

ATTRIBUTION OF PORTFOLIOS WITH CLIMATE-RELATED SIGNALS

Andrew Ang

Managing Director, BlackRock, New York City

Debarshi Basu

Managing Director, BlackRock, New York City

Marco Corsi

Managing Director, BlackRock, London

We attribute returns for portfolios constructed with climate-related signals—past and forward-looking carbon commitments; water withdrawal intensity, which falls into natural capital; and a signal of climate-related intellectual property. A key feature of the attribution is it sums to 100%, and we apply the attribution method to returns, ex ante and ex post risk, and tracking error. The decompositions without residuals better allow investors to evaluate the various contributions of these climate-related signals to risk and return, enabling more efficient and customized capital deployment. We find there is relatively low correlation among these signals, so they offer potential diversification benefits, and there are significant interactions of the climate-related signals with ex post carbon emissions.

Introduction

The transition to net zero is a topic relevant to many investors looking to mitigate the risk and take advantage of the investment opportunities associated with this critical shift. Measuring the risk and return of different approaches associated with the net-zero transition—such as current and future carbon emissions, the conservation of natural capital, and investments in new green technologies—is important for the allocation of capital, setting optimal taxes and subsidies, and assessing the real investments of governments and corporations (see Intergovernmental Panel on Climate Change 2023). But evaluating the returns and risk of different net-zero approaches can be difficult because companies may pursue more than one of these approaches simultaneously. Similarly, the majority of investors typically hold diversified portfolios anchored around a major market benchmark, and there may be several climate-related characteristics taken into account when constructing their portfolios.

Note: The views expressed herein are personal views of the authors and do not represent the views of any organization or other third party. This material is not intended to be relied upon as a forecast or research or investment advice and is not a recommendation, offer, or solicitation to buy or sell any securities or to adopt any investment strategy.

In this chapter, we attribute contributions of different climate-related signals to portfolio returns and *ex ante* and *ex post* risk.¹ The attribution method follows Moehle, Boyd, and Ang (2022), which computes Shapley (1951, 1953) attribution values in a quantitative investment context. A key feature is that the attributions sum to 100%; that is, the Shapley attribution measures the contribution of each climate-related signal such that the individual signal returns sum to the actual portfolio return. In our specific example, we decompose the risk and returns of a climate-aware portfolio that maximizes exposure to carbon emissions (both past and forward-looking commitments), water withdrawal intensity, and green R&D investments as proxied by green patents, subject to a tracking error limit relative to the MSCI World Index with sector, country, and asset-level constraints.

The Shapley attribution has several other attractive features. The attribution is symmetric: If features i and j contribute the same amount when they are added to different portfolios, then they have the same attribution. It also is linearly additive: If the contribution to feature i is added to the contribution of feature j , the attribution to the combined $(i + j)$ features is the sum of the individual contributions. In fact, Young (1985) and others show that the Shapley attribution is the *only* attribution method that satisfies all these desirable criteria.²

We find that constructing a portfolio with multiple dimensions of transition-related variables—as opposed to only carbon emissions, water withdrawal, or green patents signals taken one at a time—improves diversification. A portfolio constructed with exposure to all three climate-related characteristics generates an excess return of 63 bps per year over the benchmark MSCI World Index. The portfolio’s annualized active risk is 160 bps relative to the MSCI World universe over 1 February 2017 to 1 June 2024 (a period of 88 months). The portfolio delivers a 67% reduction in carbon emission intensity relative to the benchmark’s carbon emission intensity, with all three components of the score contributing to the reduction in emissions. It is notable that this level of reduction in carbon emissions is achieved without using an explicit decarbonization constraint in the optimization.³

A benefit of being able to compute total attribution of signals is that investors with various degrees of preferences for different sustainability approaches can use the decompositions to customize the weights of different signals—and in this case, upweight or downweight the various climate-related components. In our results, water efficiency and green patents also lead to *ex post* reductions in carbon emissions without explicitly targeting carbon emissions. In particular,

¹Note that the terms “net zero” and “transition” have a distinct meaning, especially in a regulatory context. In this chapter, we use the broader term “climate-related” to encompass climate-related goals that might not be directly included in specific net-zero frameworks. See, for example, Commission Delegated Regulation (EU) 2020/1818 of 17 July 2020: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32020R1818>.

²Shapley attribution is the only attribution method that satisfies the properties of (1) efficiency, where the individual signals sum to the actual portfolio return; (2) symmetry; (3) linear additivity, as defined in this chapter; and (4) null value, where the return is the benchmark return when no features are activated.

