Paris-aligned commodities

Taylor Dufour[1], He Ren[2], Andrew Ang[3], Filip Mena-Berlin[4], Mateo Gomez[5] and Katharina Schwaiger[6]

Abstract

We measure carbon emissions in commodities portfolios and propose a methodology to reduce them. We estimate emissions using both government/industry data and supply chain insights by applying natural language processing techniques. To reduce portfolio emissions, we tilt investments towards (away from) commodities with lower (higher) emissions per dollar value, achieving a 50% reduction in carbon intensity relative to the S&P GSCI benchmark. We show that a sustainable commodities return can closely track the benchmark while demonstrating similar inflation hedging characteristics – when including or excluding energy sector commodities.

JEL classifications: Q02, Q56, G11.

1 BlackRock, taylor.dufour@blackrock.com

2 BlackRock, he.ren@blackrock.com

3 Columbia University, aa610@columbia.edu

4 BlackRock, filip.menaberlin@blackrock.com

5 BlackRock, mateo.gomez@blackrock.com

6 BlackRock, katharina.schwaiger@blackrock.com

1. Introduction

Commodities investments have been an integral part of institutional portfolios since the early 2000s – due to legislation allowing financial institutions easier access to invest in commodities futures (Commodity Futures Modernization Act of 2000) and several influential academic studies, such as Erb and Harvey (2006) and Gorton and Rouwenhorst (2006). These studies showed that commodities futures had low or negative correlations with equities and bonds and a partial ability to hedge inflation, and that both long-only and long-short commodity investments could produce attractive risk-adjusted returns. While sizeable bodies of literature are examining sustainable investments in traditional asset classes like equities and bonds, the attention is now turning to making alternative investments more sustainable.

Integrating sustainability into commodity investing is essential because commodities underpin key industrial and consumption processes. Many commodities are treated as sectors themselves in climate forecasting models by scientific bodies such as the Intergovernmental Panel on Climate Change (IPCC), the International Energy Agency (IEA) and the Science Based Targets Initiative (SBTi). In fact, the Sectoral Decarbonisation Approach[7] explicitly models several commodity sectors with homogeneous products (eg crude oil, aluminium) or where activity is relatively homogeneous (eg wheat, beef), projecting pathways for carbon emissions, physical risk and transition risk (see, for example, Krabbe et al (2015)).

However, particular challenges to building a sustainable investment framework for commodities go from measurement to usage and impact. Commodities have complex lifecycles and, at various points in their supply chains, they can have very different carbon efficiency profiles. For example, natural gas is the most carbon-intensive fossil fuel to extract, but it is the most efficient fossil-fuel-based source of energy. Therefore, measuring emissions solely during the extraction or use phases would yield opposite insights, underscoring the importance of considering the full supply chain. Additionally, the heterogeneity of commodities – both across and within sectors – can make it difficult to standardise the units of measurement to facilitate comparisons. Finally, commodities are both inputs and outputs and, unlike equities or corporate bonds, the decision to buy or sell a commodity in a sustainable portfolio can be fraught with unintended side effects. For example, natural gas can be used as a transition fuel in emerging markets to replace coal. Excluding it entirely may bias portfolios towards developed markets with more mature renewable infrastructure and therefore unintentionally penalising developing economies. We seek to account for each challenge when building a sustainable commodities portfolio.

In this paper, we present a framework for integrating sustainability into commodity investments, broadly following the requirements set by the European Union Technical Expert Group on Sustainable Finance (EU TEG). Starting with a broad commodity universe, the framework has three steps: tilting into the commodities which have attractive carbon profiles and are best positioned for the transition, removing the worst offenders, and ensuring the portfolio targets at least a 50% carbon reduction.

To achieve this, we must estimate emissions for each commodity, which we do by both using proprietary techniques (discussed later) and pooling and aggregating

[7] For an introduction see: CDP, WRI, WWF (2015): Quick Guide to the Sectoral Decarbonization Approach.

data from a range of government, producer and industry bodies like the US Department of Agriculture (USDA), the US Department of Energy and the World Gold Council. Although these estimates are not standardised across commodities and some data are missing, they are published by independent and reputable third parties and, thus, we use them to measure the carbon profile of our portfolio.

We complement these data with emissions estimated using a proprietary approach that leverages the wealth of corporate emissions data published regularly by companies across the globe. The approach consists of the use of natural language processing (NLP) techniques on conference call transcripts, allowing us to map the companies that are most closely associated with certain commodities – i.e. organisations that are directly involved in the commodity's supply chain as either producer or consumer.

With a "supply chain map" at hand we can glean current and forward-looking insights into individual commodities. First, we estimate historical emission intensities for each commodity. We find that direct emission estimates and supply chain estimates are most highly correlated for nickel and zinc and the least correlated for lean hogs and sugar.

Next, we use the supply chain map to predict which commodities will play a significant role in a low-carbon economy. These predictions are also relevant for Paris-Aligned Benchmark (PAB) investment portfolios. To identify such commodities, we create a custom ESG dictionary and use NLP to assign each commodity an ESG score. This score reflects how often sustainability themes related to the commodity appear in corporate discussions, such as conference call transcripts. By doing this, we identify those that are seen as most sustainable and likely to be important in a greener economy.

Using these data, we then construct a sustainable commodities overlay. In our empirical work, we apply our methodology to the well-known S&P GSCI commodities index, though the approach can be applied to any benchmark (see alternatives discussed by Miffre (2012) for example). The optimisation tilts the portfolio towards commodities with the lowest supply-chain-implied emission intensities and those which are most likely to be used in greening the economy. In addition, we set as an optimisation constraint a carbon intensity reduction of 50% relative to the index, which is achievable. This reduction can be adjusted lower or higher, depending on an investor's preference for sustainability. Additional constraints, such as a 20% cap on individual holdings, can also be imposed.

