Hedging Climate Risk: a Two-Step Factor Mimicking Approach[★]

Clément Aymard[*]

Abstract

This paper proposes a flexible, data-driven approach to hedging climate risk through a two-step factor mimicking portfolio (FMP) framework. In the first step, asset exposures to climate risk (climate betas) are estimated using data at various frequencies and across multiple climate indices. In the second step, climate risk target series are projected onto the returns of beta-sorted portfolios. This approach mitigates the weak factor and short-sample problems that typically makes climate risk hedging a challenging task. Unlike prior methods that often rely on narrative-driven constructions, this framework depends entirely on past observed return comovements. Out-of-sample results demonstrate substantially improved hedging performance relative to existing FMP approaches and comparable performance to more complex, non-FMP methods, while offering greater simplicity and replicability. The framework is adaptable across asset universes and can accommodate various investment constraints.

Keywords: Climate Change, Hedging, Factor-mimicking, Sustainable Investing
JEL: G11, Q54, C58

Preliminary draft

[★]I am grateful to the Group for Research in Decision Analysis (GERAD) and Hydro-Quebec for their financial supports.

[*]HEC Montréal, Montréal, Canada. *clement.aymard@hec.ca*.

1. Introduction

Climate change poses an increasingly systemic risk to the global economy, and financial markets are beginning to reflect this reality. Investors now recognize that climate-related risks—both physical and transitional—can materially affect asset returns. Yet, hedging this new type of risk remains fundamentally challenging. Traditional insurance contracts are ill-suited to this task, as climate risk is not a rare event but a near certainty over the long run, making insurers reluctant to offer coverage. As a result, one practical alternative is self-insurance through strategic asset allocation—favoring assets that are expected to appreciate or remain resilient during periods of heightened climate risk. However, due in particular to the well-known weak factor problem—when test assets exhibit limited exposure to the hedging target (see, *e.g.*, Giglio et al. (2023b) and Pukthuanthong et al. (2022))—identifying effective hedge portfolios is particularly difficult. This paper proposes a new, flexible, data-driven approach to constructing hedge portfolios that allow to construct effective hedge portfolios.

The method builds on the traditional factor mimicking portfolio (FMP) approach pioneered by Huberman et al. (1987) (see also Breeden et al. (1989) and Lamont (2001)), which aims to construct hedge portfolios based on historical comovements between asset returns and a given risk factor. While commonly used for hedging macroeconomic risks, the standard FMP method typically projects the target series onto the returns of base assets, selecting those with the strongest estimated loadings to form the hedge portfolio[1]. When applied to climate risk, however, the FMP method often performs poorly, for two main reasons. First, as mentioned above, climate risk is a weak factor: asset returns show limited sensitivity to its fluctuations. This low signal-to-noise ratio inflates estimation error in factor exposures and degrades out-of-sample hedging performance. Second, as Alekseev et al. (2022) highlight, the FMP method tends to be "*very sensitive to the availability of time-series data, and suffers when the time series is short*". Since climate-related risks only began receiving significant investor attention in the early 2010s, the

[1] See, *e.g.*, Giglio and Xiu (2021), Giglio et al. (2022), Giglio et al. (2023a), Giglio et al. (2023b), Jurczenko and Teiletche (2022), Pukthuanthong et al. (2022).

(effective and usable) sample length is at most 25 years.[2] This paper introduces a two-step framework that addresses both of these limitations.

In the first step, I estimate the exposures of base assets to climate risk—so-called climate betas—through simple contemporaneous univariate regressions of asset returns on climate risk indices[3]. A key innovation at this stage is the use of higher-frequency data. Previous studies have been constrained to monthly or quarterly frequencies due to data limitations: for example, Engle et al. (2020) use monthly indices, while Alekseev et al. (2022) and Cao et al. (2024) rely on quarterly data from fund managers' trading decisions or earnings calls. By contrast, I exploit five climate risk indices available at the daily level—the Media Climate Change Concern (MCCC) index of Ardia et al. (2023), along with the International Summit, Global Warming, Natural Disaster, and U.S. Climate Policy indices of Faccini et al. (2023). Importantly, the only input required on the asset side is basic return data, which is readily available at high frequencies. Operating in a higher-frequency data environment offers two key advantages. First, it increases the number of observations, thereby improving the statistical accuracy of estimated exposures. Second, it allows the capture of potentially distinct temporal dynamics between asset returns and climate risk at different frequencies. This opens the door to an important empirical question addressed in the paper: whether combining climate betas estimated at different frequencies can enhance the overall hedging effectiveness.

