

Reducing carbon intensity in portfolios: Better news than you think.

February 2023

For financial adviser use only





Contents


Introduction	01
Carbon by sector	02
Choice of carbon metrics	04
The impact of reducing carbon intensity in portfolios	06
Properties of carbon-reduced portfolios: returns, sector tilts, stock selection	08
How do we interpret and use all of this?	14
Conclusion	17

Introduction

As investment managers commit to aligning their portfolios to net zero by 2050, it's more important than ever to understand the investment impact of reducing carbon exposures in portfolios.

Regulators, investor pressure, and industry collaborations¹ have led to significant improvements in corporate disclosure of greenhouse gas (GHG) emissions², so these metrics are among the most reliable and broadly-available data points for quantifying the environmental impact of listed equity portfolios. Chart 1 below shows the extent of reporting of carbon exposure over time – how it has grown as corporate responsibility, public scrutiny and regulatory requirements have evolved.

Chart 1: Proportion of firms reporting scope 1 & 2 carbon emissions³



Source: MSCI, Realindex database. Data points shown are at 30 June each year.

Against this background, quantitative portfolio construction tools provide an efficient way to adjust equity portfolios for a given carbon reduction target. In particular, combining a multi-factor risk model with a portfolio optimiser allows us to find a portfolio which achieves a target carbon reduction at the minimum possible *ex ante* tracking error.

In this article, we look at the metrication of carbon emissions for stocks and illustrate the effect of carbon reduction over time on a selection of capitalisation-weighted benchmark portfolios, in terms of risk, return, and various other characteristics such as sector exposures. This is primarily for illustration; in practice we already implement this in our Diversified Alpha strategies, and plan to implement in our value strategies in the near future. A companion paper to this one discusses the practical implementation in our value strategies in detail.

We start by noting that carbon emissions are highly concentrated in only a few sectors – utilities, materials and energy. These sectors have historically been recognised as aligned to value style investing, relative to more growth-like sectors – health care, information technology and consumer discretionary. We measure carbon exposure using current or historical values, rather than forecast values, which reflects the nature of these higher carbon emitting sectors. Growth stocks rely much more on future outcomes and so these sectors exhibit less exposure to current measures.

We then briefly discuss why we chose *carbon intensity* as our measure of GHG exposure, freely acknowledging that other better measures may exist in other specific uses. The complexity of which metric to choose is such that it almost merits a paper on its own.

Our central work here examines the impact of reducing carbon intensity by a fixed percentage (from 10% to 90%) of known cap-weighted benchmarks. Somewhat surprisingly, the tracking error induced by this, the alpha generated or lost, and the resulting differences in portfolio make-up, are small for even moderate reductions in carbon intensity. We examine how this might be due to intra-sector stock selection rather than inter-sector movement, and show that large reductions in carbon intensity do indeed generate large sector rotations. We outline some attribution of this effect as well.

1. In particular, the Task Force on Climate-related Financial Disclosures (<https://www.fsb-tcfd.org/>) and the Carbon Disclosure Project (<https://www.cdp.net/en>)
2. Following common industry practise, we loosely refer to "carbon" and "carbon emissions" to mean total greenhouse gas emissions, which include carbon dioxide and methane, among other gases. The emissions from the various greenhouse gases are converted to their carbon dioxide equivalent, based on the global warming potential (GWP) of each greenhouse gas.
3. Again note the use of "carbon emissions" when we mean GHG. Scope 1 are direct GHG emissions, while Scope 2 are indirect GHG emissions caused by purchased electricity, steam, heat and cooling (see <https://www.epa.gov/climateleadership/scope-1-and-scope-2-inventory-guidance>). Scope 3 GHG emissions arise from upstream and downstream activities due to a firm, most importantly from usage of products generated by a firm – for example, petroleum refiners and automobile manufacturers. All are converted by convention to carbon dioxide equivalent as noted above.




Carbon by sector

Where does the carbon come from? Scope 1 & 2 carbon emissions are heavily concentrated in certain industries, particularly utilities, materials, and energy. We can see this simply by summing the total Scope 1 & 2 carbon emissions by industry. Chart 3 below shows this – beyond these three sectors, very little carbon is emitted. Scope 1 emissions (primary emissions from actual production) also dominate Scope 2.

As noted in the footnote above, Scope 1 GHG emissions are *directly due* to firm operations, while Scope 2 are *indirectly caused* by purchased electricity, steam, heat and cooling. Scope 3 emissions (still poorly reported and somewhat unreliable) arise from upstream and downstream impacts of the firm's activities, across the 15 categories depicted in the schematic in Chart 2 below. Notable among Scope 3 categories is the *usage* of a firm's products, the primary driver of emissions in industries such as petroleum refining and automobile manufacturing.


