conclusion, one can observe that both methods based on the betas create an error with a higher standard deviation, and they also lead to an approximation bias. These numerical results, together with the more qualitative arguments already explained in Section 2, are convincing enough that we do not need to go any further in the comparisons between the various methods; from now on, only the cross-sectional method will be discussed.

In the next sections, we start with single factor long-short and long-only portfolios, and move towards more realistic strategies, showing how the cross-sectional method provides a useful and intuitive tool to undertand their performances.

5.2 Single factor portfolios

This section addresses the cross-sectional performance attribution for single factor portfolios, showing the role played by the cross-sectional correlations between them.

Long-short portfolios. As mentioned in Section 4, the cross-sectional performance attribution provides a perfect replication with no approximation error when applied to any linear combination of the single factor portfolios. In fact, the basic cross-sectional model of Section 4.1 is sufficient to provide a full explanation of their performance, and it is quite interesting to analyze the interpretation it provides.

The cumulative performance and factor contributions for the four single factor portfolios are shown on Figure 4. These graphs show that the performance of each single factor portfolio is explained partly by itself (the self-contribution), but also by the other factors (the cross-contributions). The existence of these cross-terms comes from the non-zero cross-sectional correlation between the single factor portfolios. From this specific example, one can conclude that low volatility and quality portfolios are mostly explained by themselves, whereas the cross-effects are much more pronounced for momentum and value. This exemplifies the way cross-sectional attribution helps shed a revealing light on the dependency structure between portfolios (see also Section 6.3 for a discussion on these cross-effects).

Long-only portfolios. One can move one step further and analyze the cross-sectional performance attribution for single factor investable (long-only) portfolios. Of course, the situation here is a slightly different, as there are now approximation errors to account for.

Consider for instance low volatility: As one can see in Figure 5, the situation here is qualitatively similar to the long-short case. The approximation error is very much under control, and the factor contributions are similar to the long-short case.

Figure 4: Factor contributions for the single factor long-short portfolios. The graphs show the cumulative alpha and factor contributions for each single factor long-short portfolio.