In the second step, I integrate these climate betas directly into the FMP framework. Specifically, I project the monthly climate risk series onto the returns of portfolios formed using beta-sorted base assets from the prior month. This reduces the dimensionality of the estimation problem, which mitigates the limited time-series data issue discussed above. For example, projecting the climate series onto thousands of individual assets would almost certainly lead to infeasability to the curse of dimensionality problem ($N > T$). By constructing a small number of portfolios (*e.g.*, $N = 3$), each composed of these thousands

[2] See also footnote 2 in Alekseev et al. (2022): "*Prior to 2010, climate risks were hardly incorporated into market prices and likely did not affect investor behavior, making all of these approaches difficult to implement.*"

[3] For alternative methods for estimating climate betas, see *e.g.*, Sautner et al. (2023b), Sautner et al. (2023a) or Li et al. (2023)

of base assets, I can circumvent the issue. Not that this approach is not in itself a novelty of this paper. In fact, it aligns perfectly with portfolio-based regression framework of Engle et al. (2020). However, it is a fundamentally different way to implement it: while Engle et al. (2020) form their portfolios by sorting assets based on green score characteristics, I use a purely data-driven characteristic—the climate betas estimated in the first step.

A central goal of this paper is to avoid narrative-driven portfolio construction and rely instead on historical observable comovement patterns, that is, remain as much data-driven as possible. Many prior studies assume a narrative-based construction that determines ex ante what "should" proxy for climate risk, typically arguing that greener assets are better hedges against climate risk. For instance, Engle et al. (2020) construct portfolios based on green ratings, drawing on models studying investor preferences for sustainability (*e.g.*, Andersson et al. (2016), Bolton and Kacperczyk (2021), Pàstor et al. (2021), Pedersen et al. (2021)). In contrast, this paper makes no ex ante assumption about which assets should hedge climate risk.

A different approach is pursued by Alekseev et al. (2022), who examine how mutual fund managers reallocate portfolios in response to idiosyncratic shifts in their beliefs about climate risk. They show that a portfolio long in stocks bought and short in stocks sold by such managers appreciates during negative aggregate climate news shocks. This method delivers strong hedging performance, outperforming traditional FMP approaches that they also implement for comparisons. More recently, De Nard et al. (2024) use an FMP framework with sustainable funds, estimating exposures based on returns projected onto a climate risk index, then optimizing the hedge portfolio with a variance minimization objective. In another study, Cao et al. (2024) construct hedge portfolios based on market reactions to climate discussions in earnings calls.

Hedging performance in this literature is generally assessed using out-of-sample correlations between hedge portfolio returns and AR(1) innovations of a given climate risk target series. For example, Engle et al. (2020) report a correlation of 0.17 using data from March 2012 to December 2016, and their Wall Street Journal Climate Change News index as hedging target. For the period between January 2015 and December 2019, Alekseev et al. (2022) find correlations of up to 0.37 with the MCCC index and similar or lower values with the four climate risk indices from Faccini et al. (2023). Using their own mimicking portfolio construction methods, they obtain correlations up to 0.18 for the MCCC index, and in the range $[0.03, 0.06]$ for the four indices of Faccini et al. (2023). In contrast to existing FMP-based methods, my two-step method generally yields higher out-of-sample correlations. Compared to more complex alternatives like that of Alekseev et al. (2022), I deliver comparable or superior performance, with the added advantage of simplicity and superior replicability—it does not rely on specific and hard-to-obtain inputs such as fund manager trades.

Furthermore, I also explore the role of the investment universe in determining hedging effectiveness. Given the weak factor problem, asset selection is critical: the goal is to find assets whose returns exhibit stable and measurable sensitivity to climate risk. Individual stocks may be too noisy, while more aggregated instruments—such as portfolios based on industries—may provide more reliable signals. To remain data-driven, I do not choose ex ante a given univrse but instead consider multiple candidates, from individual stocks (including about 2,000 assets) to General Industry Classification Standard (GICS) levels 1, 2 and 3 portfolios. I find that the individual stock universe consistently delivers superior hedging performance compared to the GICS-based universes, suggesting that selecting a sufficiently large universe is an important driver in building an effective hedge portfolio.