Chart 2: Schematic of Scope 1, 2 and 3 Emissions



Source: Corporate Value Chain (Scope 3) Accounting and Reporting Standard Supplement, page 5, Greenhouse Gas Protocol.



Chart 3: Total CO₂ emissions for MSCI ACWI firms, million tons per annum (at June 30 2022)



Source: FactSet, MSCI Carbon Metrics, Realindex database

“...carbon emissions are highly concentrated in only a few sectors – utilities, materials and energy. These sectors have historically been recognised as aligned to value style investing.”



Choice of carbon metrics

Metrics need to reflect the observation that larger firms generate more GHG in total but are not necessarily less efficient in a relative sense. If we want to avoid introducing a size bias, we need to adjust for scale or “size”, however measured. Choices might include scaling by market capitalisation, total assets or some other measure like sales revenue or gross profit.

Carbon intensity, defined as Scope 1 and 2 emissions (millions of tonnes of CO₂ equivalent) per million dollars of sales, is a common choice of metric for quantifying, and also reducing, portfolio carbon. Carbon intensity for a portfolio is computed as the weighted average of asset-level carbon intensities (WACI):

$$WACI_{\text{portfolio}} = \text{“weighted average carbon intensity”} = \sum_{i \in \{\text{assets}\}} \left(\frac{\text{dollar holdings}_i}{\text{portfolio AUM}} \right) \cdot \left(\frac{\text{scope 1 \& 2 emissions}_i}{\text{sales}_i} \right)$$

That is, carbon intensity scales emissions by sales revenue. Measuring carbon intensity in this way has a number of advantages:

Carbon intensity measures the *output efficiency* with respect to carbon, i.e. how much carbon is produced per “unit” of production, with a unit being \$1m of sales revenue. In industries with homogenous products, such as electricity generation, oil extraction, or dairy farming, it is ideal to measure carbon efficiency by the emissions generated for each unit of product created, where the unit might be a kilowatt of electricity, a barrel of oil, or a litre of milk. Measuring production as total dollar sales allows comparison within industries with a diverse set of heterogeneous products, such as pharmaceuticals, car parts, and fast-moving consumer goods.

Both numerator and denominator of carbon intensity are measured on a contemporaneous basis: we compare a fiscal year’s carbon emissions to the sales that were made in that same year. Contrast that with a market capitalisation denominator, where last year’s carbon emissions are compared to the current market cap, which itself represents the net present value (NPV) of *future* profitability.

“Carbon intensity measures the output efficiency with respect to carbon, i.e. how much carbon is produced per “unit” of production, with a unit being \$1m of sales revenue.”



Carbon intensity is agnostic to the market's valuation of a firm, instead measuring properties of the firm itself. Contrast this with a market capitalisation denominator, which causes the carbon footprint of a firm to double if the share price halves (and vice versa). The high cross-sectional correlation of carbon footprints with valuation metrics is evidence of this.

Lastly, from a societal point of view, the carbon a firm emits into the atmosphere is a negative externality that we all have to bear. This negative needs to be weighed off against the positive value that society derives from the output of a firm; arguably that is measurable by what society is collectively willing to pay for the firm's output, i.e. its total sales revenue.

These advantages explain why carbon intensity is the portfolio metric recommended by the TCFD⁴ for quantifying portfolio carbon. We use carbon intensity in this study as our primary metric, however we are well aware that it is not ideal and even can be misleading in some cases.

Examples of the limitations of carbon intensity as a metric include:

- Carbon intensity can be highly sensitive to commodity prices or foreign exchange rates, which do not reflect the underlying carbon product of the firm.
- Firms which operate late in the value chain have artificially inflated sales, which leads to lower carbon intensity when compared to firms near the beginning of the value chain.
- Higher inflation lifts sales revenue, so it artificially drops carbon intensity, all else equal.
- Scope 1 and Scope 2 emissions are not cleanly separated. The Scope 1 emissions for one firm may be part of the Scope 2 emissions for another.
- Scope 3 emissions – for example, car usage after manufacture – are not counted at all.
- Scope 1 and 2 emissions can be misleading or artificially low if their generation is outsourced to other countries or regions (e.g., from developed economies to emerging economies).

4. As above, Task Force on Climate Related Financial Disclosures (<https://www.fsb-tcfd.org/>)




The impact of reducing carbon intensity in portfolios

Now we come to a somewhat surprising result on the impact on portfolios of reducing their carbon intensity. We look at the impact on tracking error and return of *reducing carbon intensity* – compared to well-known benchmarks. See Chart 4 and Table 1 below. Data is from December 2008 to June 2022.

