

Net-Zero Investing: Harnessing the Power of Unstructured Data



NET-ZERO INVESTING: HARNESSING THE POWER OF UNSTRUCTURED DATA

Andy Moniz, CFA

Senior Vice President, Director of Responsible Investing,
Acadian Asset Management, London

Devin Nial

Senior Vice President, Portfolio Manager, ESG, Acadian Asset Management,
Boston

Matthew Picone, CFA

Senior Vice President, Portfolio Manager, ESG, Acadian Asset Management,
Sydney

Lukasz Pomorski

Senior Vice President, Portfolio Manager, Acadian Asset Management, Boston

Jerry Yu, CFA

Assistant Vice President, Associate Portfolio Manager, ESG, Acadian Asset
Management, Boston

Climate is increasingly important for investors, but to address it in an investment portfolio, one needs to overcome a significant data challenge. On the one hand, data providers try to cater to investor demand with various datasets; on the other hand, such offering is often a black box that may heavily depend on noisy historical data. This situation is of particular concern to net-zero investors, who need solutions that can be plausibly tied to companies' emission trajectories over very long periods of time. The purpose of this chapter is to explain how investors may respond to this challenge and to propose a realistic implementation that addresses it. We highlight how climate investors can leverage unstructured data through natural language processing (NLP), how they should incorporate new information that becomes available over time, and how they may deal with the uncertainty inherent in climate alignment estimates. Our example application showcases the use of NLP and unstructured data and also stresses many other design choices that, in our view, will improve net-zero solutions.

Note: The views expressed herein are those of the authors and do not necessarily reflect the views of Acadian Asset Management LLC. The views should not be considered investment advice and do not constitute or form part of any offer to issue or sell, or any solicitation of any offer to subscribe or to purchase, shares, units, or other interests in any particular investments.

Introduction

Climate considerations are increasingly important for investors, with use cases ranging from identifying potential risks and opportunities that may affect a financial portfolio to identifying targets for proxy voting and company engagement. These activities critically depend on the availability and quality of climate data; unfortunately, this is a major issue for investors. While multiple data providers offer a range of climate solutions, there are legitimate concerns about the usefulness of such data. For example, much of the data capture only historical firm behavior, but potential risk, opportunities, and engagement goals are all forward looking. This tension is particularly important for investors aiming to build net-zero-aligned portfolios. On the one hand, the idea behind net-zero investing is deceptively easy to explain: Build a portfolio of securities that are well positioned should the world economy decarbonize, potentially all the way to "net zero." On the other hand, translating this straightforward idea to an actual portfolio is exceedingly difficult because it requires investors to map company characteristics today to decades out into the future. Today, few companies can credibly claim to have achieved net zero, so building a realistic portfolio necessarily requires investors to take a stance on how issuer behavior may evolve, possibly over multiple decades. Moreover, data quality is often dubious because of both measurement problems and, perhaps even more importantly, the vagueness of corporate communications or outright greenwashing. Increasingly, many companies proclaim the desire to decarbonize and may even commit to specific targets. However, the credibility of these targets likely differs among companies, and investors today have relatively few tools to be able to assess this.

We believe that to address these challenges, investors need to increasingly rely on alternative data and on new techniques to extract actionable insights from such data. We focus primarily on textual data that may be disseminated by either the company in question or external stakeholders (e.g., nongovernmental organizations and the news media) and on the tools designed to process such data, collectively referred to as natural language processing (NLP). We explain why these data and this approach are critical for understanding firms' climate exposure and potential greenwashing by the underlying issuers. We follow up with a case study that explains in detail how one may build a measure of net-zero alignment in practice.

Our practical example illustrates an important theme that we believe all realistic climate solutions must share. There is no silver bullet to address portfolio climate needs, so investors must be prepared to use creative solutions that blend multiple data sources and techniques. The case study we present leverages NLP, but to build the overall climate measure, it also needs data that may not be directly climate related (e.g., sell-side analyst earnings forecasts) and additional statistical techniques (e.g., Bayesian updating, to update the measure as new data become available and to build not just a point estimate but also a range of possible outcomes for a given firm).

Limitations of Existing Data Solutions

Given the growing interest in climate and net-zero investing, it is not surprising that data providers have proposed a plethora of potential solutions. Unfortunately, such solutions tend to suffer from two major weaknesses: First, they usually provide only partial coverage of the investment universe, and second, they sometimes only have a tenuous relationship with the stated goal of alignment with economic outcomes far out into the future (Heal and Millner 2014; Pindyck 2017). Coverage is a perennial issue in sustainable investment, reflecting more company disclosure for large-cap issuers and for developed issuers. While intuitive, the lack of coverage is a problem for many asset owners who worry about the climate alignment of their overall portfolio and not just their, say, large-cap developed mandates. To illustrate this issue, one could survey the offering of net-zero index providers. While there are popular large-cap net-zero indexes (MSCI World Climate Paris Aligned Index, just to give one example), to the best of our knowledge, no similar small-cap indexes exist. Clearly, this situation clashes with the guidance from the Net-Zero Asset Owner Alliance that advises investors to “bring the focus of addressing the systemic risk of climate change to the entirety of investments and operations” (UN Environment Programme 2024).

