The development of an ESG attribution model

Robbert Lammers, Sam Radford and Vincent Verkerk

ABSTRACT

Environmental, social, and governance (ESG) considerations have become crucial in the investment landscape, as sustainable practices and responsible corporate behavior increasingly influence financial performance. Despite this growing importance, there is a lack of comprehensive models that integrate and attribute ESG performance, such as CO_2 emissions, leaving a gap in understanding the impact of investment decisions on ESG outcomes.

This article describes the development of an ESG attribution model. The model quantifies the influence of investment decisions on ESG factors and seeks to close the feedback loop in the investment decision-making process by combining ESG data with traditional financial metrics, enabling a holistic analysis of both financial and non-financial outcomes.

The methodology involves expanding the Brinson-Fachler method (Brinson & Fachler, 1985) and the Investment Decision Process (IDP) model (Geenen et al., 2021) to include non-financial elements such as carbon emissions and ESG scores. Case studies demonstrate the model's practical application, highlighting its potential to enhance investment decision-making and promote sustainable investing practices.

The main takeaway is that incorporating ESG factors into the investment decision-making process is essential for aligning investments with broader societal goals and improve decision-making.

INTRODUCTION

In recent years, environmental, social, and governance (ESG) considerations have gained significant prominence in the investment landscape. Investors now recognize that sustainable practices and responsible corporate behavior can impact financial performance. Consequently, integrating ESG factors into investment decisions has become a critical aspect of portfolio management (Steehouwer, 2023).

At the same time, financial institutions, once primarily focused on maximizing returns, now acknowledge that their investment choices affect the environment, society, and communities. Whether it's carbon emissions, labor practices, or board diversity, their investments wield influence.

CLOSING THE FEEDBACK LOOP: INCLUDING ESG ATTRIBUTION INTO THE INVESTMENT DECISION-MAKING PROCESS

To navigate this evolving landscape, financial institutions must gain deeper insights into the effects of the investment decisions they make. Traditionally, the investment decision-making process starts with ex-ante expectations of the future, shaping asset allocation and security selection. Performance attribution aims to complete a feedback loop by showing the ex-post realization of the portfolio's performance, allowing investors to analyze their decisions and identify missing variables.

Whilst traditional performance attribution has sought to explain portfolio outperformance in terms of returns and/or risk, investors increasingly consider factors beyond financial return (Global Sustainable Investment Alliance, 2022). For example, investors may seek to limit the carbon emissions of their portfolio, choose not to invest in regions with poor human rights

Robbert Lammers
Advisor Responsible Investment, MN

Sam Radford
Performance Consultant, Ortec Finance

Vincent Verkerk
Performance Analyst, MN

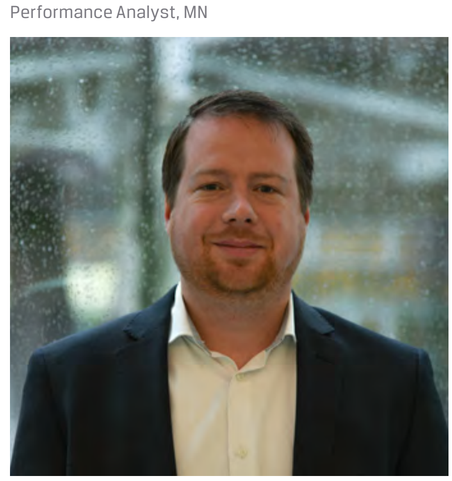

records, or avoid investing in companies which engage in child labor or have a low representation of women amongst the board.

If a portfolio is managed with the aim to keep financed emissions below the financed emissions of the benchmark, investors must analyze the impact their decisions have on both financial performance and emissions. This ensures that the feedback loop between ex-ante expectations and ex-post realization continues to enhance the investment decision-making process.

ENHANCING INVESTOR ACCOUNTABILITY: THE ROLE OF ESG ATTRIBUTION IN REGULATORY COMPLIANCE

Alongside the need for an attribution model to help improve investor decision making, regulators increasingly monitor the claims of investors relating to ESG; the EU's Sustainable Finance Disclosure Regulation mandates ESG disclosure for funds claiming to include ESG into their investment decision-making process. A robust attribution model allows investors to prove the positive impact of their decisions on targeted metrics, helping meet regulatory requirements.

EXPANDING THE MODEL: ADDING NON-FINANCIAL ELEMENTS TO AN EXISTING ATTRIBUTION MODEL

The classic Brinson-Fachler method dissects the total excess return of the portfolio relative to the benchmark into allocation and selection (and interaction) effects. This model was expanded on to create the Investment Decision Process (IDP) model, which "stacks" several Brinson-Fachler schemes. Each decision within the investment decision process should be included in the attribution and collectively explain the total excess performance of the portfolio against its primary benchmark.

INTEGRATING ESG FACTORS HAS BECOME A CRITICAL ASPECT OF PORTFOLIO MANAGEMENT

Recognizing the need to understand how ESG factors interact with investment decisions, we developed an early stage ESG attribution model. This model expands upon the existing IDP attribution model by adding a non-financial element to the model. The model can be applied to a variety of different metrics, including carbon emissions, environmental, social and governance scores, biodiversity or water usage scores.