We start by taking a range of capitalisation-weighted benchmarks, and optimally reduce carbon intensity by 10%, 20%, and more when compared to the benchmark⁵. By backtesting this process using quarterly rebalances, we then measure the *ex-post* or realised tracking error and returns of these carbon-reduced portfolios. The results of such backtests provide a frontier of (minimum) tracking error against carbon reduction. This is shown on Chart 4.

This is our central result: **for global portfolios (which have many stocks), we find surprisingly small tracking error for quite large reductions in carbon intensity.** For example a 40% reduction in carbon intensity relative to MSCI ACWI ex-AU would only have generated an *ex post* tracking error of 10bps p.a. across the whole period.⁶

Chart 4: Ex post annualised tracking error against cap weighted benchmarks (Dec 2008 to June 2022), as carbon intensity is reduced



Source: FactSet, MSCI Carbon Metrics, Realindex database

5. That is, attempting to minimise tracking error to the benchmark while applying a carbon intensity constraint.

6. We see similar results in our Value strategies – see the companion paper for details of this.



For Australia, with a smaller (and more resource oriented) market, the impact is high – for example, decreasing carbon intensity by 30% for the ASX 200 adds about 50bps of tracking error. For MSCI Emerging Markets, which has many more stocks, this is only 14bps. Results for MSCI World ex AU (developed markets only) and MSCI ACWI ex AU (developed and emerging markets) are almost identical. The tracking error impact on Australian stocks, especially small cap, is much greater than for other universes.

Table 1 extends on Chart 4. We note that in these historical backtests there is almost always a return improvement as well, when carbon intensity is reduced. While this is interesting, and tells us that a reduction in carbon would have added value in the past, we cannot necessarily extrapolate this into the future.

Table 1: Impact on active risk, active return and information ratio against cap weighted benchmarks (Dec 2008 to June 2022), as carbon intensity is reduced

	Index	10%	20%	30%	40%	50%	60%	70%	80%	90%
Active Risk (ex post, %p.a.)	ASX_200	0.10%	0.28%	0.53%	0.78%	1.11%	1.49%	1.89%	2.41%	3.14%
	ASX_SMALLS	0.21%	0.44%	0.75%	1.15%	1.59%	1.98%	2.54%	3.38%	4.94%
	MSCI_W_XAU	0.02%	0.05%	0.07%	0.11%	0.16%	0.24%	0.37%	0.58%	0.97%
	MSCI_EM	0.04%	0.08%	0.14%	0.24%	0.36%	0.54%	0.80%	1.25%	1.99%
	MSCI_ACWI_XAU	0.02%	0.04%	0.07%	0.10%	0.15%	0.23%	0.36%	0.57%	0.97%
	MSCI_W_SML	0.03%	0.06%	0.09%	0.14%	0.21%	0.33%	0.52%	0.83%	1.38%
Active Return (% p.a.)	ASX_200	0.02%	0.10%	0.24%	0.41%	0.70%	1.02%	1.34%	1.37%	0.63%
	ASX_SMALLS	-0.02%	0.02%	0.10%	0.26%	0.46%	0.87%	1.30%	2.03%	2.03%
	MSCI_W_XAU	0.01%	0.02%	0.04%	0.08%	0.14%	0.21%	0.30%	0.44%	0.83%
	MSCI_EM	0.02%	0.05%	0.07%	0.09%	0.14%	0.23%	0.41%	0.52%	0.75%
	MSCI_ACWI_XAU	0.01%	0.02%	0.05%	0.08%	0.14%	0.21%	0.32%	0.45%	0.80%
	MSCI_W_SML	0.00%	0.00%	0.00%	-0.01%	0.01%	0.06%	0.22%	0.52%	1.01%
Information Ratio	ASX_200	0.17	0.35	0.45	0.52	0.63	0.69	0.71	0.57	0.20
	ASX_SMALLS	-0.07	0.04	0.14	0.22	0.29	0.44	0.51	0.60	0.41
	MSCI_W_XAU	0.40	0.42	0.56	0.74	0.87	0.86	0.81	0.76	0.86
	MSCI_EM	0.53	0.59	0.51	0.37	0.38	0.43	0.51	0.42	0.38
	MSCI_ACWI_XAU	0.52	0.54	0.67	0.81	0.93	0.94	0.89	0.79	0.82
	MSCI_W_SML	-0.11	-0.04	-0.03	-0.06	0.04	0.19	0.43	0.63	0.74

Source: FactSet, MSCI Carbon Metrics, Axioma, Realindex database



Properties of carbon-reduced portfolios: returns, sector tilts, stock selection

The key question we have to ask here is “why does this happen?” Intuition suggests that reducing carbon intensity would have a marked effect on tracking error, and perhaps returns. We don’t really see this - Australia is the exception. Why?