The second issue is that the currently available data may be only a very noisy measure of net-zero alignment (Schneider and Kuntz-Duriseti 2002; Thiele 2020). This is partially a function of regulation. For example, the net-zero indexes, such as the one mentioned previously, reflect the EU’s minimum technical standards that prominently feature measures of carbon intensity. However, carbon intensity captures a company’s emissions today and perhaps in the near future (for a relevant analysis, see Bixby, Brixton, and Pomorski 2022), so it may not always be a good measure of emissions that are still decades away. Moreover, when data providers come up with their proprietary measures, they may use subjective or relatively opaque methodologies (Task Force on Climate-related Financial Disclosures 2020) and may struggle to demonstrate the link between them and the desired future economic outcomes. Indeed, the implied temperature scores published by data providers, often provided with decimal-point precision, suggest an unwarranted high degree of accuracy of climate forecasts (Robinson-Tillett 2022). This leads to a paradoxical situation in which we are inundated with different climate alignment data that meaningfully differ across providers, making it challenging for the asset owner to identify and justify which specific source to rely on. For example, even if an investor decides on a specific type of data (e.g., Scope 3 emissions or implied temperature scores), such data can have very low correlations between providers, potentially leading to very different investment outcomes.

Proposed Solution: Machine Learning to the Rescue

We argue that machine learning (ML) techniques offer a viable alternative to improve an investor’s situation for two overlapping reasons. First, insights about long-term climate exposure and outcomes can realistically be obtained only

from unstructured data. Second, to process unstructured data, one has little choice but to resort to ML and, in particular, to one specific subarea of these tools, NLP.

The first argument is that the data net-zero investors need are likely to be unstructured. It is probably unrealistic to expect that issuers might produce numerical data that can plausibly describe their climate exposure in, say, 2050. Even if a company does produce such an estimate or scenario, it will reflect a range of assumptions that may be specific to the given company and thus not generalizable to others. Understanding such assumptions should plausibly affect one's assessment of the company's climate exposure and alignment. For example, a company may pledge a net-zero commitment. On its own, this may seem to be a positive development, but the full assessment will likely require a careful analysis of the specific steps the company is planning to undertake, intermediate targets and milestones, current and planned future disclosures, and so on. Such diverse information will not be presented in a numerical form, and it may not even lend itself to a tabular template. Instead, it will likely be a narrative, with free-form language describing the company's ambitions.

The second argument is that to process such data at scale, it is perhaps inevitable to eventually use ML techniques. Continuing with the previous example, it is, of course, conceivable that human analysts can process information about any one issuer's net-zero commitment and arrive at an informed view about its quality and likelihood of success. Unfortunately, this model does not scale. Even large data providers may not be able to hire hundreds of analysts to assess the thousands of issuers that a large investor may hold in its portfolio. We cannot solve the coverage issue with standard statistical techniques, such as regression-type tools. As we explained previously, at least some relevant information will not be numerical, which will prevent a purely "parametric" approach. Moreover, we may have somewhat different information about each individual issuer, and we cannot resolve the problem by simply hiring hundreds of analysts. It seems unlikely that human researchers could produce data that would be comparable across a wide range, possibly thousands, of issuers. The human analyst thought process is ultimately a black box that may not easily translate between how two skilled analysts may view a given company. In our view, ML is the only realistic solution that can reliably scale and that can handle the complexity of the underlying data.

In addition to efficiently handling large volumes of unstructured data, ML could also be helpful for investors building a holistic measure that aggregates a number of climate indicators, each of which is only weakly correlated with the desired outcome. This is especially true when there are nonlinearities and interactions between various pieces of data, which we believe is likely in climate investing. Some issuers that are clearly brown today are likely to be among the most important drivers of lowering carbon emissions in the future. For example, some energy or utility companies with current high emissions may be well positioned to meet the world's future nonnegotiable energy needs; they may

also have the resources and a clear economic incentive to pursue the relevant research and development today (e.g., Cohen, Gurun, and Nguyen 2020).

Example Application: Company Decarbonization Alignment

Of course, although ML may sound good in principle, such techniques can only be beneficial when used in a carefully designed application. To illustrate one such application, we now turn to perhaps the most obvious data need net-zero investors face: predicting a company's decarbonization alignment in the future.