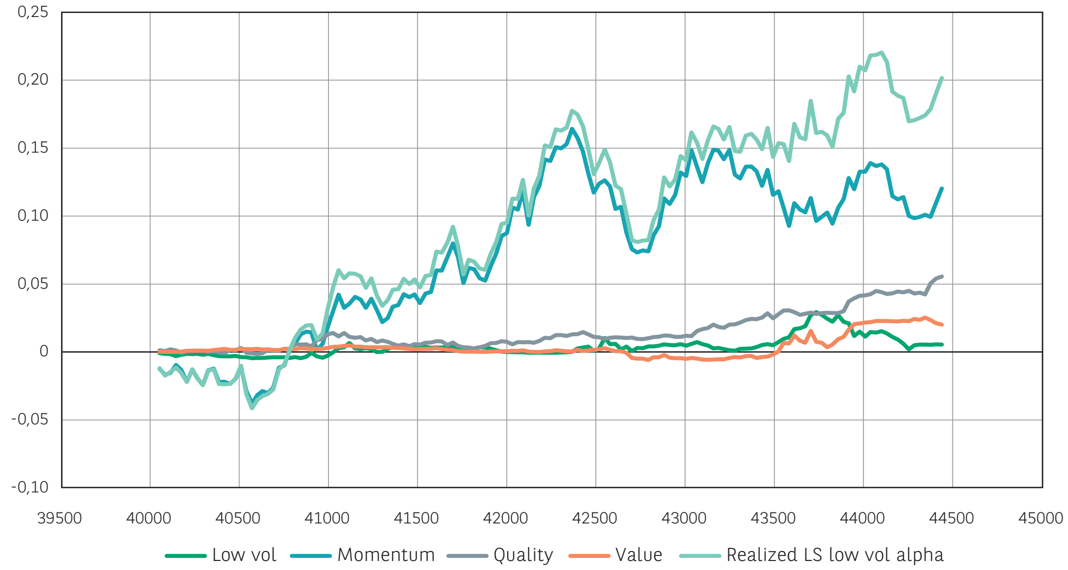

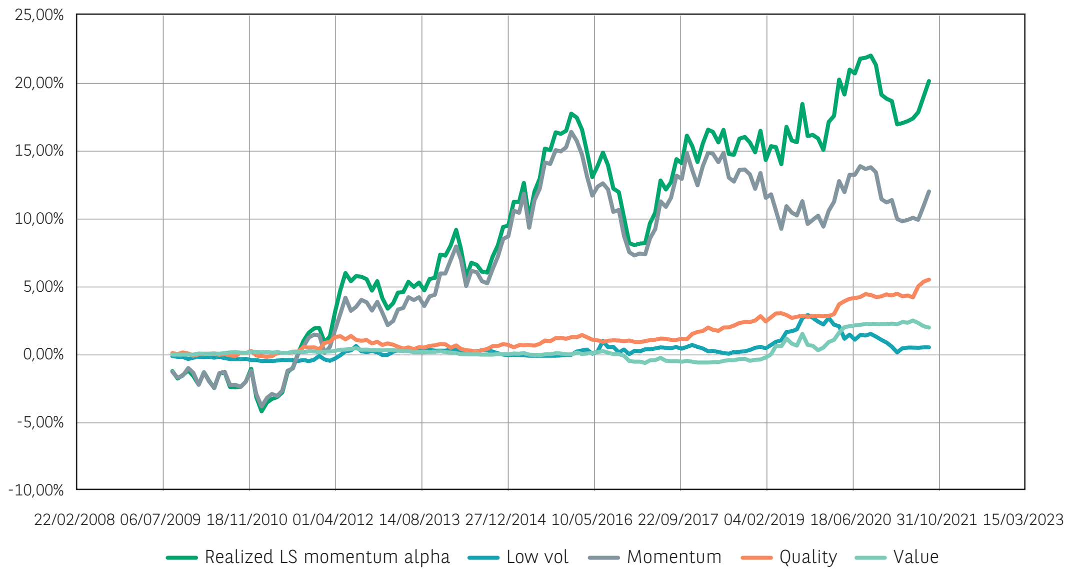

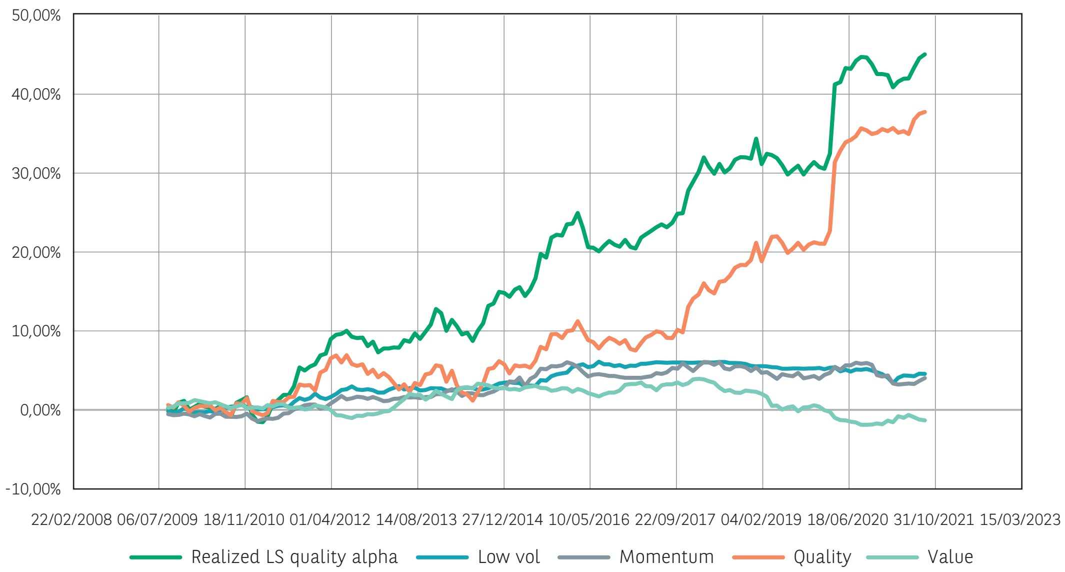
50,00%
40,00%
30,00%
20,00%
10,00%
0,00%
-10,00%
22/02/2008 06/07/2009 18/11/2010 01/04/2012 14/08/2013 27/12/2014 10/05/2016 22/09/2017 04/02/2019 18/06/2020 31/10/2021 15/03/2023
Realized LS quality alpha
Low vol
Momentum
Quality
Value

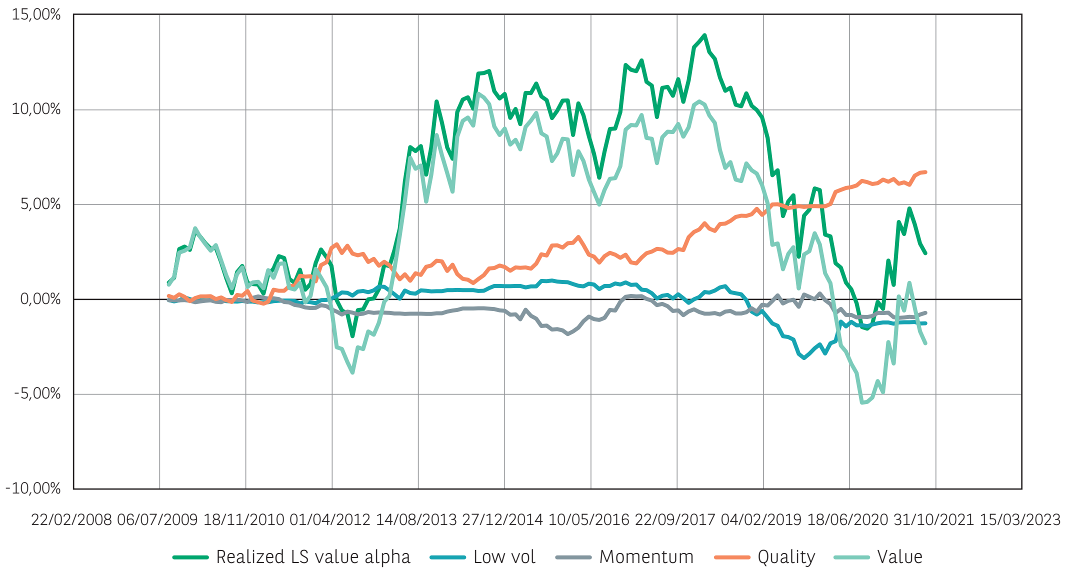
15,00%
10,00%
5,00%
0,00%
-5,00%
-10,00%
22/02/2008 06/07/2009 18/11/2010 01/04/2012 14/08/2013 27/12/2014 10/05/2016 22/09/2017 04/02/2019 18/06/2020 31/10/2021 15/03/2023
Realized LS value alpha
Low vol
Momentum
Quality
Value

Figure 5: Approximation and factor contributions for the low volatility long-only portfolio. Left: Cumulative alpha of a low volatility long-only portfolio and its cross-sectional approximation. Right: Cumulative factor contributions.

