Measuring the performance impact of ESG investing

A next gen attribution method ready for the future of sustainable investing

Zeist, September 2023
René Wijnen, Portfolio Manager Quantitative Investment Strategies Achmea IM
Ralph Sandelowsky, Senior Portfolio Manager Quantitative Investment Strategies Achmea IM
Dennis Thé, Head of Quantitative Investment Strategies Achmea IM

Investment Management

Abstract (EN)

The rise of ESG investing in recent decades has been remarkable and has changed the investment world. Meanwhile, performance attribution techniques have largely been unchanged during this period. Nowadays, the investment process isn't that one-dimensional anymore due to the rise of exclusion policies, ESG tilts, carbon reduction targets and many other ESG considerations. This complicates matters for performance attribution as it is no longer clear cut whether the performance impact of, for example, underweighting the energy sector should be attributed to the strategy or to any one of the ESG considerations, or both. We propose using Shapley values to decompose the active weights into different portfolio choices. By doing so, we reobtain transparency over which choices lead to which portfolio exposures. The results can be used in performance attribution to obtain an accurate measurement of the performance impact of strategy and ESG choices. This novel approach to performance attribution can be of great value to portfolio managers and asset owners of both active management strategies and replicating strategies of custom ESG indices.

Abstract (NL)

De opkomst van ESG-beleggen in de afgelopen decennia is opmerkelijk geweest en heeft de beleggingswereld onmiskenbaar veranderd. Ondertussen zijn performance attributietechnieken over dezelfde periode grotendeels onveranderd gebleven. Dit maakt ook dat het beleggingsproces niet meer zo eendimensionaal is vanwege de opkomst van uitsluitingsbeleid, ESG-tilting, CO_2-doelstellingen en vele andere ESG-overwegingen. Dit bemoeilijkt de performance attributie, omdat het niet langer eenduidig is of het effect van bijvoorbeeld een onderwogen energiesector moet worden toegeschreven aan de strategie of aan een van de ESG-overwegingen, of beide. Wij stellen voor om Shapley values te gebruiken om de actieve gewichten op te splitsen in verschillende portefeuillekeuzes. Hiermee verkrijgen we transparantie over welke keuzes leiden tot welke portefeuille exposures. De resultaten kunnen worden gebruikt in performance attributie om gedetailleerdere inzichten te verkrijgen van het performance impact van strategie- en ESG-keuzes. Deze nieuwe benadering van performance attributie kan van grote waarde zijn voor portefeuillebeheerders en vermogensbezitters van zowel actieve strategieën als replicerende strategieën van op maat gemaakte ESG-indices.

Table of Contents

1. Introduction

In recent years, a remarkable transformation has taken place within the realm of investment management, as investors increasingly recognise the importance of integrating non-financial factors into their decision-making processes. The rise of environmental, social, and governance (ESG) investing represents a pivotal shift towards a more holistic approach that takes into account not only financial returns but also the long-term sustainability and impact of investments.

The growth trajectory of ESG investments has been impressive, with global assets under management (AuM) in sustainable investments reaching unprecedented levels. According to industry reports, ESG-focused funds and strategies have experienced exponential growth, with double-digit annual increases in AuM observed across various regions. Institutional investors, including pension funds, sovereign wealth funds, and asset management firms, have played a pivotal role in driving this surge, as they increasingly recognise the potential for sustainable investments to generate long-term value and mitigate risks.

In their 2022 Asset and Wealth Management revolution paper, PwC finds that from 2020 to 2021, global AuM in ESG investments almost doubled from USD 9,4 trillion to USD 18,4 trillion[1]. The majority of this increase comes from Europe, where AuM almost tripled from USD 4,7 trillion to USD 12,8 trillion.

ESG investing encompasses the evaluation of environmental, social, and governance criteria when making investment decisions. The environmental criterium focuses on elements such as:

- climate change,
- resource efficiency,
- pollution, and
- biodiversity.

Social criteria encompass:

- labour standards,
- human rights,
- diversity and inclusion, and
- community relations.

Governance criteria evaluate, amongst others:

- board compensation and diversity,
- audit and financial reporting,
- business ethics and codes of conduct, and
- stakeholder management.

By integrating these non-financial factors, ESG investing seeks to align financial goals with broader environmental, societal, and governance objectives.

There are several approaches in which investors can use these non-financial factors in their portfolio construction process. The oldest and most well-known is exclusion, where an investor sells its position in financial instruments issued by a company that does not meet the investor's threshold on ESG criteria. Excluded companies cannot be bought until they are removed from the exclusion list. A second approach is engagement and voting, where an investor purchases a stake in a company and exerts influence to improve the company's scoring on these non-financial factors. A third and relatively new methodology is to use the company's scores on non-financial factors to construct a portfolio that outscores a certain benchmark. For example: a portfolio should have lower CO2 emissions or lower water usage than a benchmark[2].

[1] https://www.pwc.com/gx/en/financial-services/assets/pdf/pwc-awm-revolution-2022.pdf
[2] A portfolio's carbon footprint or water use is determined by the portfolio holdings.

In the early days of ESG investing, investors accepted that their returns could deviate from the benchmark. As the deviation was small, and investors believed that they were adding non-monetary value by implementing the ESG factors, not much attention was given to it.

As the AuM of ESG investments increased, investors began to take a more stringent approach to the companies that they invested in. By imposing tighter limits, the contribution of ESG factors to the overall return has subsequently increased. This means that a portfolio's relative return to a benchmark will be increasingly influenced by the effects of non-financial factors.

In the investment management sector, deviations from a benchmark are typically underpinned by a rigorously tested, almost scientific investment process. The portfolio management team running this investment process should be able to explain all differences based on the investment decisions that they or their models have made.

To a certain extent, this is at odds with the relative returns generated by applying non-financial factors. The returns generated by these factors face much less scrutiny and are often accepted based on the premise or belief that applying these non-financial factors adds societal value, even at the expense of financial returns.

With the growing impact of non-financial factors on relative returns, it is time to adapt the investor's performance attribution. This is done by developing a methodology that can accurately separate the effects of different financial and non-financial choices on a portfolio's composition. This paper focusses on relative weights, meaning portfolio weight deviations from a benchmark. These changes can be the result of several elements, including but not limited to:

- investment processes aimed at generating positive relative returns or reducing portfolio risk,
- normative and non-normative exclusions (can also include a best-in-class approach),
- ESG tilting to improve portfolio ESG scores. These can be generic ESG ratings or more specific ESG related scores.
- Carbon emissions reductions. Reducing the carbon footprint of the portfolio by mitigating the exposure of the portfolio's holdings to high emitting companies.

2. Problem Definition

Relative returns originate from holding either different assets, or the same assets at different weights, or a combination of both versus a benchmark. In order to estimate the real performance impact of applying non-financial factors to an investment portfolio, both sides of the equation should be taken into account: i.e. which assets are sold, and which assets are bought as a result of applying non-financial factors to the investment process.

This presents two challenges:

1) Lowering a company portfolio weight versus a benchmark, based on non-financial factors might look like the relative return can be attributed to these factors, but if the asset manager's investment process would not have included the company in the portfolio in the first place, to what extent can this relative return be attributed to the non-financial factors?
2) Once the portfolio weight of a company has been reduced, the question arises what will be bought in this company's place. If an energy company is excluded and the assets are invested in an IT company, is the relative return arising from that change due to ESG factors, or sector effects?

In order to properly answer these questions, this paper will analyse commonly used performance attribution methodologies and come up with a number of enhancements to make sure that the performance attribution measures what it intends to measure.

3. Performance attribution theory

Given that investors face an environment in which ESG integration is playing an increasingly substantial role, performance attribution should integrate methods to accurately assess the impact of ESG considerations on realised risk and return. Traditional attribution methods often fail to accurately account for ESG considerations for several reasons which will be elaborated in this next chapter.

3.1 Brinson-Fachler attribution

The Brinson-Fachler attribution, also known as Brinson attribution or Brinson-Fachler model, is a popular method used in the field of investment management to analyse the performance of an investment portfolio. It provides a framework for breaking down the sources of portfolio returns and assesses the contribution of numerous factors.

The model typically decomposes the portfolio's return into three main components: asset allocation effect (top down), security selection effect (bottom up), and interaction effect. The asset allocation effect reflects the return contribution resulting from the allocation of assets across different asset classes or investment categories, but it can also be applied to sector or regional classifications. The security selection effect measures the return contribution attributable to the selection of individual securities within each asset class. The interaction effect captures the combined impact of the interaction between asset allocation and security selection decisions. For example, when a specific sector is overweighted and within this sector, the manager picked stocks that outperformed the sectoral benchmark, these two effects amplify each other and there is a positive interaction effect.

