

Climate change investment risk: Optimal portfolio construction ahead of the transition to a lower-carbon economy*

Davide Benedetti[†] Enrico Biffis[†] Fotis Chatzimichalakis[‡]
Luciano Lilloy Fedele[‡] Ian Simm[‡]

Abstract

There is an increasing likelihood that governments of major economies will act within the next decade to reduce greenhouse gas emissions, probably by intervening in the fossil fuel markets through taxation or cap & trade mechanisms (collectively “carbon pricing”). We develop a model to capture the potential impact of carbon pricing on fossil fuel stocks, and use it to inform Bayesian portfolio construction methodologies, which are then used to create what we call Smart Carbon Portfolios. We find that investors could reduce ex-post risk by lowering the weightings of some fossil fuel stocks with corresponding higher weightings in lower-risk fossil fuel stocks and/or in the stocks of companies active in energy efficiency markets. The financial costs of such de-risking strategy are found to be statistically negligible in risk-return space. Robustness of the results is explored with alternative approaches.

Keywords: climate change, carbon pricing, Bayesian analysis, Black-Litterman model, portfolio theory.

1 Introduction

In recent years the challenge posed by climate change caused by the emission of greenhouse gases (GHGs) from human activity has risen rapidly up the agendas of policymakers worldwide (e.g., Carney, 2015; Gao, 2016). From the late 1980s, some of these decision makers have

*This version: June 4, 2019. First version: March 24, 2018. The authors are grateful to two anonymous referees, as well as to participants in several workshops and seminars, for valuable suggestions which led to a considerably improved version of the paper.

[†]Davide Benedetti (d.benedetti@imperial.ac.uk) and Enrico Biffis (e.biffis@imperial.ac.uk; corresponding author) are at Imperial College Business School, Imperial College London, South Kensington Campus, SW7 2AZ United Kingdom.

[‡]Ian Simm (i.simm@impaxam.com), Fotis Chatzimichalakis (f.chatzimichalakis@impaxam.com) and Luciano Lilloy Fedele (l.fedele@impaxam.com) are at Impax Asset Management, 30 Panton Street, London SW1Y 4AJ, United Kingdom.

expressed concern over the potential consequences of significant global temperature rise and associated climate instability. This development can be attributed to the increased robustness of statements from the scientific community on this topic, including (i) the potential magnitude and timeframes of climate change, (ii) the deepening evidence base that the change is caused by human activity rather than natural effects such as solar activity or volcanoes,¹ (iii) the scale of the potential impact on society, and (iv) the existence of tipping points leading to runaway changes that could make human existence impossible (IPCC, 2014). These factors plus others, such as pressure from “third sector” groups and indications that some leading corporations are able to supply the products and services to help mitigate climate change, have spurred policy makers into action to reduce GHGs.²

There is now widespread agreement that climate change risks can be segmented into (a) the physical risks posed for example by different weather patterns and shifting climatic belts (Physical Risk), and (b) the transition risks arising from society’s response to adapt to physical changes and to mitigate further change (Transition Risk). Those investors with large, complex portfolios (institutional investors) typically manage risk through well-established mechanisms, including asset selection, diversification, insurance and liquidity (e.g., Gründl et al., 2016). Climate change poses particular challenges for some of these investors,³ potentially rendering well-established risk management mechanisms partly or fully ineffective. On the other hand, the potential mitigating factors, including government policy responses, changes in the behaviour of individuals, and the development and deployment of technologies that reduce impact, are poorly understood (see Field, 2014; IPCC, 2015). Furthermore, the long-term availability of insurance and other risk pooling mechanisms is uncertain⁴ (e.g., IAIS, 2018). These challenges are exacerbated by the

¹In some parts of the world, weather patterns have begun to deviate from historical trends and are only statistically consistent with climate models that incorporate GHGs linked to human activity (see Anderson et al., 2016; Bader et al., 2008; Shiogama et al., 2016; Stocker, 2014).

²A notable example is represented by the Paris Agreement signed on April 22, 2016.

³Investors who care about climate change risks may do so for financial and/or non-financial reasons. The former include stranded asset risk (e.g., Caldecott et al., 2015; Simm et al., 2016), the latter reputation, ethical considerations or fiduciary duties (e.g., Krüger et al., 2018; Röell, 2019).

⁴Financial innovation has been very successful in reshaping the way in which some long-term risks, such as longevity and morbidity, can be more efficiently managed by individuals and transferred between insurance and capital market participants (e.g., Biffis and Blake, 2013, 2014), but progress in the area of climate risks has been slow (see Biffis and Chavez, 2017; Chavez et al., 2018, for some examples).

rising exposure taken by many institutional investors to real assets, i.e., those for which there is typically no liquid market for sale.

Many institutional investors who are analysing this issue are questioning whether these risks are already discounted in current valuations (Andersson et al., 2016; Daniel et al., 2016). There is growing empirical evidence that this is indeed not the case. For example, Hong et al. (2016) find that trends in droughts that are exacerbated by global warming are not priced in the valuations of food companies. Choi et al. (2018) find that, while retail investors seem to sell high-emission firms and buy low-emissions firms in response to systematically abnormal temperatures,⁵ institutional investors' portfolio decisions are instead largely unaffected. In a survey of institutional investors⁶ Krüger et al. (2018) find that respondents believe that equity valuations do not fully capture the risks of climate change, the oil sector being perceived as the most overvalued. Surprisingly, they also find that, among investors who consider stranded asset risk⁷ as being relevant for climate change risk management, only 25% believe that stranded asset risk is very high for coal producers. This all suggests that market participants may be relatively inexperienced in appreciating climate change risks and their management.

Institutional investors are exposed to a combination of Physical and Transition Risk via their holdings and are deploying different strategies to manage Transition Risk in particular (e.g., Mercer, 2015; Schrodgers, 2016). A frequently cited example is “divestment”, particularly the sale of a material percentage of energy stocks with redeployment of capital across the wider portfolio. Although reducing Transition Risk, potentially significantly, this strategy typically raises additional issues (e.g., Litterman, 2015; Bessembinder, 2016), such as a reduction in dividend streams received and in the exposure of the portfolio to energy prices, which historically have often been a positive driver of portfolio returns. A more nuanced strategy based on relative Transition Risk across a portfolio of energy stocks uses “carbon footprints” to inform portfolio

⁵A recent strand of literature links personal experience of abnormal temperatures with increased perception of climate risk in the US (see Akerlof et al., 2013; Myers et al., 2013; Zaval et al., 2014).

⁶Although the survey may arguably be biased toward investors with a high awareness of climate risk, we find the results particularly telling.

⁷Stranded asset risk is the risk that certain assets become uneconomic (firms are unable to recover their investment cost) due to Transition Risks such as taxation of fossil fuel supply or consumption (Caldecott et al., 2015; Krüger et al., 2018).

changes (e.g., Krüger et al., 2018; Andersson et al., 2016). However, the simple idea that a high carbon footprint implies a high level of financial risk has two material flaws: (i) a company's carbon footprint is generally based on current emissions data, while its equity valuation is based on an aggregated view of the company's future financial prospects, for example its cash flow, and (ii) although a company's carbon footprint may be correlated with its exposure to rising costs linked to climate change, for example taxes on fossil fuels, the footprint takes no account of the company's pricing power, i.e., its ability to pass on cost increases to its customers (e.g., Simm et al., 2016).

This paper presents an alternative approach that can be used by institutional investors to assess and manage Transition Risk, while also taking account of changing information on risk. A key benefit of our approach is that it uses some of the standard tools of investment management, with a particular focus on the potential for certain risk factors to impact a company's financial prospects. The approach is based on five main premises: (i) at the present time, institutional investors can cost-effectively manage Physical Risk through a combination of diversification and asset selection (including asset level insurance), so the analysis and recommendations are limited to Transition Risk; (ii) the modelling of Transition Risk should focus on the potential impact on future financial returns of companies operating in sectors with weak pricing power, such as Energy, Utilities and Basic Materials; (iii) an appropriate method for modelling/flexing Transition Risk is scenario analysis of the impact of Carbon Prices (e.g., World Bank, 2017) on future corporate cash flows, translated into expected risk exposure of individual stocks; (iv) the expected risk exposure for each stock represents a pricing anomaly in today's asset prices and can inform an optimisation exercise; (v) the optimised portfolio can be updated periodically to reflect new risk data, particularly the likelihood of Carbon Prices being implemented. We develop the portfolio construction framework within the standard mean-variance asset allocation model, and decide to focus on companies engaged in the exploration and production of fossil fuel assets (henceforth E&P firms). Application of the framework to a broader set of firms, as per point (ii) above, is left for further research. We use Bayesian and quasi-Bayesian approaches to build what we call Smart

Carbon Portfolios (SCPs), and show how they can allow investors to hedge against the transition to a carbon pricing state without incurring material ex-ante efficiency costs. In particular, we show how some of the proposed portfolio allocations are statistically indistinguishable ex-ante (in risk-return and portfolio weight space) from relevant portfolio benchmarks. In general, we find that naive divestment is dominated by strategies exploiting the cross-sectional heterogeneity of the E&P sector’s exposure to carbon pricing.