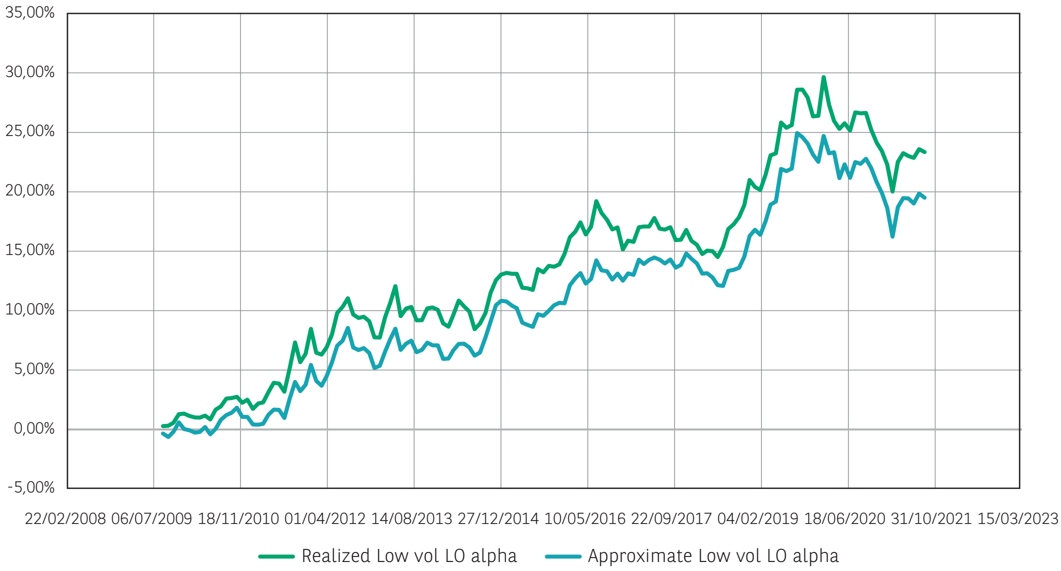

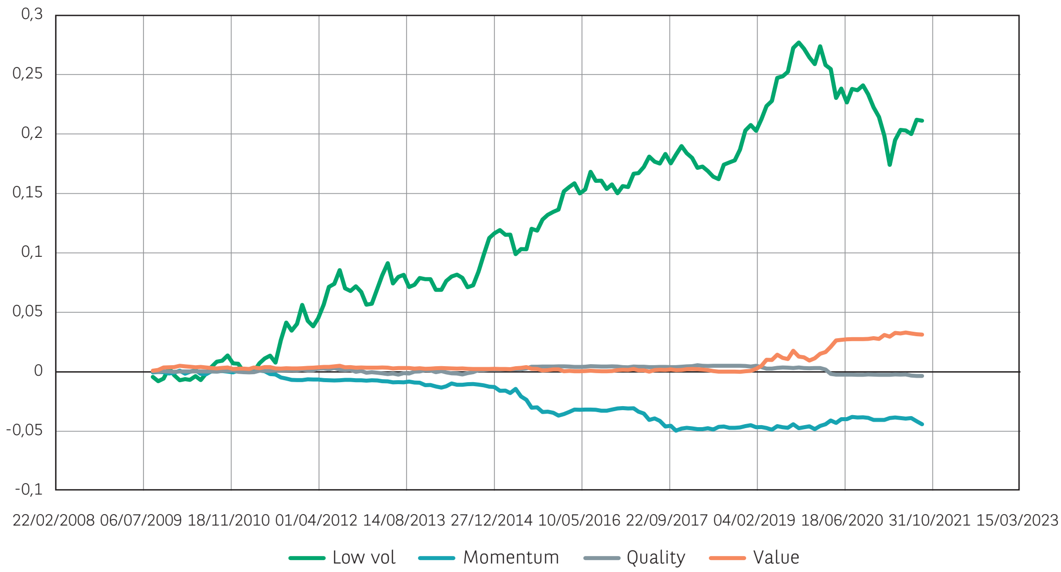

As a counter-example, and an empirical motivation for using the interaction terms in the regression, we analyze the momentum situation. The graphs in Figure 6 show that the factor-only approximation is clearly not satisfactory, whereas it becomes much more accurate if we add the interaction terms.

As a short conclusion to this analysis, one can see that the cross-sectional method provides an interpretation of the various factor contributions that accounts well for the dependency structure as measured by the tilts, but also that the basic cross-sectional regression model is not always sufficient to provide a sufficiently good global approximation, even for an investable portfolio built from a single factor. We now proceed to the analysis of multi-factor investment strategies

Figure 6: Cumulative alpha for a Momentum long-only portfolio and its approximations with or without the nonlinear terms in the cross-sectional regression.

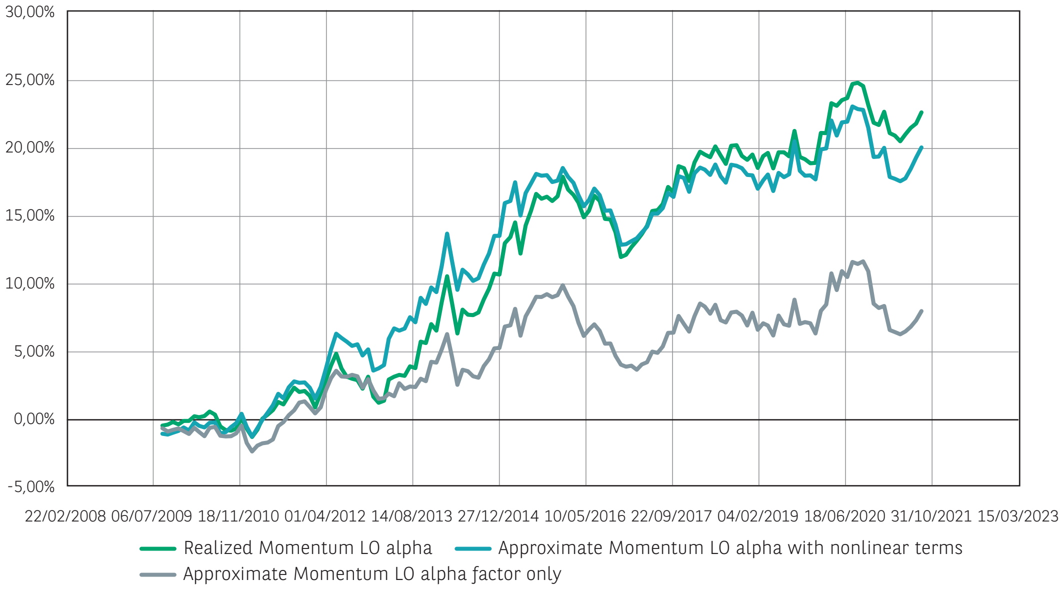

5.3 Factor contributions

Factor contribution for the theoretical long-short portfolio. As a first illustration of the cross-sectional performance attribution for a multi-factor portfolio, we show in Figure 7 the factor contributions for a theoretical long-short portfolio on the US market. This is the portfolio one would like to invest in, should the positivity and maximum weight constraints be relaxed.

Figure 7: Cumulative factor contributions for the theoretical long-short portfolio. The low volatility and quality factors contribute overall quite positively. Momentum is flat and value has contributed negatively since 2015.

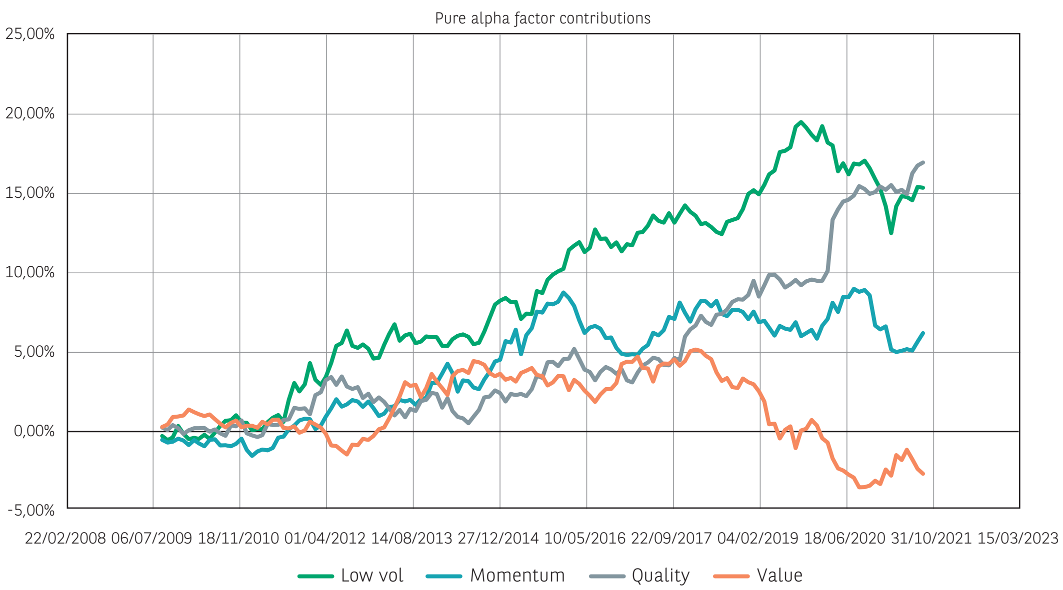

Note that, due to the cross-sectional correlation structure, and although there is no approximation error coming from the cross-sectional regression, the factor contributions are not necessarily equal to the factor performances, a fact that will be discussed in Section 6.3 below. Factor performances and contributions are nonetheless highly correlated.