The table below shows a typical Brinson-Fachler attribution, this is a stylised example with three sectors. For a more detailed description on how to compute a Brinson-Fachler attribution, see Appendix A. The table can be read as follows: the positive allocation effect of materials implies that overweighting this sector added value to the portfolio. However, the negative selection effect means that, on average, the portfolio manager selected materials stocks that performed worse than the sector benchmark. In total, the negative effects dominate, leading to a negative active return on the materials sector.

Table 1: Stylised example of typical sector-based Brinson-Fachler attribution

Sector	Portfolio Weight	Benchmark Weight	Active Weight	Portfolio Return	Benchmark Return	Allocation	Selection	Interaction	Total
Materials	25.00%	20.00%	5.00%	6.00%	8.00%	0.21%	-0.40%	-0.10%	-0.29%
Industrials	25.00%	15.00%	10.00%	7.00%	7.00%	0.33%	0.00%	0.00%	0.33%
Energy	25.00%	25.00%	0.00%	-4.00%	-2.00%	0.00%	-0.50%	0.00%	-0.50%
Financials	25.00%	40.00%	-15.00%	3.00%	4.00%	-0.04%	-0.40%	0.15%	-0.29%
Total	100.00%	100.00%	0.00%	3.00%	3.75%	0.50%	-1.30%	0.05%	-0.75%

Source: Achmea IM

The Brinson-Fachler model is a well-established method of attribution, frequently employed by investors. However, when applying this model to assess the influence of ESG integration on performance, straightforward analysis is hindered due to its inherent reliance on classification techniques. While the model effectively evaluates the impact of allocation, selection, and interaction on asset class, sector exposures, and regional allocation performance, it encounters challenges in dealing with sustainability-related data.

Certain sustainability data can be readily represented as classifications, such as binary exclusions (yes/no) for specific criteria. However, other significant sustainability metrics, like carbon emission intensity or ESG scores, do not possess inherent discrete classifications. Consequently, using the Brinson-Fachler model in such cases necessitates converting continuous variables into discrete classifications, where stocks are assigned to distinct and exclusive groups.

Using a Brinson-Fachler model to get a comprehensive understanding of the impact of multiple performance sources including sustainability factors, proves to be a profound challenge. The model's single-dimension attribution approach fails to provide a holistic view, especially given the multifaceted nature of contemporary ESG integration, which may include exclusions, CO2 reduction goals, tilting based on ESG ratings, or client-specific themes. Active management strategies might encompass even more diverse aspects.

Examining multiple individual Brinson-Fachler attributions could be attempted, but this approach falls short of revealing the complete drivers of return and fails to accurately isolate the effect of each factor. For instance, when analysing a Brinson-Fachler attribution on ESG ratings, a positive allocation effect from investing less in the worst-rated companies (CCC) may be a result of ESG tilting, but other effects may also be at play. Overrepresentation of excluded or high-emitting firms in this rating group or the manager's active strategy could also influence the observed results, indicating complicating interactions between strategic and sustainable choices. As a result, an alternative attribution method is necessary for precise measurement of these effects.

In practice, a simplified version of Brinson-Fachler is often employed to gauge the impact of exclusions on portfolio performance. By applying Brinson-Fachler on the binary classification of excluded (yes/no), one can calculate allocation effects for both excluded and non-excluded stocks, with selection effects being nullified due to the market capitalisation-based reallocation for non-excluded stocks. However, this approach only accurately reflects exclusions' performance impact in scenarios where the portfolio adheres solely to market cap weighting, and implements exclusions without any other ESG integration or active portfolio choices. For more complex situations involving additional ESG integration factors, this method fails to consider interactions adequately, leading to inaccurate measurements of exclusions' impact on performance.

For example, consider a portfolio with both a carbon reduction target and an exclusion policy for oil and gas companies. An excluded oil producer might have also been underweighted or excluded due to the carbon reduction target. This makes it difficult to attribute the full effect to exclusions alone, as that neglects the carbon reduction target's influence.

Similarly, in a scenario where both a value and a growth manager exclude oil and gas companies, the impact of exclusions would disproportionately affect the value manager compared to the growth manager, as these companies tend to be part of the value universe as opposed to the growth universe. This discrepancy arises from the value manager's likelihood of overweighting some oil and gas companies in the absence of exclusions, while the growth manager would not have selected or significantly underweighted these companies anyway. However, the simplified Brinson-Fachler approach for exclusions mentioned earlier would fail to discern these differing impacts on the two managers.

In conclusion, while the Brinson-Fachler model is a valuable tool for attribution analysis in investment contexts, its application to assess ESG integration's impact on performance necessitates a thoughtful consideration of data classification and the utilisation of alternative attribution methods to capture the complex interactions between different active choices.

3.2 Factor-based attribution

The factor-based attribution can be read as follows: the portfolio has a positive exposure to all the style factors; the value, momentum and quality factors had a positive return, and hence these exposures resulted in positive contributions to the active return. The small cap factor, on the other hand, underperformed over this period. This explains why the positive relative exposure of the portfolio resulted in a negative contribution from the small cap factor. Similarly, the over- or underweighting of certain sectors and countries in this portfolio resulted in a negative contribution.

The factor-based performance attribution overcomes two problems that the Brinson-Fachler method faced. First, it can use a multivariate view, assessing the impact of several systematic factors simultaneously. Second, it can directly use both continuous and discrete data.

Figure 1: Cumulative distribution of the MSCI ACWI carbon intensity

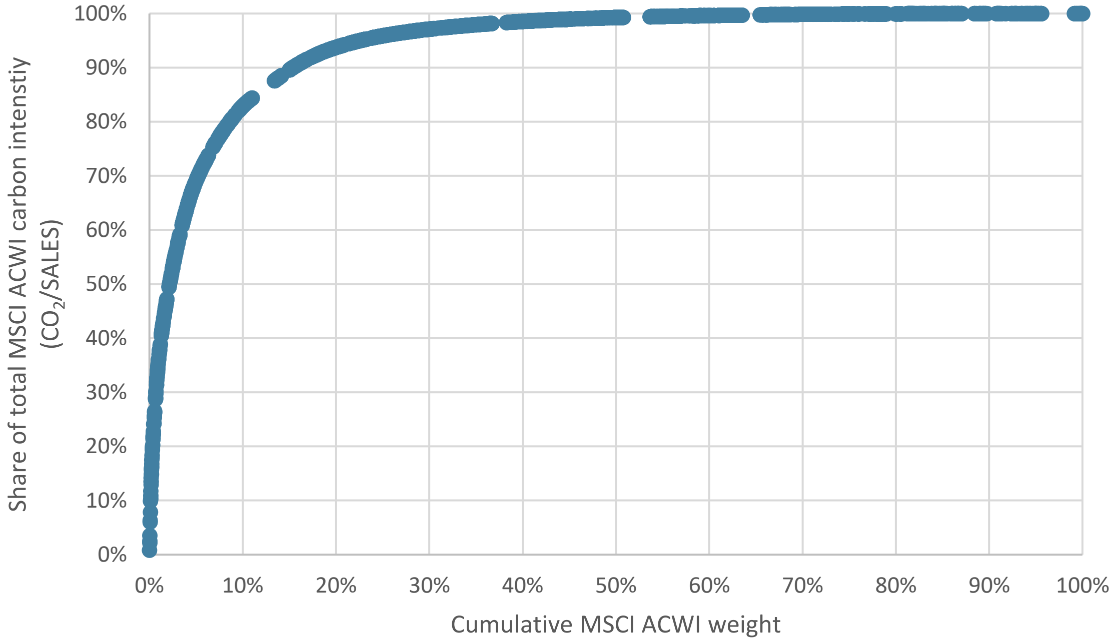

Source: Achmea IM, MSCI

A straightforward way to assess the performance impact of ESG-integration would be to add ESG related factors such as carbon emissions intensity factor. This factor should, corrected for other systematic exposures, reflect the performance of green firms (low emission intensity) versus brown firms (high emission intensity).

Initially, the idea of considering carbon emissions data seems logical. However, a significant problem arises because the data on carbon emissions is heavily skewed. The graph below illustrates how the carbon intensity (Scope 1+2 CO2 emissions/Sales) of the MSCI ACWI[3] is distributed among its constituents. In an ideal scenario where all constituents have an equal carbon intensity, the plot would form a straight line with a 45-degree slope. Instead, the chart shows that a select number of companies with high emissions dominate the index. These few companies, accounting for only 2.2% of the index weight, contribute to over half of the index's carbon intensity.