³Approaches that lower the total carbon emission intensity of a portfolio through a constraint in an optimization are taken by Bolton, Kacperczyk, and Samama (2022); Hodges, Ren, Schwaiger, and Ang (2022); and Le Guenedal and Roncalli (2022), among others.

these two climate-related signals reduce the portfolio carbon emission intensity by -18 and -7 metric tons per \$1 million sales (t/mn\$ sales), respectively, relative to the benchmark ACWI portfolio. (As expected, exposure to lowering the carbon emission scores reduces carbon emission intensity, by -39 t/mn\$ sales.) Such attribution makes it easier to understand the drivers behind a portfolio-level outcome and enables customized selection of desired climate-related characteristics to meet individuals' specific objectives.

This chapter is part of a growing literature that investigates the relationship of climate-related signals to stock returns. Some of this relationship is ambiguous: Bolton and Kacperczyk (2023) report that companies with higher carbon emissions have high excess returns, whereas Kazdin, Schwaiger, Wendt, and Ang (2021) find the opposite result. Ang, Garvey, and Schwaiger (2024) report that companies with higher profitability adjusted for carbon emissions and industry have higher returns. In contrast, Aswani, Raghunandan, and Rajgopal (2024) find there is no relation between carbon emissions and stock returns. Other studies examine climate-related variables other than carbon emissions; for example, Hsu, Li, and Tsou (2023) report that companies with higher levels of pollution are riskier and have higher returns. Of course, climate-related variables are a special case of the more general environmental, social, and governance (ESG) area. Using more than 16,000 global stocks and data from seven different ESG providers, Alves, Krüger, and van Dijk (2024) find that there is no relation between ESG ratings, regulations, or disclosure standards and stock returns. In our study, we focus on return attribution of climate-related variables in the context of an investment strategy but cannot make any statements on the relationship between returns and broader ESG scores.

The Shapley attribution we consider has not been covered in the large attribution literature in finance.⁴ Some of these studies, such as Jensen (1968), Brinson, Hood, and Beebower (1986), and Fama and French (2010), use time-series data and compute alphas relative to a benchmark. These regression-based methods are dependent on the order of variables assumed in the regression. Studies using holdings-level data, such as Grinold and Kahn (2000) and Grinold (2006), often have large residuals, which are return components not attributable to any feature. In contrast, our return decompositions are not dependent on sequential order, are residual free, and sum to 100%.

⁴There is now wide use of Shapley values in machine learning with the use of SHAP (SHapley Additive exPlanations) functions—which enable the performance gain of a predictive procedure to be attributed to different inputs of the model. See Lundberg and Lee (2017) and <https://shap.readthedocs.io/en/latest/>. There are many methods related to SHAP, including Baseline SHAP, Kernel SHAP, Tree SHAP, and Deep SHAP.

Climate-Related Portfolio Construction

In this section, we describe the climate-related variables and the portfolio construction.

Data and Signals

Our full panel dataset consists of 23,646 firm-month observations from February 2017 to June 2024 consisting of stock returns, climate-related scores, and carbon emissions.

Stock Return Universe

The universe for the portfolio is the MSCI World Index, which incorporates large- and mid-cap companies from 23 developed markets. The portfolio averages 1,626 stocks across the sample from February 2017 to June 2024.


Climate-Related Variables

For the purpose of demonstrating Shapley attribution on the portfolio constructed with climate-related characteristics, we take three signals. The first signal is *carbon emissions*, which uses historical Scope 1 and 2 emission intensity over sales from MSCI and a forward-looking commitment measure. The former represents a company's most recent Scope 1 and Scope 2 greenhouse gas (GHG) emissions measured in metric ton CO₂ equivalent, which is normalized by sales in millions of US dollars. As can be seen from **Exhibit 1**, the emission numbers exhibit a pronounced right skew, which is driven by a small number of companies with very large carbon emission intensities (see comments by Hodges, Ren, Schwaiger, and Ang 2022; Bolton and Kacperczyk 2023). We use the log transformation to remove the positive skewness, which results in the histogram reported in the right-hand plot of Exhibit 1. We Z-score and truncate this variable between $[-3, 3]$.

For future carbon commitments, we use an indicator variable of 1 or 0, which is exponentially weighted in the past, depending on whether a firm has set science-based carbon emission targets and is a signatory of the Science Based Targets initiative (SBTi). Garvey, Iyer, and Nash (2018) and Ang, Garvey, and Schwaiger (2024) show that firms with lower carbon emissions have, on average, higher profitability and efficiency metrics. In addition, Trinks, Ibikunle, Mulder, and Scholtens (2022) show that these firms also have lower systematic risk.