To examine this, we examine the return and risk trade off, resulting sector tilts and stock selection as we increase the carbon intensity constraint. Of most interest is when we look at stock selection. We pay particular attention to the attributed contributions to return and risk, and how the reduction in carbon intensity takes place within sectors rather than across them.

Return and risk


We have no reason to expect these positive historical returns to carbon reduction will persist into the future. While returns to carbon reduction have generally been positive across the entire backtest period, this is not the intent of the carbon intensity reduction study and must be viewed as luck rather than skill.

Saying this, it is likely that some of these positive returns to carbon reduction are due to a repricing of high-carbon assets, as investors have realised the social and political headwinds such stocks face. The question then remains whether the current market is correctly pricing-in these headwinds, though unfortunately that is something we will only know in the future, with hindsight, as valuations continue to evolve dynamically in response to government policies, energy and carbon prices, and of course realised corporate earnings.

There is a significant amount of time series variation in the returns to carbon intensity reduction. Market events during 2021 and 2022H1, and in particular the Russia-Ukraine war, have driven up energy prices and inflation strongly, and as a result, fossil fuel firms have performed well, leading to the weakest 12-month performance for portfolio carbon reduction since the start of the carbon data. It will come as no surprise that the positive returns are negatively correlated with the price of oil.

Chart 5 below shows the active return (or value add) from carbon intensity reduction in these backtests. Here we extract two examples of carbon intensity reduction against the MSCI ACWI ex AU benchmark – 30% (tracking error 7bps) and 50% (tracking error 15bps). A small amount of extra tracking error (from 30% reduction to 50% reduction) has a marked impact on the return outcomes. The 50% reduction case has much higher alpha early in the sample but has worse returns later, when energy prices spiked in 2022.

Chart 5: Rolling 12-month active returns for 30% and 50% carbon intensity reductions
MSCI ACWI ex AU backtests, December 2008 to June 2022



Source: FactSet, MSCI Carbon Metrics, Axioma, Realindex database

Tracking error is a different story. We have already noted the tracking errors needed to achieve carbon reductions can be surprisingly low. Taken together with the returns discussed above, in some cases this has meant IRs (Information Ratios) for carbon reduction as high as 0.7 or 0.8. But this is not the key conclusion here.



Sector tilts

It is unsurprising that large reductions in carbon intensity entail substantial down-weighting of the utilities, materials, and energy sectors. See Table 2 Panel A below. For smaller reductions in carbon intensity, tracking error remains small as well, with small rotations away from utilities and materials and into financials, industrials and consumer staples.

Sectors are of course made up of industries, two levels further down in the GICS hierarchy⁷. Table 2 Panel B shows these results⁸. The impact at industry level is similar, with most of each sector's active weight change being driven by one or two industries within it: utilities is largely driven by electric utilities, materials by chemicals and construction materials, energy by oil, gas and consumables, and financials by banks and insurance. We only see rotation out of one industry and into another for the energy sector – out of oil and gas into energy equipment. The other sectors do not see this.

Table 2: Panel A: Average sector active weights at various carbon intensity reductions,
MSCI ACWI ex AU December 2008 to June 2022

Sector	10%	20%	30%	40%	50%	60%	70%	80%	90%
Utilities	-0.10%	-0.23%	-0.39%	-0.58%	-0.82%	-1.17%	-1.68%	-2.28%	-2.44%
Materials	-0.06%	-0.15%	-0.26%	-0.41%	-0.68%	-1.17%	-1.92%	-2.96%	-4.09%
Energy	0.01%	0.03%	0.05%	0.06%	0.07%	0.01%	-0.22%	-0.89%	-3.19%
Consumer discretionary	0.01%	0.01%	0.03%	0.04%	0.06%	0.08%	0.08%	0.05%	-0.32%
Information technology	0.00%	0.01%	0.01%	0.02%	0.03%	0.05%	0.10%	0.19%	0.71%
Real estate	0.02%	0.04%	0.07%	0.11%	0.16%	0.23%	0.32%	0.44%	0.29%
Health care	0.01%	0.03%	0.05%	0.08%	0.13%	0.22%	0.37%	0.64%	1.13%
Consumer staples	0.03%	0.07%	0.13%	0.19%	0.28%	0.41%	0.57%	0.66%	-0.09%
Communication services	0.02%	0.05%	0.08%	0.13%	0.20%	0.34%	0.55%	0.88%	1.47%
Industrials	0.03%	0.07%	0.12%	0.18%	0.27%	0.49%	0.86%	1.36%	1.71%
Financials	0.03%	0.06%	0.11%	0.18%	0.29%	0.53%	0.97%	1.91%	4.84%

Source: FactSet, MSCI Carbon Metrics, Realindex database

7. See <https://www.msci.com/our-solutions/indexes/gics>

8. There are some small rounding effects which mean that the sum of industry weights do not exactly equal the sectors weights.

