



Factor-Mimicking Portfolios for Climate Risk

Gianluca De Nard, Robert F. Engle & Bryan Kelly

To cite this article: Gianluca De Nard, Robert F. Engle & Bryan Kelly (2024) Factor-Mimicking Portfolios for Climate Risk, Financial Analysts Journal, 80:3, 37-58, DOI: [10.1080/0015198X.2024.2332164](https://doi.org/10.1080/0015198X.2024.2332164)

To link to this article: <https://doi.org/10.1080/0015198X.2024.2332164>



© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 10 May 2024.



Submit your article to this journal [↗](#)



Article views: 9023



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 8 View citing articles [↗](#)

Factor-Mimicking Portfolios for Climate Risk

Gianluca De Nard , Robert F. Engle , and Bryan Kelly

Gianluca De Nard is a research fellow at NYU Stern, New York, USA; senior researcher at University of Zurich, Zurich, Switzerland; and Head of Quantitative Research at OLZ AG, Zurich, Switzerland. Robert F. Engle is Professor Emeritus of Finance, and Co-Director of the Volatility and Risk Institute at NYU Stern, New York, New York, USA, and Research Associate at NBER, Cambridge, Massachusetts, USA. Bryan Kelly is Professor of Finance at Yale University, New Haven, Connecticut, USA; Head of Machine Learning at AQR Capital Management, Greenwich, Connecticut, US; and Research Associate at NBER, Cambridge, Massachusetts, USA. Send correspondence to Robert F. Engle at rengle@stern.nyu.edu.

We propose and implement a procedure to optimally hedge climate change risk. First, we construct climate risk indices through textual analysis of newspapers. Second, we present a new approach to compute factor-mimicking portfolios to build climate risk hedge portfolios. The new mimicking portfolio approach is much more efficient than traditional sorting or maximum correlation approaches by taking into account new methodologies of estimating large-dimensional covariance matrices in short samples. In an extensive empirical out-of-sample performance test, we demonstrate the superior all-around performance delivering markedly higher and statistically significant alphas and betas with the climate risk indices.

Keywords: climate change; factor model; portfolio selection; sustainable portfolio

Disclosure: No potential conflict of interest was reported by the author(s).

PL Credits: 2.0

Introduction

Climate change poses a great risk to the global economy and ultimately to the human race. Investors recognizing this risk are demanding portfolios that recognize and optimize with respect to this risk. The financial services industry has responded with a plethora of new offerings. In April 2022, Bjoy et al. (2022) identified 860 global mutual funds and exchange-traded funds (ETFs) with a climate mandate. These funds have assets under management of \$407 billion—double the value from a year earlier. As this segment of investment products expands, it is natural to study the effectiveness of these offerings and the theory underlying their management. Of course, objectives may differ for different asset managers and investors. Morningstar now divides this universe into five buckets: low carbon, climate conscious, green bonds, climate solutions, and clean energy/tech. In most cases, a simple feature is a common component of all these categories. The portfolio over-weights stocks that are or will be prepared for climate change and underweights those that are not. The portfolio should therefore outperform the broad market in a bad climate outcome and can be considered a climate hedge portfolio. See, for example, Jurczenko and Teiletche (2022); Engle et al. (2020); Andersson, Bolton, and Samama (2016); Engle (2014); and others following the framework of Merton (1987). The actual management of these funds recognizes that there are competing objectives and many theories about how climate change will evolve and impact different firms and sectors.

We thank Heebum Lee, Olivier Ledit, Michael Wolf, and the NSF, participants at the Volatility and Risk Institute Conference at NYU Stern, for helpful comments. This work has been nominated for the Swiss Risk Award 2021.

The Norwegian Finance Initiative and the Global Risk Institute, CERES, and the UZH FAN, provided generous grant support.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Investors around the world desire products that allow them to hedge against the realizations of climate risk. Because of the long-run and nondiversifiable nature of climate risk, standard futures or insurance contracts in which one party promises to pay the other in the event of a climate disaster are difficult to implement. Indeed, no counterparty could credibly guarantee to pay claims during a climate disaster event that might materialize in many decades, in part because a bad outcome would mandate all contracts to be paid at the same time. Individual investors are therefore largely constrained to self-insure against climate risk.

In this paper, we will propose and implement an approach to forming an optimal hedge portfolio of publicly available sustainable funds. We will call this a climate-efficient factor-mimicking portfolio (CEP). Roughly, this portfolio will have both a minimum variance objective and maximum greenness criterion. The risk of these funds will be estimated by a new shrinkage estimator from De Nard (2022) that reduces the noise in large covariance matrices when samples are short. This is an important property, as the number of sustainable funds is increasing rapidly, but their track records are usually (too) short. The greenness of any fund and of the CEP will be measured by its response to new information on climate change after controlling for other standard risk factors. The balance between these two criteria will be based on out-of-sample performance of an extensive real-life empirical analysis for various climate (change) news indices.

Due to the rise in investors' awareness of the economic and financial risks of climate change and the increased demand for financial instruments to hedge these risks, there have recently been great strides in academic research on operationalizing methods for climate hedging.

Our approach is related to the "mimicking portfolio" approach of Lamont (2001), where climate risk series are projected onto a set of portfolio returns using time-series information. Since it does not take an a priori view on which assets gain or lose when climate shocks occur, it needs to learn this from assets' return performance during past climate risk realizations. The innovation of our approach is to rely on efficient long-only portfolios based on sustainable and climate-related ETFs, reducing not only exposure to climate change risks but also financial risks (portfolio standard deviation) and superior (risk-adjusted) returns.

One drawback of these mimicking portfolio approaches is that they require long time series. That is why alternative approaches focus more on long and short positions based on economic reasoning, for example, sell

coal companies and buy clean energy companies or, as suggested by Engle et al. (2020), buy (sell) companies with high (low) ESG scores. The difficulty of this approach is to use the right priors about investors' perceptions of each industry's exposures to climate risk. It only hedges climate risk if the underlying economic intuition is aligned with that of the "average investor."

A new, promising quantity-based approach from Alekseev et al. (2022) shows improved hedging performance by exploiting information on how mutual fund managers trade in response to idiosyncratic changes in their climate risk beliefs.¹ The quantity-based approach predicts the investors' capital reallocation around such idiosyncratic (climate change) belief shocks that could move equilibrium prices. Alekseev et al. (2022) show that a portfolio that is long stocks that investors tend to buy after experiencing negative idiosyncratic climate belief shocks and short stocks that investors tend to sell, appreciate in value in periods with negative aggregate climate news shocks. An advantage of the quantity-based approach is that it learns from rich cross-sectional trading responses rather than time-series price information.

Acknowledging the discussed alternative approaches, we focus on advances for the more traditional mimicking portfolio approach. Our contribution is to introduce an investible long-only climate-efficient factor-mimicking portfolio that is simple to implement, interpret, and customize. Contrary to the discussed literature, we focus on sustainable and climate-related ETFs without any short sales, as we deem our investment universe as "green." Additionally, we use (hedging) optimized portfolios rather than sorting portfolios where further (sustainability) constraints and portfolio preferences can be easily included. To tackle the curse of dimensionality (rapidly increasing number of sustainable funds, but with short track record) we use shrinkage estimators for more robust portfolio optimization. Finally, an investor needs to select only its investment universe and climate hedge target. Additional expensive and difficult-to-obtain data are not necessary.

Framework

Climate risk is the risk that changes in climate will negatively impact cash flows in the future. To some extent, the market has already priced this risk; where this risk is great, assets are priced below their intrinsic riskless value. Equivalently, expected returns must be high to compensate investors for holding assets heavily exposed to climate risk. This feature has been documented empirically by Bolton and Kacperczyk (2021).

This climate risk premium must be based on the market's view of the severity, likelihood, and urgency of climate change and the exposure of this asset to such risks. It is natural to formulate this as a climate risk factor, which is a weighted average of financial assets where the weights are based on the climate exposure of each asset. Investors unconcerned with climate risk would hold this portfolio for its higher expected return, and investors who want to hedge climate risk would short this portfolio and expect a negative risk premium and reduced returns in exchange for reduced risk. New information will induce repricing of assets. If the new information indicates that climate risk is increasing, then asset values will fall roughly in proportion to the asset exposures; therefore, the climate risk portfolio will underperform and its short counterpart—the climate hedge portfolio—will outperform.

To examine this empirically, it is necessary to have a measure of the news. This could be text-based measures from news channels or physical news such as temperature extremes. Whenever there is news that climate risks are greater than the market has priced, the climate hedge factor should appreciate in value. These changes will appear as an alpha when the climate risk factor is missing. Over time, as the market reprices with new scientific and public analyses, the market pricing should approach the scientific view and if this continues to deteriorate, the hedge portfolios should continue to appreciate. Thus, holding a climate hedge portfolio should yield negative returns in the absence of news on climate risk, but when there is news, the yield should be positive. From this perspective, the news simply accumulates. Either there is no news or there is some news and the cumulative excess returns on the climate hedge portfolio will depend upon the cumulative news. Thus, the alpha at any point in time should be associated with the news at that point in time, possibly with some lags.

A closely related theoretical analysis is presented by Pástor, Stambaugh, and Taylor (2021), wherein the driving force is differences in investor preferences for green stocks rather than differences in risk. The results instead are lower prices for green stocks and higher returns. The differences could be due to the timing of climate news as well as the characterization of green and brown stocks.

Careful analysis of the implications of climate change suggests in fact that at least two disparate factors can be defined. One is the physical risk of climate change from rising temperatures, increased storms, droughts, and wildfires and the range of physical hazards that could result. The second is the risk inherent

in the transition to a low-carbon emission economy that is a natural response to the risks of climate change. This risk is based at least in part on policy and is therefore sometimes called regulation risk, although it is more commonly called transition risk. While both factors would be expected to have positive risk premiums, the factors are often negatively correlated. If important regulations are proposed, transition risk rises and physical risk falls. However, a predicted severe hurricane season would increase both risks and a technological breakthrough in carbon sequestration would decrease both risks.

Measuring Climate Change News

A central feature of this theory is that climate hedge portfolios should appreciate when there is news that climate risk is increasing. There are many ways this news can be measured, and we have used several in this analysis, following the literature by Engle et al. (2020), Ardia et al. (2022), and Faccini, Matin, and Skiadopoulos (2021), among others. These measures have obvious shortcomings and yet are simple and easily understood. Since we are asking what news will move the market perception of climate change risk, there will always be room for improvement and it is important to understand the strengths and limitations of such measures. We feel that the series we discuss here are good compromises.

New York Times. Two measures are based on textual analysis of *The New York Times* (NYT). In Engle et al. (2020), the same methodology was applied to *The Wall Street Journal* (WSJ). The procedure follows Gentzkow, Bryan Kelly, and Taddy (2019). Each article is examined for its word frequency, and these frequencies are compared with the same body of authoritative texts on the subject of climate change. After stripping out small words (stop words) and combining words with the same root but different stems (stemming), we have a unique set of one- and two-word combinations. We count their frequency in the dictionary and in each day of the NYT. For each document in the NYT and in the dictionary, these counts are converted to “term frequency-inverse document frequency” or *tf-idf* scores. For word j , in document i , this expression is:

$$tf_{i,j} - idf_j = \frac{c_{i,j}}{n_i} \times \log\left(\frac{n}{d_j}\right), \quad (3.1)$$

where $c_{i,j}$ is the specific word count in document i , n_i is the total word count in document i , and d_j is the number of documents that contain word j . Hence, a word that is in most documents will have a weight

near zero. Rare words will be emphasized, and if they appear in document j , the $tf-idf$ can be very big. Finally, we construct our daily climate change index as the “cosine similarity” between the $tf-idf$ scores for each daily edition of NYT and the climate dictionary.