The final carbon emission signal takes 80% past carbon emissions and 20% carbon commitments. The lower weight on carbon commitments is motivated by the smaller number of firms that have made SBTi commitments to lowering future emissions. We Z-score so the variable has a mean of zero before using it in the portfolio construction process.

Exhibit 1. Log Transformation of Carbon Emissions



Notes: The histogram of the raw Scope 1 and 2 carbon emission intensity (which is normalized by sales) is plotted in the left panel. The log transformation of the raw data is plotted in the right panel.

The second signal, the natural capital signal, is *water withdrawal intensity* obtained from MSCI. The metric represents the company's reported water withdrawal (measured in cubic meters) normalized to revenues (\$ millions). As with carbon emissions, water withdrawal intensity exhibits a right skew, so we log transform and Z-score the raw data.

The final signal measures climate-related intangible capital by *green patents*, as introduced by Chan, Hogan, Schwaiger, and Ang (2020). Often, patents are the culmination of investment in research and development, and a large literature uses patents to proxy for intangible asset information (see, for example, Lee, Sun, Wang, and Zhang 2019). Green patents are patents that promote innovation consistent with the UN Sustainable Development Goals, as defined by the World Intellectual Property Organization. We follow Chan, Hogan, Schwaiger, and Ang (2020) and take the two-year rolling sum of the number of green patents owned by each company divided by market capitalization, which is then Z-scored. Green patents are a measure of intellectual property investments associated with the transition.

Finally, we further Z-score each of the three climate-related signals on a sector-by-sector basis over the MSCI World universe.

Portfolio Construction

We construct a portfolio using the climate-related scores and carbon emissions as follows. Our portfolio is long only and uses the following optimization for N portfolio weights, $\mathbf{h}_{\text{active}}$:

$$\mathbf{h}_{\text{active}} = \arg \max_{\mathbf{h}} U(\mathbf{h}) = \arg \max_{\mathbf{h}} \mathbf{h}^T \boldsymbol{\alpha} - \lambda \mathbf{h}^T \mathbf{V} \mathbf{h}, \quad (1)$$

where

$\boldsymbol{\alpha}$ is an $N \times 1$ vector that is an equal-weighted average of the three climate-related scores for each constituent of the benchmark

λ is a coefficient of risk aversion

\mathbf{V} is the variance–covariance matrix ($N \times N$) from a factor model from the Aladdin risk system (see Bass, Gladstone, and Ang 2017)

We set λ to 0.25, which corresponds to an *ex ante* tracking error between 100 bps and 150 bps of risk.

The objective function in Equation 1 maximizes the combined climate-related score of all stocks in the MSCI World Index and treats the climate-related scores as alpha components. In this formulation, we are not addressing whether there is an empirical relation between the climate-related scores and returns; the optimization exogenously assumes that the investor desires the maximum climate-related score for the portfolio subject to risk.

In addition, we assume the following constraints:

$$\begin{aligned} \mathbf{h} &\geq -\mathbf{h}_{\text{benchmark}} \\ -3.0\% &\leq \mathbf{h} \leq +3.0\%. \end{aligned}$$

$$\left| \sum_{i \in \text{Country}_j} h_i \right| \leq 2.0\%. \quad \forall \text{Country}_j \in \text{Benchmark.}$$

$$\left| \sum_{i \in \text{Sector}_k} h_i \right| \leq 2.0\%. \quad \forall \text{Sector}_k \in \text{Benchmark.}$$

Note that $\mathbf{h}_{\text{benchmark}}$ is an $N \times 1$ vector of market-cap weights in the MSCI World Index benchmark. The constraints can be interpreted as follows. The first constraint guarantees the portfolio is long only. In the second constraint, the active weight relative to the benchmark of a single security is less than or equal to 3.0%. The third and fourth constraints represent that the active country weight is limited to 2% and the maximum active sector weight is 2%, respectively.

We rebalance the portfolio on the last business day of February and August in line with the NYSE trading calendar. On the semiannual rebalance dates, we liquidate the old positions and purchase the new positions. We assume full reinvestment without any cash balances and hold these positions until the next annual rebalance date, when the process is repeated. We ignore transaction costs for our analysis for simplicity, but it is straightforward to include an additional linear term in Equation 1 to take them into account.

Finally, for the portfolio benchmark, we use a modified version of the MSCI World Index that rebalances only twice a year,⁵ in February and August. Doing so ensures that the relative performance between the portfolio and its benchmark is not affected by differences in the respective rebalancing schedules.

Shapley Attribution

We lay out an intuitive exposition of Shapley (1951, 1953) attribution using a geometric interpretation. A more general formula is in the Appendix.