Finally, this framework is practical as it offers a high degree of flexibility. Investors can tailor implementation to their needs—choosing the frequency of estimation or applying their own investment universe. Moreover, while the baseline specifications of this framework return a zero-investment long-short portfolio, the method easily accommodates long-only or other weight-related constraints.

The remainder of the paper is structured as follows. Section 2 presents the two-step methodology, data, and empirical designs. Section 3 reports in-sample results across multiple asset universes and specifications. Section 4 evaluates the hedging performance out-of-sample.

2. Methodology, Empirical Settings and Data

This section presents the two-step framework developed to construct hedge portfolios for climate risk, and outlines the empirical design and datasets employed. The methodology builds on the classical factor mimicking portfolio (FMP) approach but introduces key innovations to address its limitations in the context of climate risk—particularly the weak factor problem and limited time-series availability.

2.1. General Framework

The framework proceeds in two steps. The first step estimates climate betas, which quantify each asset's sensitivity to climate risk. These betas are obtained through simple univariate time-series regressions of excess returns on the hedging target. Formally, for each base asset $i = \{1, 2, \ldots, N\}$, I estimate:

$$r_{i,t} = \beta_i CR_t + \sum_{k=1}^{K} \lambda_{i,k} CTRL_{i,k,t} + \epsilon_{i,t} \tag{1}$$

where $r_{i,t}$ denotes demeaned excess returns of asset i, CR_t is the (mean-zero) climate risk target series, and $CTRL_{i,k,t}$ are K control variables.

Importantly, to potentially improve statistical precision and capture richer comovement structures, I consider estimating these regressions at three different frequencies, daily, weekly, and monthly, yielding three sets of estimated betas: $\boldsymbol{\hat{\beta}^{d}}$, $\boldsymbol{\hat{\beta}^{w}}$ and $\boldsymbol{\hat{\beta}^{m}}$, respectively. In the baseline specification, I include the three Fama-French factors—market (MKT), size (SMB), and value (HML)—as controls.

The second step translates these estimated betas into portfolio weights. To do so, different paths are possible. In this paper, I follow the portfolio-based projection approach of Engle et al. (2020) where the hedging target is projected onto portfolios returns that have the original base assets as components. More precisely, I employ portfolios sorted by climate betas. Specifically, I estimate:

$$CR_t = \alpha + \sum_{f \in F} \delta_{\mathcal{F}_f} \boldsymbol{Z_{t-1}^{\beta^{\mathcal{F}_f}\prime}} \boldsymbol{r_t} + \sum_{k=1}^{K} \psi_k \boldsymbol{Z_{t-1}^{CTRL_k\prime}} \boldsymbol{r_t} + \epsilon_t \tag{2}$$

where $\boldsymbol{Z^X}$ denotes a cross-sectional transformation of variable $\boldsymbol{X}$ that has a mean of zero, and $\boldsymbol{r_t}$ is the vector of asset returns in month t. Therefore, the product $\boldsymbol{Z_{t-1}^{\beta^{\mathcal{F}_f}}}'\boldsymbol{r_t}$ is a scalar that represents the return on a zero-investment portfolio that tilts toward assets with high or low climate betas estimated at daily, weekly or monthly frequency $\mathcal{F}_f \in \{d, w, m\}$. The estimated coefficient $\delta_{\mathcal{F}_f}$ captures the sensitivity of the hedging target to this portfolio. Then, the final weights for base assets at time T are recovered by multiplying the transformed value of their climate betas at time T (a N vector) with $\hat{\delta}_{\mathcal{F}_f}$ (a scalar), that is, $\boldsymbol{\hat{w}_T^{\mathcal{F}_f}} = \boldsymbol{Z_T^{\beta^{\mathcal{F}_f}}}\hat{\delta}_{\mathcal{F}_f}$, and summing across selected frequencies yields the full weight vector: $\boldsymbol{\hat{w}_T} = \sum_{f \in F} \boldsymbol{\hat{w}_T^{\mathcal{F}_f}}$. Because by construction the sum of the elements of $\boldsymbol{Z_T^{\beta^{\mathcal{F}_f}}}$ is always zero, the final weights $\boldsymbol{\hat{w}_T}$ also sum to zero. Stated differently, the hedge portfolio is a long-short zero-investment portfolio.