A second index is constructed from the same source but uses tags that are supplied by LexisNexis to each article in the daily NYT. There are several thousand tags that are constructed, and each article has one or more tags. These tags are assigned by a proprietary natural language processing procedure that may or may not correspond to the matching of word frequencies. Thus, we can interpret such tags as the result of an algorithm that is constructed for public use of the news series. We measure the proportion of news articles that include either “global warming” or “climate change” as one of the tags. Although both indices are inherently daily, we use a weekly version. The news series are reported for seven days, yet the stock market has only a five-day week. The news series, thus, is treated as constant for a week at the level from the preceding week.

The indices are plotted in Figure 1. The figure shows that the intensity of climate news coverage has steadily increased in the last two decades. In addition, the climate risk indices spike during salient climate events, such as the adoption of global climate

treaties (e.g., Paris Agreement or the Kyoto Protocol), or important global conferences to battle climate change (e.g., the 2009 UN Climate Change Conference in Copenhagen).

Wall Street Journal. Several news indices were constructed from the WSJ. These indices are based upon machine learning algorithms that use keywords to direct the search for clusters of relevant news. These are designed to reflect general climate risk, physical risk, and transition risk.

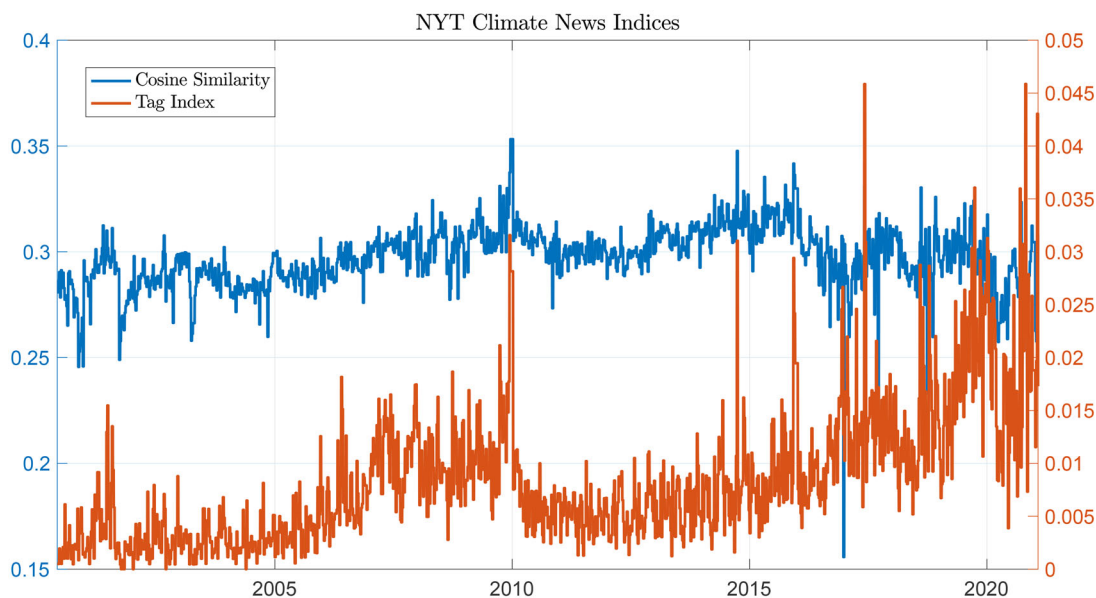
The WSJ indices are plotted in Figure A1. The figure shows also that the intensity of climate news coverage has increased in the last two decades and that the climate risk indices spike during salient climate events or important global conferences. Additionally, we see that the three WSJ indices have overall similar co-movements. In Table A1, we report the (daily) pairwise correlations among all investigated climate news indices. All the indices are positively correlated with each other, and the correlations between the WSJ versions are especially large.

Factor-Mimicking Portfolios

Selection of the Investment and Factor Universe.

We seek portfolios of underlying assets that are efficient climate hedge portfolios. We propose to focus on assets that are sustainable and

Figure 1. The New York Times Climate News Indices from June 2001 to January 2021



climate-related portfolios. For example, a selection by Morningstar of US funds includes mutual funds and ETFs that satisfy at least one of the following criteria: fossil fuel-free, low-carbon, low-environmental risk score, and sustainable-sector funds. We restrict our attention to V-LAB's 177 climate-focused funds.²

There are two general approaches to forming factor-mimicking portfolios of a non-investible factor such as climate news. The first is often called a maximum-correlation portfolio and was introduced by Huberman and Kandel (1987) and Lamont (2001). The news series is regressed on asset returns to find a portfolio with maximum correlation with the news series. This approach was employed in Engle et al. (2020), among others. The second uses a two-step method to find the sensitivity of each asset to the news series and then combines this information through a Fama and MacBeth (1973) or Lehmann and Modest (1988) cross-sectional regression to get the portfolio. These are shown in Jurczenko and Teiletche (2019, 2022) to be special cases of the general asset pricing setting which seeks a minimum-variance portfolio subject to a constraint on its correlation with the news series. Their formulation, however, does not include other investible factors and expects the underlying assets to span the universe of returns in order to form principle components that span the factor space. We provide a simple and useful extension.

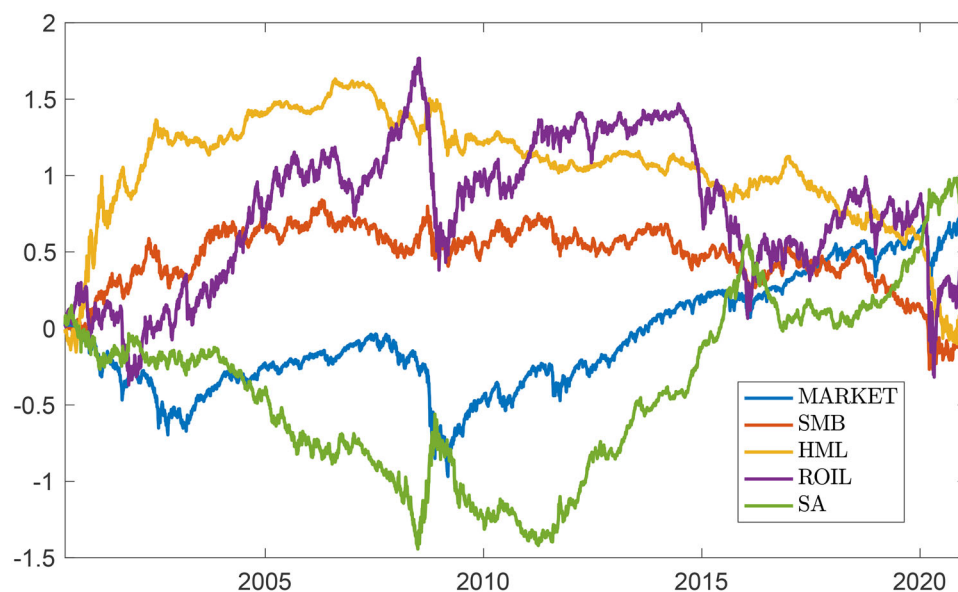
When deriving a climate factor-mimicking portfolio, we find it important—but not mandatory—to control for other well-known factors from the literature. We deem the three Fama–French factors (MARKET, HML, and SMB) as well as a stranded asset (SA) and oil short-term future factor (ROIL) the most useful for our purpose.³ Note that we are interested in *investable* risk factors. Thus, we compute the investable factors as follows:

- $\text{MARKET} := \text{VTI} - r_f$ ⁴
- $\text{HML} := \text{IWD} + \text{IWN} - (\text{IWF} + \text{IWO})$
- $\text{SMB} := \text{IWN} + \text{IWO} - (\text{IWF} + \text{IWD})$
- $\text{SA} := \text{SPY} - 0.7 * \text{KOL} - 0.3 * \text{XLE}$
- $\text{ROIL} := \text{WTI crude oil returns}$

where VTI is the Vanguard Total Stock Market Index Fund ETF; r_f is the risk-free rate; IWD is the Russell 1000 Value ETF; IWF is the Russell 1000 Growth ETF; IWN is the Russell 2000 Value ETF; IWO is the Russell 2000 Growth ETF; SPY is the SPDR S&P 500 ETF Trust; KOL is the VanEck Vectors Coal ETF; and XLE is the Energy Select Sector SPDR Fund. The investable factors are plotted in Figure 2.

The Climate Factor Signal. Basically, we are interested to find climate-related funds that mimic the behavior of an underlying climate factor. Therefore, we include our climate news series detailed in “Measuring Climate Change News.” Then,

Figure 2. Plot of the Cumulative Returns of the Various Investable Factors in Log Scale



we regress each climate-related fund, $i = 1, \dots, 177$, on the factors and the climate (change) news, CC, using a rolling regression with (one year of) daily data:

$$r_{i,t} = \hat{\beta}_i CC_t + \hat{\gamma}_i' \text{Factors}_t + \hat{\epsilon}_{i,t}, \quad (4.1)$$

where $\text{Factors} = (\text{MARKET}, \text{HML}, \text{SMB}, \text{SA}, \text{ROIL})$. Now the $\hat{\beta}_i$ for each climate-related fund is the signal upon which we want to derive our mimicking portfolio. We argue that the higher the $\hat{\beta}_i$, the better and thus the higher the fund's weight must be. However, if $\hat{\beta}_i$ is negative, we give a weight of zero to the particular fund. Consequently, we propose a long-only factor-mimicking portfolio. The reason we restrict our attention to a long-only portfolio is twofold. First, for our investment universe we select only sustainable and climate-related funds. Thus, we do not want to short sell such a fund that Morningstar and V-LAB deem to be climate-focused. Second, similar to why we consider investable (control) factors, we want to compute real-life investable portfolios without any difficult-to-implement constraints.

A simple method to generate a signal (vector) b out of the $\hat{\beta}$'s is to weight the positive $\hat{\beta}$'s:

$$b^{\hat{\beta}} := \frac{\hat{\beta}}{\sum_{i=1}^{177} \hat{\beta}_i}, \quad \tilde{\beta} := \max(\hat{\beta}, 0). \quad (4.2)$$

We also propose two alternative signals based on the squared $\hat{\beta}$'s and $\hat{\beta}$ times its t statistics:

$$b^{\hat{\beta}^2} := \frac{\tilde{\beta}^2}{\sum_{i=1}^{177} \tilde{\beta}_i^2}, \quad \tilde{\beta}^2 := (\tilde{\beta}_1^2, \dots, \tilde{\beta}_{177}^2)', \quad (4.3)$$

$$b^{t\text{-stat}} := \frac{\tilde{\beta}^{t\text{-stat}}}{\sum_{i=1}^{177} \tilde{\beta}_i^{t\text{-stat}}}, \quad \tilde{\beta}_i^{t\text{-stat}} := \frac{\tilde{\beta}_i^2}{\hat{\sigma}_i^{HC3}},$$

$$\tilde{\beta}^{t\text{-stat}} := (\tilde{\beta}_1^{t\text{-stat}}, \dots, \tilde{\beta}_{177}^{t\text{-stat}})', \quad (4.4)$$

where $\hat{\sigma}_i^{HC3}$ is the heteroskedasticity-consistent HC3 standard error of $\hat{\beta}_i$.