We work with three features, $i = 1, 2, 3$, which can be interpreted as the three climate-related signals. We denote the portfolio return as $f(\mathbf{x})$, where the vector \mathbf{x} is a configuration with all features. The benchmark MSCI ACWI return without any climate considerations is given by $\mathbf{x} = (0, 0, 0)$, with corresponding return $f(0, 0, 0)$. The portfolio return with all climate return signals is denoted by $f(1, 1, 1)$, and we denote the full configuration by $\mathbf{x} = (1, 1, 1) = \mathbf{1}$. We wish to decompose the full portfolio return, $f(1, 1, 1)$, into the three individual components.

Lifts

In Equation 2, we define the marginal contribution for feature i , or lift, as the change in performance by adding feature i :

$$f(\mathbf{x} + \mathbf{e}_i) - f(\mathbf{x}), \quad (2)$$

where \mathbf{e}_i is a vector of zeros with a 1 in the i th position. The marginal contribution depends on which features are turned on in the configuration \mathbf{x} and then adds the i th feature.

In the context of the optimization of Equation 1, the entries of 1 in the vector \mathbf{x} correspond to nonzero entries of the alpha vector, α . For example, $\mathbf{x} = (1, 0, 0)$ corresponds to having scores only for the first climate-related signal of carbon emissions turned on in the optimization. In this case, the alpha vector in Equation 1 takes the form $\alpha = (z_1 + 0 + 0)$, where z_1 represents the carbon emission scores, 0 is zero so there are no scores for the two other


⁵The MSCI World Index rebalances four times a year, on the last business day of February, May, August, and November.

climate-related signals (water withdrawal and green patents, represented by z_2 and z_3 , respectively).

Hypercube Interpretation

For n features, we visualize a hypercube with each feature corresponding to a vertex of a hypercube. For example, for three features, the axes in **Exhibit 2** correspond to three features: x_1 , x_2 , and x_3 . The origin, $(0, 0, 0)$, represents the benchmark or zero, and the full set of features, $(1, 1, 1)$, represents the actual portfolio return. The 1 entries in the vector \mathbf{x} represent the features that are turned on. For example, $\mathbf{x} = (0, 1, 0)$ represents the feature $i = 2$ turned on. The point $(1, 1, 1)$ represents the portfolio return with all climate-related features enabled.

Exhibit 2. Hypercube Interpretation of Marginal Contributions: Vertices Are Feature Configurations



With $n = 3$ features, there are six possible paths from $(0, 0, 0)$ to $(1, 1, 1)$:

$1 \rightarrow 2 \rightarrow 3$
 $1 \rightarrow 3 \rightarrow 2$
 $2 \rightarrow 1 \rightarrow 3$
 $2 \rightarrow 3 \rightarrow 1$
 $3 \rightarrow 1 \rightarrow 2$
 $3 \rightarrow 2 \rightarrow 1$

For these paths, it is understood that we always start from $(0, 0, 0)$ and then turn on the features in the order listed in each permutation.

Traveling on the edges from configuration \mathbf{x} to $\mathbf{x} + \mathbf{e}_i$ represents the lift $f(\mathbf{x} + \mathbf{e}_i) - f(\mathbf{x})$. For example, the edge from $(0, 0, 0)$ to $(1, 0, 0)$ represents adding Feature 1 starting from no features (or the origin). Then, traveling along the edge from $(1, 0, 0)$ to $(1, 1, 0)$ adds Feature 2 starting from a configuration with only Feature 1.

Marginal Contributions

We state the marginal contributions corresponding to the first feature, $i = 1$, for the six permutations:

Permutation	Marginal Contribution for $i = 1$
$1 \rightarrow 2 \rightarrow 3$	$f(1, 0, 0) - f(0, 0, 0)$
$1 \rightarrow 3 \rightarrow 2$	$f(1, 0, 0) - f(0, 0, 0)$
$2 \rightarrow 1 \rightarrow 3$	$f(1, 1, 0) - f(0, 1, 0)$
$2 \rightarrow 3 \rightarrow 1$	$f(1, 1, 1) - f(0, 1, 1)$
$3 \rightarrow 1 \rightarrow 2$	$f(1, 0, 1) - f(0, 0, 1)$
$3 \rightarrow 2 \rightarrow 1$	$f(1, 1, 1) - f(0, 1, 1)$

Take the first permutation, $1 \rightarrow 2 \rightarrow 3$. After starting at the benchmark, $(0, 0, 0)$, we turn on the first feature. The marginal contribution is then $f(1, 0, 0) - f(0, 0, 0)$. Then sequentially adding Features 2 and 3 (going from $2 \rightarrow 3$ after Feature 1 is added) no longer involves Feature 1, and the subsequent path does not further contribute to the lift of Feature 1.