We also conduct exercises where energy commodities are excluded from our portfolio. Deciding whether to include or exclude energy in a climate-focused portfolio is complex and nuanced. On the one hand, energy is an essential commodity for modern life. The energy sector also has highly desirable inflation hedging properties for investors' portfolios. In our empirical results, we show that it is possible to reduce carbon emission intensities by 50% even when energy is present.[8] On the other hand, some argue that there are ethical reasons to avoid involvement with energy commodities that contribute to further carbon emissions (see Richardson

[8] Even in climate scenarios where overall greenhouse gas emissions are significantly reduced to limit global warming, the IEA (2021) predicts that approximately one fifth of energy is still supplied by fossil fuels. Tobler et al (2021) argue that investors not participating in energy futures markets contribute to higher price volatility and that financial investors have a role to play in lowering volatility to ensure food and energy security.

(2017) for a summary). In our exercises, we therefore opt for both the inclusion and exclusion of energy in the optimisation process.

An important caveat is that our analysis assumes commodity returns as given and focuses primarily on the sustainability of commodities, except where returns factor into the risk-return trade-off in the final optimisation step. The study of commodity returns has a long history, dating back to Keynes (1923). In our approach, we treat carbon emissions as exogenous; however, in equilibrium models (like Williams and Wright (1991); Deaton and Laroque (1992)), emissions would be reflected in commodity prices. Additionally, we have not incorporated carbon futures in our analysis. Including carbon futures and other types of offsets could further reduce the carbon intensities reported in our results. As carbon markets develop, we anticipate that carbon futures will increasingly influence commodity prices (see Chevallier (2009)).

Furthermore, our analysis is based exclusively on the S&P GSCI commodity benchmark, so we cannot make broader conclusions about how a commodity sleeve should be integrated into a general asset allocation problem – especially one that incorporates sustainability. While many papers document that commodities diversify equity and bond portfolios (like the seminal Erb and Harvey (2006) and Gorton and Rouwenhorst (2006) papers), fewer papers discuss sustainability for the entire portfolio – Hodges et al (2022) being a notable exception.

The remainder of this paper is organised as follows. Section 2 introduces the two methods we use to measure commodity carbon emissions, contrasting direct emission estimates with our NLP-based supply chain estimates. Section 3 outlines the sustainable optimisation framework relative to the S&P GSCI, which reduces carbon emissions and incorporates the forward-looking green commodity signals. Section 4 demonstrates how a Paris-aligned commodities portfolio can achieve its objectives within a broader multi-asset portfolio, offering similar diversification and inflation hedging benefits – both with and without energy. Finally, section 5 concludes.

2. Commodity emissions

Net zero investment frameworks for traditional asset classes are designed to reduce emissions in the real economy. This notion is directly applicable to commodities, where direct emissions data are available to quantify the environmental impact of each commodity. To ensure consistency, we standardise direct emissions – originally reported in a variety of units, including volume, weight and energy – by denominating them in dollar terms. This approach aligns with the widely accepted sales-denominated equity emissions which express greenhouse gas (GHG) emissions relative to economic value.

2.1. Industry direct emissions

Growing efforts have been witnessed in recent years among governments, industries and companies to study and estimate the environmental impact of the commodity lifecycle. However, data collection and analysis remain highly sector specific. While sectors like energy and grains have richer coverage, data are still sparse in metals and soft commodities such as livestock, crops and grains. Even within sectors like grains and energy, it is still challenging to model emissions for derivative or co-product commodities, such as soymeal and soy oil from soybean crushing and gasoline and diesel from cracking.

We gather emission estimates for most commodities spanning the grain, energy and metals sectors from sources such as academia, governments and the sell side. Graph 2.1 presents the GHG emissions for the investment universe, showing that the energy sector is the least carbon-efficient (Graph 2.1, six right-most bars), with natural gas being the least efficient commodity. These estimates are in line with intuition, though with sparse observations and not yet standardised globally (albeit some early efforts in doing so are the GHG Protocol, the United Nations Conference on Trade and Development and the Institutional Investors Group on Climate Change). Table 2.1 summarises the data sources used. It is also worth noting that emission estimates for energy include combustion, which is the primary source of emissions in this sector.

GHG emissions — Graph 2.1

Per metric ton; number

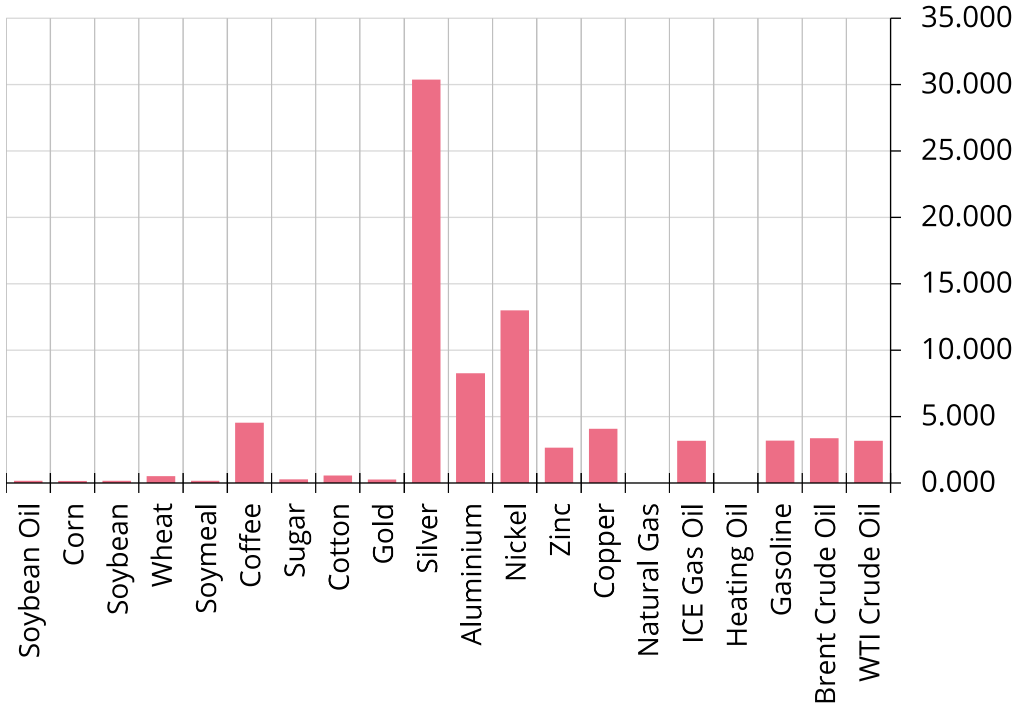

Source: see Table 2.1; BlackRock.