As we want to compute a long-only mimicking portfolio, we restrict our attention to strictly positive signals only and thus reduce the investment universe accordingly. More specifically, we consider the sub-universe of $N \leq 177$ funds with (strictly positive) signal vector b .

Note that the signal vector b on its own is already a very intuitive long-only factor-mimicking portfolio. It simply weights all funds based on their (positive) $\hat{\beta}_i$ on the climate news series we want to mimic. We denote this portfolio as our base case. Even though

the base case seems to be an intuitive start, arguably, it is suboptimal portfolio as it is inefficient. We want to improve upon the base-case portfolio by including a more sophisticated and optimized mimicking portfolio construction, resulting in an efficient portfolio by taking into account the information of the funds' covariance matrix.

In Figure 3, we plot the size of the investment universe and the actual number of funds that indeed have a positive signal (β). Note that we invest only in funds that have a complete daily return time series for a yearly rolling window with monthly rebalancing. The number of funds included in the investment universe monotonically increased from 45 in June 2001 to 177 in January 2021. The number of climate-related funds increased especially during the last five years. On the contrary, with exception of the last two years, the number of funds with a positive signal (β) did not increase significantly. In the percentage of funds with a positive signal, there was even a decrease since the financial crisis in 2008. In January 2021, almost 60% of the funds have a positive signal. The average positive β is 0.03, with an average t statistic of 0.80.

Efficient Mimicking Portfolio Construction.

The next step is to use the derived signal b to compute a constrained long-only ($w_i \geq 0$) and fully invested ($w'1 = 1$) climate factor-mimicking portfolio w . To this end, we propose the following portfolio-optimization problem:

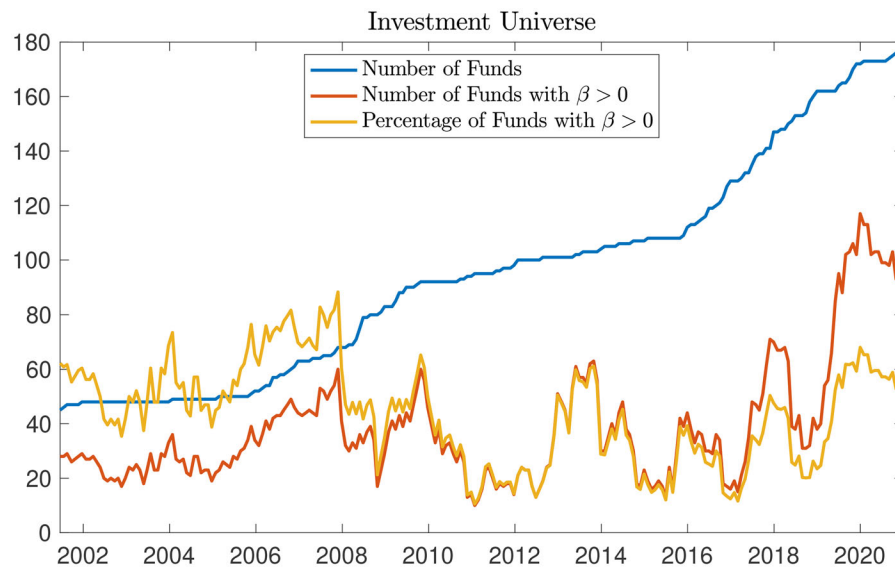
$$\min_w \quad w' \hat{\Sigma} w - \lambda w' b, \quad (4.5)$$

$$\text{s.t.} \quad w_i \geq 0 \quad \text{for all } i = 1, \dots, N \text{ and} \quad (4.6)$$

$$w'1 = 1, \quad (4.7)$$

where $\hat{\Sigma}$ is the $N \times N$ estimated covariance matrix of climate funds returns, 1 is a conformable vector of ones, and $\lambda \geq 0$ is a scaling parameter to define the form of the optimization problem. For example, if we set λ equal to zero, we are basically estimating the constrained minimum-variance portfolio without any signal information, or if we set λ to ∞ we do not take into account the covariance matrix and thus invest only in the asset(s) with the largest signal b_i . Note that the (exact) value of λ depends on the risk appetite and hedging preference of the investor: The higher (lower) the value of λ , the higher (lower) the risk appetite and hedging preference are. Arguably, the portfolio of interest should consider both the covariance matrix and the signal to be an attractive "well-diversified" and risk-controlled (climate efficient factor-)mimicking portfolio.

Figure 3. The Number of Funds in the Investment Universe and the Actual Number and Percentage of Funds in Which We Want to Invest



Notes: Funds with a positive signal and large enough time series. Numbers are for a yearly rolling window with monthly rebalancing frequency from 06/19/2001 until 01/29/2021 for the NYT Tag Climate (Change) News Index.
NYT = *The New York Times*.

In our analysis, we define the base-case scenario as $\lambda = 2$, with the identity matrix as covariance matrix estimator $\hat{\Sigma} = I$. This is an interesting scenario because the minimization function (4.5) simplifies to $w'w - 2w'b$, with the unique solution $w^* = b$. This is the very intuitive portfolio that weights its constituents according to their signal. Another special case is when we set $\lambda = 0$; thus, the constrained minimum-variance portfolio based on replacing the covariance matrix estimator with the identity matrix. This results in the equally weighted portfolio of the assets with positive signal: $w^* = 1/N$. Consequently, in the optimization problem (4.5–4.7) with $\hat{\Sigma} = I$, we derive two naive benchmarks: b weighting (base case) and equal weighting ($1/N$).

We believe that we can improve upon these naive portfolios by taking into account a more sophisticated covariance matrix estimator than just the identity matrix. However, due to the large and increasing number of sustainable and climate-related funds (N ; see Figure 3) and the short sample that is available/observable (T), this is a challenging task. In such a setting, the sample covariance matrix has a poor out-of-sample performance. Jobson and Korkie (1980) show that the sample covariance matrix suffers from high estimation error, especially when the number of available assets N is high compared to the return

time-series length T . The shrinkage estimators of Ledoit and Wolf (2022) compose a powerful class of estimators for this setting. Due to the small sample size, we suggest the (unconditional) constant-variance-covariance (CVC) linear shrinkage estimator of De Nard (2022), which shrinks the sample variances and sample covariances toward their grand mean.⁵ Recently, these shrinkage estimators were successfully extended to dynamic and factor models. Unfortunately, due to the short time-series length of the sustainable and climate-related funds, the combination of shrinkage with multivariate GARCH models as introduced by Engle, Ledoit, and Wolf (2019) and the enhancement with intraday data by De Nard et al. (2022) are suboptimal. On the contrary, the extension to factor models of De Nard, Ledoit, and Wolf (2021) is a further improvement to reduce the curse of dimensionality. Therefore, we recommend using the AFM-CVC estimator. This is an approximate factor model (AFM) in which the covariance matrix of the factors and residuals are estimated via CVC shrinkage. The proposed estimator is similar to the Dynamic Conditional Correlation - Nonlinear shrinkage approximate factor model dynamic conditional correlation-nonlinear shrinkage (AFM-DCC-NL) estimator of De Nard and Zhao (2023), but with three modifications: (i) It uses linear shrinkage instead

of nonlinear shrinkage due to the small sample, (ii) there is no (multivariate) GARCH component due to the small sample, and (iii) it uses the investable (see “Selection of the Investment and Factor Universe”) instead of the “academic” Fama–French factors. In [Appendix B](#), we give an overview of the modified approximate factor model shrinkage estimator.

Performance Measures. The first performance measure is the β of the news variable. Thereby, we regress the (return of the) climate-efficient factor-mimicking portfolio, $r_{\text{CEP},t}$, on climate news and factors using the whole dataset (without intercept):

$$r_{\text{CEP},t} = \hat{\beta} \text{CC}_t + \hat{\gamma}' \text{Factors}_t + \hat{\epsilon}_t. \quad (4.8)$$

Ideally, the estimated β should be positive and significant, and the bigger the better. This is an out-of-sample test in the sense that every month of the portfolio return is subsequent to the data used to form the portfolio, as described in “Efficient Mimicking Portfolio Construction.”

The second performance measure is the α of the mimicking portfolio. Thereby, regress the excess return of CEP on all investable factors using the whole dataset:

$$r_{\text{CEP},t} - r_{f,t} = \hat{\alpha} + \hat{\gamma}' \text{Factors}_t + \hat{\epsilon}_t. \quad (4.9)$$

Ideally, the estimated intercept, $\hat{\alpha}$, should be positive and significant, and the bigger the better. This is an out-of-sample test every month.

For the estimation of the parameters in regression 4.8 and 4.9, we use adaptive least squares (ALS) for improved inference in financial factor models, as shown by Beck, De Nard, and Wolf (2023). ALS generally leads to smaller heteroskedasticity-consistent (HC) standard errors compared to ordinary least squares (OLS), which translates into improved inference in the form of shorter confidence intervals and more powerful hypothesis tests. The ALS method “decides” between the OLS method and the weighted least squares (WLS) method based on a pretest for conditional homoskedasticity. Only if this test rejects the null—that is, if this tests detects a significant amount of conditional heteroskedasticity in the data—does one use WLS; otherwise, one uses OLS. Of course, either way, one must use corresponding HC standard errors for the inference. Note that we use HC3 to compute the test statistics.

An interesting alternative performance measure is the (Pearson’s linear) out-of-sample correlation between the CEP and the climate news index. Even

though the correlation is an intuitive performance measure to quantify the accuracy of a factor-mimicking portfolio, the goal is to derive a positively correlated factor-mimicking portfolio controlled for various factors. Therefore, in line with the validating β regression, we focus on the partial correlation between the CEP and the climate news index of interest, adjusting for the other five factors mentioned above. Thereby, we focus on daily, weekly, and monthly out-of-sample partial correlations.

Finally, we also report annualized results on the out-of-sample average portfolio return (AV), standard deviation (SD), and information ratio (IR := AV/SD).

It is clear that by optimizing these criteria, the results are no longer truly out-of-sample. A fully out-of-sample version of this portfolio is maintained and updated daily on V-LAB. In this case, the tuning parameters remain fixed and the code is fixed. News data and returns are generated daily, and the portfolio weights are updated monthly. Performance of this climate-efficient factor-mimicking portfolio (CEP) is posted daily on V-LAB.