The second permutation, $1 \rightarrow 3 \rightarrow 2$, is similar to the first permutation, $1 \rightarrow 2 \rightarrow 3$, because Feature 1 is added first and thus the contribution of Feature 1 is the same: $f(1, 0, 0) - f(0, 0, 0)$.

In the permutation $2 \rightarrow 1 \rightarrow 3$, the marginal contribution of the $i = 1$ feature is enabled after the second feature is already turned on: $\mathbf{x} = (0, 1, 0)$. Thus, in

the permutation $2 \rightarrow 1 \rightarrow 3$, the marginal contribution of the $i = 1$ feature is $f(\mathbf{x} + \mathbf{e}_i) - f(\mathbf{x}) = f(1, 1, 0) - f(0, 1, 0)$.

In the permutation $2 \rightarrow 3 \rightarrow 1$, Feature 1 is turned on last, after Features 2 and 3 are active, so the starting configuration is $\mathbf{x} = (0, 1, 1)$. In this case, the marginal contribution of the $i = 1$ feature is $f(\mathbf{x} + \mathbf{e}_i) - f(\mathbf{x}) = f(1, 1, 1) - f(0, 1, 1)$.

In the permutation $3 \rightarrow 1 \rightarrow 2$, we turn on Feature 1 after turning on Feature 3. Thus, the starting point is $\mathbf{x} = (0, 0, 1)$. The marginal contribution of Feature 1 is $f(1, 0, 1) - f(0, 0, 1)$. Feature 2's subsequent addition does not further contribute to the lift of Feature 1.

Finally, in the permutation $3 \rightarrow 2 \rightarrow 1$, we move to Feature 1 after already turning on Features 3 and 2. Thus, the marginal lift of Feature 1 is $f(1, 1, 1) - f(0, 1, 1)$.

It is important to note that each of the $f(\cdot)$ evaluations is a different optimization of Equation 1 where the α vector takes on different values depending on which features are turned on.

We can add up all the marginal contributions to Feature 1, a_1 , in each of the permutations:


$$a_1 = \frac{2}{6}[f(1, 0, 0) - f(0, 0, 0)] + \frac{1}{6}[f(1, 1, 0) - f(0, 1, 0)] + \frac{1}{6}[f(1, 1, 1) - f(0, 1, 1)] \\ + \frac{2}{6}[f(1, 0, 1) - f(0, 0, 1)].$$

We define the Shapley attribution of Feature 1 as a_1 above. There is a coefficient of 2 for the marginal contribution $f(1, 0, 0) - f(0, 0, 0)$ because two paths include the edge $(0, 0, 0)$ to $(1, 0, 0)$ on the hypercube that are the permutations $0 \rightarrow 1 \rightarrow 2 \rightarrow 3$ and $0 \rightarrow 1 \rightarrow 3 \rightarrow 2$. The terms with a 1 in the numerator contain only one edge across the six permutations. For example, only one path includes the edge $(0, 1, 0)$ to $(1, 1, 0)$ that occurs for the permutation $0 \rightarrow 2 \rightarrow 1 \rightarrow 3$.

Exhibit 3 shows the four distinct edges in the hypercube for three features that correspond to the marginal performance change for Feature 1. Note there are four edges but six permutations, so for two permutations, the marginal contribution for Feature 1 is repeated.

The Shapley attributions for the second and third features, $i = 2$ and $i = 3$, respectively, given by a_2 and a_3 , respectively, can be obtained in a similar fashion.

Exhibit 3. Hypercube Interpretation of Marginal Contributions for Feature 1




Empirical Results

We present Shapley attribution results over the period February 2017 to June 2024.

Portfolio Performance

Exhibit 4 presents the performance of the climate-related portfolio versus the benchmark from February 2017 to June 2024. We start with \$100 at the beginning of February 2017. Over the time period, the climate-related portfolio has an annualized return of 1.03% per month, compared to 0.98% per month for the benchmark, which is an outperformance of 63 bps per year. The annualized *ex post* tracking error over the sample of the climate-related portfolio versus the benchmark is 1.63%.

Exhibit 4. Climate Portfolio Performance



Attribution

Exhibit 5 reports our main results and breaks down returns, total risk (volatility), and tracking error relative to the MSCI ACWI benchmark. We also report portfolio-level metrics corresponding to the three climate-related signals: the carbon emission intensity and percentage of firms with SBTi commitments, water withdrawal scores, and green patent scores.