To assess the decarbonization alignment, we need to build a view of the company's carbon emissions at some point decades away—say, in 2050. We can then map the estimated emissions to a specific pathway and thus determine whether the firm belongs in a net-zero portfolio.

As we will show, predicting emissions will indeed involve ML and, in particular, NLP. Although these techniques will be a critical component of the resulting measure, even the most advanced ML cannot get there on its own. We need to provide additional structure and creative solutions for such tools to lead to actionable investment insights.

Structure of the Forecast

To start, we express emissions in tons as a product of the firm's expected sales and its carbon intensity:¹

$$E(\text{Emissions in tons}_{Firm\ j}^{2050}) = E(\text{Sales}_{Firm\ j}^{2050}) \times E(\text{Intensity}_{Firm\ j}^{2050}). \quad (1)$$

We rely on this identity because we believe it is more straightforward to predict these individual components than emissions in tons directly. For example, if we were to predict a company's emissions in the near future (say, in 2027 instead of 2050), we could directly use sell-side sales forecasts for the first term in the product of Equation 1. Sell-side analyst forecasts, reported in such databases as I/B/E/S, are informed predictions based on market trends, economic conditions, and company performance. For the second term of the product, expected carbon intensity in 2027, we could perhaps assume that the firm's intensity will be unchanged over such a short period of time and simply use a historical number.

It is more complicated to arrive at a forecast in 2050. For example, sell-side analyst forecasts are available for only up to five years into the future. We need to find a way to extend such forecasts for another few decades. One option is

¹Technically, the equation is an approximation: The expected value of a product does not generally equal the product of the expectations. As mentioned previously, practical solutions may require some compromises and necessary approximations.

to use solutions proposed in academic literature, such as a three-stage residual income model inspired by Gebhardt, Lee, and Swaminathan (2001):

- The first stage of the model integrates I/B/E/S sell-side analyst forecasts over the first five fiscal years (from FY0 to FY5).
- The second stage assumes that sales forecasts mean revert to a peer-group median between FY5 and FY10.
- The third stage assumes sales reach a long-run equilibrium after FY10.

Next, we need to forecast carbon intensity. Unfortunately, unlike with sales, we do not have as much guidance from academic literature on how a firm's intensity may evolve over time. We need to resort to some simplifying assumptions:

- We begin with the presumption that a company's carbon intensity will remain unchanged from its reported year-end value.
- If a company has announced a decarbonization target, however, this assumption is superseded by the target value. Since decarbonization targets are published by companies on an inconsistent basis, with differing baselines and target dates, we standardize targets and compute the expected decarbonization by the target year.


Of course, some companies with no pledges today may still pledge a decarbonization commitment at some point in the future, and some firms may change their carbon intensity over time even absent such commitments. Later, we will show how we update the distribution of intensity forecasts over time as such new data arrive.

After we forecast both sales and carbon intensity, we can return to Equation 1 and multiply the forecasts to arrive at a distribution of carbon emission forecasts across companies.

How Realistic Are Companies' Decarbonization Commitments?

Relying on a company's stated decarbonization target implicitly assumes that a company will follow through on its commitment. However, taking a commitment at face value and using it directly in our intensity forecast is probably overly optimistic. Thus, we refine this assumption and construct a proxy to assess the credibility of a company's decarbonization commitment. To do so, we will turn to ML and NLP. Specifically, at the cost of introducing some technical jargon, we fine-tune a large language model (LLM) using a supervised learning technique that teaches the model to interpret climate disclosures. Embeddings condense a huge volume of textual data within a high-dimensional vector space to encode better semantic and syntactic meaning. For instance, such phrases as "net-zero goals" and "Paris alignment" will be represented closer together in vector space than more vague terms such as "ambitions" and "pledges" will be. We illustrate this concept in **Exhibit 1** using

Exhibit 1. Mapping Company Disclosures to Climate Categories



Notes: This exhibit uses t-SNE to show a two-dimensional projection of embeddings for words and phrases. Words with similar meanings are clustered together.

t-distributed stochastic neighbor embedding (*t*-SNE), a dimensionality reduction technique designed to visualize high-dimensional data by giving each word a location within a two-dimensional map (van der Maaten and Hinton 2008). Exhibit 1 illustrates how the various words found in textual documents map to climate categories, clustering around such concepts as “emissions,” “energy transition,” or “decarbonization plans.”