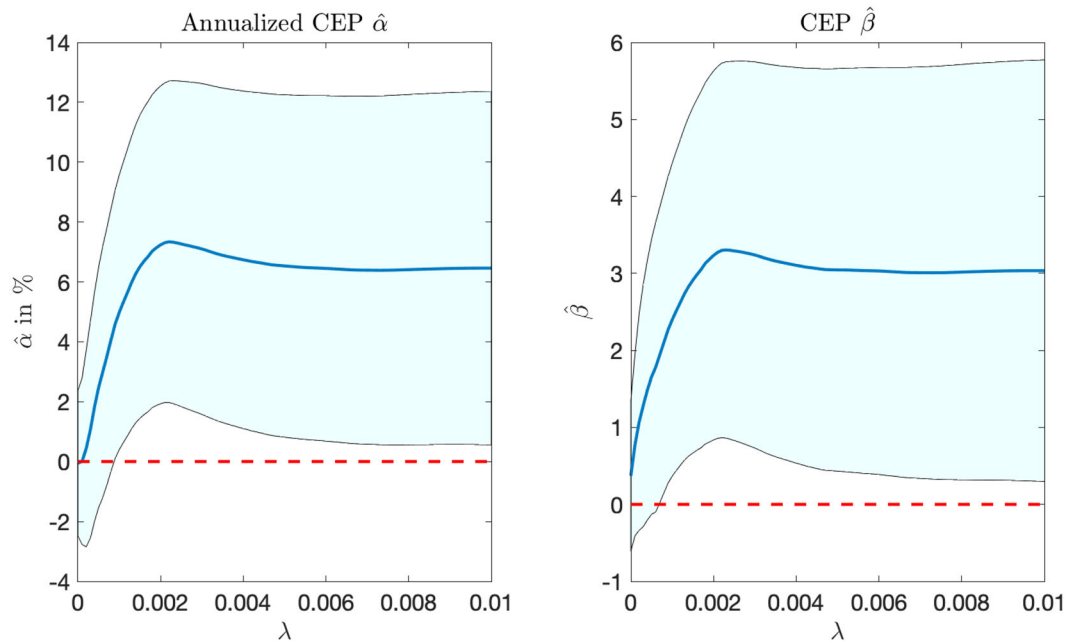
Empirical Analysis

In our empirical analysis, we use daily data with a monthly rolling window, resulting in a monthly re-estimation of the models and rebalancing of the portfolios. In this section, we present the main results of our climate-efficient factor-mimicking portfolio approach, the results for monthly rebalancing based on the $\beta \times t$ -statistics signal (4.4). The main results are based on 4,935 daily (out-of-sample) returns of the NYT Tag Index CEP from 06/19/2001 until 01/29/2021. To further the robustness of our results, in “Robustness of the Portfolio” we include other signals and climate news indices.

Main Results. To show that we have computed a well-performing climate factor-mimicking portfolio, we plot the alphas and betas for various parameterizations of our CEP optimization in [Figure 4](#).

We find that our CEP approach is very powerful and robust, as all resulting CEPs have strictly positive (out-of-sample) $\hat{\alpha}$ and $\hat{\beta}$. Interestingly, the extreme portfolio which neglects the signal by setting ($\lambda = 0$), has positive $\hat{\alpha}$ and $\hat{\beta}$. Therefore, the investment universe selection screening for funds with a positive signal is vitally important. However, the $\hat{\alpha}$ and $\hat{\beta}$ of this constrained minimum-variance portfolio are not statistically significant. To visualize all the CEPs that have statistically significant $\hat{\alpha}$ and $\hat{\beta}$, we plot their

Figure 4. The (Annualized) $\hat{\alpha}$ and $\hat{\beta}$ and Their 95% Confidence Intervals Based on HC3 Standard Errors for the Factor-Mimicking Portfolios of the NYT Climate News Based on the $\beta \times t$ -Statistics Signal



Note: NYT = The New York Times.

95% confidence bands based on HC3 standard errors. When increasing the exposure to the signal, thus λ , we observe that not only the $\hat{\beta}$ but also the $\hat{\alpha}$ of the CEP is increasing quickly and becoming statistically significant. Consequently, deviating from the minimum-variance portfolio and putting more weight on the signal are clearly beneficial. Nevertheless, putting too much weight on the signal can hurt the performance because in the limit, where $\lambda = \infty$, we focus only on the signal, resulting in a CEP investing only in the fund with the largest signal. Arguably, this is not an interesting portfolio for a climate hedging manager as it is undiversified and has a large exposure to estimation error of the signal. Finally, both extreme portfolios are suboptimal with insignificant performance measures, but there is a large spectrum of CEPs in between with large and statistically significant $\hat{\alpha}$ and $\hat{\beta}$. These CEPs find a good combination of signal exposure and variance reduction.

Due to the robustness of our CEP approach, we see that the $\hat{\alpha}$ and $\hat{\beta}$ are very stable for a “large enough” λ . Nevertheless, empirically we find an optimal and consistent region around $\lambda \approx 0.002$, with the largest and statistically significant $\hat{\alpha}$ and $\hat{\beta}$. The performance measures of this CEP and all the benchmarks are

presented in Table 1. The CEP delivers a high $\hat{\alpha}$ of 7.34% (with a t statistic of 2.68) and a high $\hat{\beta}$ of 3.30 (with a t statistic of 2.66). Thus, we have indeed derived climate(-efficient) factor-mimicking portfolios.

Now we restrict our attention to the (optimal) CEP and compare its performance with various benchmarks: the $\beta \times t$ -statistic signal-weighted portfolio (base case), the signal-only portfolio ($\lambda = \infty$), the constrained minimum-variance portfolio ($\lambda = 0$), the positive signal equally weighted portfolio ($1/N$), and the (overall) equally weighted portfolio (EW). All performance measures are summarized in Table 1.⁶

First, all included benchmarks are consistently and markedly outperformed by the proposed CEP. With exception of the $\hat{\beta}$ and monthly partial correlation of the base case, the results are not statistically significant for the benchmarks. For the equally weighted portfolio, the $\hat{\alpha}$, $\hat{\beta}$, and partial correlations are even negative. Thus, even though we focus on a sustainable and climate-related investment universe, it is a difficult task to derive a long-only climate factor-mimicking portfolio with positive and statistically significant results. The $1/N$ improves the performance by focusing only on the sub-universe with positive signal; however, the results are still bad.

Table 1. Various Annualized Performance Measures (in %) of the Monthly Re-estimated and Rebalanced Climate-Efficient Factor-Mimicking Portfolio (CEP) Based on the $\beta \times t$ -Statistics Signal and Various Benchmarks

	CEP	Base Case	$\lambda = \infty$	$\lambda = 0$	1/N	EW
$\hat{\alpha}$	7.34	2.46	5.55	0.05	-0.68	-2.32
t-Stat	2.68	1.51	1.63	0.05	-0.83	-3.14
$\hat{\beta}$	3.30	1.71	2.74	0.37	0.21	-0.40
t-Stat	2.66	2.31	1.73	0.74	0.62	-1.46
AV	14.71	9.78	13.26	5.36	6.54	4.88
SD	21.31	19.34	24.31	13.60	17.96	18.24
IR	0.69	0.51	0.55	0.39	0.36	0.27
Daily						
PCor	1.91	2.51	1.46	0.09	0.29	-0.81
t-Stat	1.34	1.76	1.02	0.07	0.21	-0.57
Weekly						
PCor	4.20	6.04	3.95	0.08	0.75	-2.38
t-Stat	1.31	1.89	1.24	0.25	0.23	-0.75
Monthly						
PCor	9.94	15.58	10.59	4.51	5.74	-1.02
t-Stat	1.50	2.36	1.60	0.68	0.87	-0.15

Notes: $\hat{\alpha}$ indicates the (estimated) alpha of the portfolio; $\hat{\beta}$ indicates the (estimated) beta of the news variable; t-Stat indicates the t statistic based on HC3 standard errors.

AV = average return; SD = standard deviation; IR = information ratio; PCor = partial out-of-sample correlation between the factor-mimicking portfolio and the climate index.

In the rows labeled $\hat{\alpha}$, $\hat{\beta}$, and IR, the largest number appears in boldface. The measures are based on 4,935 daily, respectively 987 weekly, and respectively 235 monthly (out-of-sample) returns from 06/19/2001 to 01/29/2021.

Therefore, more impressive are the large and significant $\hat{\alpha}$ (up to 7.34%) and $\hat{\beta}$ (up to 3.30) we can obtain by our CEPs.

In comparison, the naive benchmarks based on the identity matrix, $\hat{\Sigma} = I$, are not performing that well. The base case has a statistically insignificant $\hat{\alpha}$ of 2.46%, roughly one-third of the CEP, and statistically significant $\hat{\beta}$ of 1.71, almost half as large. Therefore, within our CEP approach (4.5), the identity matrix is markedly and consistently outperformed by more sophisticated covariance matrix estimators, for example, by AFM-CVC. Even though the identity matrix is a very naive covariance matrix estimator, it gives us intuitive and diversified portfolios. For example, for $\lambda = 0$, we get the (positive signal) equally weighted portfolio (1/N), and for $\lambda = 2$ we get the (positive) signal-weighted portfolio (base case). Even though the base case has positive $\hat{\alpha}$ and $\hat{\beta}$ and the 1/N has positive $\hat{\beta}$, they are suboptimal portfolios markedly outperformed by our CEP.

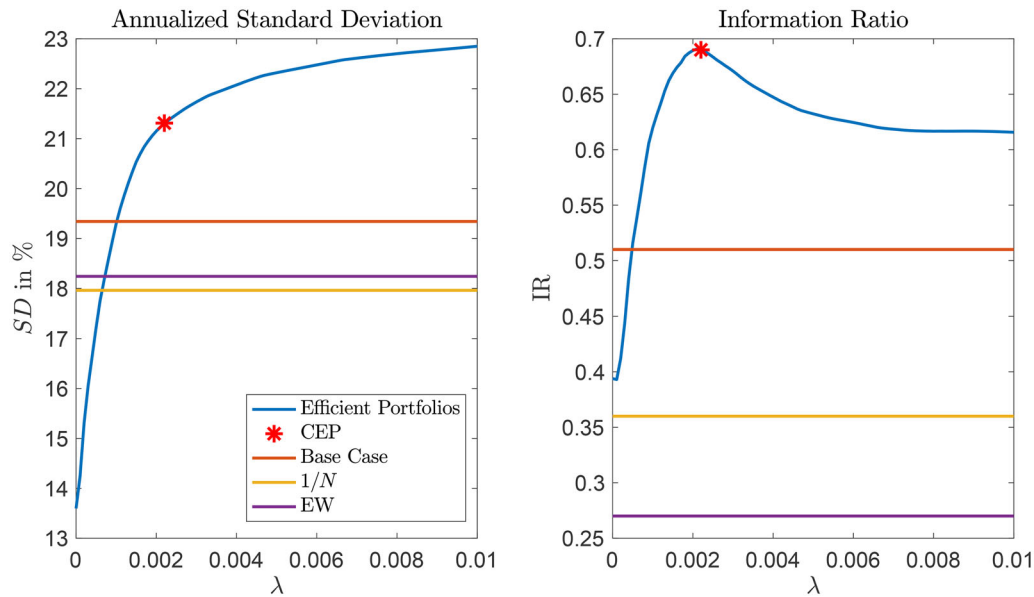
As mentioned above, there are two interesting special cases for CEP. The first is ($\lambda = \infty$) which gives a maximum weight to the signal and selects an undiversified portfolio investing only in the asset with the largest

signal at each rebalancing. This extreme portfolio is performing decently and is the best benchmark; it is also much better than the base case that weights the portfolio according to the signal size. It has high $\hat{\alpha}$ (5.55%) and $\hat{\beta}$ (2.74); however, the parameter estimates are not statistically significant and are still way below the CEP. This is due to the large estimation error of the noisy signal and no diversification effect. On the contrary, if we ignore the signal ($\lambda = 0$) and compute the long-only minimum-variance portfolio, the $\hat{\alpha}$ and $\hat{\beta}$ shrink close to zero but are still higher than those of the equally weighted naive benchmarks.