We first turn to return attributions in the first row. Over the sample period from February 2017 to June 2024, the portfolio return was 11.11% per year. We can attribute this to the carbon, water, and green patent signals, which are 14 bps, 58 bps, and -9 bps, respectively (all annualized). Starting with the benchmark return of 11.11% per year, we have

Portfolio return = Benchmark + Carbon + Water + Green patent,

or

$$11.74\% = 11.11\% + 0.14\% + 0.58\% - 0.09\%.$$

Thus, over the sample, most of the outperformance has been driven by water, whereas green patents have slightly detracted. Note that the attribution, unlike regression-based or holdings-based methods, sums to 100%.

Of the realized volatility of 16.4%, the largest contribution is the benchmark of 16.09%—as by construction, with the optimization in Equation 1 setting risk aversion and sector, country, and holdings constraints to limit deviations from the benchmark. The largest contribution to the 1.63% tracking error is from carbon (77 bps), followed by water (67 bps) and green patents (19 bps).

Exhibit 5. Shapley Attributions

	Portfolio ^a	Benchmark	Carbon Attribution	Water Attribution	Green Patent Attribution
Realized Return^b (ann.)	11.74% (63 bps)	11.11%	14 bps	58 bps	-9 bps
Realized Volatility^b (ann.)	16.41% (32 bps)	16.09%	29 bps	-9 bps	12 bps
Realized Tracking Error^b (ann.)	1.63%	—	77 bps	67 bps	19 bps
Carbon Emission Intensity^c (t/mn\$ sales)	31.38 (-63.80)	95.15	-39.06	-17.59	-7.12
Percentage of Portfolio with SBTi Approved Target^c	65.10% (22%)	43.47%	16.08%	4.32%	1.23%
Water Withdrawal Score^c	66.46%	19.84%	-7.61%	63.00%	-8.77%
Green Patent Score^c	24.67%	—	-7.11%	-9.47%	41.25%

Notes: The exhibit shows Shapley attributions applied to the portfolio realized return, volatility, tracking error, carbon emission intensity, and the percentage of portfolio with SBTi commitments, water withdrawal score, and green patent score. The return, volatility, and tracking error are annualized. ^aNumbers in parentheses indicate the difference between the portfolio and the benchmark. ^bAll figures are annualized, based on monthly return over the period February 2017 to June 2024. ^cFigures are weighted averages calculated point in time as of the end of February 2024.

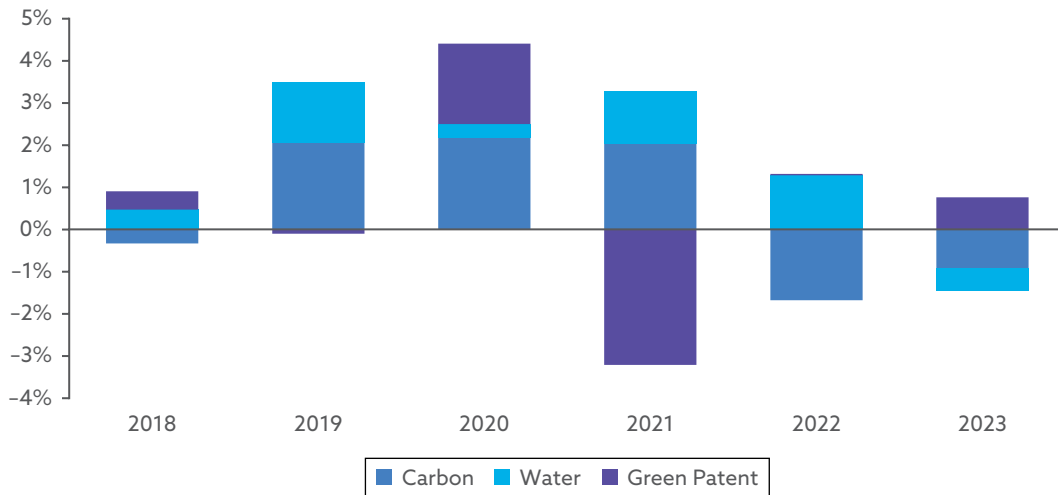
Note also that there are no risk attributions to “residual” or “idiosyncratic” components.

Of the portfolio climate-related scores, we expect each signal to have the largest contribution to the portfolio-level scores corresponding to each signal, which is evident, for example, from the fact that the largest attribution to carbon emissions is the carbon score. But there are also interesting and large cross-effects that are important for carbon emissions. The carbon emission intensity of the portfolio is 31.4 t/mn\$ sales, which represents a 67% reduction compared to the ACWI benchmark of 95.2 t/mn\$ sales. The water signal reduces carbon emission intensity by 17.6 t/mn\$ sales, and green patents reduce carbon emission intensity by 7.1 t/mn\$ sales. These reductions are on top of the reduction of 39.6 t/mn\$ sales from the carbon signal. Thus, the natural capital and green intellectual property signals also contribute to carbon emission reductions even though carbon emission is not directly captured in the definition of these signals.