As with the first step, the investor can flexibly select control characteristics for inclusion. In the baseline case, I follow Engle et al. (2020) and include cross-sectional transformations of the *SIZE* and *Book-to-Market (BM)* characteristics, computed from the distribution across base assets.[4]

While this framework draws on the structure of the FMP method of Engle et al. (2020), its implementation here departs fundamentally from their narrative-driven approach. Rather than assuming what specific characteristics (*e.g.*, green scores) *should* be the information to rely on when forming the FMP, I adopt a fully data-driven strategy that lets historical return comovements speak for themselves.

2.2. Empirical Settings and Data

The general framework described above offers considerable flexibility. In practical applications, the hedging investor has multiple degrees of freedom in how the method is implemented. This section outlines the main empirical choices made for this paper, as well as the data sources and processing steps involved.

[4] Note that these are *not* the same controls as in regressions (1) as they are specific to the universe of N base assets. Specifically, $\boldsymbol{Z_{t-1}^{SIZE^k}}$ is constructed based on distribution of market capitalization of the N base assets, and $\boldsymbol{Z_{t-1}^{BM^k}}$ is constructed based on distribution of book-to-market of the N base assets.

2.2.1. Hedging Target

A first and crucial choice in the framework is the definition of the hedging target, denoted CR_t in the previous equations. I consider five climate risk indices as potential targets: the Media Change in Climate Concerns (MCCC) index from Ardia et al. (2019), and four indices developped by Faccini et al. (2023), namely the International Summit (INTS), Global Warming (GWAR), Natural Disaster (NATD) and U.S. Climate Policy (USCP) indices. All indices are available at the daily frequency. The MCCC index spans from January 2003 to December 2023, while the four indices of Faccini et al. (2023) cover January 2000 to June 2023. Following the litterature (*e.g.*, Alekseev et al. (2022) or Engle et al. (2020)) I define each hedging target series as the AR(1) innovations from the original index levels. Figure 1 illustrates the time series these five hedging targets starting in 2010. Since climate betas are estimated at three distinct frequencies—daily, weekly, and monthly—the hedging targets are also converted to these frequencies accordingly. As expected, lower-frequency versions of the target appear visually smoother and less noisy, providing an early indication that climate betas at different frequencies may capture distinct elements of the underlying dependence structure.

2.2.2. Investment Universe

A second key empirical decision concerns the choice of hedging instruments—that is, the universe of base assets used to construct the hedge portfolio. In the context of climate risk, this choice is particularly important due to the weak factor nature of the climate signal: exposures to climate risk are likely to be small, noisy, and heterogeneous across assets. The trade-off lies between choosing a broad universe to increase the chance of capturing relevant exposures, and narrowing the universe to reduce estimation noise. For example, using individual stocks offers broad coverage and increases the likelihood of including assets with meaningful climate exposure. However, climate risk is generally not a first-order driver of stock returns, and the resulting signal may be dominated by noise. Alternatively, using more aggregated portfolios—such as industry-level groupings—can help mitigate idiosyncratic noise and may align better with economic intuition, as sectors like Energy, Utilities, and Materials are often viewed as more climate-sensitive.

To maintain a data-driven and agnostic stance, I do not commit to a single universe ex ante. Instead, I conduct the analysis across four distinct asset universes:[5] (i) STOCK, a universe of individual equities, comprising the 2,000 largest CRSP-listed stocks (share codes 10 and 11) available at each portfolio formation date, (ii) GICS1, industry portfolios based on the first level of the Global Industry Classification Standard (GICS), covering 11 broad sectors, (iii) GICS2, a finer decomposition into 27 industry subgroups, and (iv) GICS3, an even more granular classification comprising 75 industry subgroups. All the industry portfolios are constructed using a value-weighting scheme.[6]

2.2.3. Other Considerations

Several additional empirical choices might affect the performance of the hedging strategy.

Estimation windows in regressions (1) *and* (2). In the baseline specifications, climate betas are estimated using a 36-month rolling window for regression (1). For the construction of the hedging portfolio in regression (2), I use the longest available period for the in-sample analysis, and a 36-month window for the out-of-sample analyses. As robustness checks, I also consider window lengths ranging from 12 to 60 months for daily betas, 24 to 60 months for weekly betas, and 36 to 60 months for monthly betas. Similarly for regression (2), out-of-sample window lengths from 24 to 60 months are tested. See section III of the Appendix.