Table 2: Panel B: Average industry active weights at various carbon intensity reductions,
MSCI ACWI ex AU December 2008 to June 2022

		10%	30%	50%	70%	90%
Energy	Energy Equipment & Services	0.00%	0.01%	0.05%	0.26%	1.30%
	Oil, Gas & Consumable Fuels	0.01%	0.03%	0.02%	-0.48%	-4.49%
Financials	Banks	0.01%	0.05%	0.15%	0.58%	2.35%
	Capital Markets	0.00%	0.00%	0.01%	0.11%	0.69%
	Consumer Finance	0.00%	0.00%	0.00%	0.00%	0.05%
	Diversified Financial Services	0.00%	0.00%	-0.02%	-0.19%	-0.20%
	Insurance	0.01%	0.05%	0.13%	0.43%	1.84%
	Mortgage Real Estate Investment Trusts (REITs)	0.00%	0.01%	0.02%	0.04%	0.11%
	Thrifts & Mortgage Finance	0.00%	0.00%	0.01%	0.00%	0.00%
Materials	Chemicals	-0.01%	-0.07%	-0.30%	-0.96%	-1.79%
	Construction Materials	-0.05%	-0.17%	-0.27%	-0.36%	-0.38%
	Containers & Packaging	0.00%	0.01%	0.00%	-0.09%	-0.21%
	Metals & Mining	0.00%	-0.02%	-0.08%	-0.37%	-1.56%
	Paper & Forest Products	0.00%	-0.01%	-0.04%	-0.14%	-0.15%
Utilities	Electric Utilities	-0.07%	-0.33%	-0.71%	-1.20%	-1.54%
	Gas Utilities	0.01%	0.03%	0.06%	0.03%	0.04%
	Independent Power and Renewable Electricity Producers	-0.04%	-0.09%	-0.15%	-0.14%	0.06%
	Multi-Utilities	-0.01%	-0.03%	-0.08%	-0.40%	-0.92%
	Water Utilities	0.01%	0.03%	0.05%	0.04%	-0.09%

Source: FactSet, MSCI Carbon Metrics, Realindex database

Stock selection

We noted above that smaller reductions in carbon intensity see small changes in sector positions. So what is driving this? It seems most of the carbon reduction comes not from sector allocations but *from stock selection within sectors*.

We look at this in two ways:

- A return and risk decomposition across all levels of carbon intensity reduction, to show the extent of the impact of stock selection or “specific” risk and return
- A Brinson attribution drilldown into a single case (50% reduction) to partition the carbon reduction into within- and between-sector effects

Using Axioma, we decompose the sources of active return and risk for carbon intensity reductions from 10% to 90%. This tells us how much is driven by common risk factors (for example: country, style, industry) and how much is due to stock specific effects alone. Table 3 shows this.

Table 3: Annualised Returns and Risk of attributed return components for carbon reduction scenarios MSCI ACWI ex AU, Dec 2008 to Jun 2022

Panel A: Returns decomposition (%p.a.)

	c10	c20	c30	c40	c50	c60	c70	c80	c90
Portfolio	10.37%	10.38%	10.41%	10.44%	10.50%	10.57%	10.67%	10.81%	11.16%
Benchmark	10.36%	10.36%	10.36%	10.36%	10.36%	10.36%	10.36%	10.36%	10.36%
Active	0.01%	0.02%	0.05%	0.08%	0.14%	0.21%	0.32%	0.45%	0.80%
Active breakdown									
Specific	0.01%	0.02%	0.03%	0.06%	0.11%	0.18%	0.29%	0.42%	0.67%
Factor	0.00%	0.01%	0.02%	0.03%	0.03%	0.04%	0.03%	0.03%	0.13%
Factor breakdown									
Country	0.00%	0.00%	0.00%	0.00%	-0.01%	-0.01%	-0.02%	-0.04%	-0.06%
Currency	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%
Industry	0.00%	0.01%	0.02%	0.03%	0.04%	0.04%	0.03%	0.02%	0.04%
Local	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Market	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.01%
Style	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%	0.03%	0.13%