The paper is organized as follows. The next section introduces the carbon pricing scenarios methodology and discusses it within the standard mean-variance portfolio selection framework. Section 3 introduces the Bayesian approach to portfolio construction, providing in turn a parameterization that can be understood in terms of an asset manager’s views, in the spirit of the Black and Litterman (1992) approach (henceforth BL approach). Section 4 shows how the methodology can be used to build SCPs, and provides a number of examples. In section 5, we discuss further extensions and refinements to assess the robustness of our findings. Finally, section 6 concludes. An appendix provides further technical details on the portfolio construction methodology.

2 Carbon pricing scenarios approach

Consider the stylized case of two possible states of the world associated with different efficient frontiers (see figure 1 for an example). A mean-variance investor with access to a technology forecasting the likelihood of state transitions could (i) invest according to the single frontier obtained ex-ante by using the unconditional first two moments implied by the conditional moments, given the probability of state transition occurrence, or (ii) time the transition to the new state by rebalancing the portfolio only in response to a high enough probability of state transition. It is well known that in theory approach (ii) is able to exploit an investment opportunity set dominating the investment opportunity set offered by the single frontier delivered by approach (i) (see Clarke and de Silva, 1998). Ang and Bekaert (2004), for example, show how a regime switching model can be used to rebalance an investor’s portfolio when a transition to a bear market state is likely to occur. In a similar vein, Avramov and Wermers (2006) document the

out-of-sample performance of some predictability-based strategies. In practice, however, implementation of any model rests on estimation of the relevant parameters from historical data, so that considerable implementation challenges arise due to sampling error and the possibility of model mis-specification.

Let us recast now the considerations above in the context of interest: we regard the current state as being the “business as usual” state, the alternative state being instead characterized by the introduction of carbon pricing and hence referred to as “carbon pricing state”. We take the perspective of investors who think that the possibility of a switch to the carbon pricing state is not fully priced in current stock market valuations,⁸ and that the new state realization will likely lead to abrupt repricing of assets instead of smooth transition to a new equilibrium. This view is supported by the limited awareness of climate change risk among institutional investors, as documented for example in Andersson et al. (2016), as well as the empirical evidence on the inefficient pricing of risks brought on or exacerbated by climate change (see Hong et al., 2016) or the survey results presented in Krüger et al. (2018). Investors sharing this view would likely recognize the difficulty of timing the transition to the carbon pricing state, and hence prefer approach (i) to the risk of mis-timing the rebalancing of their portfolios and being exposed to a potentially disruptive price discovery process. We therefore explore the implementation of approach (i), informed by a probabilistic representation of carbon state transitions, and assess the extent of any efficiency losses, weighing them against the expected gains from monetization of climate change risk mispricing and mitigation of the risk originating from a state transition. In the next section, we present a methodology to quantify the cross-sectional impact of such a transition, conditional on our views on the timing and magnitude of carbon pricing. We then introduce a term structure of probabilities allowing us to build the relevant ex-ante efficient frontiers and understand the portfolios supporting them.

< FIGURE 1 ABOUT HERE. >

⁸As long as information on the new state is not revealed smoothly, the approach is also relevant when the possibility of a switch to the carbon pricing state is fully priced.

2.1 Oil reserves and carbon pricing state transition

For simplicity, we assume that the expected effect of all drivers of oil supply and demand, except climate change policy, is currently fully reflected in the forward price curves for oil, and that climate change risk can be captured by the analysis of a few fundamental issues illustrated in figure 3 (Simm et al., 2016). To fix ideas, we think of a one-off carbon tax introduced by governments and resulting in the transition to a new equilibrium as depicted in figure 3. There are three likely effects to consider: (i) an increase in the price of oil paid by consumers; (ii) a reduction in the wholesale price of oil received by producers; and (iii) a reduction in oil consumption, potentially rendering those assets with a high marginal cost of production “stranded” (e.g., Caldecott et al., 2015). In line with Simm et al. (2016), given the inherent volatility of the price of oil and the relatively strong lobbying power of oil producers, we consider effect (iii) to be the most significant.

To determine the impact of carbon pricing on E&P stocks we analyse the forward production curves⁹ of E&P companies within the MSCI World Index. An example from an anonymized company is depicted in figure 2. The chart reports the oil reserves, quantified in “barrel of oil equivalent” (BOE) terms, available to the firm at future dates. Different oil reserves are characterized by different break-even oil prices reflecting heterogeneous production costs. For example, deepwater and ultra-deepwater reserves can have a significant proportion of their break-even prices above \$150, whereas tight liquids reserves can have a substantial portion of their breakeven price distribution below the \$80 level. Figure 2 shows, for example, a company with 117m BOE reserves in the \$55-60 break-even bracket in year 2020. For a benchmark oil price of (say) \$60 those reserves represent 11.98% of the “profitable reserves” over the time period 2020-25.

< FIGURE 2 ABOUT HERE. >

For given forward production curve, we determine the impact of carbon pricing by using a simple supply and demand model allowing us to quantify a company’s revenue in different

⁹We are grateful to Carbon Tracker for providing this information.

carbon pricing states, where the difference in states reflects various assumptions on the timing and magnitude of the envisaged carbon tax. By introducing a term structure of probabilities for the carbon pricing state occurrence, we can then quantify the ex-ante pricing implications via the discounted cashflow approach discussed in the next section. For the examples of section 4.1, we use the 36 scenarios set out in table 1 and assume an 85% cumulative probability of global carbon price ranging between \$10 per tonne and \$60 per tonne of CO2 between 2020 and 2025. In particular, we consider the introduction of carbon pricing in any given year between 2020 and 2025. The probability of carbon pricing not being implemented is assumed to be 15%. The scenario probabilities, as well as the scenarios themselves, are subjective, and can of course be readily adjusted in response to market-implied expectations of future oil prices or any other information an investor might find valuable. We abstract away from trying to anticipate any changes in companies' production strategies, the rationale being that such plans can be adjusted only slowly due to the significant scale and complexity of operations.

< TABLE 1 ABOUT HERE. >

< FIGURE 3 ABOUT HERE. >

2.2 A simple approach to stock price impact.

Let us denote by α_j the carbon tax (in \$ per ton of CO2) that will be implemented at time T_j with probability π_j (with $j = 1, \dots, m$), the current time being $T_0 = 0$. Let us consider the global oil price and production curves, as illustrated in figure 3. Focusing on a single future date for simplicity, the introduction of a carbon tax would make the pre-tax equilibrium (P^*, V^*) switch to a post-tax equilibrium $(\bar{P}^*, \bar{V}^*)$. We compute the post-tax equilibrium point by using a first-order approximation for both the supply and the demand curve. We then use the company-specific forward production curves discussed in the previous section (see figure 2) to determine how the transition to a new equilibrium will affect a generic company i 's revenue. For ease of notation and analysis, we interpolate the information contained in forward production plans

with a smooth curve giving oil reserves volume and cost information denoted by $V_{h,j}^i$ and $C_{h,j}$, respectively, where the oil reserve index $h \in \mathcal{R}_j$ lives in a compact subset $\mathcal{R}_j$ of $\mathbb{R}$. The revenue generated by each company i 's reserve, after carbon tax α_j is introduced at time T_j , is assumed to be given by

$$\overline{REV}_{h,j}^i = \begin{cases} V_{h,j}^i(\overline{P}_j^* - \alpha_j) & \text{for } h : C_{h,j} \leq \overline{P}_j^* \\ 0 & \text{for } h : C_{h,j} > \overline{P}_j^*, \end{cases}$$

meaning that unprofitable reserves simply become “stranded”. Hence, the change in revenue for any given company due to carbon pricing in year T_j is given by

$$\Delta REV_j^i = \int_{\mathcal{R}_j} \left(\overline{REV}_{h,j}^i - REV_{h,j}^i \right) dh,$$

where we use the notation REV for pre-carbon-tax revenue. We can then consider the entire term structure of production curves, quantify the change in revenue at different horizons, and then compute the overall change in expected free cashflows, $E[\Delta FCF^i]$, by NPV-ing the future expected changes in free cash flows at different horizons:

$$E[\Delta FCF^i] = \sum_j p_j \times \Delta FCF_j^i \times (1 + \rho^i)^{-T_j},$$

where we have used a constant weighted average cost of capital (WACC) approach, with ρ^i denoting the WACC of company i (e.g., Brealey et al., 2012). Finally, we measure $E[\Delta FCF^i]$ relative to a baseline $\overline{\Delta FCF}^i$ computed via a DCF approach with constant growth rate g^i , and use the relative change, $\%E[\Delta FCF^i]$, as a proxy for the expected change in enterprise value. Using information on leverage, we recover the percentage change in equity value (market capitalization) from the formula $[(1 + \%E[\Delta FCF^i]) - DE^i] (1 - DE^i)^{-1}$, where DE^i denotes the debt-to-equity ratio of company i (e.g., Brealey et al., 2012).