A further contribution of our CEP approach is that we compute efficient portfolios and not just maximum correlation portfolios or naively weighted portfolios. We directly incorporate the covariance matrix of the funds in our optimization to control for the CEP variance. Therefore, the higher the λ , the higher the risk appetite of an investor, usually resulting in increased average return (AV), standard deviation (SD), and turnover (TO) of the portfolio. We present the relevant summary statistics in Table 1.

Additionally, in Figure 5 we plot the standard deviation and information ratio ($IR := AV/SD$) for our CEPs based on the AFM-CVC shrinkage covariance

Figure 5. The (Annualized) Out-of-Sample Standard Deviation and Information Ratio of Various NYT Climate News Factor-Mimicking Portfolios Based on the $\beta \times t$ -Statistics Signal



Notes: The efficient factor-mimicking portfolios are based on the AFM-CVC shrinkage estimator (blue), and the optimal CEP is marked with *. The numbers of the benchmarks based on the identity matrix are plotted as horizontal lines for comparison. AFM = approximate factor model; CEP = climate-efficient factor-mimicking portfolio; CVC = constant-variance-covariance; NYT = *The New York Times*.

matrix estimator (in blue) and the identity matrix as implied by the base case and the equally weighted benchmarks. All the benchmarks are markedly and consistently outperformed by more sophisticated covariance matrix estimators, for example, by AFM-CVC. AFM-CVC successfully reduces the out-of-sample standard deviation and thus increases the information ratio. In terms of SD, the outperformance of AFM-CVC over the identity matrix is (always) highly statistically significant and also economically meaningful.⁷ In terms of IR, the outperformance of AFM-CVC over the identity matrix is usually statistically significant and also economically meaningful.⁸ For example, for the CEP, thus AFM-CVC with $\lambda = 0.0022$ marked with *, the IR is 0.69 and hence much larger than all benchmarks. Consequently, the efficiency of optimizing (4.5) by the proposed CEP not only finds a markedly improved tradeoff solution of the signal-to-noise ratio, by increasing the $\hat{\alpha}$ and $\hat{\beta}$ and making them statistically significant, but also improves upon the risk-return tradeoff by generating (statistically significant) larger risk-adjusted returns; see Table 1.

As an alternative performance measure, we also present (out-of-sample) partial correlations (PCor)

between the factor-mimicking portfolios and the climate news index; see Table 1. Also in terms of partial correlation, our CEPs are very powerful and robust. Across all return frequencies, any computed CEP has strictly positive partial correlation with the climate news index. Interestingly, the partial correlation increases markedly with lower frequencies. There are two main reasons: First, aggregating daily returns to weekly or monthly returns reduces the noise extensively. Second, the climate news index is on a weekly frequency. Note that in terms of partial correlation, the base case consistently outperforms the CEPs and for monthly returns the PCor of 15.58% is even statistically significant. The equally weighted portfolios have very low or even negative PCor. This shows again that even with a sustainable and climate-related investment universe, it is difficult to compute a decent out-of-sample factor-mimicking portfolio.

Finally, when we compare the distribution of the out-of-sample performance measures, IR (Figure 5), $\hat{\alpha}$ and $\hat{\beta}$ (Figure 4), and partial correlations, we observe a similar pattern. In general, the performance can be markedly increased by increasing the weight λ toward the signal; however, a too-large λ results in a less diversified and extreme portfolio. Motivated by

Table 2. Various Annualized Performance Measures Net of 10 bps Transaction Costs of the Monthly Re-estimated and Rebalanced Climate-Efficient Factor-Mimicking Portfolio (CEP) Based on the $\beta \times t$ -Statistics Signal and Various Benchmarks

	CEP	Base Case	$\lambda = \infty$	$\lambda = 0$	1/N	EW
$\hat{\alpha}$	6.20	1.73	4.37	−0.66	−1.15	−2.35
t-Stat	2.26	1.06	1.28	−0.53	−1.41	−3.18
$\hat{\beta}$	2.95	1.50	2.40	0.19	0.07	−0.41
t-Stat	2.34	2.03	1.51	0.38	0.22	−1.49
AV	13.58	9.06	12.08	4.76	6.07	4.85
SD	21.31	19.34	24.31	13.60	17.96	18.24
IR	0.64	0.47	0.50	0.35	0.34	0.27
TO	0.95	0.60	0.99	0.50	0.40	0.03

Notes: $\hat{\alpha}$ indicates the (estimated) alpha of the portfolio; $\hat{\beta}$ indicates the (estimated) beta of the news variable; t-Stat indicates the t-statistic based on HC3 standard errors.

AV = average return; SD = standard deviation; IR = information ratio; TO = average monthly turnover.

In the rows labeled $\hat{\alpha}$, $\hat{\beta}$, and IR, the largest number appears in boldface. The measures are based on 4,935 daily, respectively 987 weekly, and respectively 235 monthly (out-of-sample) returns from 06/19/2001 to 01/29/2021.

our empirical findings, we propose to use a $\lambda \approx 0.002$ to obtain all-around large and statistically significant performance measures.

Portfolio Composition. In this section, we restrict our attention to the portfolio composition of the CEP and its benchmarks to investigate the portfolio holdings and their dynamic to obtain a successful mimicking portfolio. In the top panel of Figure A2, we plot the fund weights of the climate-efficient factor-mimicking portfolios in terms of the scaling parameter λ at the end of the sample (01/29/2021). Hence, on the very left of this panel we see the fund weights of the long-only minimum-variance portfolio that is well diversified. Now increasing λ , and thus the weight to the climate signal, we observe how most of the weights and the number of assets in the portfolio become smaller. Note that in the limit ($\lambda = \infty$), we invest only in the asset with the largest signal. The “optimal” CEP with $\lambda = 0.0022$ is almost in the middle of this top panel, indicating that we invest only in a few assets with large position in a single fund. In the bottom panel of Table A2, we list the top 10 positions of this CEP at the end of the sample and see that it invests almost 65% in *KraneShares MSCI China Environment (KGRN)*, more than 15% in *Global X Lithium & Battery Tec (LIT)*, 7% in *iShares Global Clean Energy (ICLN)*, as well as 5% in *Invesco Global Clean Energy (PBD)* and *Invesco Solar (TAN)*.

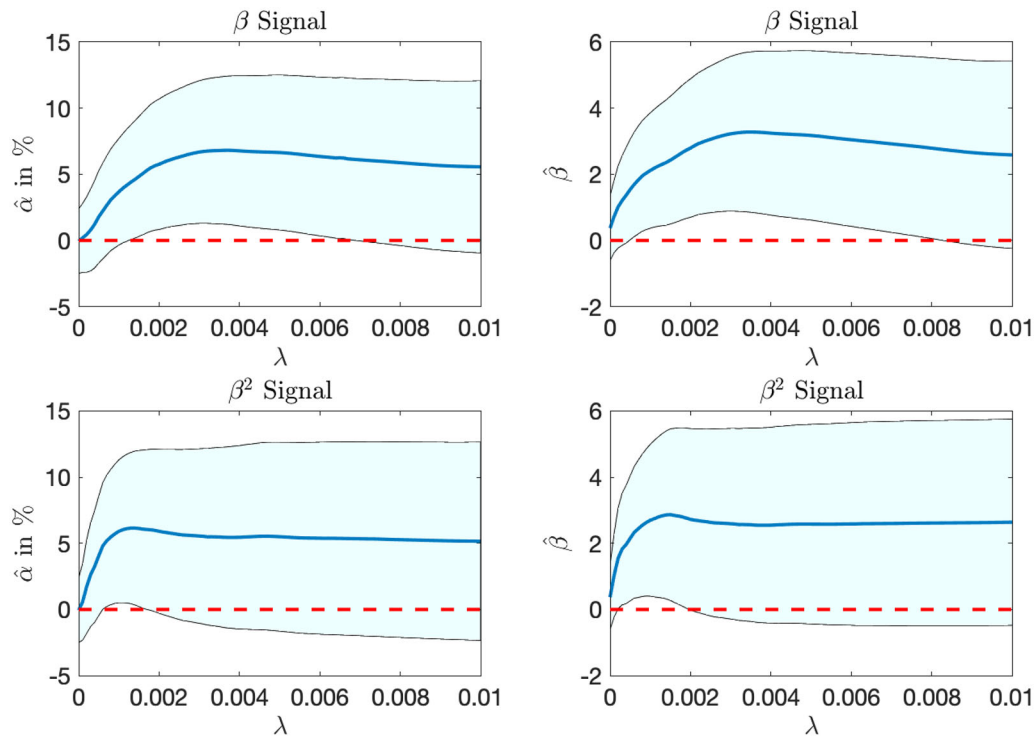
In the middle panel of Figure A2, we plot the fund weights of the “optimal” CEP over time. We see that it is erratic, with often extreme positions that can

change quickly from one rebalancing to another; thus, it has a high average monthly turnover of 95%. On the other hand, the base-case scenario is much more stable over time by investing proportional to the (positive) signal size with an average monthly turnover of 40%. Nevertheless, the CEP consistently outperforms the base case (in terms of risk-adjusted returns, $\hat{\alpha}$ and $\hat{\beta}$), indicating how important it is to compute an efficient portfolio optimizing for portfolio variance and signal exposure. Consequently, diversification is not that important for a factor-mimicking portfolio; for example, see the results for no diversification at all with $\lambda = \infty$.

In Table A2, we also show the largest average monthly holdings and the percentage of holdings larger than 1%, both over the entire sample. For the last two decades, the *New Alternatives (NALFX)* ETF has been vitally important for climate news hedging with an average portfolio position of 8.55% and being more than 14% of the rebalancings in the portfolio.

Besides examining the portfolio holdings, we also plot the sector exposures of the CEP and the base case in Figure A3. We use the V-LAB classification of the 177 climate-related funds into (i) fossil fuel free, (ii) low carbon, and (iii) low environmental risk. We observe the same pattern that the base case is more diversified and stable over time. Interesting is that for the base case, at the beginning of the sample there is a similar exposure to all three sectors, in the middle there is a clear overweight of low environmental risk, and at the end there is a clear overweight of fossil fuel-free funds (80%)

Figure 6. The (Annualized) $\hat{\alpha}$ and $\hat{\beta}$ and Their 95% Confidence Intervals Based on HC3 Standard Errors for Various Efficient Factor-Mimicking Portfolios of the NYT Climate News Based on the β and β^2 Signal



Note: NYT = The New York Times.

and no exposure anymore to low environmental risk. The CEP invests at the end of the sample only in fossil fuel-free funds, neglecting the low-carbon funds.

Robustness of the Portfolio

To further enhance the robustness of our results, we run several sensitivity checks. The results are robust also for weekly re-estimation of the models and rebalancing of the climate factor-mimicking portfolio. However, the average performance and significance of coefficients, $\hat{\alpha}$ and $\hat{\beta}$, usually decreases for higher frequencies. Therefore, we deem monthly updating the most useful for our purpose.