Shapley Attribution over Time

Exhibit 6 reports the Shapley attribution of yearly active returns. The carbon signal contributes positively from 2018 to 2021 but is negative in 2022 and 2023. The negative returns to the carbon signal are due to the Russian invasion of Ukraine in February 2022, which led to large increases in energy prices. Although the full sample attribution to green patents is slightly negative (-9 bps per year; see Exhibit 5), it provides an important source of diversification

Exhibit 6. Shapley Attribution of Annual Active Returns, 2018–2023



in certain periods—particularly in 2020. The year 2020 saw the COVID-19 shock, where after an initial sharp decline of the market in Q1 2020, there was a significant increase in growth and technology stocks that helped society function during social distancing (for further remarks, see Ang 2023). The water signal has positive returns in all years except 2023. Overall, the three climate-related signals exhibit different behavior and thus provide diversification to the full climate-related portfolio.

Conclusion

We provide a method of attribution following Shapley (1951, 1953) that exactly decomposes portfolio statistics to individual features. We apply the Shapley attribution to a climate-related portfolio that maximizes past carbon emissions and future commitments, water withdrawal intensity, and green intellectual property proxied by green patents. Over the February 2017 to June 2024 sample, the carbon and water signals positively contribute to the portfolio outperforming the MSCI ACWI benchmark, and the green patent signal slightly detracts from performance relative to the benchmark. The largest contribution to realized tracking error is from the carbon reduction signal. Interestingly, the large 67% reduction in carbon emission intensity relative to MSCI ACWI is due to all three climate-related signals, not just the signal that explicitly measures reductions in carbon emissions.

While we can measure and attribute any portfolio statistic associated with the signals or other inputs into the portfolio construction process, Shapley attribution does not make any statement on causal mechanisms. The causal relationship is often important for choosing a particular climate-related signal and also for the choice by investors of certain sustainable investment approaches. While we cannot speak to causality, proper attribution of investment performance is a useful input for verifying and measuring causal effects.

Appendix: Computation of Shapley Attribution

The Shapley attribution for feature i , a_i , is defined in Equation A.1 as

$$a_i = \frac{1}{n!} \sum_{\pi} a_{i,\pi} \quad (\text{A.1})$$

where $a_{i,\pi}$ is the marginal contribution, or lift, for permutation π for feature i defined in Equation 2. The sum in Equation A.1 is over all $n!$ permutations. In our example in the main text, there are $3! = 6$ permutations. We can interpret the six ways of transversing the hypercube from 0 to **1** as equally likely in the denominator of Equation A.1 (see Exhibits 2 and 3).

The general formula for the Shapley attribution for feature i for features $i = 1, \dots, n$ is shown in Equation A.2 as

$$a_i = \sum_{\mathbf{x} \in \chi_i} \frac{(\mathbf{1}'\mathbf{x})!(n - \mathbf{1}'\mathbf{x} - 1)!}{n!} [f(\mathbf{x} + \mathbf{e}_i) - f(\mathbf{x})] \quad (\text{A.2})$$

where $\mathbf{1}$ is an $n \times 1$ vector of ones and $\chi_i = \{\mathbf{x} \mid x_i = 0\}$ is the set of configurations without feature i .

The drawback with Shapley attributions is that there are 2^n configurations that need to be evaluated for n features, which is unwieldy for a large n . In this case, Moehle, Boyd, and Ang (2022) show that we can use a sampling procedure to evaluate Equation A.2 using a multinomial distribution.

References

Alves, R., P. Krüger, and M. van Dijk. 2024. "Drawing Up the Bill: Is ESG Related to Stock Returns Around the World?" Working paper, Rotterdam School of Management, Erasmus University. www.rsm.nl/fileadmin/Faculty-Research/Centres/EPSVC/240223WOR_Drawing_up_the_bill_-_Is_ESG_related_to_stock_returns_around_the_world_.pdf.

Ang, A. 2023. "Trends and Cycles of Style Factors in the 20th and 21st Centuries." *Journal of Portfolio Management* 49 (2): 33–56. doi:10.3905/jpm.2022.1.455.

Ang, A., G. T. Garvey, and K. Schwaiger. 2024. "A Systematic Approach to Sustainable and ESG Investing." Working paper.

Aswani, J., A. Raghunandan, and S. Rajgopal. 2024. "Are Carbon Emissions Associated with Stock Returns?" *Review of Finance* 28 (1): 75–106. doi:10.1093/rof/rfad013.