2.3 Carbon pricing impact results.

By applying the methodology outlined in the previous sections, we are able to obtain the results illustrated in figure 4, which reports the extent of repricing implied by our probabilistic scenario analysis of E&P stocks. Equivalently, it can be understood as the spreads that a suitably constructed portfolio should be able to earn over time on the mispriced stocks, should a state transition occur. The bottom panel of figure 4 disaggregates the extent of mispricing within the E&P sector. We cluster E&P stocks into eight groups (labelled from A to I), depending on the severity of their exposure to carbon pricing. In line with our discussion of figure 3, we find that companies operating assets at the higher-cost end of the production curve are more likely to experience a material reduction in output, should carbon pricing be implemented. Group E, corresponding to coal stocks, and group I, corresponding to oil sands, are both significantly exposed to carbon pricing. The degree of mispricing is negligible for group A, which is therefore not reported in the picture and ignored in the following analysis. Table 2 provides some details on total and average market capitalization of the firms within groups B to I. The findings suggest that there is considerable cross-sectional heterogeneity in exposure to carbon pricing among E&P firms, which we will exploit in the portfolio construction discussed in the next section.

< FIGURE 4 ABOUT HERE. >

< TABLE 2 ABOUT HERE. >

3 Portfolio construction methodology

We consider portfolio construction within the standard mean-variance framework. The natural way to integrate views on carbon pricing into our asset allocation is via Bayesian methods (Barry, 1974; Klein and Bawa, 1976). We consider a market with n risky assets with Normal i.i.d. returns $\{\mathbf{R}_t\}_{t=1,\dots,S}$, where $\mathbf{R}_t = (R_{1,t}, \dots, R_{n,t})^\top$ denotes the vector of random returns at each time $t \geq 0$. Let $\boldsymbol{\mu} \in \mathbb{R}^n$ and $\Sigma \in \mathbb{R}^{n \times n}$ denote the vector of mean returns and the variance-covariance matrix,

respectively. Mean-variance optimization seeks portfolio weights $\mathbf{w}^* = (w_1^*, \dots, w_n^*)^\top$ solving the following problem (e.g., Meucci, 2009; Lai et al., 2011; Roncalli, 2014):

$$\begin{cases} \min_{\mathbf{w} \in \mathcal{W}} & \frac{1}{2}\sigma_p^2 := \frac{1}{2}\mathbf{w}^\top \Sigma \mathbf{w} \\ s.t. & \mathbf{w}^\top \boldsymbol{\mu} = \bar{\mu}_p, \end{cases} \quad (3.1)$$

for fixed target level of expected portfolio return $\bar{\mu}_p > 0$ and admissible set of portfolio allocations $\mathcal{W}$. The latter can include constraints such as full investment ($\mathbf{w}^\top \mathbf{1} = 1$), as well as long-only ($\mathbf{w} \geq 0$) or capped (e.g., $w_i \in [0, c_i]$, for $c_i > 0$) allocations. The efficient frontier can be traced by letting $\bar{\mu}_p$ vary above μ_{gmvp} , the expected return of the global minimum variance portfolio.¹⁰ The problem can equivalently be formulated in terms of expected return maximization for given level of acceptable portfolio risk level σ_p . Alternatively, one can address the risk-return trade-off directly and solve the minimization problem

$$\min_{\mathbf{w} \in \mathcal{W}} \alpha \sqrt{\mathbf{w}^\top \Sigma \mathbf{w}} - (1 - \alpha) \mathbf{w}^\top \boldsymbol{\mu}, \quad (3.2)$$

where the coefficient α can be made vary between zero and one to trace the entire frontier.

An important challenge in solving the optimization problems above is that the parameters $\boldsymbol{\mu}$ and Σ are unknown and need to be estimated (see Broadie, 1993, for a discussion of true vs. estimated frontiers). An additional challenge is represented by the introduction of an asset manager's views on specific assets. A common problem arises in this case: it is not immediately obvious how to formalize a manager's views, as any changes in any stock's input parameters may result in dramatic variations in the overall portfolio allocation, due to the sensitivity of the optimizer to input parameters (Best and Grauer, 1991; Chopra and Ziemba, 2013) and the degree of comovement channeled by the variance-covariance matrix (Black and Litterman, 1992; Satchell and Scowcroft, 2000; Cheung, 2010). A tractable approach to addressing the issue of unknown parameters is Bayesian analysis, whereas the quasi-Bayesian BL approach was designed

¹⁰The structure of the global minimum variance portfolio can be derived explicitly and is given by $\mathbf{w}_{gmvp} = \Sigma^{-1} \mathbf{1} (\mathbf{1}^\top \Sigma^{-1} \mathbf{1})^{-1}$; see Lai et al. (2011), for example.

exactly to address the problem of integrating a manager’s views in portfolio construction. In the following, we will rely on a parameterization aligning the two approaches and hence addressing both sets of issues simultaneously. As a by-product, the BL approach will allow us to translate our scenarios into expected absolute/relative returns from trading strategies or “view portfolios”. Moreover, it will allow us to interpret the resulting optimal allocations in terms of departures from a benchmark, such as the CAPM equilibrium model. For robustness, we discuss refinements and extensions in section 5.

Given our working assumptions on returns, in the Bayesian approach it is convenient to assign the pair $(\boldsymbol{\mu}, \Sigma)$ a prior joint distribution of Normal Inverse Wishart ($\mathcal{NIW}$) type (e.g., Meucci, 2009; Lai et al., 2011),

$$(\boldsymbol{\mu}, \Sigma) \sim \mathcal{NIW}(\boldsymbol{\mu}_0, d_0, \Sigma_0, \nu_0). \quad (3.3)$$

After taking into account the return observations $\{\mathbf{R}_t\}_{t=0,\dots,T}$, the above leads to a posterior distribution (i.e., a time- T conditional distribution updated according to Bayes rule) which is also of $\mathcal{NIW}$ type:

$$(\boldsymbol{\mu}, \Sigma) | \mathbf{R}_1, \dots, \mathbf{R}_T \sim \mathcal{NIW}(\boldsymbol{\mu}_1, d_1, \Sigma_1, \nu_1),$$

with the parameters $\boldsymbol{\mu}_1, d_1, \Sigma_1, \nu_1$ given explicitly in the appendix. It is then natural to solve the optimization problem by using the expectations of the parameters $\boldsymbol{\mu}, \Sigma$ under the posterior distribution, which are given as follows (e.g., Lai et al., 2011):¹¹

$$\hat{\boldsymbol{\mu}}_b = \beta \boldsymbol{\mu}_0 + (1 - \beta) \hat{\boldsymbol{\mu}}, \quad (3.4)$$

$$\hat{\Sigma}_b = \gamma \Sigma_0 + (1 - \gamma) \hat{\Sigma} + \beta \gamma (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})^\top, \quad (3.5)$$

with $\beta = d_0/d_1$, $\gamma = T/\nu_1$, and where $\hat{\boldsymbol{\mu}}$ and $\hat{\Sigma}$ denote the sample return mean and variance-covariance matrix, respectively, which under our assumptions can equivalently be seen as maximum likelihood estimates of the relevant parameters. They may also be regarded as method-of-moments estimates without Normality, but with the assumption of weak stationarity of re-

¹¹This can be shown to be equivalent to shrinkage estimators, and hence mean-square error minimization.

turns. In expression (3.4), we recognize the familiar structure of shrinkage estimators, in that the Bayesian updating procedure shrinks the sample mean, $\hat{\boldsymbol{\mu}}$, towards the prior estimate, $\boldsymbol{\mu}_0$. Similar considerations apply to the variance-covariance matrix, with the addition of a contribution from the uncertainty in mean returns, as captured by the last term in expression (3.5).

The Bayesian setting offers a probabilistically coherent way of introducing our views into the specification of market parameters via the prior distribution. However, we would like now to express our views in terms of trading strategies capturing the gains expected to be delivered by our views. We adopt the BL approach, and as customary in this setting, we express our absolute or relative forecasts on individual assets via the equation

$$\boldsymbol{Q} = P\boldsymbol{\mu} + \boldsymbol{\varepsilon},$$

where $\boldsymbol{Q} \in \mathbb{R}^k$ is the vector of expert forecasts on returns, $P \in \mathbb{R}^{k \times n}$ is the pick matrix listing the trading strategies implementing our views, and $\boldsymbol{\varepsilon}$ is an error term uncorrelated with $\boldsymbol{\mu}$ and distributed according to $\boldsymbol{\varepsilon} \sim N(0, (1-\tau)\Omega)$, with $\Omega = T^{-1}P\hat{\Sigma}P^\top$ (see the appendix for details). A positive value $P_{i,j} > 0$ in P , for example, represents a long position implemented with a portfolio weight $P_{i,j}$ in asset j according to strategy i . We assume that P does not contain redundant forecasts, and in particular that it has full rank $k \leq n$. The parameter $\tau \in [0, 1]$ can be used to fine tune the confidence we have in our forecasts, with τ approaching one capturing the case in which we have full confidence in them, and τ approaching zero the case in which we have no more confidence than what market data suggest. We note that the parameterization we used is slightly different from the canonical BL model, but provides a more intuitive interpretation of the parameter τ along the lines of the sample theoretical approach of Mankert (2006). According to the BL approach, for given vector of forecasts $\boldsymbol{q}$ (i.e., conditionally on $\boldsymbol{Q} = \boldsymbol{q}$), the key parameters of the optimization problem are given as follows:

$$\boldsymbol{\mu}_{BL} = \hat{\boldsymbol{\mu}} + \tau K(\bar{\boldsymbol{q}} - \hat{\boldsymbol{\mu}}), \quad \Sigma_{BL} = \hat{\Sigma}, \quad (3.6)$$

with¹² $K := T^{-1}\hat{\Sigma}P^T \left(PT^{-1}\hat{\Sigma}P^T \right)^{-1} P$ and $\bar{\mathbf{q}}$ is the vector satisfying $P\bar{\mathbf{q}} = \mathbf{q}$, and $\bar{\mathbf{q}} = P^T(PP^T)^{-1}\mathbf{q}$, where we have used the fact that P has full rank. In line with the original contribution by Black and Litterman (1992), we have used the sample variance-covariance matrix for Σ_{BL} , but clearly other choices are possible, as will be discussed below.