Factor contribution for the investment portfolio. Figure 8 shows the factor contributions for a multi-factor equity portfolio on the US and European markets.

Figure 8: Factor contributions for the long-only portfolio. In the US market, the contributions of quality, momentum and value contributions are in line with those for the theoretical long-short portfolio. Low volatility has not contributed well in the recent period. For the European market, quality and momentum clearly have positive contributions, low volatility is flat and value has again been a negative contributor since 2015. In both markets, the additional contribution of the size factor remains very small.

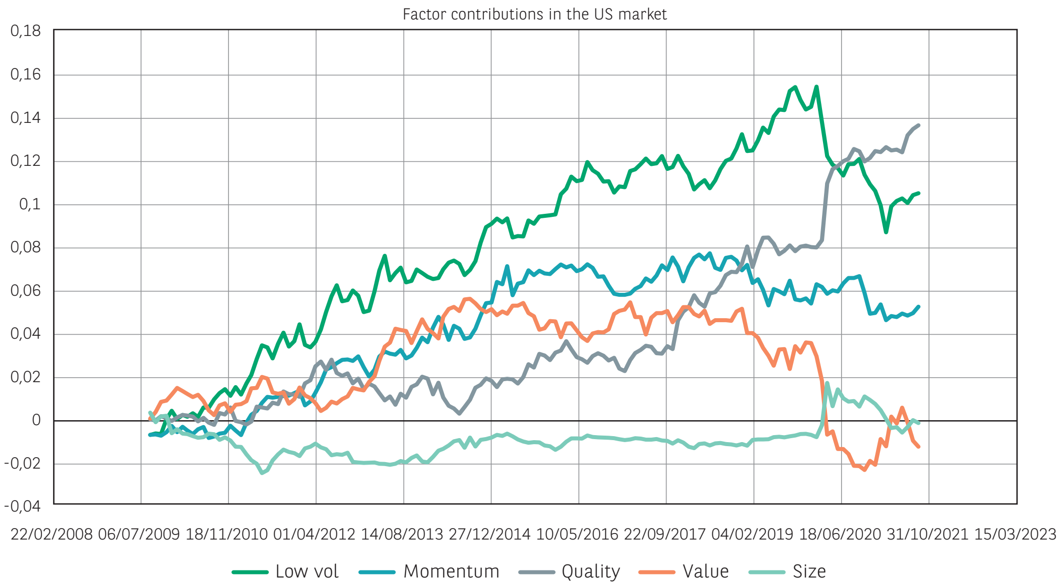

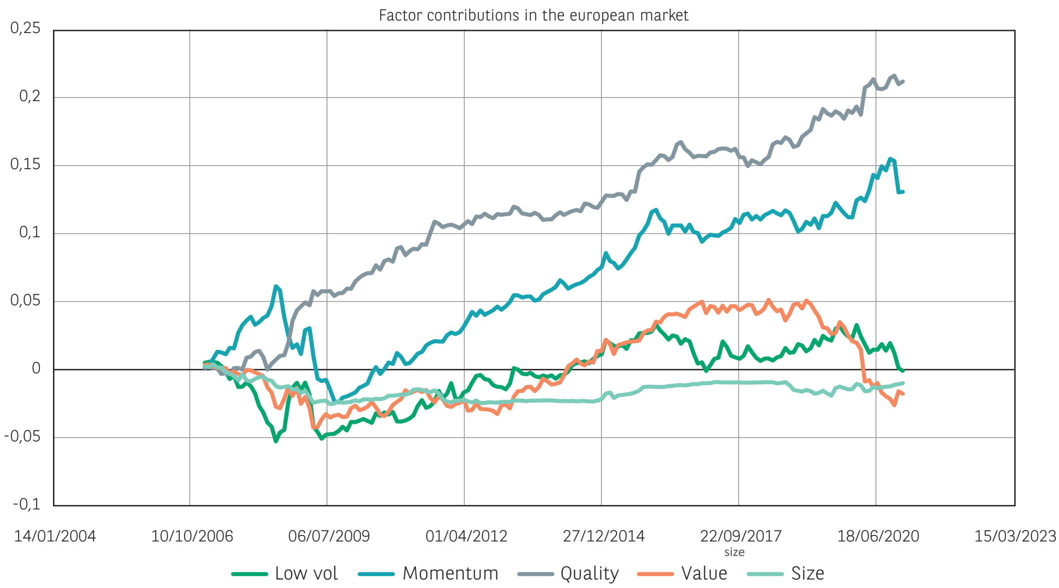

We now provide some interpretation of, and perform some sanity checks on, the proposed performance attribution. Factor investing should entail a positive exposure to the factors and therefore, contributions and pure factor performances should be qualitatively similar. For instance, it is important to understand whether negative factor contributions such as that of the value factor over the past few years are due to poor factor performances, or inadequate portfolio exposure to the factor. Hence, a significant new question naturally arises: Are the factor contributions in line with the performances of the factor portfolios?

Figures 9 and 10 provide a globally satisfactory answer to that question. The factor contributions are clearly positively correlated with the factor portfolio performances, so that, from a statistical and financial point of view, the cross-sectional performance attribution makes very good sense.

It also reflects well on the portfolio construction, and the fact that the strategy has globally positive ex post exposure to the factor performances.

Figure 9: Factor contributions and performances - US. Note the rather obvious straight regression line for all factors.

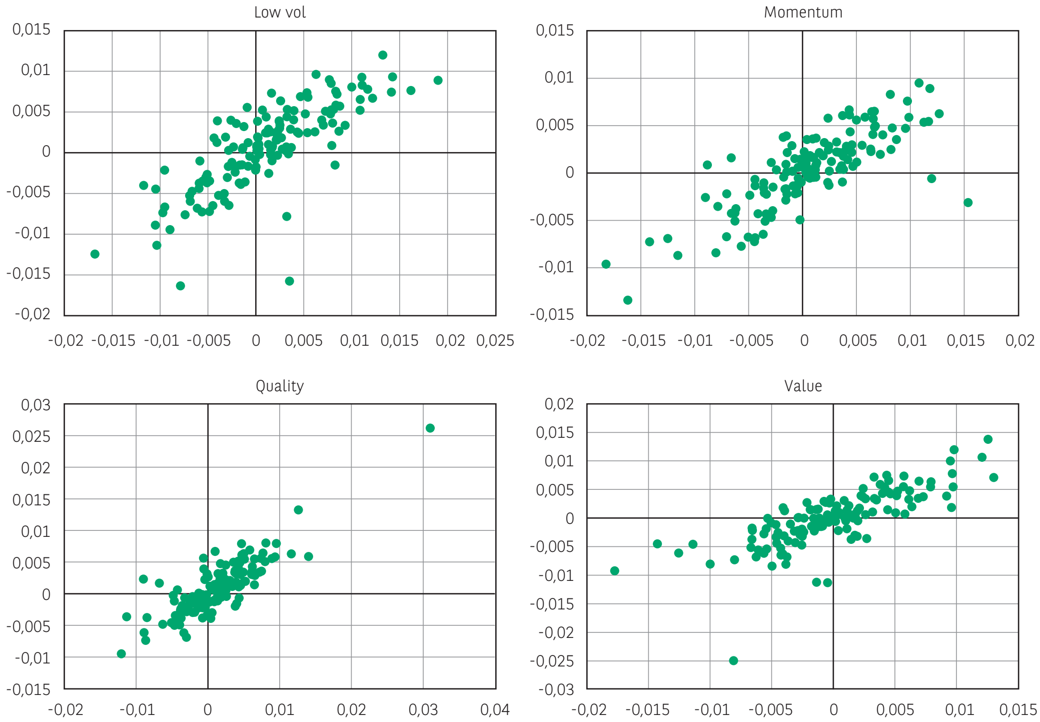

Figure 10: Factor contributions and performances - Europe. Again, a rather obvious straight regression line for all factors.