METHODOLOGY

The proposed model serves various use cases; one is discussed in detail here whilst other use cases are discussed briefly in the Future Extensions section. The first use case involves comparing a single mandate against a benchmark, considering different decisions made within the portfolio.

A MULTI-BENCHMARK APPROACH WITH THE IDP MODEL

Consider a straightforward equity portfolio that is compared against a regional equity benchmark; the **Manager Portfolio**. The portfolio manager receives an exclusion list with benchmark securities they cannot invest in and a mandate not to allocate more than x% to any single security. The exclusions are predominantly ESG driven, while the x% cap aims to increase diversification and reduce risk.

In the IDP model approach, we treat this as three distinct decisions, each associated with a different benchmark:

- **Original Benchmark:** The original client mandated benchmark
- **Exclusion Benchmark:** Derived from the Original Benchmark, applying predominantly ESG driven exclusions reallocating weight on a pro-rated basis across other securities in the benchmark
- **Capped Benchmark:** Based on the Exclusion Benchmark, capping the maximum weight of each security at x% and reallocating excess weight on a pro-rated basis across other securities in the benchmark

For attribution analysis:

- We compare the Original Benchmark to the Exclusion Benchmark, isolating the impact of the decision to exclude the securities within the exclusion list; this can be considered a pure allocation effect.
- Next, we compare the Capped Benchmark to the Exclusion Benchmark, revealing the impact of the x% weight cap.
- Finally, we assess the Manager Portfolio against the Capped Benchmark, capturing both the allocation and selection effect resulting from the manager's decisions.

This approach allows us to disentangle the impact of the mandated exclusion while ensuring a "fair" comparison between the manager and a similarly constrained benchmark. The IDP model's flexibility even allows for variations in decision order (e.g., swapping the Capped Benchmark and Exclusion Benchmark).

Figure 1
Hierarchy for comparing Manager Portfolio against Original Benchmark with multiple decisions

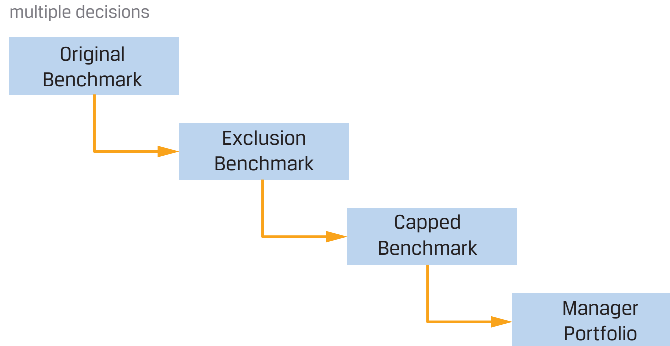

Whilst classic performance attribution explains excess financial returns, this article suggests expanding the model to also include the impact of decisions on non-financial metrics. Portfolios are often constructed considering both financial and non-financial factors, such as selecting companies with the lowest carbon emissions or excluding companies that produce products that are principally undesirable, like tobacco.

Given that investment decisions rely on both financial and non-financial factors, it is crucial to understand their effects in both domains. Measuring the impact of decisions on the intended metrics is essential to improve the decision-making process.

AGGREGATING NON-FINANCIAL METRICS

To compare aggregate non-financial factors for each of the benchmarks and the manager portfolio, various non-financial metrics were calculated by aggregating security-level data, using asset weights. This article uses five non-financial metrics:

- **Absolute financed emissions** (based on the enterprise value including cash; EVIC): the greenhouse gas emissions associated with the companies in an investor's portfolio, proportionally based on how much of these companies' activities are financed by the investor.
- **Emissions intensity** (based on revenue): the volume of greenhouse gas emissions produced per unit of revenue.
- Environmental score: a composite score derived from assessing companies' performance across various environmental issues, including carbon emissions, land use, and toxic waste.
- **Social score:** a composite score derived from assessing companies' performance across various social issue, including supply-chain labour standards, community relations, and workplace safety.
- **Governance score:** a composite score derived from assessing companies' performance across various governance issues, including board independence, renumeration, and shareholder rights.

WE SUGGEST EXPANDING ATTRIBUTION TO INCLUDE NON-FINANCIAL METRICS

Each security in the portfolio and benchmark receives a score for each of these metrics (some data was unavailable, which is discussed further in the methodology under "The Data Problem"). By combining security level weights and metrics, a total metric is calculated for each benchmark and the portfolio.

The remainder of the methodology section refers specifically to absolute financed emissions for ease, but the same methodology was applied to all other metrics. The formula used to calculate the financed emissions of the portfolio is displayed in equation one whilst the financed emissions of the Original Benchmark is shown in equation two.