The next proposition provides a parameterization allowing us to translate the BL views in terms of Bayesian priors, so as to deliver the same estimates, and hence portfolio allocations, in the two settings (see Schöttle et al., 2010).

Proposition 3.1. *Assume that carbon pricing views are captured by the parameters $P, \mathbf{q}, \tau$ in the BL approach. Let the prior distribution (3.3) in the Bayesian setup rely on the following parameterisation:*

$$d_0 = \tau(1 - \tau)^{-1}T \quad (3.7)$$

$$\boldsymbol{\mu}_0 = K\bar{\mathbf{q}} + (I - K)\hat{\boldsymbol{\mu}} \quad (3.8)$$

$$\Sigma_0 = \eta_1\hat{\Sigma} - \eta_2(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})^T, \quad (3.9)$$

where $\eta_1 = (\nu_0 - n - 1)/\nu_0$, $\eta_2 = T\tau/\nu_0$ and $\nu_0 \geq \underline{\nu}$, with $\underline{\nu}$ the lowest integer ensuring that the matrix Σ_0 is positive semidefinite. Then, the BL forecasts coincides with the Bayesian forecasts, i.e., $\hat{\boldsymbol{\mu}}_b = \hat{\boldsymbol{\mu}}_{BL}$ and $\hat{\Sigma}_b = \Sigma_{BL} = \hat{\Sigma}$.

Proof. See the appendix. □

Proposition 3.1 provides some important messages. The Bayesian approach can accommodate the manager's (absolute and relative) views on carbon pricing by expressing them in terms of trading strategies. For this to be possible, the prior distribution needs to be given by a convex combination of the baseline expected return vector $\hat{\boldsymbol{\mu}}$ and the forecast $\bar{\mathbf{q}}$ in the range of K . The resulting (posterior) mean return vector has the same components as the baseline vector in the case of securities for which we do not have any carbon pricing views. The choice of d_0 given in (3.7) supports the neat structure of expression (3.8) for $\boldsymbol{\mu}_0$, as well as for $\boldsymbol{\mu}_b$, as we have

¹²The matrix K is a projection matrix, in that $K(I - K) = 0$ and $K^2 = K$.

$\beta = \tau$ in equation (3.4). The choice of $\nu_0 \geq \underline{\nu}$ can accommodate different assumptions on the prior variance-covariance matrix, the case of ν_0 growing asymptotically large yielding the sample variance-covariance matrix as prior. Other approaches are clearly possible, given the possibility to choose (d_0, ν_0) differently.

We should now mention that, although our discussion has so far considered the parameters $(\hat{\boldsymbol{\mu}}, \hat{\Sigma})$ as simply being based on historical estimation, more refined choices are possible. Given the large number of stocks considered, in our examples we use shrinkage estimators to produce estimates of Σ , as in Ledoit and Wolf (2003a,b). We also use market-implied returns for $\hat{\boldsymbol{\mu}}$, as in the traditional BL approach, which uses the CAPM as relevant equilibrium benchmark. The optimal allocations resulting from the posterior parameters can then be formalized and interpreted in terms of benchmark portfolio tilts (Black and Litterman, 1992; Idzorek, 2005; Meucci, 2009). As we seek to benefit from the transition to a new equilibrium in which carbon pricing is (correctly) priced, it is appealing to identify explicit departures from a current-state equilibrium benchmark in the recommended portfolio allocations.

4 Smart Carbon Portfolios (SCPs)

We now provide a concrete application of our approach to the construction of portfolios incorporating the views on E&P stocks discussed in section 2. We will refer to the resulting portfolio allocations as SCPs in general, but we shall consider the following different versions in particular:

- **SCP1.** *Graded* divestment from E&P stocks in the MSCI World Index and optimal rebalancing within the E&P sector. We express views on E&P stocks only. Gradation of divestment is shaped by our views on E&P stocks in the carbon pricing state and our confidence in such views. We consider two versions of this portfolio: (a) one in which we constrain the E&P and EE sector weights to be as in the MSCI World index, in order to better understand rebalancing within the E&P sector; (b) a second case in which only the allocation to the rest of the MSCI World portfolio is constrained to the original level, in

order to better understand the rebalancing both within the E&P sector and across the E&P and EE sectors.

- **SCP2.** *Graded* divestment from E&P stocks in the MSCI World Index and optimal rebalancing within the E&P and EE sectors. The portfolio extends SCP1(b) to include *views on the EE sector*.
- **SCP3.** *Capped* divestment from E&P stocks in the MSCI World Index and optimal rebalancing within the E&P sector. We express views on E&P firms only. Capped divestment results from the introduction of constraints on the cross-section of E&P stocks to reflect additional considerations, such as the requirement to maintain a certain allocation to specific E&P firms. As an example, we force the overall allocation to E&P stocks in classes A to C to be at least as large as in the MSCI World portfolio. As we did for portfolio SCP1, we consider further constraints on the aggregate allocations to both the EE sector and the rest of the MSCI World portfolio (SCP3(a)), or just on the latter (SCP3(b)).
- **SCP4.** *Capped* divestment from E&P stocks in the MSCI World Index and optimal rebalancing within the E&P and EE sectors. The portfolio extends SCP3(b) to include *views on the EE sector*.

Portfolios SCP1-4 are constructed by using the market-implied equilibrium return vector. This allows us to see how our views translate into strategic bets on the E&P and EE sectors, tilting their allocations away from a long run equilibrium benchmark such as the MSCI World portfolio in a CAPM context. Analogous portfolios can be constructed with the historical mean vector. In constructing the SCPs, we consider the views on E&P stocks summarized in figure 4, leading to the BL mean return vector components depicted in figure 5. The equivalent assumptions for Bayesian analysis are derived from Proposition 3.1 and presented in the right panel of figure 5. All views are implemented in terms of average stock depreciation (%) across carbon pricing scenarios. We do not express a view on group A of E&P stocks, as the impact of carbon pricing is found to be negligible. For simplicity, and to narrow down our focus on the heterogeneity of carbon pricing

exposures within E&P stocks, we allow for a uniform view across the EE sector, assuming it to be undervalued by 200 or 400 basis points p.a. over the relevant trading horizon.¹³ We believe these figures to provide prudent, although indicative, baseline examples. In the robustness analysis of section 5, we will complement the allocations introduced here with simplified portfolios based on more naive divestment rules. We remark that ours is a partial equilibrium analysis, as no consideration is given to the impact of portfolio rebalancing on stock prices; we refer the reader to Röell (2019) for some interesting general equilibrium considerations.

< FIGURE 5 ABOUT HERE. >

4.1 Implementation

We provide a concrete example of the portfolio construction methodology discussed in the previous section by using stock data from the MSCI index (1,649 components). Within the index, we consider 57 E&P and 56 EE stocks on which we express our views. These are based on the analysis of E&P repricing discussed in section 2: the portfolio matrix P reports long and short allocations to stocks identified as being outperforming and underperforming relative to the benchmark, respectively, where out/under-performance is expressed in absolute terms, by annualizing the expected appreciation/depreciation of the stock over the relevant time horizon. In the case of EE stocks, the view is implemented via a market-value weighted allocation to the EE sector, as in Idzorek (2005), for example (see Satchell and Scowcroft, 2000, for the case of equally-weighted allocations). To manage dimensionality issues, as well as the noise inherent in the estimation of the variance-covariance matrix, we use a shrinkage estimator as in Ledoit and Wolf (2003a,b), with a PCA-based five factor model as shrinkage target (see also Meucci, 2009). We obtain similar results with alternative targets, such as increasing the number of factors or using the average covariance in the off-diagonal elements of the target variance-covariance matrix.

We determine optimal portfolio weights and associated efficient frontiers for the case of a mean-variance investor by using the market implied (CAPM) return vector. We repeat the

¹³This amounts to saying that our views assume both that EE stocks are undervalued and that some E&P stocks are overvalued, given the risk of transition to a carbon pricing state.

portfolio construction exercise with the BL model for different degrees of confidence in our views. In the following, we will discuss some results for the cases of $\tau = 0.25$, $\tau = 0.50$ and the $\tau = 0.99$, the latter capturing closely the case of full confidence in our views. We consider long-only portfolios in all of our examples. As illustrated in figure 6, both portfolios SCP1(b) and SCP2 result in marked divergence from the baseline frontier for high confidence in the views, but the medium confidence case leads to frontiers that are closer to the baseline. Figure 7 shows that the introduction of constraints limiting the extent of divestment from E&P stocks reduces the divergence of BL frontiers relative to the baseline. As shown in the right panels of figures 6-7, the introduction of our views on EE stocks, which reflects our expectation that the EE sector will benefit the most from carbon pricing, does not result in drastic changes in the efficient frontiers.