The LLM detects mentions of decarbonization plans in company documents. Examples include earnings call transcripts, corporate sustainability reports, and regulatory filings. The output of the LLM is a probabilistic classification that assesses the credibility of a company's decarbonization plans based on perceived alignment to the Task Force on Climate-related Financial Disclosures (TCFD) and Science Based Targets initiative (SBTi) frameworks. We refer to this as the LLM score. Intuitively, we find that companies that publish numeric information, including dates, baselines, and targets, are typically scored higher by the LLM and deemed more likely to follow through on their decarbonization commitments. In effect, the score seeks to proxy the management quality of a company through management's ability to address sustainability risks and opportunities. We illustrate this in **Exhibit 2** with example sentences for two companies.

Exhibit 2. LLM Classification for Two Hypothetical Companies

	Company A	Company B
Country	Australia	Australia
Industry	Construction materials	Construction materials
Climate target	"We're targeting to reduce our absolute Scope 1 and 2 emissions by 46% and to reduce our relevant Scope 3 emissions by 22% . . . by FY 2030."	"40% reduction in Scope 1 and Scope 2 greenhouse gas intensity by 2030."
Evidence	"Our medium-term decarbonization opportunities, which we're maturing, include optimizing our supply chain logistics and low-carbon and no-carbon alternative fuel options."	"We're pretty optimistic we're going to be able to continue to drive greater efficiencies in our operating plans."
LLM classification	High certainty of meeting the target	Low certainty of meeting the target

Such examples highlight that seemingly similar corporate pledges, such as 40% reduction in emissions, may lead to very different overall assessments based on a careful analysis of additional company disclosures. Of course, while we advocate using NLP for such analyses, we urge investors to include spot checks and "sniff tests," perhaps similar to the previous examples, where human analysts verify model output. We believe scalable, systematic processes can yield a lot of value for investors—but they should not be used sight unseen and fly purely on autopilot.

To demonstrate the benefits of using unstructured data, we perform a statistical analysis to evaluate whether the LLM score is positively correlated with independent company assessments conducted by climate experts using data from the Transition Pathway Initiative (TPI). The TPI's data underpin the Climate Action 100+ Net Zero Company Benchmark and assess performance on emission reductions, governance, and disclosure on and implementation of net-zero transition plans. As of March 2024, 151 institutional investors globally pledged their support to the TPI, representing approximately \$60 trillion in assets under management. TPI scores are available for only a small fraction of investible companies, limiting their usefulness as a comprehensive portfolio solution. Still, we believe such data could go a long way to validate and thus increase investors' comfort with other types of climate data, such as the LLM score.

Specifically, we examine whether the LLM score helps explain the TPI Management Quality score. The TPI Management Quality score consists of six levels. Levels 0 and 1 refer to companies that do not develop basic capacity to address climate risks and opportunities, lack disclosures on their carbon practices and performance, and do not integrate climate considerations into operational decision making. By contrast, Levels 4 and 5 refer to companies that develop a strategic and holistic understanding of climate risks and opportunities, with detailed and actionable transition plans that align business practices and capital expenditure decisions to their decarbonization goals.²

²See Dietz, Bienkowska, Jahn, Hastreiter, Komar, Scheer, and Sullivan (2021).

The regression specification includes three sets of variables. The first set comprises company fundamentals, including the percentage of revenue derived from the extraction of conventional and unconventional oil and gas, fossil-fuel reserves, thermal coal, and alternative energy. We further include a company's latest reported Scope 1 and 2 carbon emissions and carbon intensity. Taken together, these fundamental metrics seek to proxy exposure to carbon-related risks and opportunities as reported in a company's financial statements. The second set of variables includes a company's announced decarbonization targets. We include indicator variables equal to 1 if a company has publicly disclosed a target, if it has announced a science-based target, and if the target is approved by the SBTi and equal to zero otherwise. The final set of variables captures the comprehensiveness of a company's decarbonization plans. We include the LLM score and MSCI's Carbon Emissions Management Score.³ The latter score integrates an assessment of how aggressive any decarbonization target is, whether a company has a track record of achieving its targets, how aggressively the company has sought to use cleaner sources of energy, and carbon capture and storage/sequestration of its operational emissions. The results of the logistic regressions as of June 2024 are provided in **Exhibit 3**.

Exhibit 3. LLM Score Helps Capture Differences in TPI Management Quality Scores across Firms

	TPI Management Quality Laggards		TPI Management Quality Leaders	
	(1)	(2)	(3)	(4)
LLM		-1.086 (-4.150)***		0.426 (2.584)***
Carbon emissions	-0.161 (-1.200)	-0.094 (-0.645)	0.525 (2.752)***	0.489 (2.3103)**
Carbon intensity	-0.0019 (-0.016)	-0.067 (-0.521)	-0.390 (-4.011)***	-0.351 (-3.534)***
% Conventional oil & gas	-0.923 (-1.737)*	-0.713 (-1.392)	0.087 (-0.573)	0.043 (0.273)
% Unconventional oil & gas	-1.356 (-0.973)	-2.5794 (-0.779)	0.046 (-0.271)	0.069 (0.408)
% Thermal coal	0.203 (1.786)*	0.310 (1.733)*	-0.104 (-0.809)	-0.105 (-0.807)
% Alternative energy	0.401 (1.368)	0.513 (1.512)	0.039 (0.269)	0.089 (0.612)
Dummy ^{Carbon Underground 200}	1.233 (1.895)*	1.652 (1.931)*	0.544 (1.262)	0.468 (1.081)

(continued)

³Index source: MSCI. Copyright MSCI 2024. All rights reserved. Unpublished. Proprietary to MSCI.