Construction of $\boldsymbol{Z}$. A second empirical choice relates to the specification of the instrument vector $\boldsymbol{Z_{t-1}}$ used to weight base assets in regression (2). I consider two alternative constructions: standardized betas (z-scores) and sign-preserving betas. For standardized betas, each element of $\boldsymbol{Z^{\beta}_{t-1}}$ is computed as $\frac{\beta_{i,t-1}-\bar{\beta}_{t-1}}{\sigma(\beta_{t-1})}$. However, a potential issue with this z-score transformation is that, by assigning negative sign to the betas below the mean, it will potentially generate implicit short positions in assets that are nonetheless positively correlated with the climate risk target (that is, asset with positive climate betas). To

[5] Additional universes are considered in robustness checks; see section I of the Appendix.

[6] I could also consider portfolios sorted on environmental characteristics (*e.g.*, GHG emissions or ESG scores). I however refrain from using such portfolios as this would reintroduce narrative-driven choices into the analysis and somewhat go against the data-driven philosophy of the framework.

preserve directional exposure, I propose the sign-preserving approach where two separate portfolios are formed: one for assets with positive betas (β^+) and one for those with negative betas (β^-). For each portfolio, weights are proportional to the magnitude of the betas: $w_i^+ = \beta_i^+ / \sum_{i=1}^{N^+} \beta_i^+$ and $w_i^- = \beta_i^- / \sum_{i=1}^{N^-} \beta_i^-$. Then, the vector $\boldsymbol{Z_{t-1}}$ is constructed such that $Z_i = w_i^+$ if the asset i has a positive beta, and $Z_i = -w_i^-$ if the asset i has a negative beta. This approach maintains the sign of exposures and ensures that positively and negatively correlated assets contribute accordingly.

Portfolio Constraints. While the baseline hedging portfolio is constructed as a long-short, zero-investment portfolio, additional constraints can be introduced for practical implementation. For example, a long-only, fully invested portfolio can be formed by setting negative weights to zero and rescaling the remaining weights to sum to one. Similarly, constraints such as bounds on short positions can be imposed through appropriate truncation and normalization of the weights.

Control Variables. Finally, the choice of control variables in the regression models can influence the hedging performance. In the baseline specification, I control for the three Fama and French (1993) factors in regressions (1), and the *SIZE* and *BM* characteristics for regression (2). Additional specifications with alternative control sets are examined in section III of the Appendix.

3. In-Sample Analyses

3.1. Climate Betas

I begin by analyzing results from step (1) of the framework. The analysis covers the period from January 2010 to December 2023, using a 36-month rolling estimation window. I consider the four investment universes introduced earlier—STOCK, GICS1, GICS2, and GICS3—and five hedging targets: the AR(1) innovations of the MCCC, INTS, GWAR, NATD, and USCP indices. For each combination of universe and hedging target, I estimate climate betas at daily, weekly, and monthly frequencies, generating a time series of beta distributions spanning December 2012 to December 2023.

Figure 2 summarizes the climate beta estimates for the MCCC index. Several observations emerge. First, across all universes, most climate betas are not statistically different from zero, consistent with the weak factor problem discussed earlier: it is difficult to find base assets with strong, systematic exposure to climate risk. Nonetheless, as indicated by the blue dots in the figure, a non-negligible subset of base assets exhibits statistically significant betas, suggesting that pockets of meaningful exposure may exist—though their relevance for hedging remains an open question. Second, climate beta estimates appear to be frequency-dependent. For instance, in the STOCK universe, daily betas generally lie within the range of $[-0.02, 0.02]$, while weekly betas extend to approximately $[-0.2, 0.2]$, and monthly betas show a much wider dispersion, spanning approximately $[-1, 4]$. This scaling with frequency may suggest that lower-frequency betas capture longer-horizon, more persistent patterns of dependence between base assets and climate risk. However, at this stage, the implications for hedging performance remain to be established. To investigate further, I examine the cross-frequency correlation structure of the estimated climate betas. Specifically, for each frequency $\mathcal{F}_f$ I form a $T \times N$ matrix of estimated climate betas $\boldsymbol{\beta}^{\mathcal{F}_f}$, and compute row-wise cross-sectional correlations between frequencies—namely Daily-Weekly, Daily-Monthly, and Weekly-Monthly. Figure 3 presents the resulting distributions. Interestingly, while the correlations are generally positive, they rarely exceed 0.50 in the STOCK universe. This suggests that climate betas at different frequencies may be capturing distinct features of the climate-risk–asset-return relationship. In the industry-based universes (GICS1–3), correlations exhibit wider dispersion, and in many cases the first quartile (lower hinge) falls below zero, indicating substantial heterogeneity in frequency-specific exposures across sectors. These results tend do reinforce the idea that beta estimates at different frequencies encode different facets of the underlying dependence structure.