Panel B: Tracking Error decomposition (%p.a. *ex post*)

	c10	c20	c30	c40	c50	c60	c70	c80	c90
Portfolio	10.56%	10.56%	10.57%	10.56%	10.55%	10.56%	10.58%	10.67%	10.83%
Benchmark	10.56%	10.56%	10.56%	10.56%	10.56%	10.56%	10.56%	10.56%	10.56%
Active	0.02%	0.04%	0.07%	0.10%	0.15%	0.23%	0.35%	0.57%	0.97%
Active breakdown									
Specific	0.02%	0.04%	0.06%	0.08%	0.11%	0.17%	0.27%	0.41%	0.65%
Factor	0.01%	0.02%	0.03%	0.05%	0.07%	0.11%	0.17%	0.29%	0.56%
Factor breakdown									
Country	0.00%	0.01%	0.01%	0.02%	0.02%	0.03%	0.05%	0.08%	0.16%
Currency	0.00%	0.00%	0.01%	0.01%	0.01%	0.02%	0.04%	0.07%	0.16%
Industry	0.01%	0.02%	0.03%	0.05%	0.08%	0.12%	0.19%	0.31%	0.60%
Local	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.02%
Market	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.02%	0.02%
Style	0.00%	0.01%	0.01%	0.02%	0.04%	0.05%	0.08%	0.14%	0.25%

Source: Realindex database, Axioma



This classic attribution analysis (Table 3) on the backtested returns supports the high degree of stock selection in the carbon-reduced portfolios. The largest contributor to both risk and return is usually “specific” (i.e. idiosyncratic stock positions). After that, industry risk is the most notable, and industries are also a small positive contributor to returns. Style factors (size, growth, and so on) only start to have an impact for high levels of carbon intensity reduction.

A caveat to such a returns attribution analysis is that carbon intensity could naturally be viewed as a systematic risk factor, yet is not included in the multi-factor risk model that we used. Thus it is unsurprising when targeting reduced carbon exposure that a large proportion of measured risk exposures fall into the specific risk bucket. In future the risk (volatility) of the carbon factor may well become higher, particularly if there is high uncertainty surrounding commodity prices, fossil fuels, and government policy on emissions. In that case we would expect the returns of high carbon stocks to become more correlated to each other, through common exposure to the latent risk factor, and absent a change in the risk model it would understate portfolio risk due to mistakenly assuming uncorrelated residual returns for these climate-exposed stocks.

A final note on the choice of carbon metric. The results presented above were based on reducing carbon intensity. We might prefer to reduce carbon footprint, notwithstanding its high dependence on stock prices. Alternatively, we might insist on reducing both intensity and footprint at the same time. This is a topic for future discussion.

“The largest contributor to both risk and return is usually “specific” (i.e. idiosyncratic stock positions).”

With sectors being such a key determinant of firm-level carbon emissions, Brinson analysis of a portfolio's carbon intensity relative to a benchmark is a helpful way to understand the main sources of carbon. This second approach partitions or divides *carbon intensity* into the portions (a) due to stock selection within each sector, and (b) due to allocation between sectors. Table 4 below shows this.

Table 4: Brinson attribution of 50% carbon-reduced portfolio relative to MSCI ACWI ex AU.

(Carbon intensity decomposition, as at 30 June 2022, units are tonnes per \$m of sales)

Sector	Portfolio active weight	Portfolio carbon intensity	Benchmark carbon intensity	Portfolio contribution	Benchmark contribution	Active contribution	Allocation	Selection	Interaction	Residual
Utilities	-0.53%	468.7	1762.9	12.8	57.3	-44.5	-9.3	-42.0	6.8	0.0
Materials	-0.54%	304.1	672.3	10.1	25.9	-15.8	-3.6	-14.2	2.0	0.0
Energy	-0.03%	488.1	617.1	24.1	30.6	-6.6	-0.2	-6.4	0.0	0.0
Industrials	0.11%	72.8	122.7	7.4	12.3	-4.9	0.1	-5.0	-0.1	0.0
Consumer Discretionary	-0.08%	30.0	36.7	3.1	3.9	-0.7	0.0	-0.7	0.0	0.0
Real Estate	0.10%	70.3	96.7	2.1	2.7	-0.7	0.1	-0.7	0.0	0.0
Information Technology	0.16%	19.6	20.6	4.2	4.4	-0.2	0.0	-0.2	0.0	0.0
Financials	0.32%	18.3	19.5	2.5	2.6	-0.1	0.1	-0.2	0.0	0.0
Health Care	0.07%	17.3	18.0	2.5	2.6	-0.1	0.0	-0.1	0.0	0.0
Consumer Staples	0.21%	45.9	47.5	3.7	3.7	0.0	0.1	-0.1	0.0	0.0
Communication Services	0.21%	14.4	14.4	1.1	1.1	0.0	0.0	0.0	0.0	0.0
TOTAL	0.00%	73.5	147.1	73.5	147.1	-73.6	-12.7	-69.7	8.7	0.0

Source: FactSet, MSCI Carbon Metrics, Realindex database

Here we use the 50% carbon reduction scenario backtest in MSCI ACWI ex AU as an example. Portfolio carbon intensity at 30 June 2022 is 73.5 (tonnes of CO2 emitted per USD million dollar of sales), compared with 147.1 for the benchmark. This is a reduction of 73.6 at this point in time (50% as indicated).