Equation 1
Financed emissions of the Manager Portfolio

$$\varepsilon^{P} := \sum_{i=1}^{n} w_i^{P} \times \varepsilon_i$$

Equation 2
Financed emissions of the Original Benchmark

$$\varepsilon^{OB} := \sum_{i=1}^{n} w_i^{OB} \times \varepsilon_i$$

Where:

- ε^{P} is the financed emissions of the Manager Portfolio
- ε^{OB} is the financed emissions of the Original Benchmark
- w_i^{P} is the weight of security i in the Manager Portfolio
- w_i^{OB} is the weight of security i in the Original Benchmark
- ε_i is the financed emissions of a security i
- n is the total number of securities in the portfolio or benchmark, respectively

The financed emissions of the Exclusion Benchmark (ε^{EB}) and of the Capped Benchmark (ε^{CB}) are calculated using the same formula as equation two but adjusted to use the weight of securities in the respective benchmark. The analysis is done annually and for w_i^{P} and w_i^{OB} the average weight of the security throughout the year is used. Average weights are used because ideally, we would prefer to use daily emissions data combined with daily weights, however, emissions data is currently only available on an annual basis and therefore using the annual average weight is the closest methodology available.

The total outperformance in terms of financed emissions can then be shown using equation three, whilst this can also be split into smaller parts as shown in equations four to seven.

Equation 3
Total financed emissions excess performance

$$\varepsilon^{XS} := \varepsilon^{P} - \varepsilon^{OB}$$

Equation four represents the total impact of exclusions, calculated as the difference in absolute financed emissions between the Exclusion Benchmark and the Original Benchmark (in the Results section this is shown as "Due to exclusions").

Equation 4
Excess financed emissions attributable to the Exclusion Decision

$$\varepsilon^{XS}_{Excl} := \varepsilon^{EB} - \varepsilon^{OB}$$

Equation five represents the total impact of the security level weight cap, calculated as the difference in absolute financed emissions between the Capped Benchmark and the Exclusion Benchmark (in the results section this is shown as "Due to cap").

Equation 5
Excess financed emissions attributable to the Cap Decision

$$\varepsilon^{XS}_{Man} := \varepsilon^{P} - \varepsilon^{CB}$$

Equation six represents the total impact of manager decisions, calculated as the difference in absolute financed emissions between the Manager Portfolio and the Capped Benchmark (in the results section this is shown as "Due to manager choices").

Equation 6
Excess financed emissions attributable to the Manager's Decisions

$$\varepsilon^{XS}_{Man} := \varepsilon^{P} - \varepsilon^{CB}$$

Equation 7
Financed emissions Excess Performance Decision Decomposition

$$\varepsilon^{XS} := \varepsilon^{XS}_{Excl} + \varepsilon^{XS}_{Cap} + \varepsilon^{XS}_{Man}$$

MEASURING THE IMPACT OF EXCLUSION FOR DIFFERENT REASONS

Within the Exclusion Benchmark, securities can be excluded for various reasons. Besides assessing the overall portfolio-level and security-level impacts of each exclusion, it's valuable to understand the impact at the exclusion reason level. For example, excluding companies that derive a high percentage of revenue from energy production using thermal coal is expected to reduce the portfolio's financed emissions. In evaluating an investment decision, it is desirable to understand the actual impact of the decision on the metrics that were intended to be impacted.

IT IS DESIRABLE TO UNDERSTAND THE IMPACT OF DECISIONS ON METRICS THEY WERE INTENDED TO IMPACT.

To understand each exclusion's impact, each security is assigned a single exclusion reason. While this prevents considering a security as excluded for multiple reasons, it ensures that contributions sum up appropriately. One methodology considered was to create a series of exclusion benchmarks, with each benchmark excluding more securities than the previous one. For example, moving from the Original Benchmark to the Exclusion Benchmark could be split into several decisions:

- Original Benchmark
  - Original Benchmark ex Coal
  - Original Benchmark ex Coal and Oil & Gas
  - Original Benchmark ex Coal and Oil & Gas and Low % of Women on the board
- Exclusion Benchmark

However, this methodology was tested and found to exaggerate the effects of the later decisions due to the re-scaling of weights within each benchmark.

Therefore, if it is desirable to be able to show the impact of each reason for exclusion (rather than just a total exclusion effect) it is required to use a different methodology than creating a series of benchmarks. Instead, a single Exclusion Benchmark containing all exclusions (where exclusions are mutually exclusive, and a security can only be excluded for one reason) is constructed and compared against the Manager Benchmark.

Each security in the exclusion list is tagged with a specific exclusion reason. The weight of each exclusion reason in the exclusion list can then be calculated using the formula in equation eight.

Equation 8
The total weight of Exclusion Reason y in the Original Benchmark

$$w^{OB}_{y} := \sum_{i=1}^{n} w^{OB}_{y_i}$$

Where:

- w^{OB}_{y} is the weight of the securities excluded for reason y in Exclusion Benchmark relative to the Original Benchmark
- $w^{OB}_{y_i}$ is the weight of a security i which is excluded for reason y from the Exclusion Benchmark relative to the Original Benchmark

The total contribution to Original Benchmark of securities excluded for each reason can then be calculated using the formula in Equation nine.