3.2. Second-step Estimation Results

With climate betas in hand, I now turn to the second step of the framework and estimate regression (2). This step maps climate betas into weights used to construct the hedging portfolio. To examine the role of frequency, I consider seven models using different combinations of betas: three models using a single frequency (daily, weekly, or monthly), three models using two frequencies, and one model using all three frequencies. In all cases, the instrument vectors $\boldsymbol{Z}_{t-1}^{\beta^{\mathcal{F}_f}}$ are constructed using the sign-preserving method described earlier. As in Step 1, the dependent variables are AR(1) innovations of the five climate indices. The estimation period extends from January 2013 to December 2023 for the MCCC index, and from January 2013 to June 2023 for the four indices of Faccini et al. (2023). Control variables $\boldsymbol{Z}_{t-1}^{\boldsymbol{CTRL_k}}$ include size and book-to-market characteristics, standardized using z-scores.[7]

Table 1 reports the estimated coefficients for the MCCC hedging target; results for other indices are presented in section II of the Appendix. Several findings are notable. Across all universes, the estimated coefficients $\delta_{\mathcal{F}_f}$ for daily and weekly climate betas are consistently negative and frequently statistically significant. For instance, in the STOCK universe, model (1)—which includes only daily betas—yields an estimate of -1.57, significant at 10% level. In model (4), which includes both daily and weekly betas, the coefficients are -0.57 and -2.44, respectively. In contrast, the estimated coefficients for monthly betas are always positive and typically significant. These results seem to suggest a frequency-dependent reversal in the sign of optimal weights: base assets with *negative* daily or weekly betas should be *overweighted*, while those with *positive* monthly betas should also be *overweighted*. In other words, high-frequency betas exhibit a "reversal" pattern, whereas low-frequency betas follow a "momentum" pattern.

4. Out-of-sample results

Ultimately, the effectiveness of any hedging strategy must be judged by its out-of-sample performance—how well it mitigates realized climate risk after the portfolio is formed. To assess this, I implement a 36-month rolling window procedure for estimating regression (2), starting in December 2015. Specifically, the first estimation corresponds to what the hedging investor would do at the end of December 2015: she would form her portfolio based on past observations from January 2013 to December 2015, and hold

[7] Control characteristics are based on the constituents of each universe. For STOCK, controls are computed across the 2,000 stocks. For the GICS-based universes, controls are computed from the underlying securities forming the portfolios. See Appendix III.

it during the following month. Then, at the end of January 2016, she would form her portfolio based on past observations from February 2013 to January 2016, and hold it during the following month. The process is repeated until the sample is exhausted, that is, November 2023. To facilitate comparison with the study of Alekseev et al. (2022), I also consider a shorter evaluation window ending in December 2019.[8]

As for the full-sample analysis, I consider seven models based on varying combinations of climate betas (daily, weekly, monthly), and the five hedging targets derived from AR(1) innovations of the MCCC, INTS, GWAR, NATD, and USCP indices. The instrument vectors $\boldsymbol{Z}$ are constructed using the sign-preserving approach, and the control variables are standardized z-scores for the size and book-to-market characteristics.

To measure hedging performance, I rely on Pearson correlations between monthly hedge portfolio returns and the realized monthly values of the hedging targets. Table 2 reports these out-of-sample correlations. Focusing first on the 2016–2019 period—approximately aligned to the evaluation period in Alekseev et al. (2022)—performance under my two-step framework is broadly comparable to theirs. For instance, with respect to the MCCC target, their best-performing specification (*"Pooled: All Shocks"*, see their table 9) yields a correlation of 0.37, while my framework produces an average correlation of 0.36 across the seven model variants. For the NATD target, their correlation is 0.06 versus an average of 0.10 in my framework. For the USCP target, both approaches yield near-zero correlations. For the GWAR and INTS targets, however, my approach underperforms their best performing method. On the other hand, when benchmarked against their factor-mimicking most sophisticated method—*"Lasso Reg: All-Industries + FF"*—the two-step framework introduced here often delivers superior performance. Specifically, for MCCC and NATD, Alekseev et al. (2022) report correlations of -0.16 and -0.07, respectively, whereas my method returns 0.36 and 0.10.