< FIGURE 6 ABOUT HERE. >

< FIGURE 7 ABOUT HERE. >

As is well known, closeness of portfolios in mean-standard-deviation space can be misleading, as it provides no reliable indication of closeness in portfolio weight space. We therefore now look at the composition of our optimal portfolios by fixing a target expected return coinciding with the one of the MSCI World portfolio. The top panels of figure 8 and table 3 show that portfolios SCP1(b) and SCP2 result in greater allocations to the E&P sector, which are magnified in the case of capped divestment supporting portfolios SCP3(b) and SCP4. The bottom panels of figure 8 and table 3 show that a positive view on EE stocks translates into a greater allocation to the sector only when a cap on a specific group of E&P stocks is introduced. As illustrated in table 4, a closer examination of the portfolio weights reveals that most of the rebalancing takes place within the E&P sector itself. For example, portfolio SCP1(a) shows that class D stocks are the main beneficiaries of rebalancing when the sectoral allocations to E&P, EE and rest of the MSCI World stocks are kept fixed. SCP1(b) shows that, as soon as rebalancing across the E&P and EE sectors is allowed for, the allocation to E&P stocks actually increases at the expenses of the EE sector, group D being again the main beneficiary. The effect is considerably reduced

in portfolios SCP3(a,b), when divestment from classes A to C is capped. On the other hand, imposing a cap on the allocation to group D yields a significant rebalancing of funds toward the EE sector; see portfolios SCP1(b) and SCP2 in panel (B) of tables 3-4. These results all demonstrate the importance of exploiting the cross-sectional heterogeneity of the E&P sector when constructing an SCP, as simple E&P divestment or naive confidence in the outperformance of the EE sector are dominated by the optimal portfolio allocations illustrated here.

< FIGURE 8 ABOUT HERE. >

< TABLE 3 ABOUT HERE. >

< TABLE 4 ABOUT HERE. >

5 Robustness and extensions

Our modelling approach has so far focused on mitigating the risk arising from a transition to carbon pricing. We now discuss more in detail ex-ante efficiency and differences in portfolio weight space. The results obtained for portfolios SCP1-4 demonstrate that low and medium confidence BL portfolios result in portfolios close to the baseline efficient frontier. Portfolios are more markedly distant from the baseline frontier in the case of full confidence in our views. We now argue that efficiency losses can be small even in the latter case, once estimation risk is explicitly taken into account. As was amply discussed earlier, the solutions to the mean-variance optimization problem are sensitive to estimation error in the input parameters (Best and Grauer, 1991; Chopra and Ziemba, 2013). A common approach to understanding this issue and producing allocations robust to estimation error is the portfolio resampling procedure introduced by Michaud (1998). This is essentially a form of bootstrapping producing statistically equivalent portfolio weights, which can be averaged for different target expected return levels to produce a resampled frontier (e.g., Lai et al., 2011). Resampled frontiers reduce the gap between the baseline frontier and the BL frontiers for the SCPs. Alternatively, one can interpret portfolios

on the resampled frontier as proxying for a BL frontier resulting from an intermediate confidence level in the portfolio manager’s views (as given by some $\tau \in (0.25, 0.99)$, for example). Portfolio resampling confirms that SCPs are associated with statistically negligible efficiency losses.

A useful by-product of the resampling procedure is that we can use the resampled portfolio weights to test for the statistical distance of different allocations in portfolio weight space (e.g., Scherer, 2002). A possible metric is given by $(\mathbf{w} - \mathbf{w}_{target})^\top \hat{\Lambda}^{-1}(\mathbf{w} - \mathbf{w}_{target})$, with $\hat{\Lambda}$ the variance-covariance matrix estimated from the resampled portfolio weights; under our assumptions on returns, the resulting test statistic is χ^2 -distributed with degrees of freedom equal to the number of assets (e.g., Scherer, 2002). However, in the presence of constraints invalidating the normality assumption, the following alternative metric is more suitable and is equivalent to the squared tracking error: $(\mathbf{w} - \mathbf{w}_{target})^\top \hat{\Sigma}(\mathbf{w} - \mathbf{w}_{target})$, (see Scherer, 2002; Michaud and Michaud, 2003). We find that portfolios SCP1-4 are statistically distinguishable (at the 95% confidence level) from the MSCI World target, but not among themselves.

In portfolios SCP3-4, we relied on portfolio weights constraints to introduce additional considerations to those channelled by BL views, such as the desire to limit divestment from some E&P stocks. In mean-variance optimization, however, constraining portfolio weights can go much further, as it can be shown to be equivalent to shrinking moment estimators and hence to mitigate the impact of estimation error on portfolio construction (e.g., De Miguel et al., 2009). We could therefore ask whether direct constraints on our E&P stock allocations can mimick the SCPs obtained via the full-fledged Bayesian/BL procedure. To answer this question, we consider simple alternative approaches to portfolio construction by introducing stylized portfolio constraints. Although a wider range of implementations could be considered, we limit ourselves to the following ones:

- **F-REB.** Full divestment from E&P stocks and optimal rebalancing of the index targeting the expected return on the MSCI World portfolio.
- **G-REB.** Graded divestment from E&P stocks and optimal rebalancing of the index targeting the expected return on the MSCI World portfolio. Graded divestment is obtained by

letting the optimizer invest in E&P stocks with portfolio weights allowed to range between zero and the current benchmark allocation (MSCI World).

- **F-EE.** Full divestment from E&P stocks and optimal reallocation toward EE firms within the benchmark portfolio (MSCI World).
- **C-EE.** Capped divestment from E&P stocks and optimal reallocation toward EE names in the benchmark portfolio (MSCI World). Capped divestment is implemented as follows: the E&P stocks with lower exposure to carbon pricing (classes A, B and C) must have allocations on aggregate at least as high as in the benchmark portfolio; the portfolio is optimally divested from the other E&P groups without imposing further constraints.

We assess whether any of the portfolios above can be used as a reasonable approximation for some of our SCPs. In figure 8 and tables 3-4, we report the allocations supporting the above portfolios and compare them with the SCPs. The results show that portfolios F-REB and F-EE do not exploit the rich cross-sectional information on the exposure of the E&P sector to carbon pricing risk. Portfolio G-REB does a better job of picking up group D’s relatively lower exposure to carbon pricing. In comparison with portfolios SCP3 and SCP4, capped divestment misleads portfolio C-EE into allocating to group B and EE stocks, while ignoring group D. We find that portfolios F-REB and F-EE are statistically different from our SCPs, whereas portfolios G-REB and C-EE are closer (according to the test statistics introduced above). The results show that only a judicious choice of constraints can lead to good proxies for some SCPs, but clearly use of the actual SCPs offers a considerably more robust and coherent way to reflect our views on carbon pricing in portfolio construction.

6 Conclusion

In recent years, climate change caused by anthropogenic GHG emissions has risen rapidly up the agendas of policymakers worldwide. A key factor driving this change is the development of more accurate climate models linking weather patterns deviation from historical trends and

anthropogenic GHG emissions. As a result, there is an increasing likelihood that governments of major economies will act within the next decade to reduce these GHG emissions. Carbon pricing (e.g., taxation or cap & trade mechanism) is the most likely mechanism through which governments will seek to achieve this outcome. Industries with weak pricing power and high emission intensity are likely to experience a more significant impact from the introduction of carbon pricing than other industries. We identified Energy, Utilities and Basic Materials as industries with such characteristics, and decided to focus on companies engaged in the exploration and production of fossil fuel assets (E&P stocks). The limited awareness of climate change risk among institutional investors, as well as the empirical evidence on the inefficient pricing of risks brought on by climate change, support the view that such risk is not adequately priced in current stock market valuations. We proposed a methodology that can be used by institutional investors to assess and manage this transition risk, while also taking account of changing information on risk. A key benefit of our approach is that it uses some of the standard tools of investment management, with a particular focus on the potential risk factors to impact a company's financial prospects. The approach captures the potential impact of carbon pricing on fossil fuel stocks, and uses it to inform Bayesian and Black-Litterman type portfolio construction methodologies, which are then used to create what we call Smart Carbon Portfolios (SCPs). We find that investors could reduce ex-post risk by lowering the weightings of some fossil fuel stocks with corresponding higher weightings in lower-risk fossil fuel stocks and/or in the stocks of companies active in energy efficiency markets. For robustness, we discussed portfolio resampling to show how the methodology reduces the risk of a transition to a carbon pricing scenario while sacrificing virtually no ex-ante financial gains. The use of ad hoc constraints on allocations and/or divestment was also considered as a simpler alternative to fully fledged Bayesian analysis. The extension of our analysis to encompass other sectors, such as Utilities and Basic Materials, is left for further research.