Data sources for commodity emissions Table 2.1

Commodity	Data Source
Wheat, Soybean, Corn, Lean Hog, Live Cattle	USDA, University of Arkansas
Oil, Gasoline, Diesel, Natural Gas	US Department of Energy
Copper, Aluminium, Zinc, Nickel	Industry Associations (Copper Alliance, Aluminium Association, International Zinc Association, Nickel Institute)
Silver	Goldman Sachs
Gold	World Gold Council
Cotton	Cotton Incorporated, University of Arkansas
Sugar	Seabra et al., USDA
Cocoa	World Cocoa Foundation
Coffee	Giraldi-Daz et al.

Source: BlackRock.

2.2. Supply chain map

While direct emissions data are available for most of our investment universe, estimation methods are still nascent and lack global standardisation. To address these gaps, we introduce a proprietary approach which leverages corporate emissions data from supply chain companies to estimate indirect scope 1 and 2 GHG emissions, using the extensive ESG data sets that are available for corporations.

Our approach employs an NLP technique to analyse corporate conference call transcripts and compute stock-level commodity scores. These scores reflect the number of times a commodity term is mentioned in a corporate conference call transcript divided by the total number of words used during the conference call. For each commodity in our list (see Graph 2.1), we then identify the 30 companies with the highest score and aggregate those company-level emission intensities as a proxy of the carbon intensity for that particular commodity. We use a simple aggregation methodology of equal weights. The procedure is intuitive, capturing a wide range of relationships, from major producers like Monsanto (agriculture) and Rio Tinto (metals), to more nuanced consumers like restaurants and grocery stores for agricultural commodities. This indirect metric is also more holistic than direct emissions in that it has the potential to capture the full value chain for a given commodity.

2.3. Scaling by dollar

Next, we need to calculate carbon emission intensities. The typical approach uses scientific metrics defining GHG emissions per unit of value added – or "GEVA", which is defined by Randers (2012) and endorsed by various climate science bodies, including the SBTi, Carbon Disclosure Project (CDP) and the UN Global Compact. In equities and corporate bonds, carbon emission intensities are defined by GHG equivalents divided by a dollar quantity like sales or market capitalisation. Sales are computed by MSCI (see Frankel et al (2015)) and market capitalisation is recommended by the EU TEG. In order to make meaningful economic comparison across commodities, we standardise indirect emissions by price. This approach aligns with the widely accepted sales-denominated emissions in equities and the enterprise-value-weighted emission intensities required under the EU TEG framework. It expresses GHG emissions in terms of economic cost; a method employed by Bloomberg (Ghia et al (2021)) and others has also used emissions per dollar for commodities.

To make an economically sensible comparison of carbon emissions across commodities, we define carbon intensities to be GHG emissions divided by the market value of each commodity. This is equivalent to setting intensity targets based on tonnes of carbon dioxide (CO2) equivalents (tCO2d) divided by US dollars – or a value-added contribution of GHG. We find the highest-intensity commodities to be natural gas, Brent, West Texas Intermediate (WTI) and Gasoil and the lowest-intensity commodities to be gold, silver and cocoa (see Graph 2.2).

Carbon emission intensity Graph 2.2

Per USD value; number

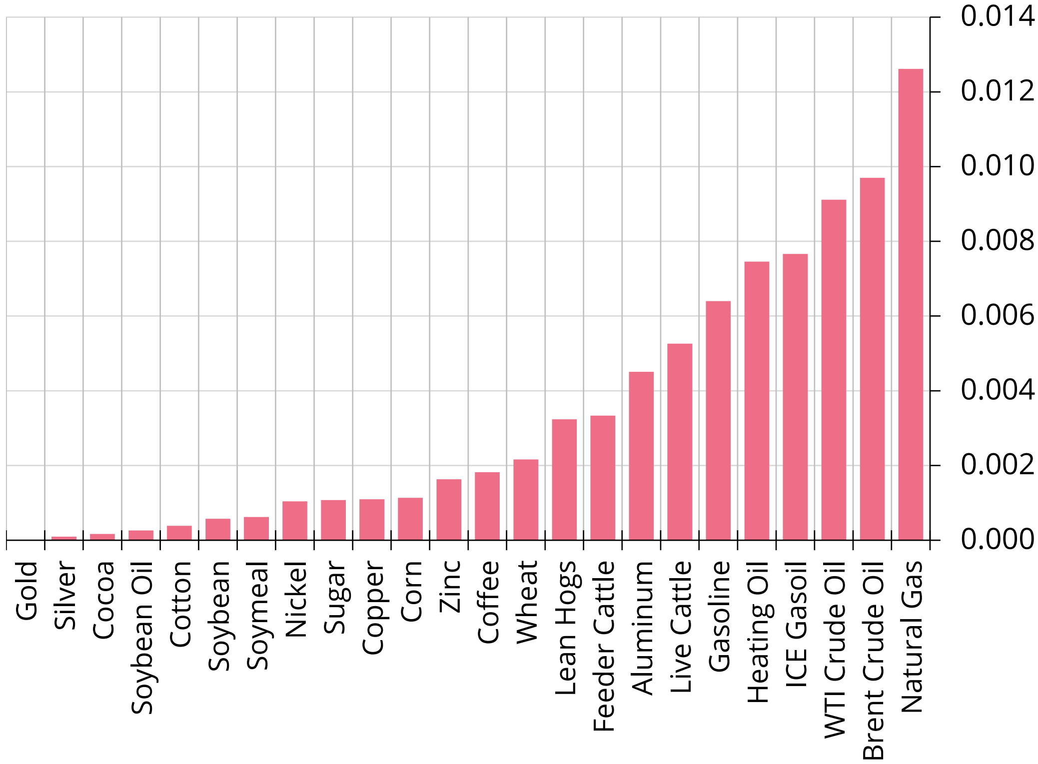

Source: see Table 2.1; BlackRock.