The Brinson analysis shows that:

- More than 90% of the carbon intensity reduction comes from the three sectors: utilities (44.5 out of the 73.6), materials (15.8 out of the 73.6) and energy (6.6 out of the 73.6)
- Of this, the vast majority is due to *stock selection rather than sector allocation* within utilities (42.0), materials (14.2) and energy (6.4). Very little of the reduction is based on sector level allocation.

The stock selection within utilities reduces this sector's carbon intensity from 1,763 tonnes/\$m to 467 tonnes/\$m, a reduction of almost 75%, whereas in the healthcare or IT sectors the reduction is only around 5%.

Note that as carbon intensity is reduced, we gradually sell down rather than out of stocks. Stock weights are reduced slowly (as carbon intensity is reduced) until the point of exclusion, rather than being excluded immediately.⁹

9. For the MSCI ACWI ex AU:

At 30% reduction, only one third of underweights are exclusions

At 50% reduction, approximately one half of underweights are exclusions

At 70% reduction, as many as two thirds of underweights are exclusions




How do we interpret and use all of this?

What is it about carbon intensity that allows us to rotate between stocks within sectors to reduce it, rather than making large changes to sector weights? And can we understand how this might occur in the more general setting that includes alpha forecasts? We think that there are three primary points to make here:

1. There are broad ranges of carbon intensities within each sector - many stocks with little or no carbon intensity and a few stocks with a great deal. As an example, Chart 6 below shows the carbon intensity of the stocks in the utilities sector at Dec 2022. Only a few stocks have very large carbon intensity, so switching these for lower carbon intensity stocks within this sector will have a dramatic effect. This holds true for other sectors as well, especially materials and energy.
2. Sectors are important drivers of risk, and tracking error, in portfolios. So to keep tracking error low against the cap-weighted benchmark when reducing carbon intensity, the easiest way is to *switch between stocks in the same sector*. But how easy is this?

Chart 7 overleaf shows this by asking the following question: if we keep sector weights exactly the same while reducing carbon intensity by excluding the worst stocks in that sector, what percentage of the sector needs to be excluded? For low levels of carbon intensity, only a small percentage of stocks need to be omitted. For example, in the energy sector we can reduce carbon intensity by 30% by excluding 11% of stocks (by weight), and for other sectors the proportion is lower (typically 5-10% of stocks). For larger reductions, the percentage increases.


Chart 6: Carbon intensity by stock in the utilities sector, MSCI ACWI ex AU, Dec 2022



Source: FactSet, MSCI Carbon Metrics, Realindex database

“Only a few stocks have very large carbon intensity, so switching these for lower carbon intensity stocks within this sector will have a dramatic effect.”

Chart 7: Percentage of stocks to exclude so as to reduce carbon intensity AND retain sector weight, MSCI ACWI ex AU, Dec 2022



Source: FactSet, MSCI Carbon Metrics, Realindex database

3. In practice, when we consider alphas and carbon intensity together, we would wish to combine their trade-off in a single step. Said another way, this amounts to us showing that dropping high carbon intensity stocks in favour of lower carbon intensity has little effect on portfolio alpha.

Let's use an example. Build an optimal portfolio¹⁰, and then drop the top 10 highest carbon intensity stocks in favour of the next best 10 alpha stocks. Carbon intensity drops sharply but alpha exposure hardly moves at all. Note that we deliberately avoid constraining sectors here to concentrate on the alpha/carbon intensity trade-off.

Chart 8 shows the scatterplot of alphas – blue markers are included stocks, grey is excluded stocks, red crosses are the 10 dropped high carbon intensity stocks and green pluses are the 10 next best alpha stocks that are added. Note the log scale for carbon intensity. Table 5 shows that carbon intensity drops sharply but alpha exposure drops but only a small amount.

Table 5: Alpha and carbon intensity change after 10 stock switch

	Average alpha	Average carbon intensity
Starting portfolio	1.282	207.4
“Switched” portfolio	1.246	117.4

Source: FactSet, MSCI Carbon Metrics, Realindex database

Chart 8: Scatterplot of stocks: alpha against carbon intensity Global (ex-Aus) Universe, Dec 2022



Source: FactSet, MSCI Carbon Metrics, Realindex database

10. For illustrative purposes here the “optimal” portfolios is a simple one: it is the maximum expected return portfolio that consists of 250 equally-weighted positions.