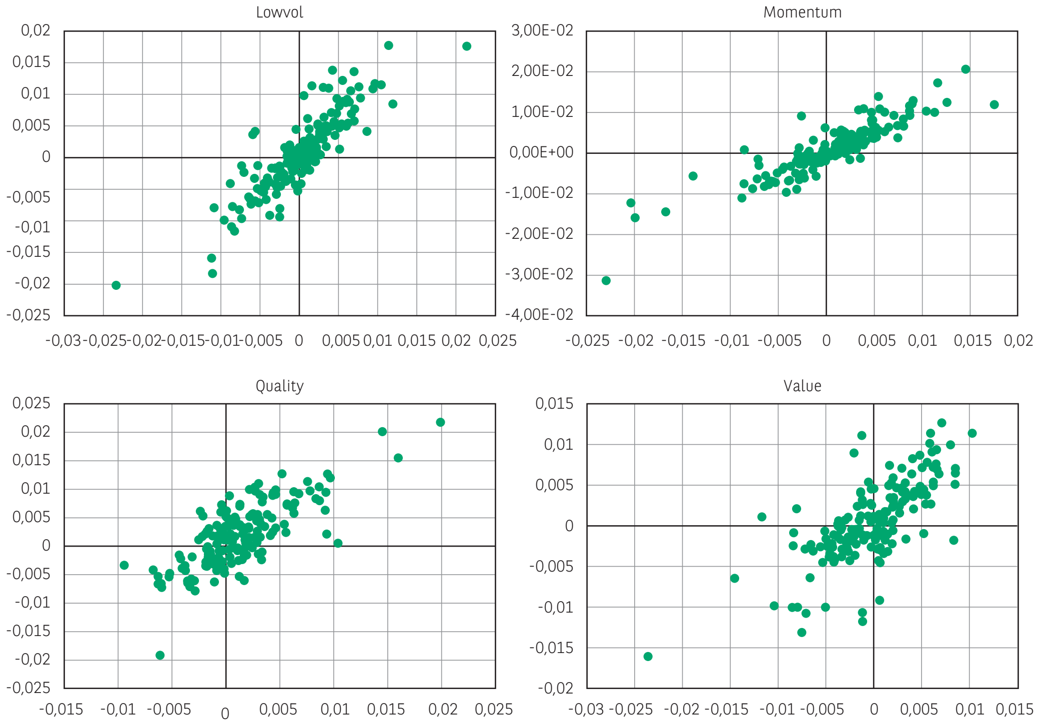

6. ROBUSTNESS OF THE CROSS-SECTIONAL ATTRIBUTION

In this section, we challenge the robustness of the cross-sectional methodology by changing some of our working hypotheses and analyzing the consequences of the cross-sectional attribution.

6.1 Different portfolio constructions

Naive portfolio. A simple thought experiment is useful in understanding both the relevance and the limitations of the cross-sectional approach to performance attribution. Consider an admissible, i.e., investable, portfolio satisfying all the constraints, but built from a set of expected returns that are not necessarily factor-driven. The question is to understand whether one can expect a useful and accurate performance attribution of such a portfolio, based on the factors. As a specific example, we analyze the already encountered low volatility long-only portfolio performance attribution, but this time, without using the low volatility factor portfolio as an explanatory variable in the cross-sectional regression. One can ask oneself whether the performance attribution of such a portfolio continues to make sense, and Figure 11 provides a clearly negative answer to this question: The approximation error becomes huge, and the approximated portfolio looks nowhere near the realized one.

Figure 11: Naive factor-based performance attribution. The long-only low volatility factor portfolio is poorly approximated by the three remaining factors (to be compared with Figure 5.2).

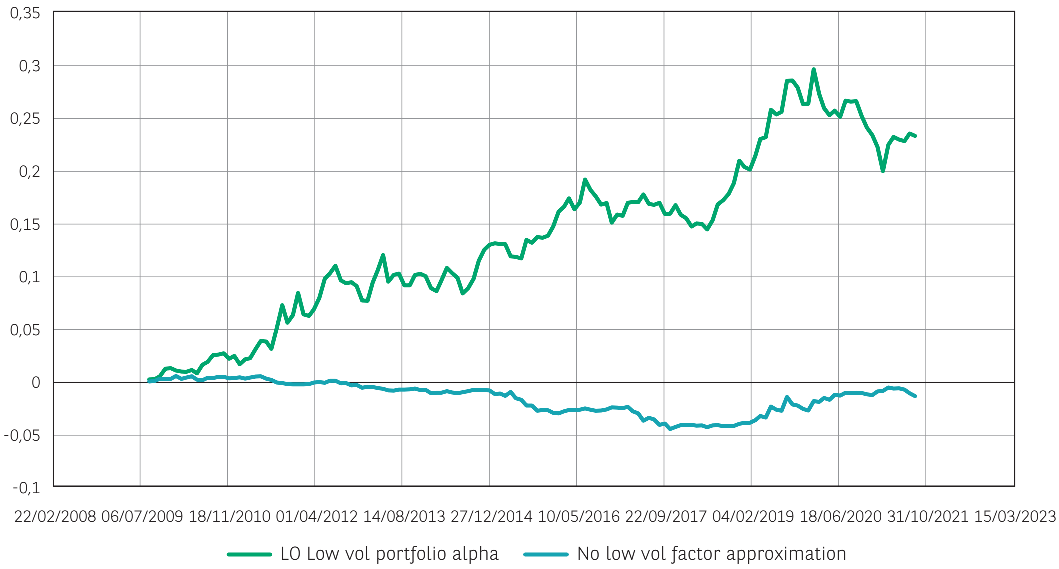

We cannot insist enough on the fact – and this is one of the limitations of the method – that only a portfolio built from factors and that somehow "looks like them", can be successfully interpreted via this cross-sectional approach. Once understood, this constraint is not a restriction in itself, only an indication that one should tread carefully when considering applying the cross-sectional method to an arbitrary portfolio.