3. Portfolio construction

3.1. Optimisation

We use mean-variance optimisation to maximise returns, minimise risk and simultaneously meet climate objectives. We denote the commodity weights of the benchmark index as h_{bmk}. From these benchmark weights, we infer implied alphas,

$\alpha_{implied}$. Under the assumption of a mean-variance representative agent, following Black and Litterman (1991) and others:

$$\alpha_{implied} \propto V h_{bmk}, \tag{1}$$

where V is the covariance matrix of the benchmark commodity returns.

To incorporate our environmental targets, we set up a new optimisation problem that uses implied alphas from the benchmark, while adding two constraints: one to upweight ESG awareness (ESG) and another to reduce carbon emissions ($CO2$). This is:

$$\max_{h} \alpha_{implied}^{\top} h - \lambda h^{\top} V\, h, \tag{2}$$

such that

$$A_{ESG} h \geq LB_{ESG} \times A_{ESGi} h_{bmk} \tag{3}$$

$$A_{CO2} h \leq UB_{CO2} \times A_{CO2} h_{bmk} \tag{4}$$

where h are the holdings of the environmentally optimised portfolio such that $h - h_{bmk}$ reflects the active environmental tilts relative to the commodity benchmark. Equation (2) is a standard mean-variance optimisation, with risk aversion λ and the implied alphas from the benchmark (see equation (1)).

The benchmark used is the S&P GSCI index, with historical data starting in December 2009. The ESG and carbon emission data are estimated using the methodologies discussed earlier. The optimisation is re-run on a monthly basis, resulting in portfolio weights which meet the ESG and carbon emission constraints relative to the benchmark. We call this portfolio the "Paris-Aligned commodities portfolio". Before conducting any optimisation, an important decision to make in our portfolio design is whether to include or exclude energy commodities. This is discussed in the next subsection.

3.2. Energy: to exclude or include?

Energy commodities are critical to modern society, and the International Energy Agency (IEA (2021)) estimates a hypothetical decline of 75% in oil usage – representing a substantial but not complete reduction in society's reliance on fossil fuels – to limit global warming well below 2°C by 2050. Additionally, including these commodities can be viewed as philosophically consistent with EU TEG guidelines for equity portfolios which specify equal or higher holdings in high-climate-impact sectors. In the context of the multi-asset portfolios, oil has been shown to have attractive inflation hedging properties (see Erb and Harvey (2006); Gorton and Rouwenhorst (2006)).

Conversely, divestment from energy commodities reflects the broader societal shift away from fossil fuels. Arguments to divest from energy include investors' legal responsibility to combat climate change, the need to avoid complicity in the fossil fuel industry and the potential effectiveness of financial incentives in encouraging climate-friendly business practices (Richardson (2017)). Moreover, while the EU TEG cites the importance of investor advocacy when discouraging equity sector exclusions, such advocacy is not relevant to commodity futures investing.

We explore both potential options and contrast the portfolio weights; the weights of a Paris-aligned commodities portfolio including or excluding energy commodities are shown in Graphs 3.1 and 3.2, respectively, compared with the S&P

GSCI index. We incorporate a constraint to reduce emissions by 50% and are able to meet that constraint in both portfolios.

We find that in the sustainable portfolio including energy (Graph 3.2), there are no holdings in Gasoil, Brent or WTI – which are naturally left out of the portfolio excluding energy (Graph 3.1). However, slightly larger-than-benchmark weights are assigned to heating oil and natural gas. It is also interesting that excluding energy tilts the climate-aligned portfolio into metals (aluminium, copper and zinc), whereas including energy results in tilts towards precious metals like gold and silver.

Benchmark and portfolio weights without energy | Graph 3.1

Weights; %

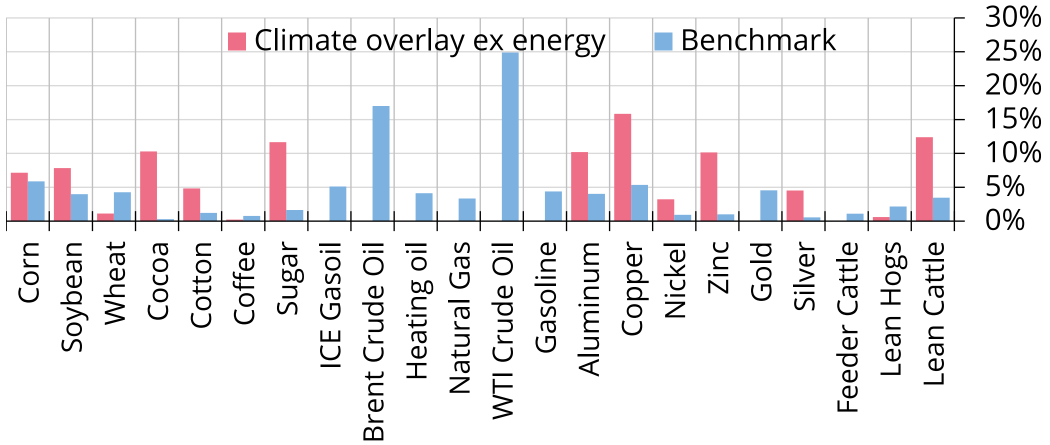

Source: BlackRock calculations. Data as of end-April 2024.