Equation 9
The contribution of an exclusion reason to the Original Benchmark's financed emissions

$$CB^{OB}_{y} := \sum_{i=1}^{n} w^{OB}_{y_i} \times \varepsilon_i$$

Where:

- CB^{OB}_{y} is the contribution to the Original Benchmark's financed emissions of all securities excluded for a specific reason

Due to the re-scaling of weights in the Exclusion benchmark, the attribution effect of excluding each set of securities cannot simply be equal to the contribution to the Original Benchmark. In other words, the non-excluded securities have their weights increased (as the weight of excluded securities is assigned to the allowed securities on a pro-rated basis) and so the Exclusion Benchmark financed emissions is not simply equal to the contribution of non-excluded securities in the Original Benchmark but will be increased by an amount equal to the impact of assigning this weight to the allowed securities.

Equation 10
The financed emissions contribution in the Original Benchmark of all securities that are not excluded from the Exclusion Benchmark

$$CB^{OB}_{NonExc} := \sum_{i=1}^{n} w^{OB}_{NonExc} \times \varepsilon_i$$

Where:
- CB^{OB}_{NonExc} is the contribution to the Original Benchmark's financed emissions of all securities that are not excluded for any reason

Therefore, to calculate the attribution effect of each exclusion reason, the formula shown in equation eleven is used:

Equation 11
The attribution effect in the IDP model for excluding securities for a given reason

$$\varepsilon^{XS}_{y} := w^{OB}_{y} \times CB^{OB}_{NonExc} - CB^{OB}_{y}$$

Where:
- ε^{XS}_{y} is the attributable change in financed emissions due to excluding securities for reason y

THE DATA PROBLEM

Data for traditional performance attribution (holdings, transactions, exchange rates etc.) is readily available and verifiable, allowing reliable performance attributions. However, non-financial data needed for our proposed model often lacks completeness and quality. For example, company level emissions data is not always reported, and when it is, it may lack reliability or consistency. New regulations, like the EU's Corporate Sustainability Reporting Directive, will compel companies to produce more and better data. Over time, we expect both the completeness and quality of this data to improve. In the meantime, we must use of "low quality" data and in some cases use proxies to estimate metrics for some companies.

Non-financial data is also often subjective and measured relative to peers. This means that if all securities improve equally, their scores may remain unchanged due to the relativistic methodology.

USING A PROXY FOR MISSING DATA

We propose a solution to the problem of data coverage, while acknowledging the availability of other options. For the purposes of this article, when data was missing for a security, a proxy was constructed using the following hierarchy:
1. Average score of other companies in the same sector
2. Average score of other companies in the universe (benchmark)

Using imperfect estimates may skew results, as especially poor performers have an incentive not to report this data. A data quality score has been calculated as part of the result to inform on the reliability of the underlying data. The reported score combines a quality score for available data (scored 0-100) with a score for estimated data (always scored 0). Excluding companies that do not report the necessary data, results in a portfolio with higher data quality score and more reliable results.

Ultimately, the results convey the correct story. They should however be interpreted as estimates rather than precise results. One possible extension, that this article does not discuss further, is to build confidence intervals around the results based on the data quality of the input data to each calculation.

RESULTS

Data for two global equity portfolios was used over a three-year period (2021-2023), benchmarked against a standard developed markets equity index. Both portfolios had an ESG driven exclusion list, and a financially driven x% security weight cap as explained in the methodology. While sharing the same Original Benchmark, each portfolio had its own separate Exclusions Benchmark and Capped Benchmark based on their own exclusion list. All non-financial data (carbon emissions, environmental/social/governance scores etc.) are sourced from major global ESG data providers.

Both portfolios are constructed by applying a series of exclusions to the benchmark and imposing a weight cap on individual securities. The key divergence lies in the details of the investment philosophy of the two portfolios; Portfolio One includes more extensive exclusion rules, whilst Portfolio Two emphasizes engagement. This results in different sets of excluded securities. This can be noted in the results as we see larger contributions from exclusions in Portfolio One versus Portfolio Two.

ATTRIBUTING ABSOLUTE FINANCED EMISSIONS

Table 3 presents the 2021 results for Portfolio One. The top row shows the results for Original Benchmark (in this case a Global Equity Benchmark) whilst each of the columns show:
- **Return:** The annual return of the global equity benchmark
- **Scope 1 $tCO_2e/EVIC$:** Total absolute financed scope 1 carbon emissions attributable to the benchmark. Scope 1 emissions are the direct emissions caused by a company's own activities. For example, the emissions of a company's fleet of trucks. Both the scope 1 and 2 emissions include the CO_2 equivalents methane, nitrogen oxides, hydrofluorocarbons, perfluorocarbons and sulphur hexafluoride. This scope 1 $tCO_2e/EVIC$ is calculated by first calculating a benchmark "absolute financed scope 1 emissions" (where financed emissions are based on enterprise value including cash) using the formula shown in equation two. This is then multiplied by the total market value of the portfolio to calculate the total tonnes of carbon attributable assuming all money was invested in the Original Benchmark.
- **Scope 2 $tCO_2e/EVIC$:** The same as above but using scope 2 emissions instead of scope 1 emissions. Scope 2 emissions

are the indirect emissions, for example from electricity purchased by a company. They occur at distant power plants but are company's responsibility.