Transaction Costs. An important disadvantage of the CEP is the discussed large monthly turnover and often low diversification compared to (most of) the benchmarks. Even though we empirically show that it is important to have sometimes large and dynamic exposures to single funds due to optimal

hedging goals, arguably, the large turnover should be penalized due to investor preferences and costs. Consequently, we run a sensitivity analysis for performance numbers net of various transaction costs. In Table 2, we report the usual performance measures, taking into account 10 bps of transaction costs and report additionally the monthly average turnover (TO). The 10 bps are a conservative estimate of transaction costs because the 177 sustainable V-LAB funds are in general very large and liquid.

From Table 2, we see that the transaction costs of 10 bps reduce the $\hat{\alpha}$ from 7.34 to 6.20 and the $\hat{\beta}$ from 3.30 to 2.95. Due to the large turnover, the reduction is substantial and larger compared to the benchmarks, but the coefficients are still positive as well as economically important and statistically significant. Even net of transaction costs, there is superior all-around performance of CEP compared to all the benchmarks. Note that for CEP, we would need (unrealistically) large transaction costs of 25 bps to get insignificant but still positive $\hat{\alpha}$ and $\hat{\beta}$.

Table 3. Various Annualized Performance Measures (in %) of the Monthly Re-estimated and Rebalanced Climate-Efficient Factor-Mimicking Portfolio (CEP) Based on the β Signal ($\lambda = 0.0035$) as Well as β^2 Signal ($\lambda = 0.0015$)

	β signal				β^2 signal		
	CEP	Base Case	$\lambda = \infty$		CEP	Base Case	$\lambda = \infty$
$\hat{\alpha}$	6.79	1.33	4.28	Daily Weekly Monthly	6.12	2.92	4.28
t-Stat	2.40	0.93	1.08		2.03	1.37	1.08
$\hat{\beta}$	3.28	1.21	2.61		2.87	1.88	2.61
t-Stat	2.65	1.96	1.58		2.17	2.07	1.58
AV	13.93	8.64	12.00		13.36	10.42	12.00
SD	21.46	19.17	26.42		22.27	20.97	26.42
IR	0.65	0.45	0.45		0.60	0.50	0.45
PCor	2.14	2.10	1.64		1.55	2.02	1.64
t-Stat	1.50	1.47	1.15		1.09	1.42	1.15
PCor	4.88	5.34	4.97		4.18	5.47	4.97
t-Stat	1.53	1.67	1.56		1.31	1.71	1.56
PCor	9.27	14.20	4.71		6.81	13.38	4.71
t-Stat	1.40	2.15	0.71		1.03	2.03	0.71

Notes: $\hat{\alpha}$ indicates the (estimated) alpha of the portfolio; $\hat{\beta}$ indicates the (estimated) beta of the news variable; t-Stat indicates the t-statistic based on HC3 standard errors.

AV = average return; SD = standard deviation; IR = information ratio; PCor = partial out-of-sample correlation between the factor-mimicking portfolio and the climate index.

For both signals, in the rows labeled $\hat{\alpha}$, $\hat{\beta}$, and IR, the largest number appears in boldface. The measures are based on 4,935 daily, respectively 987 weekly, and respectively 235 monthly (out-of-sample) returns from 06/19/2001 to 01/29/2021.

Alternative Signals. To assess how robust and general our CEP approach is, we focus now on alternative (simpler) signals. In the empirical analysis, we used the $\beta \times t$ -statistic defined in (4.4). The idea is to multiply the β with its t statistic to take into account not only the size but also the power of the signal. Alternatively, one could use directly the β or the squared β as a signal to overweight the actual size. The problem, especially of the latter, is that large(r) (absolute) signal numbers often have also large(r) estimation error and are prone to over-fitting and “error maximization” in optimization problem (4.5). This error maximization is comparable with the estimation error problem of the covariance matrix. Michaud’s (1989) explanation is that the most extreme sample covariance matrix coefficients, or here signals, tend to be extreme not because this is necessarily true but because they contain an extreme amount of error. Consequently, a Markowitz (1952) signal-variance investor interested to solve (4.5) takes the highest bets on the unreliable extreme coefficients of the beta and the sample covariance matrix for “optimal” signal exposure and risk control, which

is counterproductive. This is why we suggest using a shrinkage estimator of the sample covariance matrix (AFM-CVC) and to weight the beta coefficients by their power (t statistics) to reduce estimation error.

To show the benefit of weighting the signals by their power, we present the results of the β and β^2 signals in Figure 6 and Table 3. First, the results show that our CEP approach is robust in terms of the signal choice as we observe similar results and patterns. Also for the alternative signals, we find a consistent outperformance when taking into account the AFM-CVC covariance matrix estimator with overall positive $\hat{\alpha}$, $\hat{\beta}$, and partial correlation. Nevertheless, we see some intuitive differences between the signals. For example, the number of significant $\hat{\alpha}$ and $\hat{\beta}$ reduces drastically for the β and β^2 signal compared to the $\beta \times t$ -statistic signal. Furthermore, not only does the region of statistically significant CEPs shrink (in terms of λ) but also the level of $\hat{\alpha}$ and $\hat{\beta}$ reduces consistently and markedly and, in general, the risk-adjusted returns also get smaller. Thus, controlling for the power of the β signal is clearly beneficial by avoiding very large positions in unreliably high β funds.

Table 4. Various Annualized Performance Measures (in %) of the Monthly Re-estimated and Rebalanced WSJ General and Physical Climate Factor-Mimicking Portfolio Based on the $\beta \times t$ -Statistics Signal

	WSJ General			WSJ Physical		
	CEP	Base Case	$\lambda = \infty$	CEP	Base Case	$\lambda = \infty$
$\hat{\alpha}$	8.19	1.71	4.92	6.15	1.33	4.60
t-Stat	3.19	1.24	1.50	2.69	0.92	1.47
$\hat{\beta}$	6.58	2.00	3.48	7.32	2.48	4.30
t-Stat	2.43	1.40	1.07	1.60	1.04	0.77
AV	14.66	8.66	11.95	12.50	8.44	11.81
SD	20.19	18.40	23.20	19.47	18.86	23.21
IR	0.73	0.47	0.51	0.64	0.45	0.51

Notes: $\hat{\alpha}$ indicates the (estimated) alpha of the portfolio; $\hat{\beta}$ indicates the (estimated) beta of the news variable; t-Stat indicates the t-statistic based on HC3 standard errors.

AV = average return; SD = standard deviation; IR = information ratio.

For both climate news indices, in the rows labeled $\hat{\alpha}$, $\hat{\beta}$, and IR, the largest number appears in boldface. The measures are based on 4,893 daily, respectively 978 weekly, and respectively 233 monthly (out-of-sample) returns from 06/19/2001 to 11/30/2020.

Note that the results are very robust, such that for every signal investigated the CEP approach systematically outperforms all the benchmarks, delivering a statistically significant factor-mimicking portfolios with better (risk-adjusted) return profile. Again, only in terms of (monthly) out-of-sample partial correlations, the base-case scenario is hard to beat. In terms of “optimal” parametrization, we observe that the β -CEP λ of 0.0035 is larger than the β^2 -CEP λ of 0.0015. Thus, the β^2 -CEP gives more weight to the minimum-variance portfolio, which is intuitive as it does not want to put too high a weight on the very noisy signal.

Alternative Climate News Indices. To assess how robust and general our CEP approach is, we focus now on alternative climate news indices introduced in “Measuring Climate Change News.” Therefore, we also replicated the empirical analysis for the WSJ indices and find overall similar results. In Table 4, we report the most important summary statistics.

For the alternative WSJ indices, we find a consistent outperformance when taking into account the AFM-CVC covariance matrix estimator with overall positive $\hat{\alpha}$, $\hat{\beta}$, and partial correlations. Interesting is that the level of $\hat{\alpha}$ and $\hat{\beta}$ is in general higher for the WSJ indices compared to NYT; however, the estimation error in the coefficient is also larger, resulting in a smaller region of statistically significant performance measures. For example, the “optimal” WSJ physical CEP with $\lambda = 0.0011$, presented in Table 4, has an impressive but insignificant $\hat{\beta}$ of 7.32. Nonetheless,

the “optimal” WSJ general CEP with $\lambda = 0.0020$, also presented in Table 4, has the largest and statistically significant $\hat{\alpha}$ of 8.19%, $\hat{\beta}$ of 6.58, and IR of 0.73, across all indices and signals. Thereby, the $\hat{\beta}$ of the WSJ general CEP almost doubles compared to the NYT CEP.

Note that the results are very robust, such that for every investigated signal and index the CEP approach systematically outperforms all the benchmarks, delivering (in general) statistically significant factor-mimicking portfolios with better (risk-adjusted) return profiles.

Conclusion

Earth’s climate is changing, but uncertainty around the trajectory and the economic consequences of climate change is substantial. As a result, investors around the world desire products that allow them to hedge against the realizations of climate risk. In this article, we provide a rigorous and efficient methodology for constructing portfolios that hedge against risks that are otherwise difficult to insure.

We demonstrate how an efficient mimicking portfolio approach can be successful in hedging climate (change) risk across a number of out-of-sample performance tests. A central feature is that climate hedge portfolios should appreciate when there is news that climate risk is increasing. For this purpose, we first estimate climate risk by (daily) textual analysis of *The New York Times* and *The Wall Street Journal*, following the procedure of Gentzkow, Bryan

Kelly, and Taddy (2019) and Engle et al. (2020). Second, we define a climate-efficient factor-mimicking portfolio (CEP) approach that takes into account not only the betas of the assets with the climate risk news (maximum-correlation approach) but also their risk by the estimation of their covariance matrix. Thereby, we focus on long-only portfolios based on V-LAB's 177 climate-focused funds. Note that the (real-life) performance of the NYT CEPs are posted daily on <https://vlab.stern.nyu.edu/climate>.

Taken together, we show that our CEP approach is very powerful and robust. It has superior all-around performance against a variety of benchmarks, delivering markedly higher and statistically significant alphas

and betas with the climate news indices. Additionally, due to the optimization considering the covariance matrix, it returns an improved maximum correlation versus variance tradeoff with statistically significant larger information ratios.

Disclaimer

The views and opinions expressed herein are those of the authors; do not necessarily reflect the views of AQR Capital Management or OLZ AG, its affiliates, or its employees; do not constitute an offer, solicitation of an offer, or any advice or recommendation to purchase any securities or other financial instruments; and may not be construed as such.