As a consequence, relying solely on the pure carbon emission factor results in a heavy dependence on these few outliers, leading to an inaccurate representation of a systematic source. The resulting pure carbon emission factor would mostly represent the performance of a few high emitting companies instead of a broader systematic source of return. This risks that non-carbon related returns are misunderstood as carbon emissions related returns.

To address this issue, one approach is to normalise the carbon data, similar to how style factors are commonly handled. While this approach would overcome the reliance on a few outliers, it still faces another drawback. The normalised factor might not accurately reflect the actual impact of a carbon reduction target on a portfolio's composition. In practice, implementing a carbon reduction strategy in a portfolio would likely involve eliminating the highest emitters to minimise tracking error. For all other stocks, the restriction would have limited or no significant impact. However, this would not be adequately reflected in the exposure and contribution from a normalised carbon emission intensity factor.

To summarise, using non-normalised data would lead to inaccuracies due to heavy reliance on outliers. On the other hand, utilising normalised data would poorly reflect the real portfolio implications, still resulting in inaccuracies. Finding the right balance between these approaches is crucial for an accurate assessment of a carbon emission reduction strategy in a portfolio.

Even when the previous point is neglected, or the problem is solved, the resulting attribution does not strictly separate the ESG choices from other strategic choices, or separate one ESG choice from another. As an example, consider the portfolio of a value manager with an ESG restriction (improving ESG rating compared to benchmark). This manager is likely to have less value exposure as he would have had in the case where he would not have

[3] MSCI All Countries World Index

been imposed with the ESG restriction. However, this effect is not clear from the attribution, since there is only one value exposure/contribution, and no granularity as to which portfolio choice added or removed value exposure.

In summary, the factor-based attribution method is an improvement over the Brinson-Fachler approach in many ways, and it starts to give some idea of the impact that ESG integration has on performance. However, it still lacks the option to actually understand which choices lead to the eventual portfolio and how to separate one choice from the other. This insight becomes increasingly important when the investment profession shifts from a financial-only focus to a more integrated approach where financial and non-financial factors are both considered. An asset owner should know if the performance of a manger is due to the manager’s skill, or due to the ESG choices that were imposed upon the manager. Moreover, even if the manager is not making active choices, a better understanding of the impact of different ESG considerations is still required to be fully transparent about the potential costs or benefits from each of the implemented ESG considerations.

4. Methodology

The previous chapter focused on existing attribution methods and their shortcomings in measuring the impact of integrating ESG factors in an investment process. The main problem is that traditional attribution methods focus on explaining portfolio performance by looking at the portfolio exposures. As an example, consider a portfolio with a underweight to the energy sector, while energy is outperforming the broader benchmark. Traditional attribution methods are perfectly capable in measuring the impact of this underweight on return, providing answers to “what” happened. However, these methods lack the ability to provide a more thorough understanding of “why” this underweight existed in the first place. Is it due to the portfolio manager’s strategy, due to exclusions, ESG rating enhancement or carbon reduction? A solution must be capable of separating portfolio exposures into different choices to more accurately measure the effect of different portfolio choices such as strategy, exclusions, ESG tilting and carbon reduction on performance. The following chapter will shed more light on the proposed solution. The figure below visualises the attribution problem described here.

Figure 2: Schematic visualisation of the attribution problem of ESG performance impact

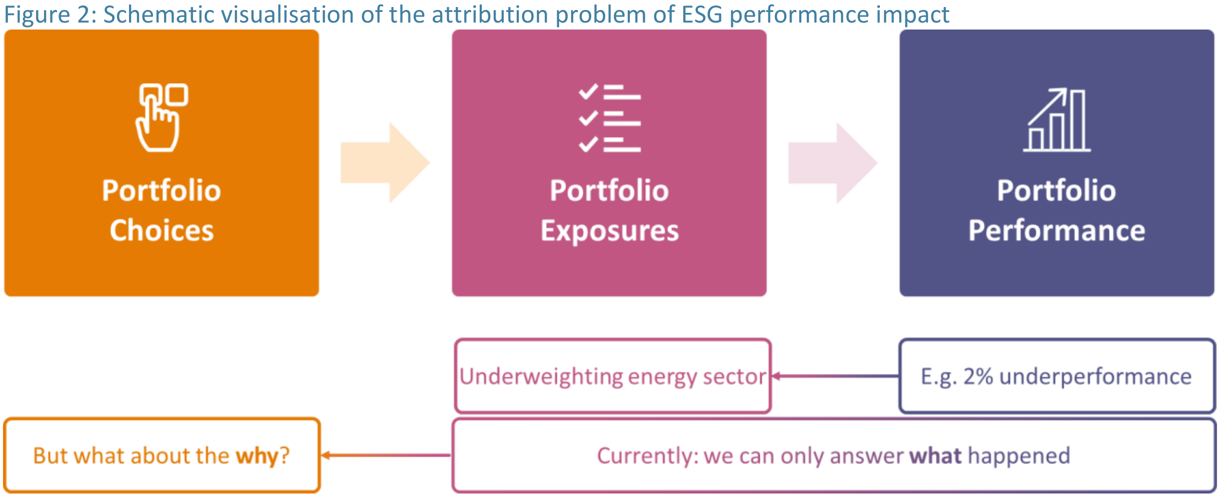

4.1 Introduction of Shapley values

Shapley values are a concept in cooperative game theory that describes how the payoff of a game can be fairly allocated to each player based on their input. The allocation is based on marginal value, i.e. how much did each player add to obtain a certain outcome. Intuitively, players who added more value should receive a bigger part of the payoff. Shapley values have found useful applications beyond game theory in various fields, most notably in machine learning, for obtaining interpretability of model outcomes. They provide a way to allocate the contributions of individual factors or features in a model to the overall outcome, providing understanding of the relative importance of these factors.

Understanding Shapley values is particularly useful when interpreting complex machine learning models, such as neural networks. By applying Shapley values to a portfolio construction process, valuable insights are gained into how different features or factors impact the model's predictions, or, in this case, the portfolio weights. This leads to understanding and interpretation of the impact generated by both financial and non-financial factors on the portfolio. Next follows a basic example of how Shapley values are calculated to get a better idea of the concept before it is applied to portfolio construction.

4.2 Example: Calculating Shapley values

Consider a scenario where three planes: Plane A, Plane B, and Plane C, need to split the costs of constructing a runway. Each plane requires a different length of runway, and the construction costs should be allocated fairly to the different planes.

The required runway lengths for each plane are as follows:

- Plane A requires 500 meters of the runway.
- Plane B requires 300 meters of the runway.
- Plane C requires 200 meters of the runway.

To calculate the Shapley values for plane A, one needs the marginal value of Plane A when added to all possible permutations, which are listed below:

Table 2: Example of Shapley value calculation: plane A

Permutation	Set before A is added	Required length before	Set after A is added	Required length after	Marginal Value
$\{A,B,C\}$	$\{\emptyset\}$	0	$\{A\}$	500	500
$\{A,C,B\}$	$\{\emptyset\}$	0	$\{A\}$	500	500
$\{B,A,C\}$	$\{B\}$	300	$\{B,A\}$	500	200
$\{B,C,A\}$	$\{B,C\}$	300	$\{B,C,A\}$	500	200
$\{C,A,B\}$	$\{C\}$	200	$\{C,A\}$	500	300
$\{C,B,A\}$	$\{C,B\}$	300	$\{C,B,A\}$	500	200

The average marginal value of A is 316.67. Similarly, the following computation can be made for Plane B.

Table 3: Example of Shapley value calculation: plane B

Permutation	Set before B is added	Required length before	Set after B is added	Required length after	Marginal Value
$\{A,B,C\}$	$\{A\}$	500	$\{A,B\}$	500	0
$\{A,C,B\}$	$\{A,C\}$	500	$\{A,C,B\}$	500	0
$\{B,A,C\}$	$\{\emptyset\}$	0	$\{B\}$	300	300
$\{B,C,A\}$	$\{\emptyset\}$	0	$\{B\}$	300	300
$\{C,A,B\}$	$\{C,A\}$	500	$\{C,A,B\}$	500	0
$\{C,B,A\}$	$\{C\}$	200	$\{C,B\}$	300	100

The average marginal value of B is 116.67. And finally, the following computation can be computed for Plane C.