- **Scope 1+2 tCO_2e/EVIC:** The sum of the previous two columns.

Table 2
Absolute financed emissions attribution for Portfolio One in 2021

	Return	Scope 1 tCO_2e/EV	Scope 2 tCO_2e/EV	Scope 1 & 2 CO_2e/EV
Global Equity Benchmark	**31.07%**	**536,341**	**129,526**	**665,867**
Regulatory	-0.07%	-29,480	-2,575	-32,054
Inappropriate activities	0.32%	12,022	-4,377	7,645
ESG score	0.17%	-19,469	712	-18,757
Environmental Score	-0.01%	-68,192	-5,126	-73,319
Operational Exclusions	0.06%	19,092	4,072	23,164
Financial score	1.13%	-41,947	-8,272	-50,218
Exclusion Benchmark	32.65%	408,368	113,959	522,327
Capped Benchmark	29.83%	482,994	132,720	615,714
Portfolio One	**29.34%**	**406,909**	**138,933**	**545,842**
Total Outperformance	***-1.72%***	***-129,432***	***9,407***	***-120,024***
Due to exclusions	***1.59%***	***-127,973***	***-15,567***	***-143,539***
Due to cap	***-2.83%***	***74,626***	***18,761***	***93,387***
Due to manager choices	***-0.48%***	***-76,085***	***6,213***	***-69,872***

The Exclusion Benchmark, Capped Benchmark and Portfolio rows display identical metrics, but calculated using the constituents of each benchmark or portfolio, as outlined in the methodology section. The rows between Original Benchmark and Exclusion Benchmark show the contribution of each exclusion reason. The sum of these figures explains the difference between the Original Benchmark and Exclusion Benchmark, as described in Equation nine.

Additionally, the lower three rows of the table show the total contribution of each major decision, as shown in Equation seven.

FINANCIAL RETURNS

In terms of financial returns the calculated figures show the following:

- The Portfolio underperformed the Original Benchmark in financial terms by 1.72%
- The exclusions contributed positively, increasing returns by 1.59%
- The cap on the maximum weight of individual securities and allocation decisions made by the manager reduced returns by 2.83% and 0.48%, respectively.

ABSOLUTE FINANCED EMISSIONS

Interestingly, the Portfolio achieved significantly lower emissions than the benchmark. The exclusion related to the "Environmental Score" was crucial in achieving these lower emissions. This exclusion method screens companies based on their performance on certain environmental metrics and is applied to sectors with the largest climate impact. This result aligns with our expectations regarding the implications of this investment decision.

What was less expected, however, is the impact of capping the weight of individual securities on total financed emissions. The calculation reveals that capping the weight of individual stocks offsets 65% of the emissions reduction from exclusions. This outcome can be attributed to the underweighting of technology stocks compared to the benchmark. Technology companies tend to have relatively low scope 1 and 2 emissions compared to other sectors.

ATTRIBUTING EMISSIONS INTENSITY

Similar calculations can be made for attributing the emissions intensity. Table 3 shows the same analysis, but this time emissions intensity (tons CO_2 per million Euros of revenue) is used instead of absolute financed emissions. Unlike the data in table 2, this output is a relative measure instead of absolute, making it useful for comparing different portfolios. The Scope 1 & 2 emissions intensity comparison provides insight in the short-term difference between the two investment strategies.

Table 3
Emissions intensity attribution for Portfolio One and Two in 2022

	Portfolio One	Portfolio Two	Difference
	Scope 1 & 2 tCO_2e intensity	Scope 1 & 2 tCO_2e intensity	Scope 1 & 2 tCO_2e intensity
Global Equity Benchmark	**143.52**	**143.52**	**0.00**
Regulatory	-3.42	-2.88	-0.54
Inappropriate activities	1.78	2.07	-0.29
ESG score	0.48	1.83	-1.35
Environmental Score	-23.74	-20.32	-3.42
Operational Exclusions	5.34	2.51	2.83
Financial score	-11.00	-6.57	-4.43
Exclusion Benchmark	112.95	120.15	-7.20
Capped Benchmark	134.87	141.22	-6.35
Portfolio	**134.79**	**140.26**	**-5.47**
Total Outperformance	***-8.73***	***-3.26***	***-5.47***
Due to exclusions	***-30.57***	***-23.36***	***-7.21***
Due to cap	***21.92***	***21.06***	***0.86***
Due to manager choices	***-0.08***	***-0.96***	***0.88***

Interestingly, table 3 reveals that one of the most significant individual differences in emission intensity between Portfolio One and Two results from excluding companies based on their financial score. In Portfolio One, this exclusion reduces the weighted average CO_2e intensity by 11.00, while, in Portfolio Two, it declines by 6.57 tCO_2e intensity. This 4.43 tCO_2e intensity difference between the two portfolios, accounts for approximately two-thirds of the total discrepancy (7.20 tCO_2e intensity) between the Global Equity Benchmark and the Exclusion Benchmark.