Technical aside:

A more sophisticated version of this example recoups almost two-thirds of the (fairly small) alpha that was lost from the first approach. As this is done internally within portfolio optimisers, we can treat the carbon reduction as a linear constraint and introduce a Lagrangian which creates a penalty to alpha for carbon intensity – sometimes termed a shadow price. In this case the shadow price needed to match the carbon reduction of the example above is a penalty of 0.42 to alpha for each extra 1,000 tonnes per million of carbon intensity. The dotted green line in Chart 8b below shows the resulting line of equivalence (curved due to the use of a log scale on the x-axis). Note the new set of stocks that are dropped.

Chart 8b: A more efficient version of chart 8, using a shadow price of carbon to penalise alpha



FactSet, MSCI Carbon Metrics, Realindex database

Table 5b: Alpha and carbon intensity change after 10 stock switch

	Average alpha	Average carbon intensity
Starting portfolio	1.282	207.4
"Switched" 10 stocks (as above)	1.246	117.4
Improved version using shadow price	1.270	117.4

Source: FactSet, MSCI Carbon Metrics, Realindex database



Conclusion

This analysis looks at the impact of reducing carbon intensity on various universes of stocks, with the goal of seeing how much tracking error is created by this reduction, and where it comes from.

Firstly, we argue for the merits of carbon intensity (carbon emissions divided by sales) as a primary metric to focus on when reducing carbon, in the absence of anything better, and noting its widespread adoption.

We then find that significant carbon reductions can be achieved at surprisingly low tracking errors in global portfolios. To do this, we use a multi factor risk model and optimiser combined with Scope 1 & 2 emissions data. Most interesting is that much of the resulting carbon reductions, and portfolio risk, arise from stock positions within sectors, rather than sector or style tilts.

Having said that, we see that exposures to higher-carbon sectors are reduced somewhat, particularly utilities, particularly when large reductions are targeted. These results suggest it is the idiosyncratic nature of carbon emissions which allow such low tracking error when the available investment universe is large.

Historical returns have generally been positive to such an approach, but this is probably incidental – it is certainly not the motivation for this study. Note also that with the Ukraine-Russia war causing global energy supply pressures that have pushed up fossil fuel prices sharply, recent "alpha" has been negative, again incidentally.

In summary, a desire for moderate reductions in carbon intensity can exploit the fact that optimal portfolios seem to require only a small move towards lower carbon stocks, not a sector level rotation, resulting in a rather small impact on performance outcomes (low tracking error against a portfolio without carbon intensity reduction). Substantial reductions in portfolio carbon – a strong move towards lower carbon stocks – cause higher tracking error, with the attendant risk to long-term performance outcomes.

In this article we have quantified the tracking errors required for a given carbon reduction, and also shown the associated historical returns. The approach described above can create an implementable carbon reduction overlay to an existing model portfolio. In practice, we would generally add additional constraints as safeguards into such a process, both around individual holdings and around sector and factor exposures. In some cases it may be more efficient – from both a tracking error and transaction cost perspective – to integrate such a carbon reduction constraint into the wider portfolio construction process, in which case the results above provide some guidance on the expected impact.

The practical implementation of this work in Realindex portfolios is underway. We expect to be communicating the results and suggested process changes through a directed research paper to clients in the near future.

Important Information

This material has been prepared and issued by First Sentier Investors (Australia) IM Ltd (ABN 89 114 194 311, AFSL 289017) (FSI AIM, Realindex), which forms part of First Sentier Investors, a global asset management business. First Sentier Investors is ultimately owned by Mitsubishi UFJ Financial Group, Inc (MUFG), a global financial group. A copy of the Financial Services Guide for FSI AIM is available from First Sentier Investors on its website.

This material contains general information only. It is not intended to provide you with financial product advice and does not take into account your objectives, financial situation or needs. Before making an investment decision you should consider, with a financial advisor, whether this information is appropriate in light of your investment needs, objectives and financial situation.

Any opinions expressed in this material are the opinions of the individual author at the time of publication only and are subject to change without notice. Such opinions: (i) are not a recommendation to hold, purchase or sell a particular financial product; (ii) may not include all of the information needed to make an investment decision in relation to such a financial product; and (iii) may substantially differ from other individual authors within First Sentier Investors.

To the extent permitted by law, no liability is accepted by MUFG, Realindex nor their affiliates for any loss or damage as a result of any reliance on this material. This material contains, or is based upon, information that Realindex believes to be accurate and reliable, however neither MUFG, Realindex nor their respective affiliates offer any warranty that it contains no factual errors. No part of this material may be reproduced or transmitted in any form or by any means without the prior written consent of Realindex.

Any performance information has been calculated using exit prices after taking into account all ongoing fees and assuming reinvestment of distributions. No allowance has been made for taxation. Past performance is not indicative of future performance.

Copyright © First Sentier Investors, 2023

All rights reserved.