Editor's Note

Submitted 13 October 2023

Accepted 15 March 2024 by Luis Garcia-Feijóo, CFA, CIPM

Notes

1. They focus on two types of idiosyncratic belief shocks: (i) instances when fund advisors experience local extreme heat events that are known to shift climate change beliefs and (ii) instances when fund managers change the language in shareholder disclosures to express concerns about climate risks.
2. Webpage: vlab.stern.nyu.edu.
3. Alternatively, one could also include further momentum or profitability ETFs. In practice, the choice of factors will be up to the researcher and will depend upon the factors relevant to the investor.
4. Before VTI is available, we take the SPDR S&P 500 ETF Trust, $SPY - r_f$.
5. Alternatively, one could use the analytical nonlinear shrinkage formula of Ledoit and Wolf (2020). However, De Nard (2022) and Ledoit and Wolf (2017) have found that linear shrinkage, especially CVC, outperforms nonlinear shrinkage for smaller sample ($T \leq 252$) and investment universes ($N \leq 100$).
6. We also analyze the value-weighted benchmark, measured by asset under management of the ETFs, and the buy and hold strategy of the largest ETF: American Funds Fundamental Investors (ANCFX). As these benchmarks have similar (poor) out-of-sample performance as EW, respectively $1/N$, and due to readability, we do not report the results here. Nevertheless, the results are available upon request.
7. Note that to assess the goodness of a covariance matrix estimator, one needs to compare only the SD of the minimum-variance portfolio, thus 13.60 for $\lambda = 0$, as it is a "clean" problem without estimation of the expected returns or signals. To test the outperformance of the CEPs based on AFM-CVC, a two-sided p -value for the null hypothesis of equal standard deviations is obtained by the prewhitened HAC p_W method described in Ledoit and Wolf (2011, Section 3.1). As the out-of-sample size is very large at 4,935, there is no need to use the computationally more involved bootstrap method described in Ledoit and Wolf (2011, Section 3.2), which is preferred for small sample sizes.
8. A two-sided p -value for the null hypothesis of equal information ratios is obtained by the prewhitened HAC p_W method described in Ledoit and Wolf (2008, Section 3.1). Since the out-of-sample size is very large at 4,935, there is no need to use the computationally more expensive bootstrap method described in Ledoit and Wolf (2008, Section 3.2), which is preferred for small sample sizes.

References

- Alekseev, G., S. Giglio, Q. Maingi, J. Selgrad, and J. Stroebe. 2022. "A Quantity-Based Approach to Constructing Climate Risk Hedge Portfolios." National Bureau of Economic Research, Working Paper w30703.
- Andersson, M., P. Bolton, and F. Samama. 2016. "Hedging Climate Risk." *Financial Analysts Journal* 72 (3): 13–32.
- Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht. 2022. "Climate Change Concerns and the Performance of Green versus Brown Stocks." National Bank of Belgium, Working Paper No. 395.
- Beck, E., G. De Nard, and M. Wolf. 2023. "Improved Inference in Financial Factor Models." *International Review of Economics & Finance* 86: 364–379. doi:10.1016/j.iref.2023.03.009.

- Bjoy, H., B. Wang, A. Stankiewicz, and A. Alladi. 2022. *Investing in Times of Climate Change*. Morningstar Manager Research.
- Bolton, P., and M. Kacperczyk. 2021. "Do Investors Care about Carbon Risk." *Journal of Financial Economics* 142 (2): 517–549. doi:10.1016/j.jfneco.2021.05.008.
- De Nard, G. 2022. "Oops! I Shrunk the Sample Covariance Matrix Again: Blockbuster Meets Shrinkage." *Journal of Financial Econometrics* 20 (4): 569–611. doi:10.1093/jfnec/nbaa020.
- De Nard, G., R. F. Engle, O. Ledoit, and M. Wolf. 2022. "Large Dynamic Covariance Matrices: Enhancements Based on Intraday Data." *Journal of Banking & Finance* 138: 106426. doi:10.1016/j.jbankfin.2022.106426.
- De Nard, G., O. Ledoit, and M. Wolf. 2021. "Factor Models for Portfolio Selection in Large Dimensions: The Good, the Better and the Ugly." *Journal of Financial Econometrics* 19 (2): 236–257. doi:10.1093/jfnec/nby033.
- De Nard, G., and Z. Zhao. 2023. "Using, Taming or Avoiding the Factor Zoo? A Double-Shrinkage Estimator for Covariance Matrices." *Journal of Empirical Finance* 72: 23–35. doi:10.1016/j.jempfin.2023.02.003.
- Engle, R. F. 2014. "A Financial Approach to Environmental Risk." *Resources for the Future*, Featurette 187.
- Engle, R. F., S. Giglio, B. Kelly, L. Heebum, and J. Stroebel. 2020. "Hedging Climate Change News." *The Review of Financial Studies* 33 (3): 1184–1216. doi:10.1093/rfs/hhz072.
- Engle, R. F., O. Ledoit, and M. Wolf. 2019. "Large Dynamic Covariance Matrices." *Journal of Business & Economic Statistics* 37 (2): 363–375. doi:10.1080/07350015.2017.1345683.
- Faccini, R., R. Matin, and G. Skiadopoulos. 2021. "Are Climate Change Risks Priced in the Us Stock Market?." Danmarks Nationalbank, Working Paper No. 169.
- Fama, E., and J. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy* 81 (3): 607–636. doi:10.1086/260061.
- Gentzkow, M., B. Bryan Kelly, and M. Taddy. 2019. "Text as Data." *Journal of Economic Literature* 57 (3): 535–574. doi:10.1257/jel.20181020.
- Huberman, G., and S. Kandel. 1987. "Mean-Variance Spanning." *Journal of Finance* 42 (2): 873–888.
- Jobson, J. D., and B. M. Korkie. 1980. "Estimation for Markowitz Efficient Portfolios." *Journal of the American Statistical Association* 75 (371): 544–554. doi:10.2307/2287643.
- Jurczenko, E., and J. Teiletche. 2019. "Macro Factor-Mimicking Portfolios." Working Paper. doi:10.2139/ssrn.3363598.
- Jurczenko, E., and J. Teiletche. 2022. "Hedging Climate Risks: A Cross-Asset Approach." In *Climate Investing: New Strategies and Implementation Challenges*, edited by E. Jurczenko. Wiley-ISTE.
- Lamont, O. 2001. "Economic Tracking Portfolios." *Journal of Econometrics* 105 (1): 161–184. doi:10.1016/S0304-4076(01)00074-4.
- Ledoit, O., and M. Wolf. 2008. "Robust Performance Hypothesis Testing with the Sharpe Ratio." *Journal of Empirical Finance* 15 (5): 850–859. doi:10.1016/j.jempfin.2008.03.002.
- Ledoit, O., and M. Wolf. 2011. "Robust Performance Hypothesis Testing with the Variance." *Wilmott Magazine* 2011 (55): 86–89. doi:10.1002/wilm.10036.
- Ledoit, O., and M. Wolf. 2017. "Nonlinear Shrinkage of the Covariance Matrix for Portfolio Selection: Markowitz Meets Goldilocks." *The Review of Financial Studies* 30 (12): 4349–4388. doi:10.1093/rfs/hhx052.
- Ledoit, O., and M. Wolf. 2020. "Analytical Nonlinear Shrinkage of Large-Dimensional Covariance Matrices." *Annals of Statistics* 48: 3043–3065.
- Ledoit, O., and M. Wolf. 2022. "The Power of (Non-)Linear Shrinking: A Review and Guide to Covariance Matrix Estimation." *Journal of Financial Econometrics* 20 (1): 187–218. doi:10.1093/jfnec/nbaa007.
- Lehmann, B. N., and D. M. Modest. 1988. "The Empirical Foundations of the Arbitrage Pricing Theory." *Journal of Financial Economics* 21 (2): 213–254. doi:10.1016/0304-405X(88)90061-X.
- Markowitz, H. 1952. "Portfolio Selection." *The Journal of Finance* 7 (1): 77–91. doi:10.2307/2975974.
- Merton, R. C. 1987. "A Simple Model of Capital Market Equilibrium with Incomplete Information." *The Journal of Finance* 42 (3): 483–510. doi:10.2307/2328367.
- Michaud, R. 1989. "The Markowitz Optimization Enigma: Is Optimized Optimal?" *Financial Analysts Journal* 45 (1): 31–42. doi:10.2469/faj.v45.n1.31.
- Pástor, L., R. F. Stambaugh, and L. A. Taylor. 2021. "Sustainable Investing in Equilibrium." *Journal of Financial Economics* 142 (2): 550–571. doi:10.1016/j.jfneco.2020.12.011.

Appendix A. Additional Figures and Tables

Figure A1. The WSJ Climate News Indices for General, Physical, and Transition Risk

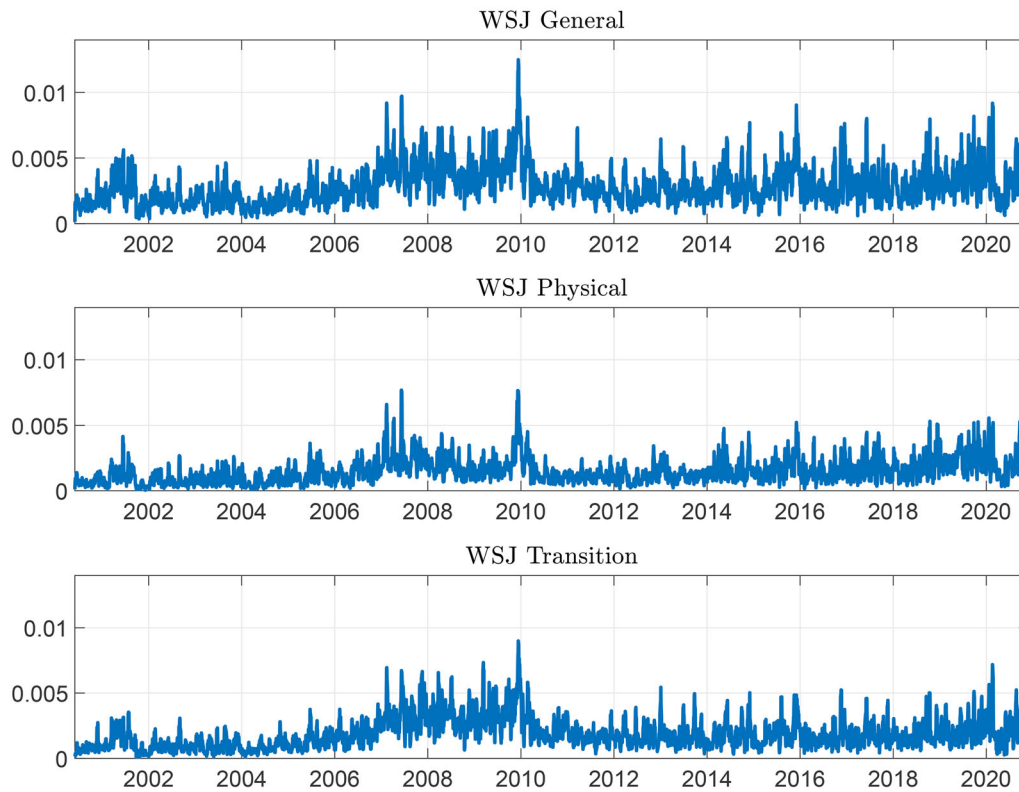


Figure A2. Top: The Fund Weights of the Climate-Efficient Factor-Mimicking Portfolios in Terms of the Scaling Parameter λ at the End of the Sample (01/29/2021); Middle: The Fund Weights of the CEP ($\lambda = 0.0022$) over Time; Bottom: The Fund Weights of the Base Case over Time