Table 4: Example of Shapley value calculation: plane C

Permutation	Set before C is added	Required length before	Set after C is added	Required length after	Marginal Value
$\{A,B,C\}$	$\{A,B\}$	500	$\{A,B,C\}$	500	0
$\{A,C,B\}$	$\{A\}$	500	$\{A,C\}$	500	0
$\{B,A,C\}$	$\{B,A\}$	500	$\{B,A,C\}$	500	0
$\{B,C,A\}$	$\{B\}$	300	$\{B,C\}$	300	0
$\{C,A,B\}$	$\{\emptyset\}$	0	$\{C\}$	200	200
$\{C,B,A\}$	$\{\emptyset\}$	0	$\{C\}$	200	200

The average marginal value of C is 66.67. Note that the sum of the average marginal values of A, B and C adds up to exactly 500. These outcomes can be used to compute the following cost allocation.

Table 5: Example of Shapley value calculation: allocation results

	Plane A	Plane B	Plane C	Total
Average marginal value	316.67	116.67	66.67	500.00
% cost allocation	63.33%	23.33%	13.33%	100.00%

4.3 Applications in portfolio construction

The example above is a simple example of the idea behind Shapley values and how to compute them. This application can be extended to portfolio construction. In the previous example, there were three features and the function for the required length of the runway was:

$$Required\ Length = \max(S)\ for\ each\ possible\ set\ S$$

The portfolio construction can also be seen as a function with several inputs. Take for example:

$$AW = f(Alpha, Exclusions, ESG\ tilting, CO_2\ reduction)$$

Where all the inputs are binary variables where 1 means that a portfolio input is considered (activated) and 0 means that is not considered (deactivated). One special case is the benchmark, i.e. all active weights are zero:

$$0 = f(0,0,0,0)$$

The actual portfolio is the case where all features are activated:

$$AW^* = f(1,1,1,1)$$

Finally, the pure strategy, i.e. without any ESG considerations can be defined as:

$$\widehat{AW} = f(1,0,0,0)$$

This example uses alpha and ESG considerations, but other portfolio inputs such as tracking error or country and sector limits could also be considered. Furthermore, the number of variables can be adjusted in order to fit the investment process used. The computational cost of calculating the Shapley values increases exponentially with each additional variable. However, there are methods to reduce the additional computation cost. For example, by grouping variables (simultaneously assess the impact of sector and country limits) or by imposing a certain hierarchy in the way features are added[4]. Another possibility would be to simplify the function, in this case the portfolio construction, by a computationally lighter alternative that gives comparable results. All these methods are lower in accuracy but that might be worth it for the additional insights they provide.

The main objective of this paper is to improve the existing performance attribution to gain valuable insights into the impact of ESG integration, or other factors, on portfolio performance. In order to do so, all possible permutations of the identified portfolio variables should be computed. This allows for the calculation of the Shapley values for the active weights of all stocks in the investment universe. These permutations are stated in the table below.

Table 6: Full set of permutations for Shapley Calculation

Permutation	Alpha	Exclusions	ESG	CO_2
Actual strategy	1	1	1	1
P1	1	1	1	0
P2	1	1	0	1
P3	1	0	1	1
P4	1	1	0	0
P5	1	0	1	0
P6	1	0	0	1
Pure Strategy	1	0	0	0

The introduction of a turnover constraint on the portfolio complicates matters. This complication arises out of the fact that the current portfolio, where all factors are implemented, is taken as a starting point. In permutation three from the table above, where the Exclusions factor is not part of the investment process, it is likely that a substantial part of the turnover budget is used to buy back shares that are excluded from the current portfolio. By doing so, it is allocating less of the turnover budget to the other factors and thereby is giving an inaccurate view of how a scenario without exclusions would have looked like. Appendix C contains a detailed description of how to correct for turnover in computing the different permutations.

After calculating all the required permutations the Shapley values can be computed to explain the difference between the actual strategy and the pure strategy. This results in a decomposition of the active weight for each stock:

[4] Note that in this example, the hierarchy of alpha first is imposed, thereby applying Shapley to explain the difference between the pure and actual strategy. Alternatively, no hierarchy can be applied if the difference between benchmark and actual strategy is of interest.

Figure 3: Stylised example of Shapley values for decomposition of active weight

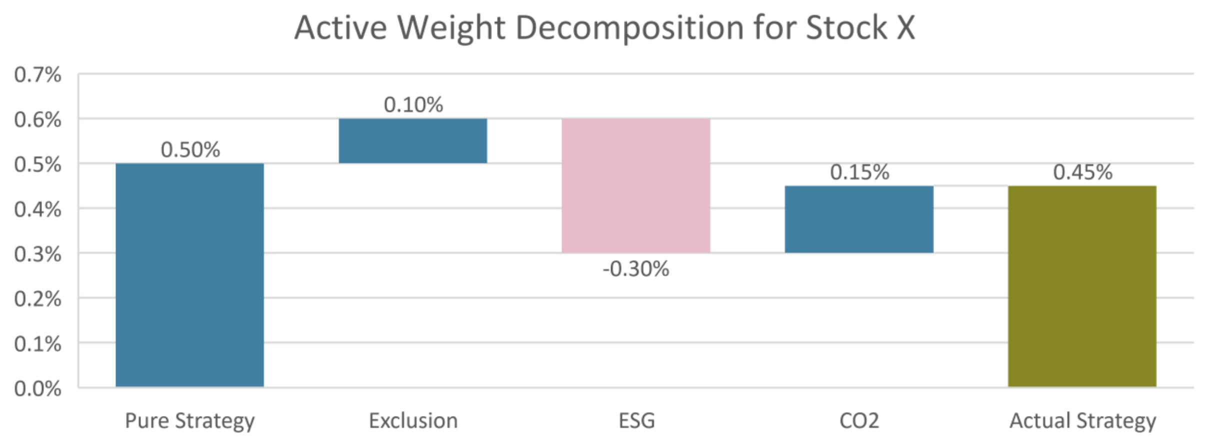

Instead of having one set of active weights for the entire portfolio, the active weights in this example are decomposed in four separate sets of active weights. For all of these individual sets, the sum over all stocks in the universe equals zero by definition. The sum of the decomposed weights is exactly equal to the actual strategy's active weights, also by definition. Hence the actual strategy can be seen as a pure strategy with an ESG overlay strategy, in this case consisting of three individual ESG sub strategies.

These active weights can then be used in a traditional Brinson-Fachler or factor-based attribution model to get insights into how different ESG considerations impacted portfolio performance. In the next chapter, these additional sets of active weights will be applied in a factor-based attribution model as this should give the best insights in the performance impact of ESG integration.

Note that many applications for Shapley values can be considered. First of all, the active weight decomposition can improve transparency of the portfolio construction step in the investment process. Such insights can also be used to assess how different restrictions (potentially also non-ESG related) behave or interact with each other and in case of active management, with the strategy. Another application could be risk attribution or tracking error decomposition and many more applications can be thought of.

5. Shapley Attribution – Results and Interpretation

In this chapter, a factor-based attribution analysis is conducted on the decomposed active weights, instead of the full active weights. This provides a more granular view of the impact of different portfolio components on performance. The Shapely attribution is applied to the Achmea IM Global Multi Factor Equity fund (GMF) for the period starting at 31-12-2021 and ending at 30-06-2023. Below is a list of the most important properties of GMF:

- Benchmark: MSCI World Index
- Alpha signal: Multi factor of value (sector-neutral), momentum and quality
- Max. active country and sector limit: 5%
- Max. tracking error: 5%
- ESG integration:
    - Exclusions: Controversial weapons producers, normative violators[5], tobacco producers, fossil fuel companies[6]
    - ESG enhancement, improvement relative to benchmark on:
        - MSCI ESG score
        - Sector-neutral thematic score for Climate theme[7]
        - Sector-neutral thematic score for Environmental theme[8]
    - Carbon reduction:
        - Until June 2022: 20% reduction of carbon emission intensity (CO2/SALES) relative to benchmark
        - From June 2022 onwards: pathway reduction to net-zero in 2050. Reduction in 2021 of 30% from 2020 levels. Then, for each year a 7% year-on-year reduction.

Applying the different sets of active weights to a factor-based attribution model leads to a multi-dimensional attribution, because active weights are no longer one dimensional. The table below summarises the result.