Notably, financial score-based exclusions have a greater impact than environmental score-based exclusions. This underscores the importance of integrating ESG performance attribution in the investment decision-making process, as seemingly unrelated decisions can significantly affect the portfolio's overall ESG performance.

ATTRIBUTING DATA QUALITY

Data quality is estimated by combining the reported trust of each data point with the total amount of missing data. Data was missing for 2.80% of the weight in the Original Benchmark but only for 0.84% of Portfolio One. Alongside this, the Data Quality Score reflects the estimated trust in the data for each row, considering factors like the amount of missing data and the source reliability (with estimates or unreliable sources reducing the trust score).

Table 4
Data quality attribution for Portfolio One in 2022

	Scope 1 & 2 tCO_2e/EV	Data Quality Score	Total weight missing data
Global Equity Benchmark	665,867	75.35%	2.80%
Regulatory	-32,054	-0.07%	0.01%
Inappropriate activities	7,645	0.38%	0.01%
ESG score	-18,757	3.81%	-0.72%
Environmental Score	-73,319	0.03%	-0.01%
Operational Exclusions	23,164	0.85%	-1.15%
Financial score	-50,218	0.68%	-0.08%
Exclusion Benchmark	522,327	**81.04%**	**0.85%**
Capped Benchmark	615,714	**78.97%**	**1.01%**
Portfolio	**545,842**	**78.94%**	**0.84%**

ATTRIBUTING ESG SCORES

Similar to emissions data, the methodology can be applied to ESG scores as well. Table 5 shows the impact of the exclusion policies on ESG-scores in the portfolio.

Table 5
ESG score attribution for Portfolio One in 2022

	Environment	Social	Governance
Global Equity Benchmark	**6.52**	**5.38**	**4.74**
Regulatory	0.01	0.01	0.00
Inappropriate activities	0.04	-0.01	0.07
ESG score	-0.01	0.10	0.16
Environmental Score	0.02	-0.01	-0.01
Operational Exclusions	-0.05	0.01	0.02
Financial score	0.03	0.03	0.04
Exclusion Benchmark	6.56	5.51	5.00
Capped Benchmark	6.57	5.69	4.96
Portfolio	**7.03**	**6.09**	**5.32**
Total Outperformance	***0.51***	***0.71***	***0.58***
Due to exclusions	***0.04***	***0.13***	***0.26***
Due to cap	***0.01***	***0.18***	***-0.04***
Due to manager choices	***0.46***	***0.40***	***0.36***

This table shows that the ESG score-based exclusions mainly improve the Social and Governance score of the exclusion benchmark but have less impact on the Environment score. This may be because companies with poor environmental scores were excluded under Inappropriate activities. Ultimately, the choices made by the manager have had a more positive effect on the ESG scores than the exclusion policies during this period.

FUTURE EXTENSIONS

DECOMPOSING CHANGES IN NON-FINANCIAL METRICS OVER TIME

The attribution calculations in this article focus on the impact of investment decisions on an annual basis. However, comparing results across different years would be a valuable addition to investor's toolkit.

For the calculations in this article, daily portfolio data was used to determine the weights of securities in both the benchmark and portfolio. An average weight for the year was then used to aggregate emissions data annually, with only the yearly totals shown in the data tables of the previous chapter. Increasing the frequency and extending the time period of these calculations would allow for a more detailed analysis of trends, as illustrated in figure 2.

Figure 2
Attributing financed emissions over time

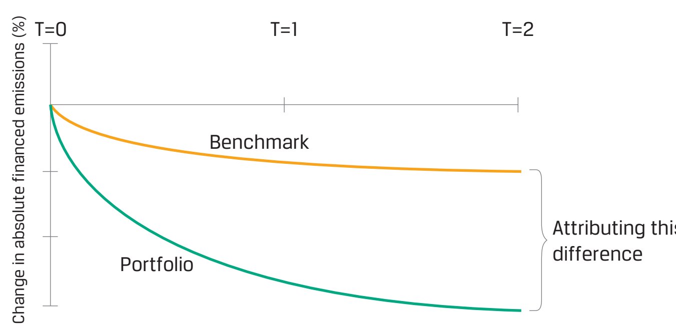

Whilst we can calculate these effects, annual calculations are not frequent enough for quality trend analysis. Using a fully automated and systemized solution and more regular data points (e.g. monthly or quarterly) would provide higher quality results for the discussed decomposition.

Conducting such an analysis over multiple years involves several challenges. Firstly, when decomposing changes in ESG metrics over time, additional factors such as the changes in data quality, as well as the effects of inflation, must be considered. The next paragraphs explore a few of these additional considerations in more detail.