References

- Akerlof, K., Maibach, E. W., Fitzgerald, D., Cedenro, A. Y., and Neuman, A. (2013). Do people “personally experience” global warming, and if so how, and does it matter? *Global Environmental Change*, 23(1):81–91.
- Anderson, T. R., Hawkins, E., and Jones, P. D. (2016). CO₂, the greenhouse effect and global warming: From the pioneering work of Arrhenius and Calendar to today’s earth system models. *Endeavour*, 40(3):178–187.
- Andersson, M., Bolton, P., and Samama, F. (2016). Hedging climate risk. *Financial Analysts Journal*, 72(3):13–32.
- Ang, A. and Bekaert, G. (2004). How regimes affect asset allocation. *Financial Analysts Journal*, 60(2):86–99.
- Avramov, D. and Wermers, R. (2006). Investing in mutual funds when returns are predictable. *Journal of Financial Economics*, pages 339–377.
- Bader, D. C., Covey, C., Gutowski Jr, W. J., Held, I. M., Miller, R. L., Tokmakian, R. T., Zhang, M. H., et al. (2008). Climate models: An assessment of strengths and limitations. *Geological and Atmospheric Sciences Reports*.
- Barry, C. B. (1974). Portfolio analysis under uncertain means, variances, and covariances. *The Journal of Finance*, 29(2):515–522.
- Bessembinder, H. (2016). Frictional costs of fossil fuel divestment. Working Paper, Arizona State University.
- Best, M. J. and Grauer, R. R. (1991). On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *The Review of Financial Studies*, 4(2):315–342.
- Biffis, E. and Blake, D. (2013). Informed intermediation of longevity exposures. *Journal of Risk and Insurance*, 80(3):559–584.
- Biffis, E. and Blake, D. (2014). Keeping some skin in the game: How to start a capital market in longevity risk transfers. *North American Actuarial Journal*, 18(1):14–21.
- Biffis, E. and Chavez, E. (2017). Satellite data and machine learning for weather risk management and food security. *Risk Analysis*, 37(8):1508–1521.
- Black, F. and Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43.
- Brealey, R. A., Myers, S. C., Allen, F., and Mohanty, P. (2012). *Principles of corporate finance*. Tata McGraw-Hill Education.
- Broadie, M. (1993). Computing efficient frontiers using estimated parameters. *Annals of Operations Research*, 45:21–58.
- Caldecott, B., Dericks, G., and Mitchell, J. (2015). Stranded assets and subcritical coal. Smith School of Enterprise and the Environment, University of Oxford.
- Carney, M. (2015). Breaking the tragedy of the horizon - climate change and financial stability. Speech given by Mark Carney, Governor of the Bank of England and Chairman of the Financial Stability Board.
- Chavez, E., Christiaensen, L., Shafir, E., Biffis, E., Ghil, M., Picard, P., Louaas, A., Toda,

- A., Nartey, C., Barron, K., Makuch, Z., Madani, K., Conway, G., Thompson, W., Cousin, E., Lucarini, V., Ramachandran, M., Xiong, W., Vandecasteele, J., Tarnavsky, E., Gonzalez-Pelaez, A., Cabrera, E., and Kilian, M. (2018). Rethinking agricultural supply chains. Working Paper, Brevan Howard Centre for Financial Analysis, Imperial College Business School, London.
- Cheung, W. (2010). The Black–Litterman model explained. *Journal of Asset Management*, 11(4):229–243.
- Choi, D., Gao, Z., and Jiang, W. (2018). Attention to global warming. Working Paper, The Chinese University of Hong Kong.
- Chopra, V. K. and Ziemba, W. T. (2013). The effect of errors in means, variances, and covariances on optimal portfolio choice. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, pages 365–373.
- Clarke, R. G. and de Silva, H. (1998). State-dependent asset allocation. *The Journal of Portfolio Management*, 24(2):57–64.
- Daniel, K. D., Litterman, R. B., and Wagner, G. (2016). Applying asset pricing theory to calibrate the price of climate risk. Working Paper No. 22795, National Bureau of Economic Research.
- De Miguel, V., Garlappi, L., Nogales, F. J., and Uppal, R. (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5):798–812.
- Field, C. B. (2014). *Climate change 2014–Impacts, adaptation and vulnerability: Regional aspects*. Cambridge University Press.
- Gao, Y. (2016). China’s response to climate change issues after paris climate change conference. *Advances in Climate Change Research*, 7(4):235–240.
- Gründl, H., Gal, J., et al. (2016). The evolution of insurer portfolio investment strategies for long-term investing. *OECD Journal: Financial Market Trends*, pages 1–55.
- Hong, H. G., Li, F. W., and Xu, J. (2016). Climate risks and market efficiency. *Journal of Econometrics*. forthcoming.
- IAIS (2018). Issues Paper on Climate Change Risks to the Insurance Sector. International Association of Insurance Supervisors.
- Idzorek, T. (2005). A step-by-step guide to the Black-Litterman model. Technical report.
- IPCC (2014). Climate change 2014: Impacts, adaptation, and vulnerability. part B: Regional aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Barros, V.R., C.B. Field, D.J. Dokken, M.D. Mastrandrea, K.J. Mach, T.E. Bilir, M. Chatterjee, K.L. Ebi, Y.O. Estrada, R.C. Genova, B. Girma, E.S. Kissel, A.N. Levy, S. MacCracken, P.R. Mastrandrea, and L.L.White (eds.). Cambridge University Press.
- IPCC (2015). *Climate change 2014: Mitigation of climate change*, volume 3. Cambridge University Press.
- Klein, R. W. and Bawa, V. S. (1976). The effect of estimation risk on optimal portfolio choice. *Journal of Financial Economics*, 3(3):215–231.
- Krüger, P., Sautner, Z., and Starks, L. (2018). The importance of climate risks for institutional investors. Research Paper Series N. 18-58, Swiss Finance Institute.
- Lai, T. L., Xing, H., and Chen, Z. (2011). Mean-variance portfolio optimization when means and covariances are unknown. *The Annals of Applied Statistics*, 5(2A):798–823.

- Ledoit, O. and Wolf, M. (2003a). Honey, I shrunk the sample covariance matrix. UPF Economics and Business Working Paper.
- Ledoit, O. and Wolf, M. (2003b). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5):603–621.
- Litterman, R. (2015). David Swensen on the fossil fuel divestment debate. *Financial Analysts Journal*, 71(3):11.
- Mankert, C. (2006). The Black-Litterman model - mathematical and behavioral finance approaches towards its use in practice. Licentiate Thesis. Royal Institute of Technology, Stockholm.
- Mercer (2015). Investing in a time of climate change. Mercer, London.
- Meucci, A. (2009). *Risk and asset allocation*. Springer.
- Michaud, R. (1998). Efficient asset management: A practical guide to stock portfolio management and asset allocation. *Financial Management Association, Survey and Synthesis Series*. HBS Press, Boston, MA.
- Michaud, R. and Michaud, R. (2003). “Portfolio resampling: Review and critique”: A comment. *Financial Analysts Journal*, 59(3):17–17.
- Myers, T. A., Maibach, E. W., Roser-Renouf, C., Akerlof, K., and Leiserowitz, A. A. (2013). The relationship between personal experience and belief in the reality of global warming. *Nature Climate Change*, 3(4):343.
- Röell, A. (2019). Divestment, portfolio choice and asset prices: An equilibrium approach. Columbia University’s School of International and Public Affairs, Imperial College Business School Working Paper.
- Roncalli, T. (2014). *Introduction to Risk Parity and Budgeting*. Chapman & Hall, CRC Press.
- Satchell, S. and Scowcroft, A. (2000). A demystification of the black-litterman model: Managing quantitative and traditional portfolio construction. *Journal of Asset Management*, 1(2):138–150.
- Scherer, B. (2002). Portfolio resampling: Review and critique. *Financial Analysts Journal*, 58(6):98–109.
- Schöttle, K., Werner, R., and Zagst, R. (2010). Comparison and robustification of Bayes and Black-litterman models. *Mathematical Methods of Operations Research*, 71:453–475.
- Schroders (2016). Climate change: Redefining the risks. Schroders, London.
- Shiogama, H., Stone, D., Emori, S., Takahashi, K., Mori, S., Maeda, A., Ishizaki, Y., and Allen, M. (2016). Predicting future uncertainty constraints on global warming projections. *Scientific Reports*, 6:18903.
- Simm, I., Winter, J., Chatzimichalakis, F., and Biffis, E. (2016). Carbon risk for investors: Building a “smart carbon” portfolio. IMPAX Asset Management and Imperial College Business School.
- Stocker, T. (2014). *Climate change 2013: the physical science basis: Working group contribution to the Fifth assessment report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
- World Bank (2017). State and trends of carbon pricing 2017. Technical report. Washington, DC.
- Zaval, L., Keenan, E. A., Johnson, E. J., and Weber, E. U. (2014). How warm days increase belief in global warming. *Nature Climate Change*, 4(2):143.

Appendix

A Technical details

A.1 Bayesian approach.