Table 7: GMF Shapley attribution over 2021 / 2023H1 period in bps

Factor	Strategy	Exclusions	ESG	CO_2	Total
Style Total	992.2	-45.3	-27.6	-9.1	910.2
Sector Total	-175.6	-56.8	2.3	-14.7	-244.8
Country Total	-133.6	-26.3	21.8	6.0	-132.1
Stock-Specific Total	10.3	-70.8	9.7	-25.3	-76.1
Grand Total	693.2	-199.2	6.3	-43.1	457.2

In the most right column of the table above shows the attribution results of the actual implemented strategy. These are exactly the same as in the example used for the traditional factor-based attribution, except that the different style factors are aggregated. The other columns display the factor- based results of the decomposed active weights and can be interpreted in the same way as with the normal factor-based attribution. For example, the choices in active weights that were made from the pure alpha strategy resulted in an exposure to the style factors that contributed 992.2 bps to the relative return. The bottom row shows the impact of the different active weight drivers. For example, in this period, exclusions had a negative attribution of 199.2 bps to the relative return of the portfolio.

This table has several advantages over traditional factor-based attribution. First, it clearly separates the choices made from a strategic or financial point of view from those from an ESG or sustainability perspective. This differentiation can be used to assess how a manager performed and whether ESG considerations had an impact

[5] UN Global Compact, OECD guidelines and UN Guiding principles on business and human rights.
[6] Companies with a exposure larger than 5% to coal, tar sands, shale gas and oil, Arctic gas and oil..
[7] Custom made
[8] Custom made

on performance. Second, it also gives insights into why ESG considerations contributed the way they did. For example, by looking at the ESG enhancements, one notices that the sector contribution is negligible. This is to be expected as ESG enhancements were integrated mostly sector neutral. By showing potential systematic sources in ESG performance impact, portfolio managers can review if the strategy is working as expected or perhaps requires adjustment.

Another way of displaying these attribution results is to show the total effects of the different portfolio choices over time, as depicted in figure 4 below. By plotting the attribution over time, it helps to see if there were certain events that might have led to the bulk of relative performance. In the case below one notices that during 2021, the impact of ESG considerations was negligible. However, since the Russian invasion of Ukraine and the sharp increase in energy prices thereafter, exclusions started to hurt more over the following months. As many weapon manufactures and fossil fuel companies are excluded, and both saw their profits and share prices rise sharply during this period, exclusions have had a negative impact on performance.

Figure 4: Shapley attribution GMF over time

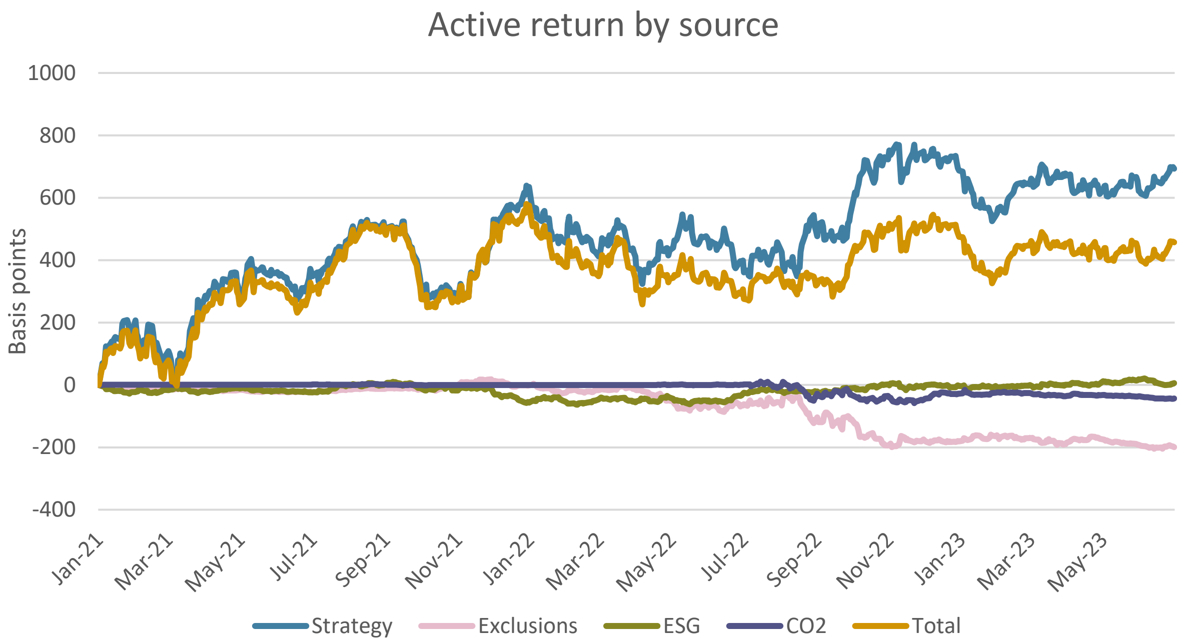

At first, the effect of exclusions on performance was limited. To some extent, this was because of the momentum strategy. Before the invasion started, momentum was positioned towards IT and health care and away from materials, energy and utilities which contain more of the excluded companies. This means that in the pure strategy, exclusions were mostly underweighted at the start of the invasion, which limited the impact of exclusions. However, as the materials, energy and utility sectors started to outperform, the momentum factor was shifting focus towards materials, energy, and utilities and away from IT and health care. In this scenario, exclusions started to hurt more because the pure strategy would have overweighted these firms.

At this point, it is interesting to reflect how the Shapley attribution differs from the more standard approaches. As discussed before, exclusion impact is often estimated by reweighting the index without the excluded names by market capitalisation, effectively redistributing the excluded index weights proportional to the size of a company. While this straightforward approach can serve as a proxy, it is missing out on possible interaction effects. To get a better idea of this error, the graph below compares both methods.

Figure 5: Exclusion performance impact comparison for GMF over 2021 / 2023H1

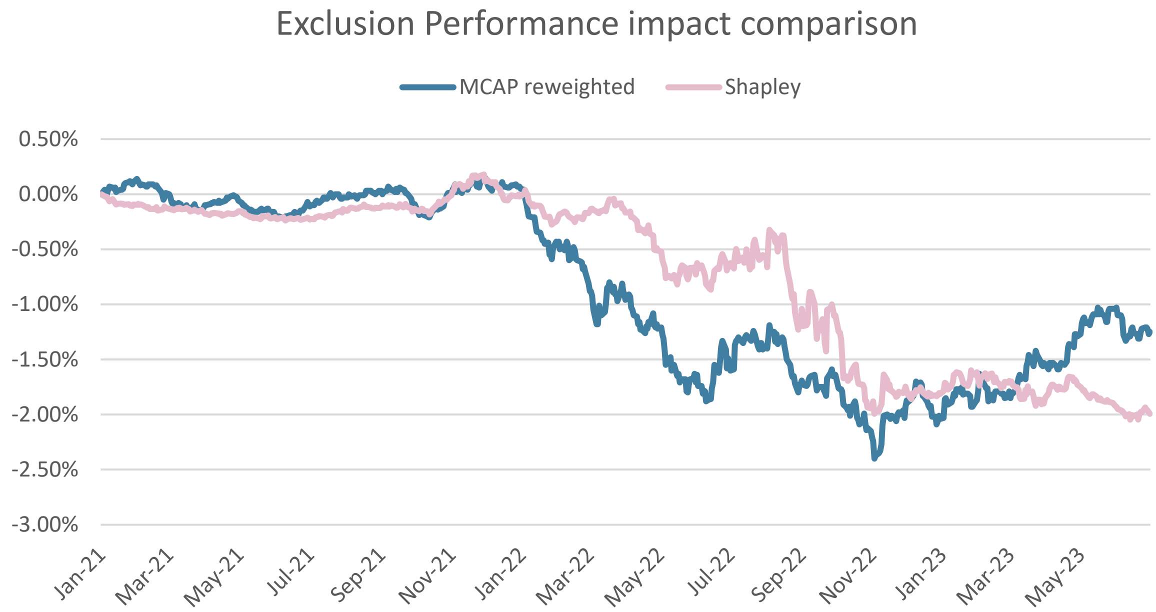

While both approaches show a somewhat similar impact, substantial deviations can arise due to the fact that interactions of exclusions with strategy or other ESG considerations are ignored with the market capitalisation reweighted approach. Following the Russian invasion, it first overestimated the exclusion impact since strategy and exclusions were actually aligned, leading to double counting. Later on, however, when the factors switched sectors, the strategy and exclusions starting misaligning. In this case, the exclusion impact was underestimated by the market capitalisation reweighting approach as it failed to account for the strategy's opportunity loss by not being able to overweight firms on the exclusion list.

Another interesting view is to see which sources where most dominant in the determination of the total active return[9]. The plot below shows how the dominance of this investment strategy shifted away from roughly 75% in 2021 to around 60% in the most recent year. This too can help to assess the performance of a manager. Especially if a manager has a limited tracking error budget, too little strategy dominance might be a reason to reconsider the tracking error limit to ensure that the strategy remains able to achieve its financial goals.