Benchmark and portfolio weights with energy | Graph 3.2

Weights; %

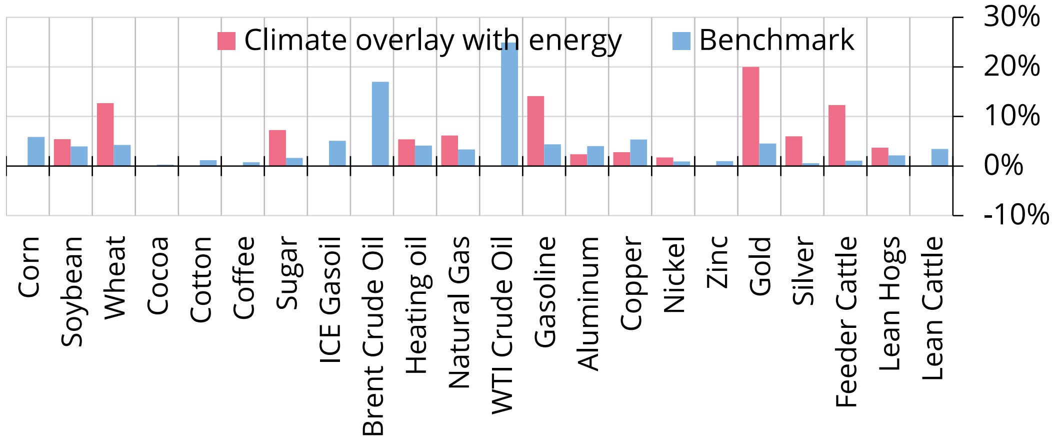

Source: BlackRock calculations. Data as of end-April 2024.

Next, we present the cumulative performance of the Paris-aligned commodity portfolios and their benchmark in Graph 3.3. Importantly, as we will discuss further in the next section, the returns of both Paris-aligned portfolios – with and without energy – are highly correlated with the GSCI index, with correlations exceeding 95%. This indicates that the resulting commodity allocation would retain its diversification and inflation hedging characteristics in a multi-asset portfolio context (see, for example, Rouwenhorst and Tang (2012)).

Benchmark and portfolio performance
Cumulative

Graph 3.3

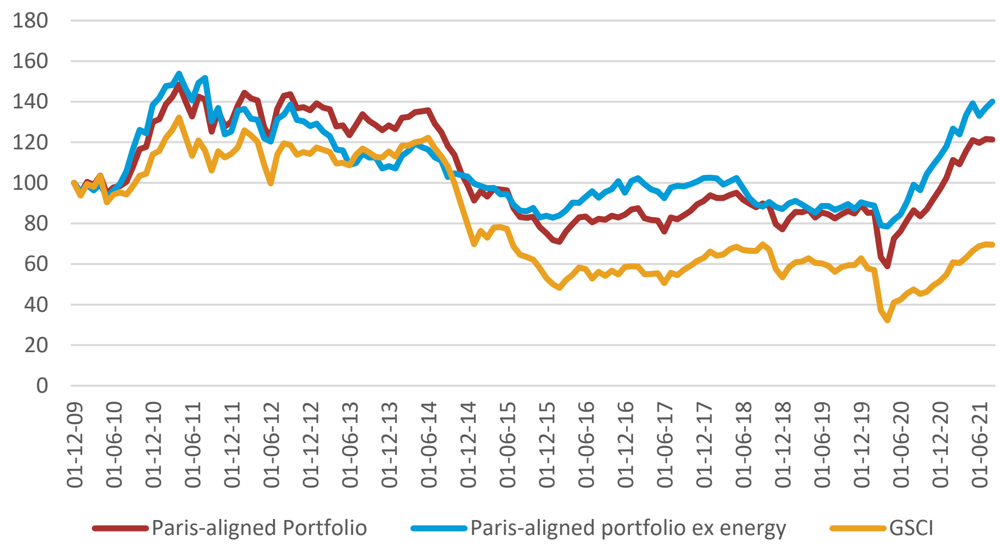

Source: BlackRock calculations. Data as of end-August 2021

4. The role of Paris-aligned commodities in a multi-asset portfolio

In this section, we integrate the Paris-aligned commodities portfolio with other typical Paris-aligned asset class sleeves: equities, sovereign and corporate bonds, and real estate (see Hodges et al (2022) for a multi-asset example, and Kaul et al (2022) and Schwaiger et al (2023) for sovereign bonds in particular). To do so, we must build a multi-asset portfolio, whose sub-asset class sleeves are all Paris-aligned. For the commodities sleeve, we apply the methodology reviewed in section 3, excluding energy commodities.[9]

Our EU PAB-compliant building blocks can be used to construct any multi-asset portfolio. For illustrative purposes, we focus on a balanced multi-asset portfolio built using macroeconomic factors (see Chen et al (1986)) following Bass et al (2017) and

9 Nonetheless, the results are very similar when using the portfolio including energy, given the high correlations between the two, as discussed in the previous section.

Fergis et al (2019). The hypothetical multi-asset portfolio allocations are shown in Table 4.1, along with standard asset class benchmarks. For simplicity we assume those weights to being constant over time.

Portfolio weights and benchmarks — Table 4.1

	Weight	Benchmark
Sovereign Bonds	57%	Custom Equal Weighted across DM and EM
Equities	19%	MSCI World, MSCI World Small Cap and MSCI Emerging Markets
Corporate Bonds	17%	Bloomberg Barclays US IG and US HY Indexes
Commodities	6%	S&P GSCI
Real Estate	2%	FTSE EPRA Nareit Developed Index

Source: BlackRock.

Graphs 4.1 and 4.2 show the cumulative returns of the resulting Paris-aligned multi-asset portfolio. Notably, the Paris-aligned sovereign bond portfolio also has a 10% allocation to green bonds, in addition to the Paris-aligned nominal bond and real bond climate overlay.