We characterize our joint *prior distribution* of return parameters $\boldsymbol{\mu}$, Σ as a Normal Inverse Wishart ($\mathcal{NIW}$). In particular, we assume Σ to have an Inverse Wishart ($\mathcal{IW}$) distribution and the conditional distribution of $\boldsymbol{\mu}$, given Σ , to be Normal ($\mathcal{N}$):

$$(\boldsymbol{\mu}, \Sigma) \sim \mathcal{NIW}(\boldsymbol{\mu}_0, d_0, \Sigma_0, \nu_0), \quad \boldsymbol{\mu}|\Sigma \sim \mathcal{N}(\boldsymbol{\mu}_0, d_0^{-1}\Sigma), \quad \Sigma \sim \mathcal{IW}(\nu_0\Sigma_0, \nu_0 + n + 1),$$

where we have that $E[\boldsymbol{\mu}|\Sigma] = \boldsymbol{\mu}_0$ and $cov(\boldsymbol{\mu}|\Sigma) = d_0^{-1}\Sigma$ (see: Meucci, 2009; Schöttle et al., 2010; Lai et al., 2011). Given market return observations $\mathbf{R}$, we can compute the sample mean return and variance-covariance as follows:

$$\hat{\boldsymbol{\mu}} = \frac{1}{T} \sum_{t=1}^T \mathbf{R}_t, \quad \hat{\Sigma} = \frac{1}{T} \sum_{t=1}^T (\mathbf{R}_t - \hat{\boldsymbol{\mu}})(\mathbf{R}_t - \hat{\boldsymbol{\mu}})^\top.$$

Given our assumptions, we have that the posterior distribution is again of $\mathcal{NIW}$ type. In particular, denoting by $\mathcal{F}_T$ the information generated by the time series of returns $\{\mathbf{R}_t\}_{t=1}^T$, we have $\boldsymbol{\mu}, \Sigma|\mathcal{F}_T \sim \mathcal{NIW}(\boldsymbol{\mu}_1, d_1, \Sigma_1, \nu_1)$, with parameters given by:

$$\begin{aligned} \boldsymbol{\mu}_1 &= \beta\boldsymbol{\mu}_0 + (1 - \beta)\hat{\boldsymbol{\mu}}, \\ \Sigma_1 &= \gamma_0\hat{\Sigma} + (1 - \gamma_0)\Sigma_0 + \beta\gamma_0(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})^\top, \\ \beta &= d_0/d_1, \quad \gamma_0 = \nu_0/\nu_1, \quad d_1 = d_0 + T, \quad \nu_1 = \nu_0 + T. \end{aligned}$$

The posterior estimates for the relevant market parameteres are therefore as follows (e.g., Meucci, 2009; Schöttle et al., 2010; Lai et al., 2011):

$$\boldsymbol{\mu}_b = E[\boldsymbol{\mu}|\mathcal{F}_T] = \boldsymbol{\mu}_1, \quad \hat{\Sigma}_b = \frac{\nu_1}{\nu_1 - n - 1} \Sigma_1,$$

where we note that $\hat{\Sigma}_b$ goes to $\hat{\Sigma}$ as ν_0 grows asymptotically large.

A.2 Black-Litterman approach.

Consider a mean return vector characterized by the distribution $\mathcal{N}(\tilde{\boldsymbol{\mu}}, \tau\tilde{\Sigma})$, for some $\tau \in [0, 1]$, where $\tilde{\boldsymbol{\mu}}$ can be based on the sample mean $\hat{\boldsymbol{\mu}}$ or, as in Black and Litterman (1992), on the expected return vector implied by an equilibrium model such as the CAPM, in which case we set $\tilde{\boldsymbol{\mu}} = 2\lambda\Sigma\mathbf{w}_{mkt}$, with $\mathbf{w}_{mkt}$ the market capitalization weights and λ a suitable risk aversion parameter. A useful choice for $\tilde{\Sigma}$ is $T^{-1}\hat{\Sigma}$ (e.g., Mankert, 2006; Schöttle et al., 2010). Given a set of views or expert forecasts on different assets, one can expect to monetize them by defining a view portfolio matrix P yielding the return $\mathbf{Q} = P\boldsymbol{\mu} + \boldsymbol{\varepsilon}$, which is random not only because of

the unknown parameter $\boldsymbol{\mu}$, but also because of potential errors in the portfolio manager's views. The error term $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, (1 - \tau)\Omega)$ is assumed to be uncorrelated with $\boldsymbol{\mu}$, and results in $\mathbf{Q}$ being distributed according to $\mathcal{N}(P\boldsymbol{\mu}, (1 - \tau)\Omega)$. We note that we have used a slightly different parameterization from the canonical BL approach (e.g., Mankert, 2006; Schöttle et al., 2010), so that we can use τ to express our confidence in the prior information (historical or market-implied) relative to the expert views: for τ approaching one, we have full confidence in the views; for τ approaching zero, we have full confidence in the prior. Finally, we set $\Omega = P\tilde{\Sigma}P^\top$ (e.g., Idzorek, 2005; Mankert, 2006; Schöttle et al., 2010), which simply means that¹⁴ $\boldsymbol{\varepsilon}$ is exposed to the same uncertainty as $\boldsymbol{\mu}$. One can then show that the conditional mean return is given by (e.g., Meucci, 2009):

$$\boldsymbol{\mu}_{BL} = E[\boldsymbol{\mu}|\mathbf{Q}] = \tilde{\boldsymbol{\mu}} + \tau\tilde{\Sigma}P^\top \left(P\tilde{\Sigma}P^\top\right)^{-1} (\mathbf{q} - P\tilde{\boldsymbol{\mu}}).$$

A.3 Proof of Proposition 3.1.

We consider the BL model of the previous section with $\tilde{\boldsymbol{\mu}} = \hat{\boldsymbol{\mu}}$. The Bayesian setting can handle unknown variance-covariance matrices comfortably, but this is more complex in the case of the BL methodology. We therefore assume $\tilde{\Sigma}$ to be known and rely on the sample variance-covariance estimate. To prove the proposition, let us first note that for each forecast $\mathbf{q}$ we can always write $\mathbf{q} = P\bar{\mathbf{q}}$ for some $\bar{\mathbf{q}} \in \mathbb{R}^k$, or equivalently $\bar{\mathbf{q}} = P^\top (PP^\top)^{-1} \mathbf{q}$, as P is assumed to have full rank. We can then write $\boldsymbol{\mu}_{BL}$ as follows:

$$\begin{aligned}\boldsymbol{\mu}_{BL} &= \hat{\boldsymbol{\mu}} + \tau\hat{\Sigma}P^\top \left(P\hat{\Sigma}P^\top\right)^{-1} P(\bar{\mathbf{q}} - \hat{\boldsymbol{\mu}}) \\ &= \hat{\boldsymbol{\mu}} + \tau K(\bar{\mathbf{q}} - \hat{\boldsymbol{\mu}}) = (I - \tau K)\hat{\boldsymbol{\mu}} + \tau K\bar{\mathbf{q}},\end{aligned}$$

where we have denoted by K the matrix $\hat{\Sigma}P^\top(P\hat{\Sigma}P^\top)^{-1}P$. Setting $d_0 = \tau(1 - \tau)^{-1}T$ as in (3.7), expression (3.4) becomes $\boldsymbol{\mu}_b = \tau\boldsymbol{\mu}_0 + (1 - \beta)\tilde{\boldsymbol{\mu}}$. By imposing the equality $\boldsymbol{\mu}_{BL} = \boldsymbol{\mu}_b$, we can solve for the vector $\boldsymbol{\mu}_0$ and obtain $\boldsymbol{\mu}_0 = K\bar{\mathbf{q}} + (I - K)\hat{\boldsymbol{\mu}}$. By imposing now the equality $\Sigma_{BL} = \Sigma_b$, and using again $d_0 = \tau(1 - \tau)^{-1}T$, we can solve for Σ_0 and obtain:

$$\Sigma_0 = \hat{\Sigma} - \frac{n+1}{\nu_0}\hat{\Sigma} - \frac{S\tau}{\nu_0}(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})^\top.$$

The matrix on the left-hand side is positive semidefinite. Hence, there must be a minimum integer value $\underline{\nu}$ for ν_0 such that the right-hand side is always a positive semidefinite matrix for $\nu_0 \geq \underline{\nu}$. So far we have made the BL approach use historical data as a benchmark, which is more amenable to a straightforward comparison with the Bayesian approach. Replacing $\tilde{\boldsymbol{\mu}}$ with $2\lambda\Sigma\mathbf{w}_{mkt}$ instead, we could rewrite the results for the case in which the benchmark is an equilibrium model such as the CAPM.

¹⁴As P is assumed to have full rank, we can always write $\boldsymbol{\varepsilon} = P\boldsymbol{\eta}$ for some $\boldsymbol{\eta} \in \mathbb{R}^n$. The least square choice for $\boldsymbol{\eta}$ results in the representation $\boldsymbol{\eta} = P^\top (PP^\top)^{-1}$, from which we obtain $\mathbf{Q} = P(\boldsymbol{\mu} + \boldsymbol{\eta})$ and hence $\Omega = P\Lambda P^\top$ for some matrix Λ . The choice $\Lambda = \tilde{\Sigma}$ seems natural, as long as we believe that the uncertainty in $\boldsymbol{\varepsilon}$ is driven by the same factors affecting that of $\boldsymbol{\mu}$.