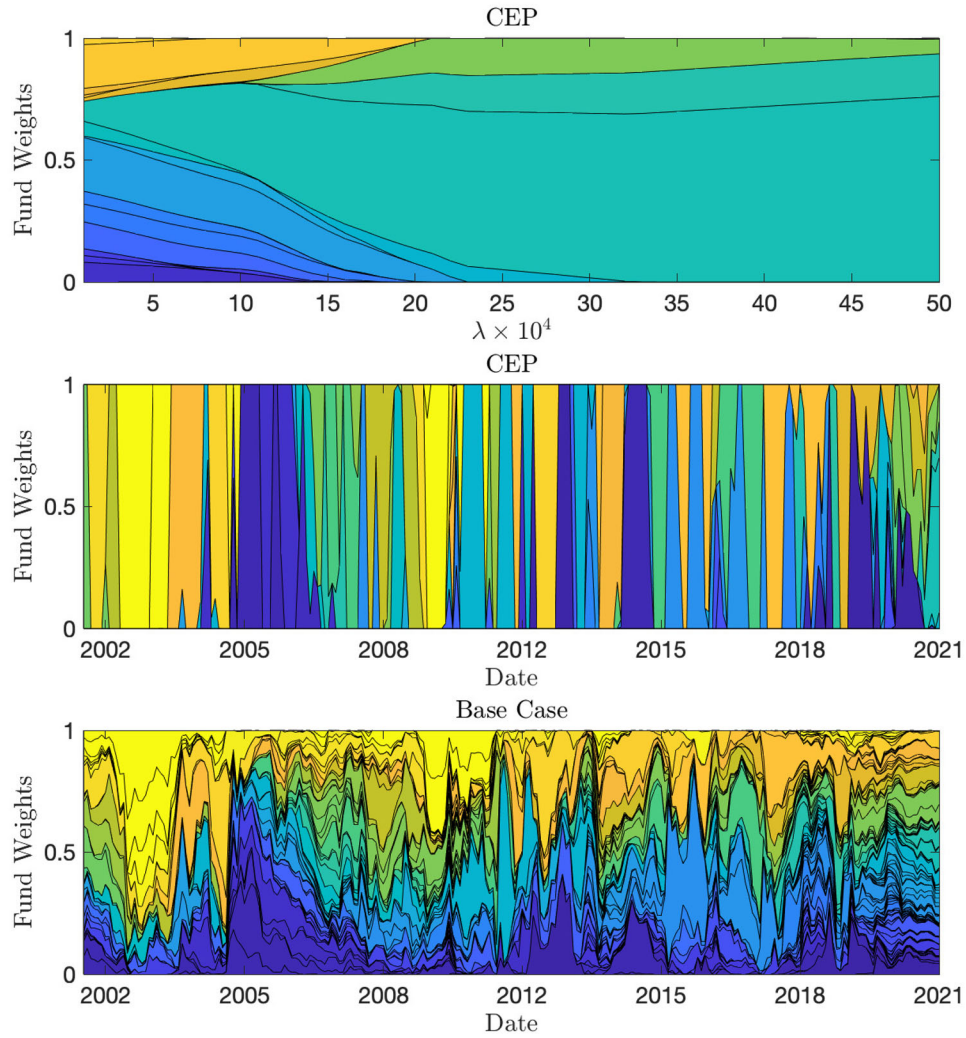


Figure A3. Top: The Sector Weights of the CEP ($\lambda = 0.0022$); Bottom: The Sector Weights of the Base Case over Time

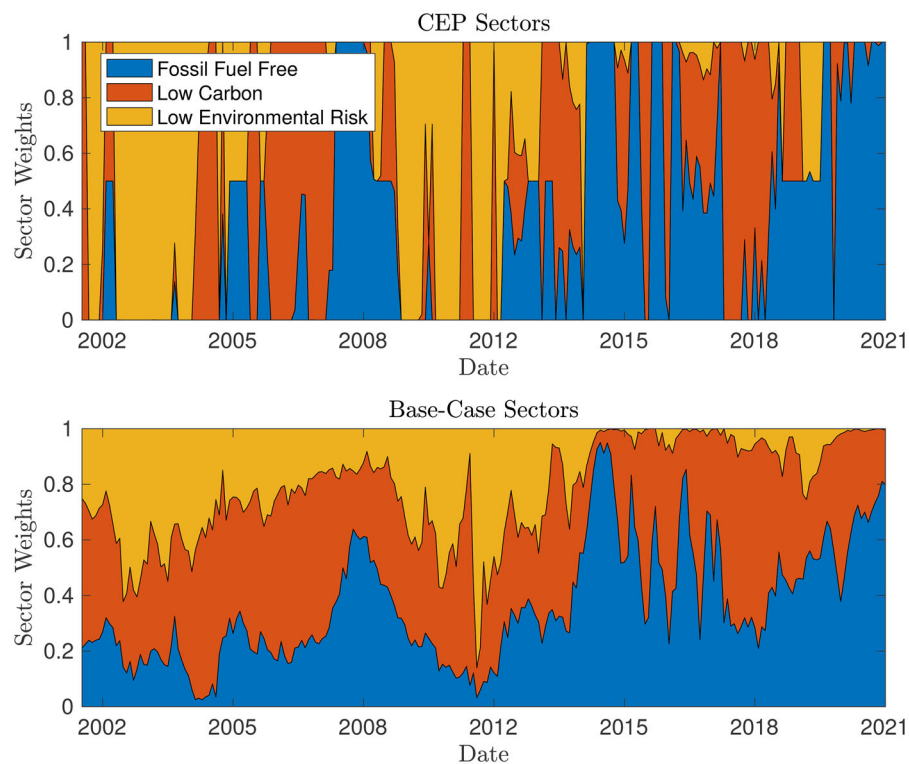


Table A1. (Daily) Pairwise Correlations between All NYT and WSJ Climate News Indices

	NYT Cosine	NYT Tag	WSJ General	WSJ Physical	WSJ Transition
NYT Cosine	1.00	0.32	0.19	0.17	0.18
NYT Tag	–	1.00	0.25	0.27	0.23
WSJ General	–	–	1.00	0.76	0.87
WSJ Physical	–	–	–	1.00	0.54
WSJ Transition	–	–	–	–	1.00

Table A2. Top 10 Funds of the CEP ($\lambda = 0.0022$)

Top Funds in CEP		
#	Fund	\bar{w}_i
1	New Alternatives (NALFX)	8.55%
2	Hotchkis & Wiley Value Opps A (HWAAX)	6.34%
3	AMG Yacktman Focused Fund (YAFFX)	6.22%
4	Eventide Healthcare & Life Sciences (ETIHX)	5.92%
5	AB Sustainable Intl Thematic A (AWPAX)	5.22%
6	Touchstone Global ESG Equity Fund (TEQAX)	3.39%
7	Akre Focus Fund Retail (AKREX)	3.21%
8	First Trust Global Wind Energy (FAN)	3.18%
9	S&P 500 ex-Energy (SPXE)	3.12%
10	Aberdeen Global Equity Impact Instl (JETIX)	3.12%
#	Fund	$w_i > 1\%$
1	New Alternatives (NALFX)	14.04%
2	Hotchkis & Wiley Value Opps A (HWAAX)	8.94%
3	AMG Yacktman Focused Fund (YAFFX)	8.94%
4	Invesco WilderHill Clean Energy (PBW)	8.51%
5	Eventide Healthcare & Life Sciences (ETIHX)	7.66%
6	VanEck Environmental Services (EVX)	7.23%
7	AB Sustainable Intl Thematic A (AWPAX)	6.81%
8	S&P 500 ex-Energy (SPXE)	6.38%
9	Global X Lithium & Battery Tech (LIT)	5.96%
10	Renewable Energy Producers (RNRG)	5.96%
#	Fund	$w_{i,T}$
1	KraneShares MSCI China Environment (KGRN)	64.58%
2	Global X Lithium & Battery Tec (LIT)	15.35%
3	iShares Global Clean Energy (ICLN)	7.08%
4	Invesco Global Clean Energy (PBD)	5.55%
5	Invesco Solar (TAN)	5.28%
6	New Alternatives (NALFX)	1.54%
7	First Trust Global Wind Energy (FAN)	0.54%
8	Goldman Sachs ESG Emerging Markets (GEBSX)	0.03%
9	First Trust EIP Carbon Impact (ECLN)	0.02%
10	iShares MSCI Global Sustainable Development Goals (SDG)	0.01%

Notes: The upper panel shows the largest average (monthly) holdings over time, \bar{w}_i . The middle panel shows the percentage of holdings larger than 1% over time, $w_i > 1\%$. The weights of the top 10 ETFs in the end-of-sample CEP are in the bottom panel, $w_{i,T}$. The measures are based on the 235 monthly CEP holdings from 06/19/2001 to 01/29/2021.

CEP = climate-efficient factor-mimicking portfolio; ETF = exchange-traded fund.

Appendix B.

Covariance Matrix Estimator

The proposed covariance matrix estimator of asset returns is based on a static factor model structure, as explained in De Nard, Ledoit, and Wolf (2021):

$$\hat{\Sigma}_r := \hat{B}\hat{\Sigma}_f\hat{B}' + \hat{\Sigma}_\epsilon, \quad (\text{B.1})$$

where \hat{B} is a $K \times N$ matrix whose i th column is the vector $\hat{\beta}_i$, $\hat{\Sigma}_f$ is the $K \times K$ covariance matrix of factors, and $\hat{\Sigma}_\epsilon$ is the $N \times N$ covariance matrix of residuals, assuming that, for every asset $i = 1, \dots, N$,

$$r_{i,t} = \hat{\alpha}_i + \hat{\beta}_i' r_{\text{Factors}} + \hat{\epsilon}_{i,t}. \quad (\text{B.2})$$

An exact factor model (EFM) assumes in addition that Σ_ϵ is a diagonal matrix. In contrast, an approximate factor model (AFM) only assumes that Σ_ϵ is matrix with bounded L^1 or L^2 norm. De Nard, Ledoit, and Wolf (2021) and De Nard and Zhao (2023) show that the EFM assumption is too strict and suggest to use a nonlinear shrinkage estimator in conjunction with a multivariate GARCH model, namely the

DCC-NL model of Engle, Ledoit, and Wolf (2019). However, due to the small sample size, DCC and nonlinear shrinkage are suboptimal in this case. This is why we use the (linear) constant-variance-covariance (CVC) shrinkage estimator of De Nard (2022) instead, which is known to perform well in this challenging and small-sample setting. The intuition of CVC is to shrink the sample variances s_i^2 and sample covariances s_{ij} toward their grand mean:

$$\hat{\Sigma}_{CVC}^* = \hat{\delta}^* (\bar{s}^2 I + \bar{s}_{ij} J) + (1 - \hat{\delta}^*) S, \quad (\text{B.3})$$

where $\hat{\delta}^*$ is the optimal shrinkage intensity derived in De Nard (2022), S is the sample

covariance matrix, and $J := \mathbf{1}\mathbf{1}' - I$ is the off-diagonal matrix.

Note that we use CVC (B.3) to estimate both Σ_f and Σ_e , following the suggestion and findings of De Nard and Zhao (2023).

Another modification of our covariance matrix estimator is that we use investable risk factors, as explained in “Selection of the Investment and Factor Universe” instead of (academic) Fama–French factors or latent/unobservable factors. Finally, we use adaptive least squares regressions to estimate (B.2); see Beck, De Nard, and Wolf (2023).