Performance of Paris-aligned multi-asset portfolio and benchmark — Graph 4.1

Cumulative

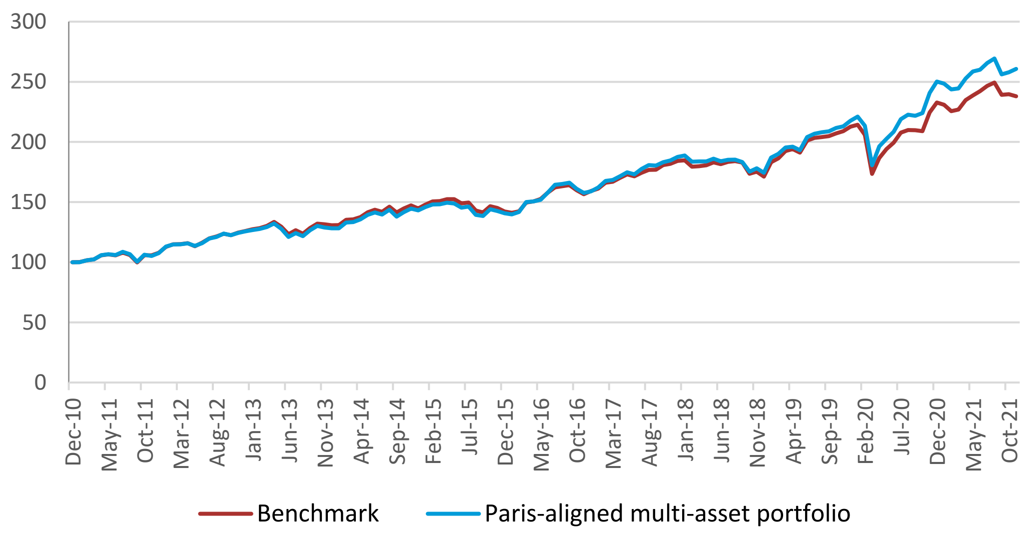

Sources: Bloomberg, BlackRock

Rolling tracking error of Paris-aligned multi-asset portfolio versus benchmark | Graph 4.2

Ex post, %

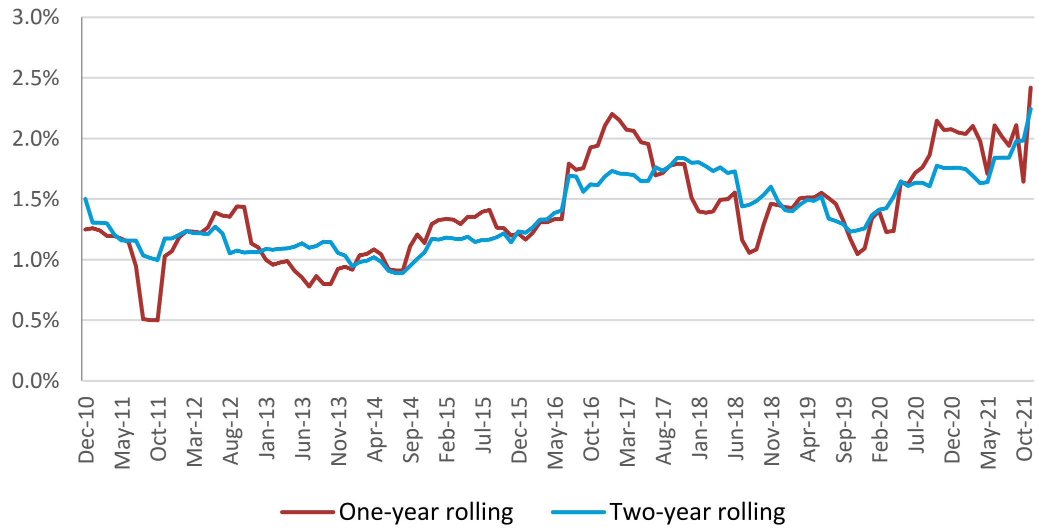

Sources: Bloomberg, BlackRock

The Paris-aligned multi-asset portfolio delivers a risk and return profile comparable to the benchmark portfolio: the mean and standard deviation of the Paris-aligned version are 9.1% and 9.7% per year, respectively, compared with the mean and standard deviation of 8.3% and 9.1% per year, respectively, for the benchmark. As shown in Graph 4.1, the Paris-aligned portfolio has even, at times, outperformed the benchmark. This is especially the case from 2019 to 2021. Generally, as illustrated in Graph 4.2, the Paris-aligned portfolio tracks its benchmark closely, with a realised tracking error of 1.8% on average.

Finally, we examine the decarbonisation pathway of the Paris-aligned multi-asset portfolio. Each asset class embeds a relative decarbonisation target (compared with its benchmark) and a self-decarbonisation target (reducing its carbon intensity over time). We aggregate the decarbonisation targets across asset classes and simulate the decarbonisation pathway into the future. The simulation assumes that the multi-asset allocation remains constant over time starting in January 2010 and that the self-decarbonisation targets are binding – that is, the benchmark does not decarbonise faster than the weighted average self-decarbonisation rate of the Paris-aligned portfolio.

Graph 4.3 presents the results, where a value of 100 represents the carbon emissions of the benchmark portfolios. There is an initial average relative decarbonisation rate of 27% and an average self-decarbonisation rate of 2.8% leading to an overall reduction of 45% by 2030 and 55% by 2040 (Graph 4.3, red line). We then compare this with the decarbonisation pathways of Paris-aligned listed equity and corporate bond portfolios as well as a 100% allocation into the Paris-aligned sovereign bond portfolio (Graph 4.3, blue and yellow lines).