Finally, as more periods and factors like inflation adjustments are included, the complexity of the calculation increases. To ensure these calculations are robust and efficient, specialised software will need to be developed, enabling investors to perform these calculations with relative ease.

DATA QUALITY AND DATA COVERAGE EFFECTS

This article made use of various data proxies in order to address the problem of data coverage, while acknowledging the availability of other options. Over time, it is expected that the availability as well as the quality of ESG data will improve. When comparing the total absolute financed emissions of a portfolio from one year to the next, three important consideration arises:

- If more companies publish data, the calculation will rely less on estimated data. All other things being equal, the total financed emissions of the portfolio will change purely as a result of the increased use of published data instead of proxied or estimated data.
- As more data becomes available, the overall data coverage (both published and estimated) will increase. All other things being equal, this will result in an increase of the total financed emissions.
- The methodology used to measure emissions may change over time as technology improves.

INFLATIONARY IMPACTS

When combining ESG metrics with financial measures, it is essential to account for the impact of inflation to ensure a fair comparison over time. Emissions metrics, for example, are calculated by dividing the total emissions attributable to a company by a financial denominator, such as enterprise value or revenue. Metrics like enterprise value and revenue are used to approximate the emissions produced per unit of productivity by the company. Using revenue or enterprise value as a proxy enables the measurement of productivity consistently across different sectors.

In high inflationary environments, it is crucial to recognize that emissions metrics may decrease solely due to inflation. Therefore, adjusting for inflation is essential to accurately assess changes in metrics like financed emissions or emissions intensity.

Table 6 illustrates an extreme example of inflation effecting the emissions intensity of a company. Consider a scenario where a company's total emissions remain constant at ten thousand tons CO_2e over five years, whilst revenue fluctuates during the same period. Nominally, the emissions intensity (measured as emissions per unit of revenue) appears to decline from twenty in 2019 to just over fifteen in 2023.

However, it is important to recognize this revenue growth is partly due to inflation. The inflationary growth figure shows the expected growth in revenue based on the inflation rate; if inflationary growth is stripped out and the emissions intensity is calculated on a real basis, the company's emissions intensity does not decrease to fifteen. Instead, it consistently hovers around nineteen throughout the entire period.

If a decomposition is done over time, inflationary increases in revenue could appear as a fall in a company's emissions intensity (tCO_2e/EV). In the most extreme cases, as demonstrated in table 6, total emissions (tCO_2e) could remain static whilst inflation-driven revenue growth significantly reduces the emission intensity. Stripping out inflation shows that the "real" emissions intensity only decreases slightly. This example, although extreme and hypothetical, underscores the importance of adjusting for inflation to avoid falsely reporting reductions in emissions intensity during high inflation periods.

OVERALL DECOMPOSITION

Once data coverage and inflation effects are separated, a more granular decomposition of changes over time is possible. This article proposes four major effects, with two further subdivided. The list of effects is likely to evolve over time and vary depending on the metric being used for the analysis.

Trading impacts emissions intensity over time. Selling high-emissions intensity companies and buying low-emissions intensity ones will naturally reduce emissions intensity over time. Weights also change based on the financial performance of each company. These effects can be decomposed as one large trading or allocation effect, or broken down into a new investments effect, divestments effect and weight/allocation change effect.

Security level metrics also change over time. Changes in emissions intensity can result from changes in absolute emissions (numerator) or revenue/enterprise value (denominator). Revenue changes can be driven by local revenue changes, exchange rates, or both. For example, a company may have both stable absolute emissions and revenue in terms of USD, leading to flat emissions intensity in USD terms, but have a rising or falling emissions intensity in EUR terms. These effects can once again be represented as a single effect or broken down into more granular effects.

Table 6
Example of the potential impact of inflation on a stock's emissions intensity

Year	Total Emissions (tCO_2e)	Revenue ($m)	Emissions Intensity (tCO_2e/€100)	Inflation	Inflationary Revenue Growth	"Real" Emissions Intensity (tCO_2e/€100)
2019	10000	500	20.00			
2020	10000	525	19.05	2.00%	10.00	19.42
2021	10000	575	17.39	5.00%	26.25	18.56
2022	10000	600	16.67	10.00%	57.50	19.75
2023	10000	650	15.38	6.00%	36.00	19.22

To gain a clear understanding of this change over time it is therefore necessary to calculate (at least) four major effects, two of which can be broken down further:

- Data coverage effect
- Inflation effect
- Trading effects
  - New investments
  - Complete divestments
  - Changes in weights/allocation
- Security level effects
  - Changes in absolute emissions
  - Changes in revenue/enterprise value due to exchange rate changes
  - Changes in revenue/enterprise value (after stripping out inflation)

This decomposition would allow a deeper level of analysis when comparing the strategies of exclusion and engagement. For an investor employing mostly exclusion rules, it would be expected that the trading effect is dominant, whilst investors who opt for active engagement would expect to see the security level effects increase over time.