Exhibit 3. LLM Score Helps Capture Differences in TPI Management Quality Scores across Firms (continued)

	TPI Management Quality Laggards		TPI Management Quality Leaders	
	(1)	(2)	(3)	(4)
Dummy ^{Numeric target}	-0.594 (-1.095)	-0.573 (-1.047)	2.026 (3.081)***	1.985 (2.990)***
Dummy ^{SBTi approved}	-1.417 (-2.117)**	-1.052 (-2.061)**	1.699 (6.072)***	1.65 (5.848)***
Dummy ^{SBTi commitment}	-1.332 (-1.419)	-1.019 (-1.268)	0.2909 (0.854)	0.249 (0.727)
MSCI Carbon Management	-0.432 (-1.621)	-0.366 (-1.581)	0.129 (1.252)	0.089 (0.841)
Sector fixed effects	Yes	Yes	Yes	Yes
Pseudo R^2	0.274	0.332	0.235	0.345
N	528	528	528	528

Notes: This exhibit reports the results of a logistic regression in which the dependent variable is the TPI Management Quality score. The dependent variable in columns 1 and 2 is an indicator variable equal to 1 if a company has a TPI Management Quality score of 0 or 1 and is equal to 0 otherwise. The dependent variable in columns 3 and 4 is an indicator variable equal to 1 if a company has a TPI score of 4 or 5 and is equal to 0 otherwise. An intercept term is included in the regression, although it is not displayed given space limitations. "LLM" represents the output of a probabilistic text classification derived from an LLM that scores the perceived credibility of a company's decarbonization plans. "Carbon emissions" represents the cross-sectional Z-score of a company's Scope 1 and 2 carbon emissions. "Carbon intensity" is the region- and industry-relative Z-score of a company's carbon intensity. "% Conventional oil & gas" is the percentage of revenue a company derives from conventional oil and gas. "% Unconventional oil & gas" is the percentage of revenue a company derives from unconventional oil and gas. "% Thermal coal" is the percentage of revenue derived from the mining of thermal coal, including lignite, bituminous, anthracite, and steam coal. "% Alternative energy" is the percentage of revenue derived from renewable energy sources. Dummy^{Carbon Underground 200} is an indicator equal to 1 if a company is on the Carbon Underground 200 list; the list identifies the top 100 coal and the top 100 oil and gas public companies ranked by the potential carbon emission content of their reported reserves. Dummy^{Numeric target} is an indicator variable equal to 1 if a company has disclosed its target percentage reduction in its carbon emissions and is equal to 0 otherwise. Dummy^{SBTi approved} is an indicator variable equal to 1 if a company has had its target approved by the SBTi. Dummy^{SBTi commitment} is an indicator variable equal to 1 if a company has committed to setting science-based targets. MSCI Carbon Management is MSCI's assessment of how aggressive a decarbonization target is, whether a company has a track record of achieving its targets, and how aggressively it has sought to use cleaner sources of energy. For each variable, we report corresponding z-values, where ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. The sample period is June 2024.

Source: Carbon and revenue data are sourced from MSCI.

Index source: MSCI. Copyright MSCI 2024. All rights reserved. Unpublished. Proprietary to MSCI.

Columns 1 and 2 in Exhibit 3 provide the results of a logistic regression where the dependent variable is an indicator variable equal to 1 if a company has a TPI Management Quality score of 0 or 1. We observe that climate laggards are more likely to derive revenue from thermal coal and appear on the Carbon Underground 200 list, consistent with the view that such companies may hold stranded assets. Column 2 includes the LLM score and shows a highly significant, negative coefficient, which means the lower the LLM score, the more likely the company is to be considered a climate laggard. In columns 3 and 4, the dependent variable is changed to an indicator variable equal to 1 if a company has a TPI Management Quality score of 4 or 5. Companies are more likely to be categorized by the TPI as a climate leader if they have lower carbon intensities than peers and have a target approved by the SBTi. Column 4 shows

that the LLM score is highly statistically significant, showing that the higher the LLM score, the more likely the firm is to be considered a climate leader.