Bass, R., S. Gladstone, and A. Ang. 2017. "Total Portfolio Factor, Not Just Asset, Allocation." *Journal of Portfolio Management* 43 (5): 38–53. doi:10.3905/jpm.2017.43.5.038.

Bolton, P., and M. Kacperczyk. 2023. "Global Pricing of Carbon-Transition Risk." *Journal of Finance* 78 (6): 3677–3754. <https://onlinelibrary.wiley.com/doi/10.1111/jofi.13272>.

Bolton, P., M. Kacperczyk, and F. Samama. 2022. "Net-Zero Carbon Portfolio Alignment." *Financial Analysts Journal* 78 (2): 19–33. doi:10.1080/0015198X.2022.2033105.

Brinson, G. P., L. R. Hood, and G. L. Beebower. 1986. "Determinants of Portfolio Performance." *Financial Analysts Journal* 42 (4): 39–44. www.jstor.org/stable/4478947.

Chan, Y., K. Hogan, K. Schwaiger, and A. Ang. 2020. "ESG in Factors." *Journal of Impact and ESG Investing* 1 (1): 26–45. doi:10.3905/jesg.2020.1.1.026.

Fama, E. F., and K. R. French. 2010. "Luck versus Skill in the Cross-Section of Mutual Fund Returns." *Journal of Finance* 65 (5): 1915–47. doi:10.1111/j.1540-6261.2010.01598.x.

Garvey, G. T., M. Iyer, and J. Nash. 2018. "Carbon Footprint and Productivity: Does the 'E' in ESG Capture Efficiency as Well as Environment?" *Journal of Investment Management* 16 (1): 59–69.

Grinold, R. C. 2006. "Attribution." *Journal of Portfolio Management* 32 (2): 9–22. doi:10.3905/jpm.2006.611799.

Grinold, R. C., and R. N. Kahn. 2000. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*. New York: McGraw Hill.

Hodges, P., H. Ren, K. Schwaiger, and A. Ang. 2022. "Net-Zero Investing for Multi-Asset Portfolios Seeking to Satisfy Paris-Aligned Benchmark Requirements with Climate Alpha Signals." *Journal of Portfolio Management* 48 (4): 33–58. doi:10.3905/jpm.2022.1.334.

Hsu, P.-H., K. Li, and C.-Y. Tsou. 2023. "The Pollution Premium." *Journal of Finance* 78 (3): 1343–92. doi:10.1111/jofi.13217.

Intergovernmental Panel on Climate Change. 2023. *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Geneva. www.ipcc.ch/report/ar6/syr/.

Jensen, M. C. 1968. "The Performance of Mutual Funds in the Period 1945–1964." *Journal of Finance* 23 (2): 389–416. doi:10.1111/j.1540-6261.1968.tb00815.x.

Kazdin, J., K. Schwaiger, V.-S. Wendt, and A. Ang. 2021. "Climate Alpha with Predictors Also Improving Company Efficiency." *Journal of Impact and ESG Investing* 2 (2): 35–56. doi:10.3905/jesg.2021.1.030.

Le Guenedal, T., and T. Roncalli. 2022. "Portfolio Construction with Climate Risk Measures." Working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3999971.

Lee, C. M. C., S. T. Sun, R. Wang, and R. Zhang. 2019. "Technological Links and Predictable Returns." *Journal of Financial Economics* 132 (3): 76–96. doi:10.1016/j.jfineco.2018.11.008.

Lundberg, S. M., and S.-I. Lee. 2017. "A Unified Approach to Interpreting Model Predictions." From the 31st Conference on Neural Information Processing Systems (NIPS 2017). https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf.

Moehle, N., S. Boyd, and A. Ang. 2022. "Portfolio Performance Attribution via Shapley Value." *Journal of Investment Management* 20 (3): 33–52.

Shapley, L. S. 1951. "Notes on the n-Person Game-II: The Value of an n-Person Game." RAND Corporation. www.rand.org/pubs/research_memoranda/RM0670.html.

Shapley, L. S. 1953. "A Value for n-Person Games." In *Contributions to the Theory of Games (AM 28)*, vol. II, edited by H. W. Kuhn and A. W. Tucker, 307–18. Princeton, NJ: Princeton University Press. www.jstor.org/stable/j.ctt1b9x1zv.24.

Trinks, A., G. Ibikunle, M. Mulder, and B. Scholtens. 2022. "Carbon Intensity and the Cost of Equity Capital." *Energy Journal* 43 (2): 181–214. doi:10.5547/01956574.43.2.atr.

Young, H. P. 1985. "Monotonic Solutions of Cooperative Games." *International Journal of Game Theory* 14 (2): 65–72. <https://link.springer.com/article/10.1007/BF01769885>.