CONSIDERING EXTRA DECISIONS AND METRICS

This article focuses on a single portfolio use case, where the excess carbon emissions of one portfolio compared to the benchmark are broken down into three primary effects:

- ESG driven exclusions
- The capping of the maximum weight of individual securities
- Allocation and selection decisions made by the portfolio manager

The IDP approach (Geenen et al., 2021) considers various top-down investment decisions inherent in institutional investors' investment process. While many are primarily financially driven, they also impact non-financial metrics. A logical extension of the proposed model is to consider all decisions, financial or non-financial, and assess their impact on both performance aspects.

Consider an institutional investor with a strategic asset allocation (SAA) updated annually and a tactical asset allocation (TAA) updated quarterly. Within each asset class in the SAA and TAA, further allocation decisions taking place – for instance, allocating between developed markets and emerging markets equities. Additionally, exclusions are applied within each asset class (potentially with different exclusion rules), each targeting different objectives. For example, one would expect a different effect from the exclusion of oil and gas producers than from excluding companies with high water usage or high impact to biodiversity.

To comprehensively evaluate the performance of portfolios and benchmarks, non-financial metrics should be included in performance measurement especially when institutional clients aim to integrate ESG in their investment decisions. Combining financial and non-financial data give insights into how the client's fund overall performance compares not only in financial terms but also in non-financial aspects relative to the SAA (the primary fund benchmark). If a decision is taken to increase an allocation towards Equity, it is important to understand the impact of this decision not just on financial performance but also on carbon emissions as a decision taken for one reason may often had an adverse impact on another objective of the portfolio.

POSSIBLE FURTHER IMPROVEMENTS RELATED TO DATA QUALITY AND COMPLETENESS

Another additional improvement could involve combining country- and sector-specific data, considering that different regulations apply across countries and sectors. For instance, Europe tends to have stricter environmental laws than the United States, making a US-based coal company an imperfect proxy for a European-based one.

When evaluating results over time, it's essential to consider data quality and completeness, the estimation of data quality can be seen in the Results section under "Attributing Data Quality". Improved data completeness reduces reliance on proxied/estimated data. All other things being equal, this change in data coverage impacts the financed emissions of the portfolio. To address this, it is required to split out this impact into a separate "data coverage effect". The calculation of this effect is further discussed in the future extensions section under "Data Quality and Data Coverage Effects".

CONCLUSION

Integrating ESG factors into investment decision-making is increasingly important from both regulatory and responsible investing perspectives. As this trend is likely to continue, tools have to be developed to gain deeper insights into the effects of investment decision beyond traditional performance metrics. In this article, we propose an ESG attribution model as a tool to offer a more holistic overview of the effects of investment decisions, whether driven by financial or ESG considerations.

We demonstrate that a standard performance attribution model can be used to attribute ESG metrics to specific investment-decisions. Initially, the model is used to show the effects of decision within one year. This enables investors to, for example, see how decisions can have unintended consequences to ESG metrics, such as financed emissions. This underscores the importance of including these types of attribution calculations within the investment decision-making process.

An important next step for an ESG attribution calculation is to analyze the ESG performance of the portfolio over time and attribute the performance compared to the benchmark to different investment decisions. Expanding the calculation over a longer period introduces new challenges, such as the need to account for inflation. As these calculations become increasingly complex, specialized systems are needed to perform them. However, as with financial performance we believe that the ability to analyze the impacts of decisions over a long period of time is vital. For emissions related goals in particular, the decisions made now are often made with the view to having a positive impact a decade or more into the future. Tracking the impact over a long period is therefore vital to gain a full understanding of the impact of our investment decisions.

Furthermore, this article acknowledges issues regarding data quality and coverage, meaning the results should be viewed as estimates. These data issues become increasingly relevant when expanding the model to attribute over longer periods. The article discusses several possible extensions to account for the impact of changes in data quality and coverage over time. As more investors analyze ESG decisions, the demand for complete and high quality data will increase, ultimately improving overall data quality and coverage.

NOTES

The methodology discussed on pages 19 and 20 for Absolute Financed Emissions aligns with the approach for calculating financed emissions as prescribed by the Partnership for Carbon Accounting Financials (PCAF). It is based on the ownership principle, where an investor's ownership percentage in a company corresponds to the percentage of emissions they have financed.

Literature

- Brinson, G.P., and N. Fachler, 1985, Measuring Non-US Equity Portfolio Performance, Journal of Portfolio Management, Spring, pp. 73-76
- Geenen, J., E. van de Burgt, M. Heemskerk, and M. Heerema, 2021, Decision-Based Evaluation of the Performance of a Hierarchically Structured Investment Process, Journal of Performance Measurement, Spring, pp. 47-62
- Global Sustainable Investment Alliance, 2022, Global Sustainable Investment Review, Report on https://www.gsi-alliance.org/wp-content/uploads/2023/12/GSIA-Report-2022.pdf
- Steehouwer, H., 2023, From climate risk analysis to investment decision making, Report on https://www.ortecfinance.com/en/insights/whitepaper-and-report/from-climate-risk-analysis-to-investment-decision-making