Factor-based portfolio. Consider a multi-factor investment strategy based on a different set of expected returns. For instance, instead of the more involved reverse optimization we advocate, one could simply introduce expected returns that are explicit functions of the aggregate scores. In that case, after a straightforward modification of the non-linear terms introduced in Section 3, Figure 12 shows that the cross-sectional performance attribution continues to provide a neat approximation.

Figure 12: The global approximation for a portfolio with a different set of expected returns.

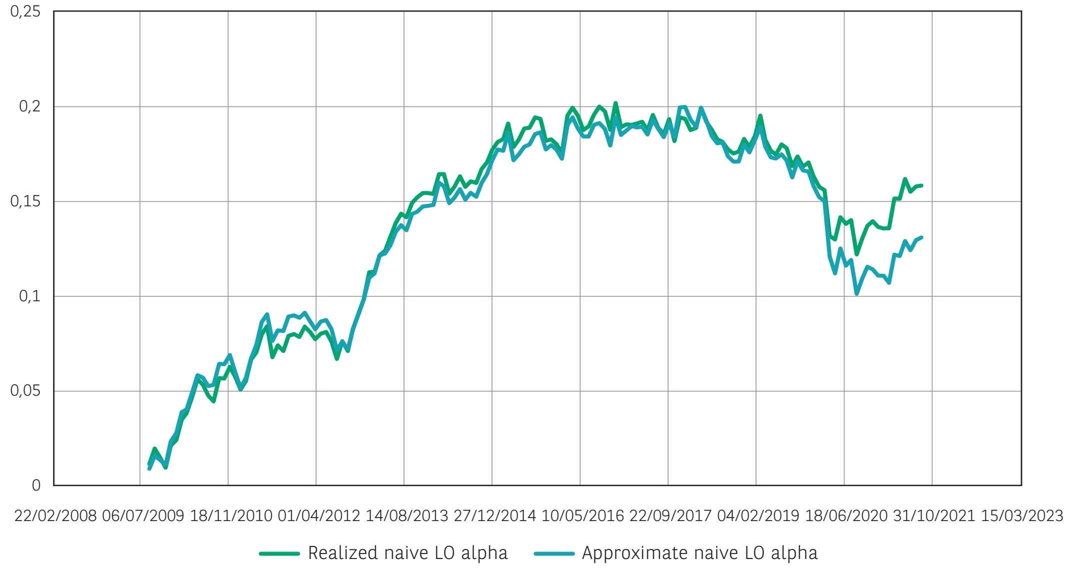

6.2 Enhancing the cross-sectional model

Changes in the factor contributions. When moving from the basic cross-sectional model to the one we actually use, we added some regressors that were justified by the portfolio optimization objective function and constraints. It is important to evaluate the extent to which the performance attribution is modified by these new terms. The first graph in Figure 13 corresponds to the naive, factor-only attribution, the second (already shown above) to the more sophisticated approach we advocate.

A look at the graphs reveals some strong qualitatively similarities: Adding the size and the interaction terms does not modify the aspect of the factor contributions; it does, however, decrease the approximation error: The standard deviation of the monthly error drops form 0.68% to 0.65% thanks to the size factor and then, more significantly, from 0.65% to 0.51% thanks to the interaction terms.

Better econometric model. Another direction for improvement could be to enrich the regression model with more general factors, e.g., country or industry dummies. We performed such an analysis by adding macro-sector indicatrices to the regression model, thereby more than doubling the number of explanatory variables. Specifically, we ran the comparison between two models: The first using the vector of betas and five factors (including high size) as regressors; the other using the betas, the five factors and six industrial sector indicator functions as regressors. As it turns out, and due to the specific portfolio construction for multi- factor strategies, we do not see any significant improvement of the global approximation, even though the more general regression itself is obviously a better econometric model for the universe of stock returns: The approximation without sectors (resp. with sectors) leads to a monthly error with 0.65% (resp. 0.70%) standard deviation and 0.03% (resp. -0.06%) mean. Furthermore, we see in Figure 14 that the sector contributions remain globally very small compared to the factor contributions:

Figure 13: Comparing factor contributions with or without interaction terms. Factor contributions in both approaches remain qualitatively very similar, with realized correlation over 85%.