[9] Dominance in active return is, for each source, determined as:

$$Dominance\ of\ source\ j,t = \frac{|x_{j,t}|}{\sum_i |x_{i,t}|}$$

where $|x_{i,t}|$ is the absolute active return of source i at time t.

Figure 6: Active return of GMF explained by source

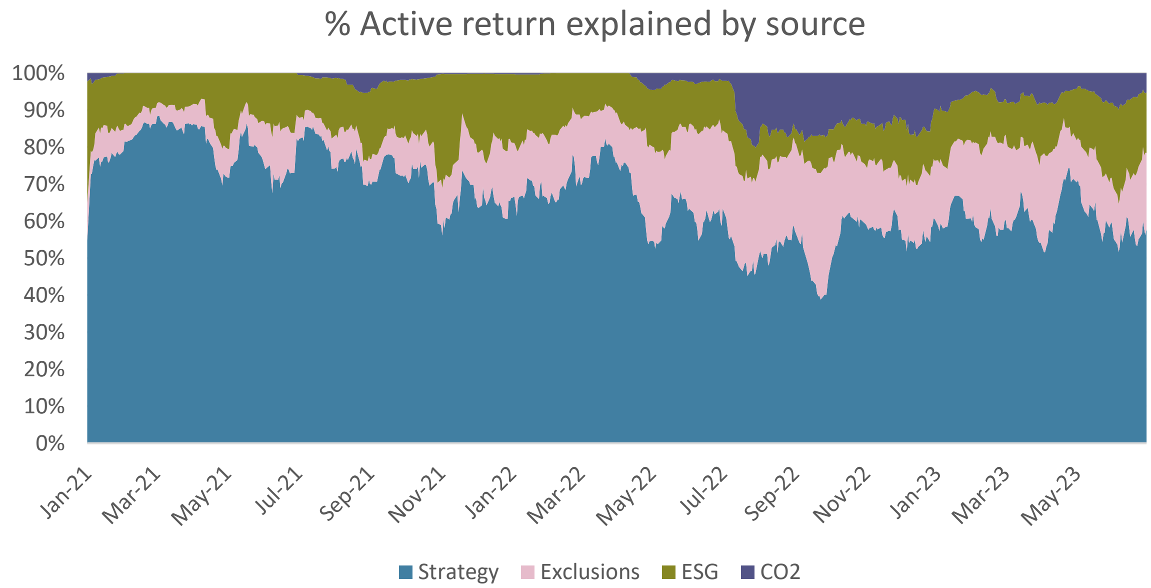

Another interesting effect visible from this figure is the CO2 performance impact. Before the Russian invasion, the momentum and quality factors were positioned to mostly low emitting sectors such as health care, communication services and IT. However, following the resulting energy crisis, the momentum factor shifted focus and the strategy shifted weight towards the more high emitting sectors, i.e., the pure strategy was no longer green enough by itself to satisfy the carbon reduction targets. At that point, the carbon reduction target started to make an impact as can be seen from the graph above.

6. Explaining performance differences of custom ESG indices

In recent years, many asset owners have chosen for a more passive implementation of their portfolio. This shift was largely the result of an increased focus on cost efficiency combined with a period of disappointing excess returns of many active strategies. Other advances in investments such as the development of exchange traded funds (ETFs) also facilitated this shift. This transition coexisted with the rise of ESG integration as discussed in the introduction of this paper. As a result, passive management started to embed more choices into it. Whereas traditionally, passive management often implied replicating a market capitalisation weighted index, nowadays, passive management often focusses on a more rule-based strategy in which deviations of market cap weighting are becoming the rule rather than the exception.

The investment space now offers unlimited custom indices that deviate from the traditional market cap weighting scheme. Many custom indices start from a traditional market cap weighted parent index and try to improve on certain characteristics while limiting the total deviation from the parent index. Such characteristics can be financial, like indices with a focus on high dividend paying stocks. But in the recent decade there has been a tremendous increase in the development of custom ESG indices in which ESG ratings or other ESG related metrics such as carbon footprint were the characteristics that were enhanced. This was accompanied by European regulation that gave guidelines on certain ESG indices, such as the EU Climate Transition Benchmarks Regulations act[10]. This regulation gave rise to the Paris Aligned and Climate Transition Aligned indices.

Intuitively, it might seem that performance attribution is redundant for a passive implementation of a custom (ESG) index. However, in most cases, the deviation of the performance of a custom index compared to its parent index is still of major interest. This performance difference provides information on the costs or benefit of pursuing the strategy of the custom index. Moreover, in most asset allocation studies, traditional, market cap weighted indices are used to determine strategic and tactical asset allocation. When the actual implementation is deviating from these traditional indices, performance differences arise that demand an explanation when evaluating the asset allocation.

At first, the performance difference between the custom and the parent index feels like a sufficient explanation of the custom strategy and hence, no advanced performance attribution is required as is the case for active management. However, when one takes a closer look, one notices that many of the custom indices, especially ESG related ones, contain many choices which can have different goals and are often accompanied with tracking error minimisations with respect to the parent index. When one only assesses the total performance difference between the custom and parent index, one does not obtain the required insights to evaluate what the performance impact of each choice was.

As an example, consider the popular MSCI World Climate Paris Aligned Index. It is a custom index based on the traditionally market cap weighted parent index, the MSCI World Index. The Paris Aligned index follows the requirements set by the EU and adds several other ESG related choices. Some are normative exclusions, for example tobacco companies. Others are carbon reduction targets to move towards a net-zero portfolio. The index also tries to reduce climate risk exposures and wants to exploit climate transition opportunities. Next to all of these ESG related choices, it also adds several restrictions to avoid concentration risk and tries to limit factor exposures and minimise tracking error relative to the parent index.

In total, a lot of active choices are made that result in deviations from the parent index. These deviations, in turn, result in excess returns which deserve to be understood and explained. However, by simply comparing the total return of the MSCI World Climate Paris Aligned index with its parent index, too little information is obtained about the performance impact of the different choices that were made, making it impossible to provide a proper explanation.

[10] https://finance.ec.europa.eu/regulation-and-supervision/financial-services-legislation/implementing-and-delegated-acts/eu-climate-transition-benchmarks-regulation_en

The figure below shows the cumulative index performance of both indices since inception of the MSCI World Climate Paris Aligned Index. During this period a underperformance of -4.6% (-1.6% annualised) was realised.

Figure 7: Comparison of cumulative index performance of MSCI ACWI vs parent index (MSCI World)

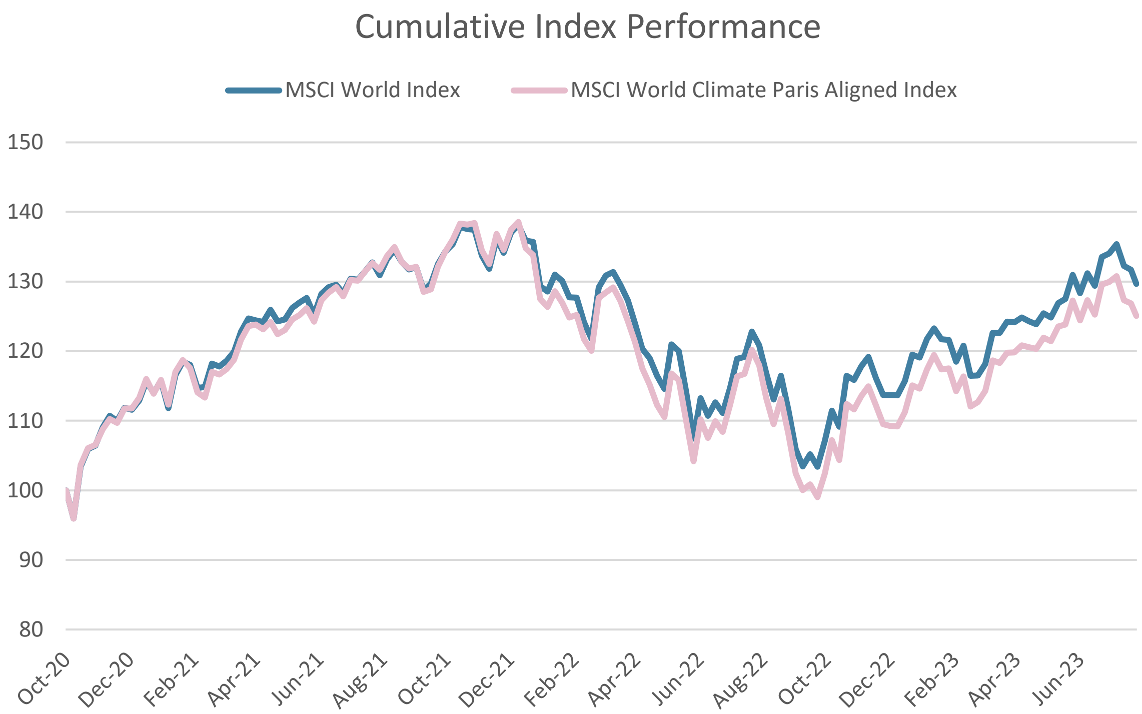

Using a traditional sector based Brinson-Fachler attribution shows that underperformance is largely explained by two sources:

- The allocation effect of the energy sector, -3.4%
- The selection effect within the industrials sector, -2.1%

While the Brinson-Fachler attribution gives some insights as to which exposures are responsible for this performance, it cannot be attributed to the different choices that were made in the portfolio construction process. For example, it is not clear whether the selection effect in the industrials sector was the result of the exclusions, increasing exposure on transition opportunities, reducing exposure to climate risks, reducing the carbon footprint, the tracking error minimisation, or, most likely, a combination of all. From the current attribution methods, one can only assess the total combined impact. This lack of information also means that it is impossible to evaluate whether climate risk reduction has paid off, or what the potential costs were of the exclusion policy.

Given the potentially substantial performance impact that custom (ESG) indices can have compared to their parent index, a more granular performance attribution is required. Luckily, due to the systematic nature of the index construction methodologies, the Shapley attribution could also be applied to custom indices to obtain these insights.

7. Conclusion

This paper started with the notion of the growing importance of ESG considerations in portfolio construction. Besides the traditional financial factors, non-financial factors such as exclusions, ESG tilting and carbon emission reduction are playing a more decisive role in portfolio construction. This implies that the impact of these non-financial factors on performance is increasing. However, investors and asset owners are struggling to accurately measure the impact of separate ESG considerations.

The paper assessed the first generation of attribution models, the Brinson-Fachler attribution, and investigated the second generation of attribution models, the factor-based attribution. Both attribution methods are limited in assessing the impact of ESG considerations, especially when multiple ESG choices are made and/or applied simultaneously with an active investment strategy. Most notably, both models are based on outputs, or portfolio exposures, but lack additional insights into which inputs lead to those outputs or exposures.

To overcome these issues in performance attribution, the technique of Shapley values is applied to a portfolio construction process. This technique provides a decomposition of the active weights of a portfolio. Not only does this improve the understanding of how the portfolio construction works, it also offers the option to add another dimension to the traditional attribution models. This development extends the factor-based models and can be referred to as Shapley attribution.

Shapley attribution successfully separates the impact of different portfolio choices, both financial and non-financial. In doing so, it accurately accounts for interactions between factors. The insights from this attribution method can be used to obtain a better understanding of the impact of ESG integration on performance. In doing so, it can also aid to more accurately evaluate an active manager’s performance by separating the impact of strategy from ESG considerations.

But the use is not limited to active management alone. With a steady rise in assets under management of both passive investment approaches and ESG investments, passive implementation of custom ESG indices can result in substantial deviation from the parent index performance. Given that many custom ESG indices contain many different design choices for different reasons, it is impossible for current attribution techniques to assess these choices separately. Here too, Shapley attribution could be used to improve the understanding of the custom index performance relative to its parent index.

This paper proposes Shapley attribution to shed more light on the impact of ESG integration, but its use is not limed to just ESG integration. Other portfolio restrictions such as tracking error limits, or country and sector limits can be investigated as well. While the Shapley attribution is computationally and data technically more demanding, it offers considerable insights that are otherwise missed out on. Given the growing trend of ESG integration, Shapley based attribution might well be the foundation for the next generation of performance attribution models.

8. Appendix A – Brinson-Fachler Attribution

Next follows a brief instruction and example on how to compute a Brinson-Fachler attribution. This is a stylised example which uses three sectors:

1. Gather portfolio and benchmark weights ($W_{p,i}$ and $W_{B,i}$) and portfolio and benchmark returns ($R_{p,i}$ and $R_{B,i}$) for each sector i. Also compute the total portfolio and benchmark return (R_p and R_B).
2. For each sector i, calculate the allocation effect as:
$$(W_{p,i} - W_{B,i}) * (R_{B,i} - R_B)$$
3. For each sector i, calculate the selection effect as:
$$(W_{B,i}) * (R_{P,i} - R_{B,i})$$
4. For each sector i, calculate the interaction effect as:
$$(W_{p,i} - W_{B,i}) * (R_{P,i} - R_{B,i})$$

Table 8: Stylised example of typical sector-based Brinson-Fachler attribution

Sector	Portfolio Weight	Benchmark Weight	Active Weight	Portfolio Return	Benchmark Return	Allocation	Selection	Interaction	Total
Materials	25.00%	20.00%	5.00%	6.00%	8.00%	0.21%	-0.40%	-0.10%	-0.29%
Industrials	25.00%	15.00%	10.00%	7.00%	7.00%	0.33%	0.00%	0.00%	0.33%
Energy	25.00%	25.00%	0.00%	-4.00%	-2.00%	0.00%	-0.50%	0.00%	-0.50%
Financials	25.00%	40.00%	-15.00%	3.00%	4.00%	-0.04%	-0.40%	0.15%	-0.29%
Total	100.00%	100.00%	0.00%	3.00%	3.75%	0.50%	-1.30%	0.05%	-0.75%

9. Appendix B – Factor-Based Attribution

Next follows a brief instruction and example on how to compute a factor-based attribution. This process follows these steps:

1. Define factor scores for all stocks within the investment universe, ensuring that the factors are normalised to have a mean of zero. These factors can encompass various style factors, as well as sector, country, or currency classifications. The resulting matrix, denoted as *S*, containing *N* rows representing each stock, and *K* columns for each systematic factor.
2. Multiply the factor scores matrix *S* with the portfolio active weight vector, AW, to obtain the portfolio factor exposures, X:

$$X = AW' * S$$

3. Compute the pure factor returns, which represent the returns of factor mimicking portfolios (FMPs). The primary objective of these FMPs is to exhibit maximum correlation with the original factor while neutralising any exposure to other factors. Using FMPs can isolate returns and avoid, for example, that the return of the value factor is partially explainable by quality exposure. To compute FMPs, orthogonalise the matrix S and multiply it with a vector of stock returns, R:

$$FMP = ((S' * S)^{-1} * S') * R$$

4. To obtain the factor contributions, C, multiply the portfolio factor exposure, X, with the pure factor returns or FMP returns, FMP:

$$C = X * FMP$$

5. Finally, calculate the stock-specific contribution, referred to as "STSP," which represents the portion of relative return that cannot be attributed to a systematic factor. This is done by subtracting all systematic contributions from the total active return:

$$STSP = AW' * R - \mathbf{1}' * C$$

where, $\mathbf{1}$, represents a vector of ones with length equal to the number of systematic factors, K.

10. Appendix C – Turnover adjustments

This appendix first describes why and how turnover constraints complicate a fair calculation of all permutations and then describes viable solutions that can be used.

When a strategy has a turnover limit, running permutations with certain ESG considerations inactivated might run into a problem. This is because the current portfolio, which includes all the ESG considerations, is taken as a starting point. This starting point, however, might not accurately reflect what the previous portfolio would have looked like under a different set of ESG considerations. The starting point does not matter if turnover is not considered, but if it is considered, it complicates matters.

Take exclusions as an example. If a substantial part of the benchmark is excluded, say 10%, the portfolio permutation without exclusions will likely buy back part of that 10% that was excluded in the previous portfolio. The exact amount will be depending on other portfolio properties such as the alpha signal of the excluded companies or the tracking error limit. Quite realistically, this would lead to a trade that is predominantly buying back exclusions, instead of giving an idea of what the portfolio would have looked like in an alternative scenario where no exclusions would have been considered.

To illustrate the challenge of a turnover constraint on correctly calculating what the portfolio permutations would have looked like, consider the graph below.