B Tables and figures

		\$10	\$20	\$30	\$40	\$50	\$60
Year	2020	4%	4%	3%	2%	1%	1%
	2021	3%	3%	3%	3%	1%	1%
	2022	2%	3%	4%	3%	1%	1%
	2023	1%	2%	4%	4%	2%	1%
	2024	1%	1%	3%	4%	3%	2%
	2025	1%	1%	2%	4%	4%	2%

Table 1: Carbon pricing scenarios probability table: unconditional probabilities by timing of introduction (calendar year) and magnitude (\$ per ton of CO₂), under the assumption that the probability of carbon pricing not being introduced is 15%.

Group	Total market cap (\$m)	Average market cap (\$m)	Energy reserves (BOE m)
B	1,140,097.57	190,016.26	12781.97
C	284,577.02	142,288.51	9607.33
D	535,240.68	17,265.83	1329.37
E	4,166.60	4,166.60	
F	68,940.73	68,940.73	5038.00
G	12,891.96	12,891.96	2152.00
H	224,187.63	28,023.45	1608.85
I	113,864.92	16,266.42	4001.90
Total	2,383,967.11		36519.43

Table 2: Breakdown of E&P stocks by groups B to I. Group E stands for Coal producers, group I for Oil Sands. Total and average market capitalization (in million USD), as well as energy reserves (in million BOE), are as of May 8, 2019.

	MSCI	SCP1(a)	SCP1(b)	SCP2	SCP3(a)	SCP3(b)	SCP4	F-REB	G-REB	F-EE	C-EE
(A)											
E&P	5.42%	5.42%	6.2%	6.2%	5.42%	7.29%	7.29%	0.00%	1.65%	0.00%	3.13%
EE	2.95%	2.95%	2.17%	2.17%	2.95%	1.08%	1.08%	2.95%	2.95%	8.37%	5.24%
MSCI R.	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	97.05%	95.4%	91.63%	91.63%
(B)											
E&P	5.42%	5.42%	2.92%	2.92%	5.42%	5.52%	5.52%	0.00%	1.65%	0.00%	3.13%
EE	2.95%	2.95%	5.45%	5.45%	2.95%	2.85%	2.85%	2.95%	2.95%	8.37%	5.24%
MSCI R.	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	97.05%	95.4%	91.63%	91.63%

Table 3: Portfolio weights for different portfolios targeting the expected return for the MSCI World portfolio (12.97%, last three years sampling period). The weights are grouped into three sectors: E&P, EE, and rest of the MSCI World portfolio (MSCI R.). SCPs are built with the BL methodology and parameter $\tau = 0.25$. In panel (A), we report the results for the case of an undervaluation view on the EE sector of 2% p.a.. In panel (B), we report the results for the case in which the allocation to the D group of E&P stocks is capped at the MSCI World portfolio level (see table 4 for details on allocations within the E&P sector). Figures in **bold** format indicate an increase in the allocation relative to the baseline MSCI World portfolio.

	MSCI	SCP1(a)	SCP1(b)	SCP2	SCP3(a)	SCP3(b)	SCP4	F-REB	G-REB	F-EE	C-EE
(A)											
B	2.42%	0.00%	0.00%	0.00%	2.83%	2.81%	2.81%	0.00%	0.00%	0.00%	3.04%
C	0.71%	0.00%	0.00%	0.00%	0.29%	0.31%	0.31%	0.00%	0.00%	0.00%	0.09%
D	0.99%	4.25%	4.96%	4.95%	1.51%	3.16%	3.15%	0.00%	1.13%	0.00%	0.00%
F	0.2%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
G	0.05%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
H	0.68%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Coal	0.01%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Oil Sands	0.36%	1.17%	1.25%	1.25%	0.78%	1.01%	1.01%	0.00%	0.52%	0.00%	0.00%
EE	2.95%	2.95%	2.17%	2.17%	2.95%	1.08%	1.08%	2.95%	2.95%	8.37%	5.24%
MSCI R.	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	97.05%	95.4%	91.63%	91.63%
(B)											
B	2.42%	1.35%	0.00%	0.00%	2.83%	2.83%	2.83%	0.00%	0.00%	0.00%	3.04%
C	0.71%	0.03%	0.00%	0.00%	0.30%	0.30%	0.30%	0.00%	0.00%	0.00%	0.09%
D	0.99%	0.99%	0.99%	0.99%	0.99%	0.99%	0.99%	0.00%	1.13%	0.00%	0.00%
F	0.20%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
G	0.05%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
H	0.68%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Coal	0.01%	0.31%	0.06%	0.06%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Oil Sands	0.36%	2.75%	1.86%	1.87%	1.30%	1.40%	1.41%	0.00%	0.52%	0.00%	0.00%
EE	2.95%	2.95%	5.45%	5.45%	2.95%	2.85%	2.85%	2.95%	2.95%	8.37%	5.24%
MSCI R.	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	91.63%	97.05%	95.4%	91.63%	91.63%

Table 4: Portfolio weights for different portfolios targeting the expected return for the MSCI World portfolio (12.97%, last three years sampling period). The SCPs are built with the BL methodology and parameter $\tau = 0.25$. In addition to the SCPs, the table reports also portfolios F/G-REB and F/C-EE introduced in section 5. In panel (A), we report the results for the case of an undervaluation view on the EE sector of 4% p.a.. In panel (B), we report the results for the case in which the allocation to the D group of E&P stocks is capped at the MSCI World portfolio level. Figures in **bold** format indicate an increase in the allocation relative to the baseline MSCI World portfolio.




Figure 1: Efficient frontiers in the current (business as usual) state and in the carbon pricing state. In addition to the state-dependent frontiers, we plot the ex-ante frontier obtained by applying the methodology of section 3.




Figure 2: Forward production data for an anonymized entity labelled “Company 1”. The chart depicts the oil reserves, quantified in barrel of oil equivalent (BOE) terms, available in different future years between 2020 and 2025. Different oil reserves are characterized by different break-even oil prices. For example, deepwater and ultra-deepwater reserves can have a significant proportion of their break-even prices above \$150, whereas tight liquids reserves can have a substantial portion of their breakeven price distribution below a level of \$80. The particular example depicted here shows that Company 1 has most of its oil reserves with a break-even price below \$60, and that, for example, reserves with a \$15-20 break-even price bracket are expected to decrease during the period 2020-25.




Figure 3: Simple supply-demand diagram for a fossil fuel sector (e.g., crude oil) taken from Simm et al. (2016). The x -axis reports assets with different volumes and marginal costs of production: to fix ideas, Middle Eastern oil is likely to be to the left of the curve, while Arctic deep sea reserves will likely be to the right. Drawing a demand curve gives the volume consumed by the industry, V_0 , and the market price P_0 which is both paid by consumers and received by suppliers. Considering now one type of carbon pricing, a tax T imposed by the government on the industry, raises the supply curve, which now intersects with the demand curve at a new point. The volume consumed drops to V_1 , while consumers pay P_1 but suppliers receive only P_2 . Assets represented by bars between V_0 and V_1 become stranded. See Simm et al. (2016) for further details.




Figure 4: Mispricing of EE and E&P stocks (left panel), and cross-sectional disaggregation of views within the E&P sector (right panel). E&P stocks are clustered into eight groups, depending on the severity of their exposure to carbon pricing. Groups are labeled from A to I. Group A is not reported as the impact of carbon pricing is found to be negligible. Group E corresponds to coal stocks, group I to Oil Sands. Mispricing is expressed in terms of average stock depreciation (% p.a.) across carbon pricing scenarios.




Figure 5: Black-Litterman (BL) expected return vector μ_{BL} for confidence levels $\tau = 0.25$ and $\tau = 0.99$, with components clustered into the main E&P classes, the EE sector and the rest of the MSCI World index. By expression (3.8), the case of τ approaching one can be understood as yielding the Bayesian prior μ_0 . As confidence in the views increases, the BL adjustment becomes more material for classes Coal and Oil Sands in particular.




Figure 6: Baseline efficient frontier (plug-in estimators) and BL frontiers for portfolios SCP1(b) (left panel) and SCP2 (right panel) in the cases of medium ($\tau = 0.5$) and high ($\tau = 0.99$) confidence. The case of $\tau = 0.25$ gives a frontier very close to the baseline.




Figure 7: Baseline efficient frontier (plug-in estimators) and BL frontiers for portfolios SCP3(b) (left panel) and SCP4 (right panel) in the cases of medium ($\tau = 0.5$) and high ($\tau = 0.99$) confidence. The case of $\tau = 0.25$ gives a frontier very close to the baseline.




Figure 8: Portfolio weights supporting SCP1-4, as well as portfolios F-REB, G-REB, F-EE, and C-EE, for a target expected return coinciding with the one delivered by the MSCI World Index (last three year sampling period). Allocations are clustered into E&P stocks, the EE sector, and the rest of the MSCI World portfolio. The top panel reports the case with an undervaluation view on the EE sector of 4% p.a., whereas the bottom panel reports results for the same case but with the allocation to group D of E&P stocks (31 firms in total) capped at the current MSCI World allocation level. Views on E&P stocks are as per figure 4. The top panel consider the case of $\tau = 0.99$, the bottom panel the case of $\tau = 0.25$.