The Paris-aligned equity and corporate bond portfolios have the highest relative and self-decarbonisation rates at 50% and 7%, respectively – by construction, as these are the requirements specified for these asset classes by the EU TEG. In the equity and corporate bond asset classes, the methodology leads to a portfolio-wide reduction in emissions of 75% in 2030 and almost 90% in 2040. In contrast, the Paris-aligned sovereign bond strategy maintains a steady 14% reduction relative to its benchmark throughout the period.

Decarbonisation pathways — Graph 4.3

100 = original emissions

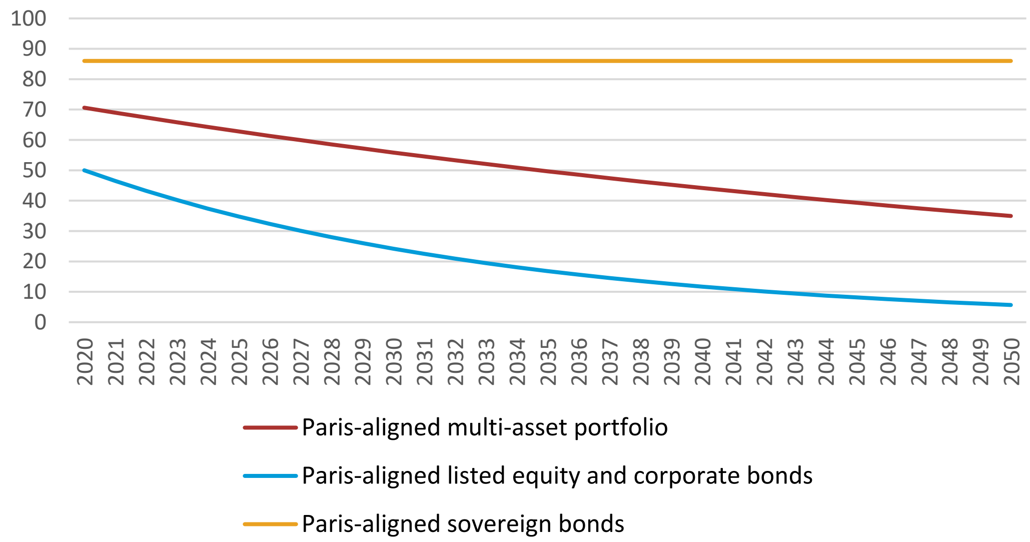

Sources: Bloomberg, BlackRock

To complete our analysis, we explore the inflation hedging properties of commodities in a Paris-aligned context. While we have already observed that the Paris-aligned multi-asset portfolio retains most of its risk and return characteristics compared with the benchmark portfolio, we now focus on additional benefits, such as its ability to hedge against inflation.

We evaluate inflation hedging using two measures: realised US monthly consumer price index (CPI) and US 10-year break-even inflation rates. Graphs 4.4 and 4.5 show the five-year rolling correlations of the commodity portfolio's return with month-on-month realised US CPI and US break-even inflation changes. The former measures the linkage of commodity portfolios to CPI baskets whereas the latter focuses on the changes of expectations, ie inflation surprises. We compare the Paris-Aligned commodity portfolios with the GSCI and GSCI Lite indexes. The correlations show that the commodities portfolio offers comparable inflation hedging properties to more traditional commodity indexes when it is Paris-Aligned (both with and without energy).