A statistically significant relationship between the LLM score and the TPI Management Quality score points to the ability of an LLM to assess the credibility of companies' decarbonization plans, thereby codifying the perceptions of climate experts. Taken together, the regression results are consistent with the idea that the LLM score captures additional information beyond the company fundamental data and numeric disclosure targets.

Importantly, the LLM score is not meant to replace TPI measures. These measures are noisy themselves and may not reflect all relevant information about a given issuer. They do, however, capture some relevant information. Exhibit 3 suggests that the LLM score also incorporates such information, as reflected in both the statistical significance of the estimates and in the increase in the R^2 when we incorporate LLM: The R^2 for the laggards increases by about 20% of its level, and that of the leaders increases by about 47% of its level.

Bayesian Approach: Updating the Distribution over Time

With any data analysis, we must recognize that the underlying companies and their environment change over time and adjust our forecasts accordingly. Perhaps the most straightforward approach would be to recompute the forecasts, as explained previously, every time the underlying data changes. This approach is substandard, if only because data are noisy and any given snapshot may lead to erroneous inferences about a given company. This may be because of both outright mistakes in the data and potential greenwashing or other strategic manipulation by the company—or even because of transient economy-wide shocks. For example, corporate emissions were depressed in 2020 because of COVID-19, but it would have been a mistake to assume the 2020 reported figures are the optimal predictor of future emissions. Indeed, emissions reverted to the long-term historical average soon thereafter.

We can do better by gradually updating our forecasts as more data become available. To formalize this intuition, we use a Bayesian approach, which allows us not only to effectively update our forecasts over time but also to model the inherent uncertainty associated with companies' decarbonization trajectories. In general, Bayesian inference offers a framework to incorporate prior knowledge, such as historical data and expert opinions, with new evidence. These inputs may be combined to provide a probabilistic assessment of a company's decarbonization trajectory. One of the major advantages of Bayesian inference is that it offers not just point estimates but also confidence intervals for parameters. This probabilistic aspect may enable investors to assess risks more comprehensively.

There are three essential components underlying Bayesian statistics (for an overview, see van de Schoot, Kaplan, Denissen, Asendorpf, Neyer, and van Aken 2014). The first is the background knowledge on the parameters of the model—that is, all knowledge captured by the prior distribution, such as a normal

distribution, before seeing the data.⁴ The choice of prior reflects how much information we have before data collection and how accurate we believe the information to be. The variance of the prior distribution reflects our uncertainty about the population parameter. A smaller variance implies greater confidence that the prior mean reflects the population mean. In other words, the prior distribution represents the current state of knowledge or current description of uncertainty about the model parameters prior to data being observed. The second key component is information about the data. It is the observed evidence (i.e., the sample distribution) expressed in terms of the likelihood function of the data given the parameters. The third component is based on combining the first two components, known as the posterior distribution, and reflects one's updated knowledge, balancing prior knowledge with observed data. We describe these three components of the model in turn.

Prior Distribution

At the outset of the analysis, it is perhaps easiest to start with a diffuse (uninformed) prior and then adjust it given historical information. In other words, the analyst would use such historical information to compute the emission forecasts as described earlier without imposing any first-principles restriction on the outcome. For analytical ease, we chose to model the log ratio of a company's 2030 emissions to its latest annual emissions with a normal prior distribution. These priors can approximate the diffuse case when we assume they have a large variance. Thus, we allow for a wide range of possible outcomes before we see the data.

Sample Distribution

The sample distribution is derived from the company's realized carbon emission trajectory. As companies report their actual emissions over time, these data are used to construct the empirical distribution of observed emissions and capture a company's operational changes, market conditions, and policy impacts. We assume that the log ratio of a company's realized emissions to its latest annual emissions also follows a normal distribution. We use a statistical time-series ARIMA (autoregressive integrated moving average) model to compute a forecast for each company's carbon emission trajectory to 2030 and obtain the mean forecast and standard error.⁵


Posterior Distribution

The prior and sample distributions are combined to form the posterior distribution, providing an updated belief on a company's decarbonization alignment.

⁴We model the log ratio of a company's 2030 emissions to its latest annual emissions with a normal distribution, which is equivalent to modeling the ratio with a log normal distribution. This distribution can accommodate all possible values of a company's 2030 emissions.

⁵A company's carbon emission trajectory is modeled as the log ratio of a company's future annual emissions to its latest annual emissions.

Exhibit 4. Bayesian Updating of Carbon Emission Forecasts



When a parameter can be modeled by a prior normal distribution, Bayesian statistics show that the sample dataset from the same process can be used to update the prior to obtain a posterior normal distribution. The weighting of the two distributions is determined by their relative variances, reflecting the confidence in the prior information versus the realized data.

Exhibit 4 shows a schematic depicting the overall estimation process.