Figure 8: Visualisation of the turnover constraint problem

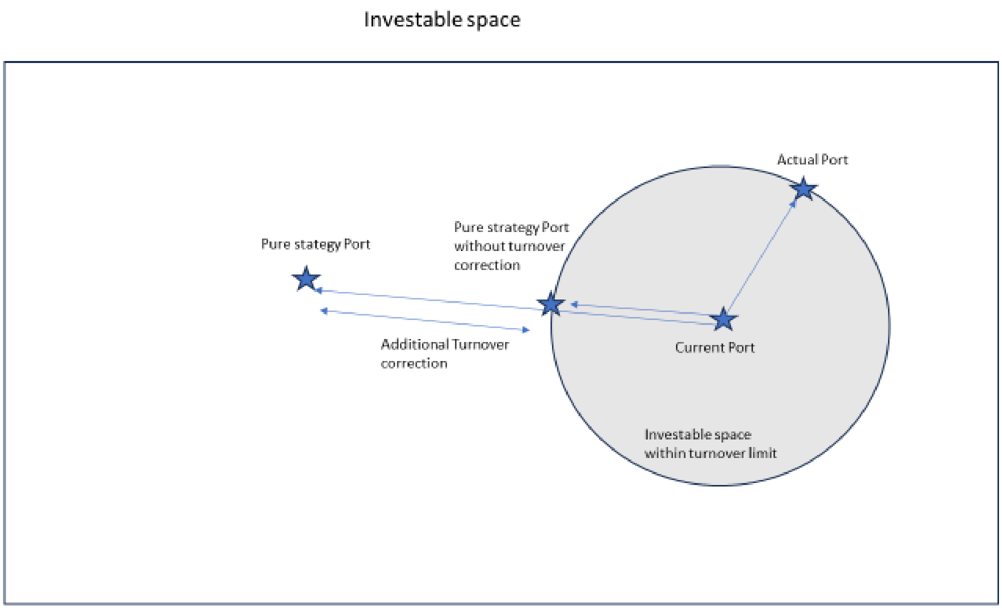

This graph is a visual representation of the investable space of a portfolio. Note that this example is stylised, the axes have no obvious meaning. Two portfolios close together on the graph imply similar portfolios, i.e. similar active weights. All portfolios within the outer square satisfy the constraints of the least restricted portfolio. When dealing with a turnover constraint, the rebalanced portfolio can only move a certain amount of distance from the current portfolio, as represented by the grey circle around the current portfolio. The problem that arises is that the different portfolio permutations face the same turnover constraint. Take the pure strategy for example, which has no ESG considerations. The pure strategy portfolio without a correction to the turnover constraint, ends up within the grey circle and has used a significant part of its turnover budget to buy back previously excluded stocks, instead of maximising factor exposure. This is a poor reflection of what the "true" pure strategy would have looked like, because it would not have had to deal with previously excluded stocks and hence could have allocated more turnover to maximising factor exposure and thus end up differently.

An appropriate correction to the turnover budget should be given to these permutations to reflect the additional turnover that is needed to correct for deactivated ESG considerations that were still active in the determination of the current portfolio. If this is neglected, the portfolio permutations will seem to be closer to the actual portfolio due to the turnover constraint, while in reality, they would be further away from the actual portfolio. In terms of Shapley values, this would imply an underestimation of the impact of the ESG considerations on the determination of the actual active weights.

The next part describes how this turnover correction is determined. To make appropriate adjustments for this issue, several options are investigated:

1. Adding the turnover constraint as a separate factor to the attribution analysis
2. Using the benchmark as the current portfolio and take the current active weight as turnover constraint
3. Adjusting the turnover constraint for the impact each separate factor has on the actual strategy

If the turnover constraint is added as a separate factor, the Shapley attribution will show the extent to which this constraint is determining the active weight of a stock. While this looks like a step in the right direction, it opens the door to further extensions of the model. Because, why should country, sector or stock limits not be separate factors? With every factor that is added, the matrix listing the permutations grows exponentially. As the matrix uses active weights of each stock in the universe, it increases the computational load significantly. At some point, this reaches the limits of what computationally can be achieved. Moreover, restrictions such as turnover can be seen as part of the strategy and therefore there would be no need to sperate the impact from the turnover limit from other strategic choices. For these reasons, this approach is considered suboptimal.

The second solution would be to take a level the playing field by taking a similar starting point for the portfolio permutations and giving a comparable amount of turnover that the actual strategy has. The following steps can be taken in the setup of the parameters for the different permutations:

1. Calculate the Active Share of the current portfolio
2. Adjust the turnover constraint to this number
3. Run the permutations using benchmark constituents and weights as the current portfolio

The issue with this approach is that the active weights of the actual portfolio are the result of multiple optimisations through time. During each period, the actual portfolio is often restricted by the turnover constraint during rebalances and hence only executes part of the intended trades at each rebalancing and so the actual strategy moves slowly towards its unlimited version. By allowing the permutations to use the total turnover in a single optimisation, the model will come up with portfolio compositions that realistically would not have been possible if a turnover constraint was used at each separate rebalancing. This can be verified by giving the actual strategy the active share as turnover constraint, this leads to a significantly different portfolio from the actual strategy that has been computed over time with a more stringent turnover limit.

A third solution would be to account for the impact that the factors have on the active weights of the actual strategy. In doing so, there is a correction for the likely transactions that will occur once a factor is deactivated, without limiting the transactions as a result of the other factors. This requires an estimation of how much of the actual strategy's active weight is determined by each factor.

The following approach captures which part of the positioning of the actual strategy is determined by each source of interest. This is done by running a regression over all stocks i, of the actual active weights of the portfolio on the alpha signal and the ESG elements:

$$Active_i = \beta_1 * Alpha_i + \beta_2 * EXCLUSIONS_i + \beta_3 * ESG_i + \beta_4 * CO2_i$$

In this regression the dummy variable of exclusions is multiplied with benchmark weight. Also, all features are normalised, which sets the intercept to zero since the average active weight is zero by definition. The coefficients of this regression give an idea how active weight is built up on average. Even though for individual stocks this might be a somewhat poor estimate of active weight decomposition, on the portfolio level it gives a good and unbiased estimate which is sufficient for the goal of adjusting the turnover limit at the portfolio level.

The following calculation computes how much additional turnover is allowed for each permutation.

$$Turnover\ addition = f(Exclusions, ESG, CO2) =$$
$$= \sum_i |\, \mathbb{1}_{Exclusions} * \beta_2 * EXCLUSIONS_i + \mathbb{1}_{ESG} * \beta_3 * ESG_i + \mathbb{1}_{CO2} * \beta_4 * CO2_i |$$

Where, $\mathbb{1}$ represents the indicator function, which is one if the specified ESG consideration is activated and is zero otherwise.

For portfolios where the goal is tracking error minimisation or minimum variance instead of alpha generation, the alpha signal can be adjusted to features that are better in estimating such a portfolio's build up. Pairwise correlation and/or (idiosyncratic) volatility, obtained from the risk model that is used in the portfolio construction process could be of assistance.

Disclaimer

Achmea Investment Management B.V. ("Achmea IM") has compiled the information contained in this document with the utmost care. This information is only intended for qualified investors and/or professional investors as defined in the Dutch Act on Financial Supervision ("Wft"). This document is or constitutes (i) no offer or invitation to purchase, sell or trade financial instruments (within the meaning of Article 1:1 Wft), (ii) no recommendation or investment advice (within the meaning of Article 1:1 Wft), (iii) no research and (vi) no legal or tax advice. Please consult, where appropriate, a legal adviser or tax consultant. Achmea IM advises you not to base your (investment) decision entirely on the information contained in this document. Achmea IM will not be liable for any loss or damage that may result from such an (investment) decision.

The information in this document is (partly) based on information which Achmea IM has acquired from information sources that are considered reliable. Achmea IM does not guarantee the reliability of the information acquired from those sources. The information in this document is for information purposes only and is not intended to confer any legal rights on you. All information reflects a particular point in time, unless explicitly stated otherwise. The provision of this document after the original publication date is no guarantee that the information contained herein is still accurate and complete on such later date. Achmea IM has the right to change this information without notice. Any reference to returns on investments is for information and clarification purposes only. In doing so, Achmea IM does not make a forecast for the return on or share price movements of those investments. The value of your investment may fluctuate and results achieved in the past are no guarantee for the future. The information in this document is based on assumptions made by Achmea IM. It is possible that other assumptions will be made than the assumptions made by Achmea IM. The conclusions drawn by Achmea IM are therefore not necessarily correct and/or complete.

Published by:

Achmea Investment Management B.V., with its registered office and place of business in Zeist (Trade Register no. 18059537), The Netherlands. Achmea IM is registered with the Netherlands Authority for the Financial Markets ("AFM") in Amsterdam (AFM registration 15001209) as manager of alternative investment funds as defined in Article 2:65a of the Wft and is authorised under this licence to provide investment services as listed in the register maintained by the AFM.