Correlation to US CPI | Graph 4.4
Rolling five years

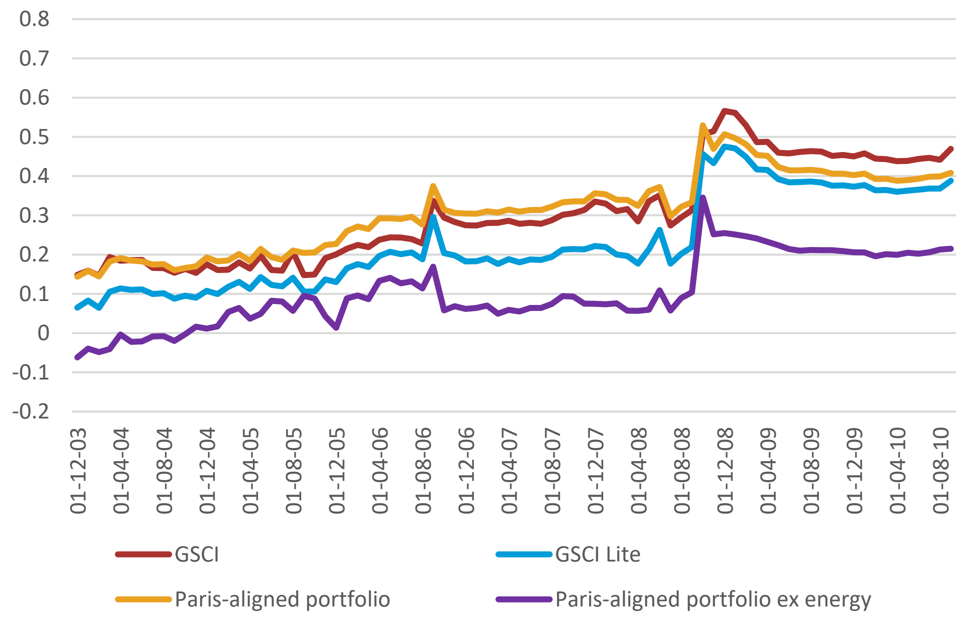

Sources: Bloomberg, BlackRock

Correlation to US breakeven inflation | Graph 4.5
Rolling five years

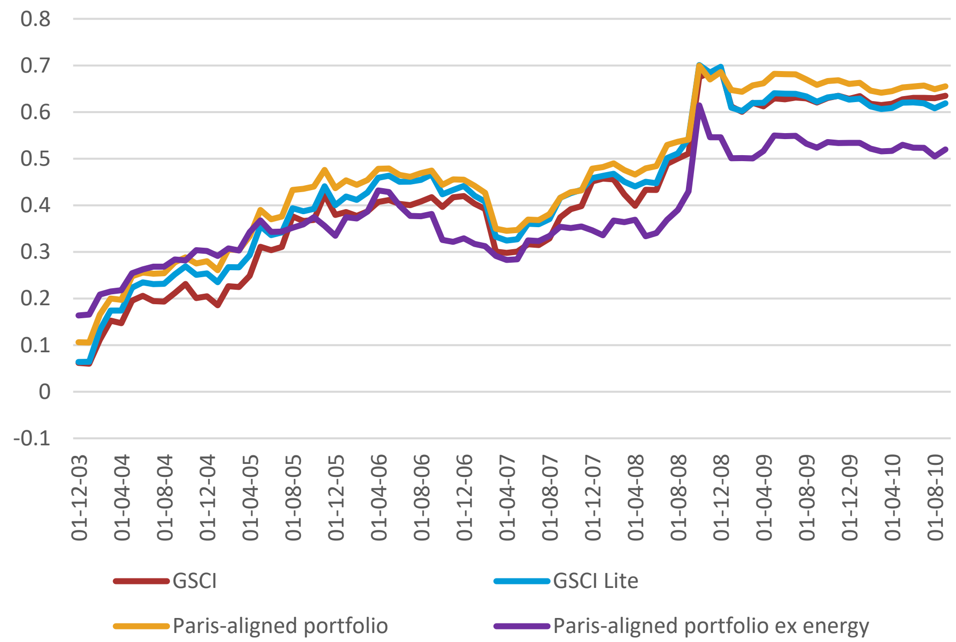

Sources: Bloomberg, BlackRock

Concluding remarks

Commodities play an important role in multi-asset portfolios due to their diversification benefits. Investors who seek to invest in a Paris-aligned manner may also want commodities exposure that is aligned with net zero principles. In this paper, we present a range of metrics that enable the measurement of transition readiness, including carbon emissions, and a framework to build Paris-aligned commodities portfolios. These can be applied in contexts that include and exclude the energy sector. They exhibit excess performance statistics as well as inflation hedging benefits compared with the generic commodity benchmark in our exercises.

References

Bass, R, S Gladstone and A Ang (2017): "Total portfolio factor, not just asset, allocation", *Journal of Portfolio Management*, vol 43, no 5, pp 38–53.

Black, F and R Litterman (1991): "Asset allocation: combining investor views with market equilibrium", *Journal of Fixed Income*, vol 1, no 2, pp 7–18.

Chen, N-F, R Roll and S Ross (1986): "Economic forces and the stock market", *Journal of Business*, vol 59, no 3, pp 383–403.

Chevallier, J (2009): "Carbon futures and macroeconomic risk factors: a view from the EU ETS", *Energy Economics*, vol 31, no 4, pp 614–25.

Deaton, A and G Laroque (1992): "On the behavior of commodity prices", *Review of Economic Studies*, vol 59, no 1, pp 1–23.

Erb, C and C Harvey (2006): "The strategic and tactical value of commodity futures", *Financial Analysts Journal*, vol 62, no 2, pp 69–97.

EU Technical Expert Group on Sustainable Finance (2019): *TEG final report on climate benchmarks and benchmarks' ESG disclosures,* September.

Fergis, K, K Gallagher, P Hodges and K Hogan (2019): "Defensive factor timing", *Journal of Portfolio Management*, vol 45, no 3, pp 50–68.

Frankel, K, M Shakwipee and K Nishikawa (2015): "Carbon footprinting 101", *MSCI Issue Brief*, September.

Ghia, K, A J Lindeman and M Zhang (2021): "ESG comes to town", *Global Commodities Applied Research Digest*, JP Morgan.

Gorton, G and K G Rouwenhorst (2006): "Facts and fantasies about commodity futures", *Financial Analysts Journal*, vol 62, no 2, pp 47–68.

Hodges, P, H Ren, K Schwaiger and A Ang (2022): "Net-zero investing for multi-asset portfolios satisfying Paris-aligned benchmark requirements with climate alpha signals", *Journal of Portfolio Management Multi-Asset Special Issue*, vol 48, no 4, pp 33–58.

International Energy Agency (IEA) (2021): *Net zero by 2050: a roadmap for the global energy sector,* www.iea.org/reports/net-zero-by-2050.

Kaul, K, K Schwaiger, M Si and A Ang (2022): "Sustainable alpha in sovereign and corporate bonds", *Journal of Investment Management*, vol 20, no 2, pp 30–50.

Krabbe, O, G Linthorst, K Blok, W Crijns-Graus, D van Vuuren, N Hohne, P Faria, N Aden and A C Pineda (2015): "Aligning corporate greenhouse-gas emissions targets with climate goals", *Nature Climate Change*, vol 5, pp 1057–60.

Miffre, J (2012): "Comparing first, second and third generation commodity indices", *SSRN Working Paper*.

Randers, J (2012): "Greenhouse gas emissions per unit of value added ('GEVA') – a corporate guide to voluntary climate action", *Energy Policy*, vol 48, pp 46–55.

Richardson, B (2017): "Divesting from climate change: the road to influence", *Law & Policy*, vol 39, no 4, pp 325–48.

Rouwenhorst, K G and K Tang (2012): "Commodity investing", *Annual Review of Financial Economics*, vol 4, pp 447–67.

Schwaiger, K, R Ali and A Ang (2023): "Net zero sovereign bond portfolios", *Evolving practices in public investment management: proceedings from the Eighth Public Investors Conference*, Bank for International Settlements, October, pp 167–84.

Tobler, A, M Hasse and P Sigg (2021): "Integrating sustainability into commodity investing", in *Handbook on sustainable investments: background information and practical examples for institutional asset owners*, CFA Institute Research Foundation and Swiss Sustainable Finance, chapter 18, pp 129–33.

Williams, J and B Wright (1991): *Storage and commodity markets*. Cambridge University Press.