Results: Expected Decarbonization in 2030

In this section, we outline the merits of the Bayesian framework for portfolio climate analytics. In particular, we show how investors can quantify portfolio alignment to the socioeconomic pathways of the Intergovernmental Panel on Climate Change (IPCC). The five shared socioeconomic pathways (SSPs), described in the IPCC's (2021) "Sixth Assessment Report," outline representations of an uncertain future. The pathways range from a "Taking the Green Road" scenario, in which CO₂ emissions decline drastically to carbon neutrality by 2050 and are negative in the second half of the century (SSP1-1.9), to a fossil-fueled development ("Taking the Highway") scenario, in which CO₂ emissions continue to rise sharply to twice current levels in 2050 and more than three times current levels in 2100 (SSP5-8.5).

Exhibit 5 illustrates the resulting posterior probability distributions for three major benchmarks: the MSCI All Country World Index (ACWI), MSCI ACWI Climate Transition, and MSCI ACWI Paris-Aligned.⁶ For each benchmark, we plot the distribution of the forecasted change in emissions. The vertical lines represent the decarbonization rates implied by each IPCC SSP. The SSP1-1.9 line implies the greatest reduction in carbon emissions, and the SSP5-8.5 line implies an increase in carbon emissions.

Exhibit 5 is based on an idea similar to the well-known MSCI Implied Temperature Rise metric. The key difference is that Exhibit 5 also gives investors information about the likely range of outcomes and allows them to quantify the risk that the portfolio might miss its climate objectives, rather than merely providing a point forecast. This is critical given the inherent uncertainty

Exhibit 5. Probability Distribution of Expected Decarbonization by 2030



Notes: The exhibit displays the posterior distribution for the MSCI ACWI, ACWI Climate Transition, and ACWI Paris-Aligned indexes as of June 2024. The vertical lines indicate the decarbonization rates under each IPCC SSP. SSP1-1.9 is the IPCC's most optimistic scenario, in which global CO₂ emissions are cut to net zero around 2050, with warming reaching 1.5°C and then stabilizing to around 1.4°C by the end of the century. SSP1-2.6 is the next-best scenario, in which global CO₂ emissions are cut severely, reaching net zero after 2050. Temperatures stabilize at around 1.8°C higher by the end of the century. SSP2-4.5 is the "middle-of-the-road" scenario; CO₂ emissions start to fall mid-century but do not reach net zero by 2100, and temperatures rise 2.7°C by the end of the century. Under the SSP3-7.0 scenario, CO₂ emissions approximately double from current levels by 2100, with average temperatures rising by 3.6°C by the end of the century. The SSP5-8.5 scenario is a future to avoid at all costs: Current CO₂ emissions levels double by 2050 with economic growth fueled by exploiting fossil fuels. By 2100, the average global temperature is 4.4°C higher. The exhibit was created using the methodology described in this chapter and then bootstrapping by simulating individual securities' decarbonization paths from each security's posterior distribution.

Index source: MSCI. Copyright MSCI 2024. All rights reserved. Unpublished. Proprietary to MSCI.

⁶Index source: MSCI. Copyright MSCI 2024. All rights reserved. Unpublished. Proprietary to MSCI.

associated with climate analysis. From a top-down perspective, this includes uncertainty regarding the future direction of government and regulatory policies, technological innovation, and how consumer preferences may evolve. From a bottom-up perspective, our approach considers ongoing uncertainty associated with companies' decarbonization trajectories and willingness to follow through on their plans.

As an example application of this framework, by integrating the area under the probability distribution, we can infer alignment to a given SSP scenario. For example, Exhibit 5 shows that the core benchmark (MSCI ACWI)⁷ clearly misses the mark for net-zero alignment (SSP1-1.9). The area in the left tail of the distribution up to the SSP1-1.9 vertical threshold indicates the probability that the benchmark is net-zero aligned, which is about 0.1. This suggests that this popular benchmark is highly likely to miss the climate goal of net-zero investors because the individual portfolio companies are unlikely to decarbonize promptly enough for the index to be net-zero aligned. It is more likely that the index will be aligned with the SSP2-4.5, "middle-of-the-road" scenario, but even here, we see only even odds of achieving that outcome (Exhibit 5 implies a probability of 0.46). In contrast, the two climate-oriented versions of the index, Climate Transition and especially Paris-Aligned, have a much more attractive net-zero alignment. The probability of meeting SSP1-1.9 is 0.37 for the former and 0.61 for the latter, with obviously an even higher probability of aligning with at least the SSP2-4.5 scenario (0.67 for Climate Transition and 0.81 for Paris-Aligned).