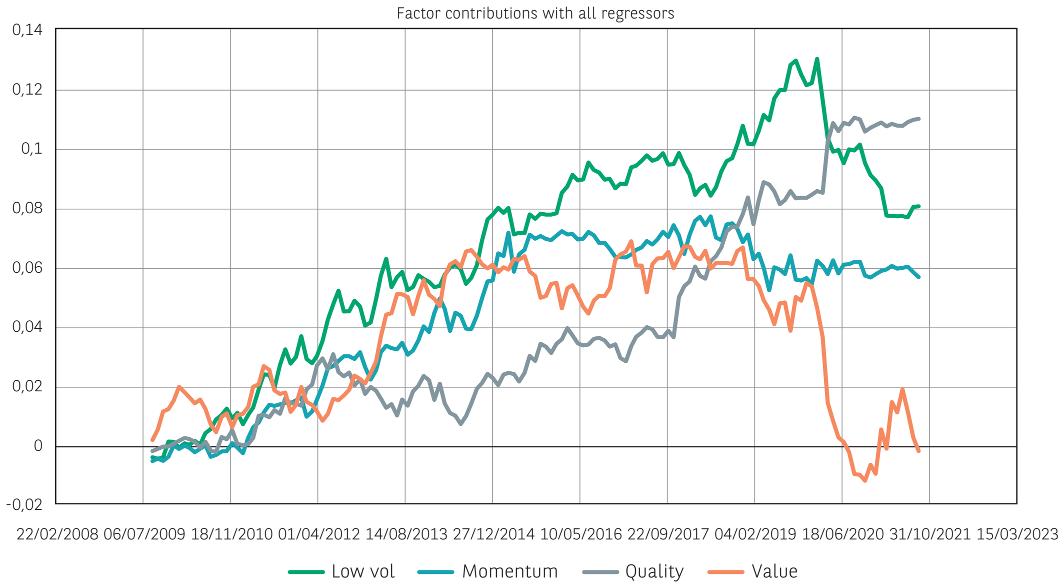

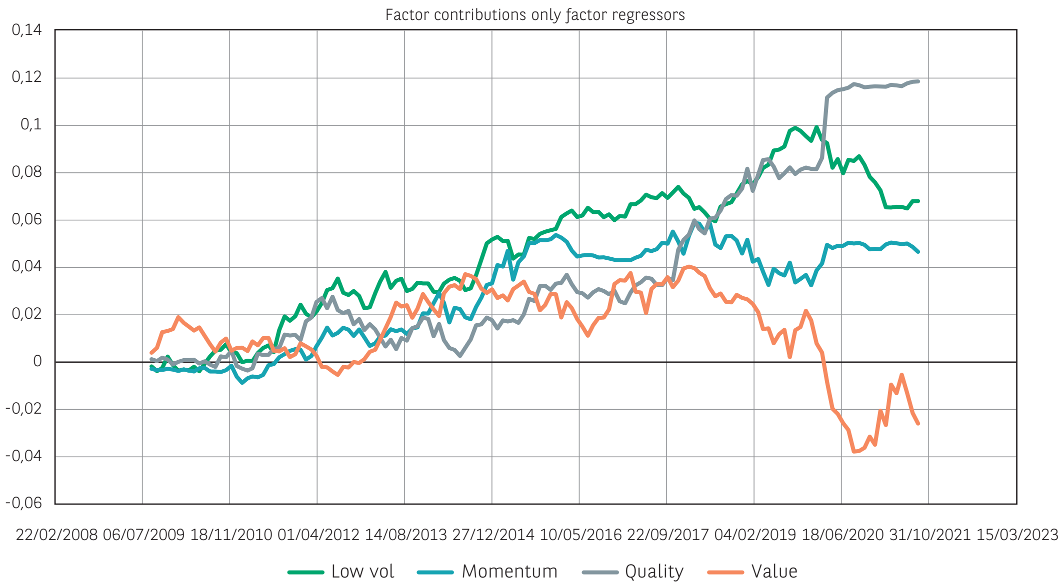

Figure 14: Cumulative factor and sector contributions. The better econometric model does not improve the accuracy of the approximation, nor does it help understand the portfolio performances.

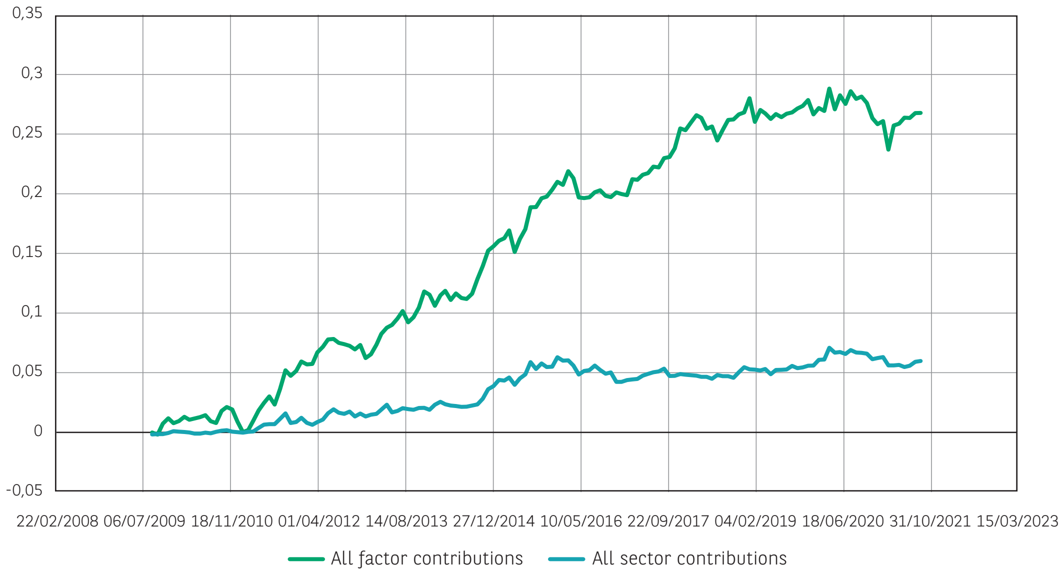

As a conclusion, we can say that improving the regression using terms that are in some sense local, i.e., that "resemble" the portfolio, greatly helps improve the approximation, whereas using more general terms that improve the econometric model does not.

6.3 The cross-selectional regression: return or portfolio vew?

This section is a short discussion on two different cross-sectional approximations corresponding to what we term the return and portfolio views. It is somewhat more theoretical than the rest of this paper, but we feel it sheds a revealing light on some easily confusing pitfalls.

Let us return to the performance attribution of the theoretical long-short portfolio with weights

$$A = \sum_{i=1}^{4} \alpha_k z_k$$

Figure 15 shows the factor contributions for the theoretical long-short portfolio and, for comparison purposes, the realized performances of each single factor portfolio.

Figure 15: Factor contributions and performances for the theoretical long-short portfolio. Profiles are very similar but not identical.