Conclusion

Climate investing and, in particular, net-zero investing are a complex but also fascinating challenge for investors. Unlike with historical carbon emissions, no company-reported, broadly comparable measures exist that could capture a firm's net-zero alignment decades from now. Instead, companies are likely to report different information, frequently in a narrative form. To process such information and to inform their broader portfolios, investors have little choice but to use ML and, in particular, NLP.

Moreover, there is no single "silver bullet" source of net-zero data, so investors must be prepared to combine different datasets and various statistical techniques in their net-zero strategies. And even then, investors will face substantial uncertainty around the estimates they produce. We believe portfolio applications should reflect this uncertainty and rely not just on our best estimate (best guess) but also on the range of possible outcomes around it—for example, through Bayesian updating. Our realistic case study showcases NLP and also highlights other important components of a holistic net-zero solution.

We conclude that while climate investing may be both art and science, there is already plenty of science investors should rely on when building net-zero portfolios.

⁷Index source: MSCI. Copyright MSCI 2024. All rights reserved. Unpublished. Proprietary to MSCI.

References

Bixby, S., A. Brixton, and L. Pomorski. 2022. "Looking Forward with Historical Carbon Data." In *Climate Investing: New Strategies and Implementation Challenges*, edited by E. Jurczenko, 29–47. Wiley and ISTE. [doi:10.1002/9781394192373](https://doi.org/10.1002/9781394192373).

Cohen, L., U. G. Gurun, and Q. H. Nguyen. 2020. "The ESG-Innovation Disconnect: Evidence from Green Patenting." NBER Working Paper 27990 (October). [doi:10.3386/w27990](https://doi.org/10.3386/w27990).

Dietz, S., B. Bienkowska, V. Jahn, N. Hastreiter, V. Komar, A. Scheer, and R. Sullivan. 2021. "TPI's Methodology Report: Management Quality and Carbon Performance, Version 4.0." Transition Pathway Initiative (November). www.transitionpathwayinitiative.org/publications/uploads/2021-methodology-report-management-quality-and-carbon-performance-version-4-0.

Gebhardt, W. R., C. M. C. Lee, and B. Swaminathan. 2001. "Toward an Implied Cost of Capital." *Journal of Accounting Research* 39 (1): 135–76. [doi:10.1111/1475-679X.00007](https://doi.org/10.1111/1475-679X.00007).

Heal, G., and A. Millner. 2014. "Uncertainty and Decision Making in Climate Change Economics." *Review of Environmental Economics and Policy* 8 (1): 120–37. [doi:10.1093/reep/ret023](https://doi.org/10.1093/reep/ret023).

IPCC. 2021. "Sixth Assessment Report." www.ipcc.ch/assessment-report/ar6/.

Pindyck, R. S. 2017. "The Use and Misuse of Models for Climate Policy." *Review of Environmental Economics and Policy* 11 (1): 100–14. [doi:10.1093/reep/rew012](https://doi.org/10.1093/reep/rew012).

Robinson-Tillett, S. 2022. "MSCI Makes Implied Temperature Scores Less Precise after Pushback." *Responsible Investor* (1 August). www.responsible-investor.com/msci-makes-implied-temperature-scores-less-precise-after-pushback/.

Schneider, S. H., and K. Kuntz-Duriseti. 2002. "Uncertainty and Climate Change Policy." In *Climate Change Policy: A Survey*, edited by S. H. Schneider, A. Rosencranz, and J. O. Niles, 53–88. Washington, DC: Island Press.

Task Force on Climate-related Financial Disclosures. 2020. "Forward-Looking Financial Sector Metrics: Consultation" (October). www.fsb.org/uploads/P291020-4.pdf.

Thiele, L. P. 2020. "Integrating Political and Technological Uncertainty into Robust Climate Policy." *Climatic Change* 163 (1): 521–38. [doi:10.1007/s10584-020-02853-9](https://doi.org/10.1007/s10584-020-02853-9).

UN Environment Programme. 2024. "To Serve Net-Zero Clients, Asset Managers Need a Systematic Approach to Climate." Press release (6 February). [www.unepfi.org/industries/to-serve-netzero-clients-asset-managers-need-a-systematic-approach-to-climate/](http://www.unepfi.org/industries/to-serve-net-zero-clients-asset-managers-need-a-systematic-approach-to-climate/).

van de Schoot, R., D. Kaplan, J. Denissen, J. B. Asendorpf, F. J. Neyer, and M. A. G. van Aken. 2014. "A Gentle Introduction to Bayesian Analysis: Applications to Developmental Research." *Child Development* 85 (3): 842–60. www.ncbi.nlm.nih.gov/pmc/articles/PMC4158865/.

van der Maaten, L. J. P., and G. E. Hinton. 2008. "Visualizing Data Using t-SNE." *Journal of Machine Learning Research* 9 (November): 2579–605.