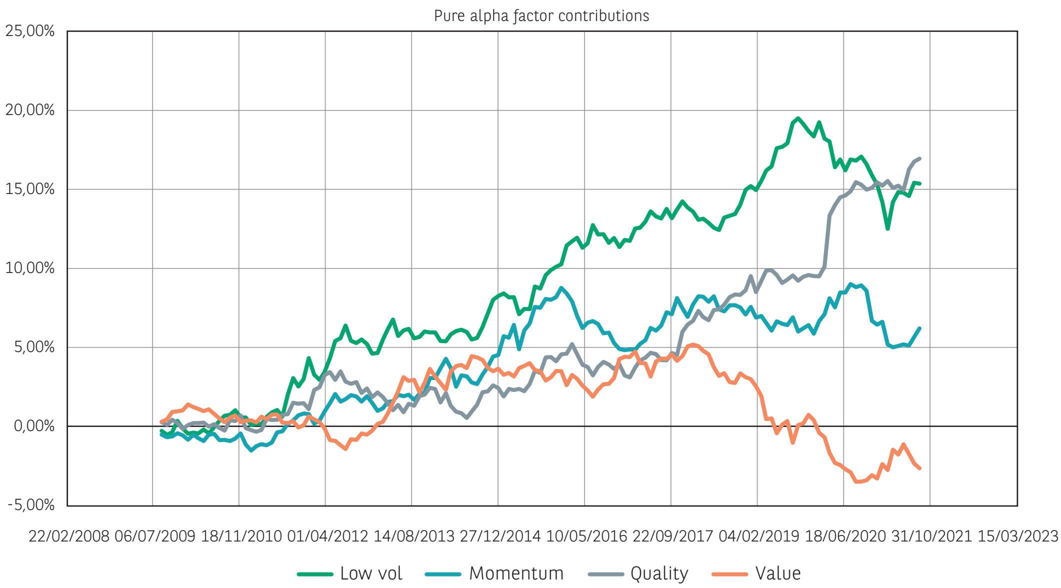

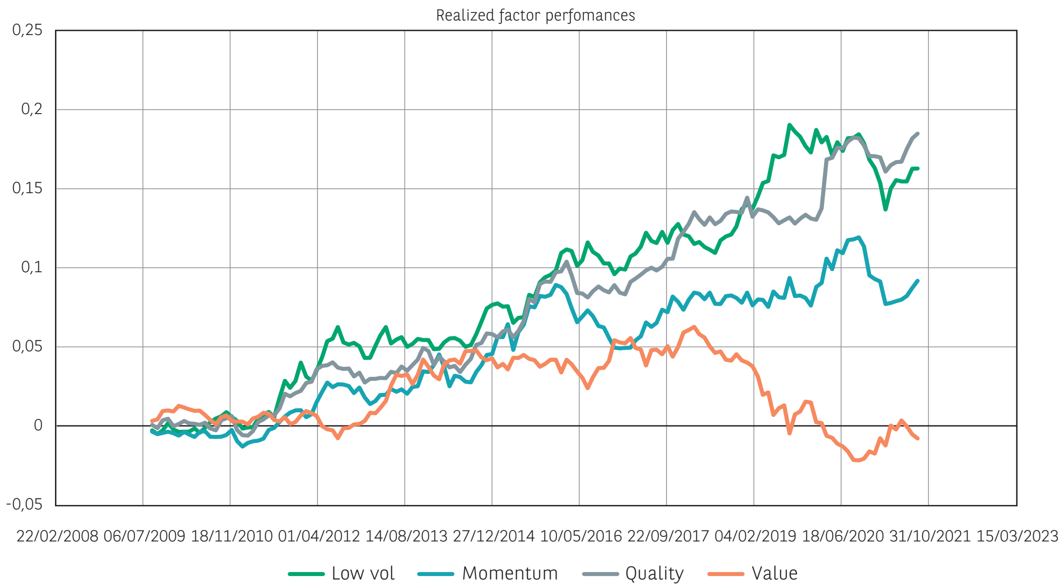

Clearly, the two graphs, although quite similar, are different. The contribution-performance correlations for the four factors are high (around 90%) but they are not equal to 100%.

This might seem a bit surprising at first, since the theoretical long-short portfolio is perfectly explained in the cross-sectional approach, see Section 4.1. The point here is that the theoretical long-short portfolio's tilts are not equal to the a_k, even after normalization; this property would only hold true if the single factor portfolios were cross-sectionally pairwise orthogonal.

This leads us to reflect on two mathematically equivalent approaches. In the simple cross-sectional regression with only the K = 4 factors as in (1), the approximate portfolio returns are given by $\boldsymbol{w}.\hat{R}$, where $\hat{R} = \sum_{k=1}^{k} \lambda_k z_k$ are the approximate asset returns. An alternate, seemingly natural, route one could follow would be to linearly regress the portfolio weights – rather than the universe returns – onto the same factor portfolios. Denote by $\widehat{W}$ these approximate portfolio weights: $\widehat{W} = \sum_{k=1}^{k} \gamma_k z_k$

Thanks to the elementary properties of orthogonal projections, there follow the identities

$$\widehat{W}.R = W.\hat{R} = \widehat{W}.\hat{R}.$$

In other words, it is mathematically equivalent to project the returns and compute the approximate portfolio return, or to project the portfolio weights and then scalar-multiply them with the realized asset returns: Both projections will lead to perfectly identical approximations. The difference lies in the interpretation of the factor contributions. Since the factor portfolios are not orthogonal, choosing one projection over the other amounts to choosing between a column or row-based representation, namely, writing the approximate portfolio returns as

$\widehat{W}.\hat{R} = \sum_{k=1}^{K} \gamma_k \left(\sum_{p=1}^{K} C_{kp}\lambda_p\right)$ or with $\widehat{W}.\hat{R} = \sum_{k=1}^{K} \lambda_k \left(\sum_{p=1}^{K} C_{pk}\gamma_p\right)$

$$C_{pk} = C_{kp} = \left(z_p . z_k\right)$$

The former expression gives the investment portfolio return as a combination of the elementary long-short portfolios returns, whereas the latter expresses it as a linear combination of the portfolio tilts. Because of the dependency structure between the single factor portfolios, we find the return-based cross-sectional regression much more satisfactory, as it sheds some specific light on the interaction between factors. In fact, consider, as in Section 5.2, one of the single factor portfolios, and apply both methods to its performance attribution. In the portfolio view, there is only one term that does not vanish, so that the whole performance of a particular portfolio is fully explained by itself, whereas in the cross-sectional approach, the dependency structure comes into play and all the factors contribute to the performance of any one of them.

7. IN CONCLUSION

The cross-sectional approach to performance attribution is a simple and flexible tool providing an intuitive and accurate explanation for portfolio returns in the context of multi-factor investing. A full-fledged, exhaustive analysis has shown that a suitable extension of the naive approach found e.g., in [11], can be used to analyze the performances of multi-factor investment strategies, and provided some major improvements. In particular, our analysis highlights the importance of incorporating non-linear terms to improve the overall accuracy of the approximation. Such terms appear quite naturally during the portfolio construction process; they are, to some extent, a measure of "how different" the investment portfolio is from the theoretical long-short portfolio, and the accuracy of the performance attribution process for multi-factor investment strategies relies quite heavily on them. We have provided some concrete examples of performance attribution for multi-factor investment strategies in the US and European markets, and have also analyzed in detail the possible extensions and shortcomings of the cross-sectional methodology. We think this approach should become a standard tool in the performance attribution of multi-factor investment strategies, as it provides a faithful and intuitive description of the portfolio performance that directly relates to the building blocks used during the portfolio construction process.

ACKNOWLEDGEMENTS

The authors would like to thank their colleagues Z. Amghar and J. Czhyr, who initiated the cross-sectional performance attribution project, for numerous fruitful discussions, their relevant and useful suggestions. They are especially grateful to R. Leote de Carvalho for his careful readings, comments and suggestions, which greatly helped improve this paper.

REFERENCES

[1] C. Bacon. Performance attribution: history and progress. 2019.

[2] Dimitris Bertsimas, Vishal Gupta, and Ioannis Ch. Paschalidis. Inverse optimization: A new perspective on the black-litterman model. Operations Research, 60(6): 1389–1403, 2012.

[3] G. P. Brinson, L. R. Hood, and G. L. Beebower. Determinants of portfolio performances. Financial analysts journal, 1995.

[4] Mark M. Carhart. On persistence in mutual fund performance. The Journal of Finance, 52(1): 57–82, 1997.

[5] K. Daniel and S. Titman. Evidence on the characteristics of cross-sectional variation in stock returns. The Journal of Finance, 52(1): 1– 33, 1997.

[6] E. F. Fama and K. R. French. Common risk factors in the returns on stocks and bonds. Journal of financial economics, 3, 1993.

[7] Robert A. Haugen and Nardin L. Baker. Commonality in the determinants of expected stock returns. Journal of Financial Economics, 41(3): 401–439, 1996.

[8] G. He and R. Litterman. The intuition behind Black-Litterman model portfolios. 2002.

[9] C. Heatter, G. Charles, and W. Yi. A four-factor performance attribution model for equity portfolios. Journal of Performance Measurement, 9, 2004.

[10] N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. The Journal of Finance, 48(1): 65–91, 1993.

[11] R. Litterman and Goldman Sachs Quantitative resources group. Modern Investment Management: An Equilibrium Approach. Wiley finance, 2003.

www INVESTORS' corner

November 2021 - Design: Creative Services BNPP AM - CC112021

BNP PARIBAS
ASSET MANAGEMENT

The sustainable investor for a changing world