These comparisons underscore two key advantages of the two-step framework proposed here: (i) superior performance relative to other factor-mimicking portfolio approaches,

[8]Note that Alekseev et al. (2022) use data from 2015 to 2019. Due to my reliance on a 36-month estimation window for both Step (1) and Step (2), my earliest possible evaluation period begins in January 2016. In section III of the Appendix, I examine results based on a 24-month window for 2015 as a robustness check.

and (ii) competitive performance relative to more sophisticated, non-FMP methods, but with significantly greater transparency, simplicity, and replicability, since it avoids proprietary or hard-to-obtain data sources.

Extending the evaluation period to December 2023 yields similar qualitative patterns, albeit with somewhat reduced magnitudes. The decline in performance may be partly attributable to structural changes in asset return dynamics or climate sentiment during this period—most notably, the COVID-19 shock in 2020, which may have introduced noise or disrupted the usual link between climate risks and financial markets.

Lastly, comparing performance across investment universes reveals a clear result: individual stocks outperform industry-level portfolios in this framework. Specifically, the STOCK universe consistently delivers higher out-of-sample correlations, whereas the GICS-based universes often yield low or even negative correlations. This finding highlights the importance of selecting a sufficiently high number of base assets. Portfolios built from broad industry aggregates appear too coarse to capture the nuanced exposures necessary for effective climate risk hedging.

References

Alekseev, G., Giglio, S., Maingi, Q., Selgrad, J., Stroebel, J., 2022. A quantity-based approach to constructing climate risk hedge portfolios. Working paper. doi:10.3386/w30703.

Andersson, M., Bolton, P., Samama, F., 2016. Hedging climate risk. Financial Analyst Journal 72, 13–32.

Ardia, D., Bluteau, K., Boudt, K., 2019. Questioning the news about economic growth: Sparse forecasting using thousands of news-based sentiment values. International Journal of Forecasting 35, 1370–1386. doi:10.1016/j.ijforecast.2018.10.010.

Ardia, D., Bluteau, K., Boudt, K., Inghelbrecht, K., 2023. Climate change concerns and the performance of green vs. brown stocks. Management Science 69, 7607–7632.

Bolton, P., Kacperczyk, M., 2021. Do investors care about carbon risk. Journal of Financial Economics 142, 517–549.

Breeden, D., Gibbons, M., , Litzenberger, R., 1989. Empirical tests of the consumption-oriented CAPM. Journal of Finance 44, 231–262.

Cao, Y., Liu, M., Zhang, R.X., 2024. The risk sharing value of disclosure: A real-time market response approach to hedge climate change risk. Working paper. doi:10.2139/ssrn.4847914.

De Nard, G., Engle, R.F., Kelly, B., 2024. Factor-mimicking portfolios for climate risk. Financial Analyst Journal 80, 37–58.

Engle, R.F., Giglio, S., Kelly, B., Lee, H., Stroebel, J., 2020. Hedging climate change news. Review of Financial Studies 33, 1184–1216.

Faccini, R., Matin, R., Skiadopoulos, G., 2023. Dissecting climate risks: Are they reflected in stock prices? Journal of Banking and Finance 155, 106948.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Giglio, S., Kelly, B., Xiu, D., 2022. Factor models, machine learning, and asset pricing. Annual Review of Financial Economics 14, 337–68.

Giglio, S., Xiu, D., 2021. Asset pricing with omitted factors. Journal of Political Economy 128, 1947–1990.

Giglio, S., Xiu, D., Zhang, D., 2023a. Prediction when factors are weak. Working paper. doi:10.2139/ssrn.4397485.

Giglio, S., Xiu, D., Zhang, D., 2023b. Test assets and weak factors. Journal of Finance (Forthcoming) doi:10.2139/ssrn.3768081.

Huberman, G., Kandel, S., Stambaugh, R., 1987. Mimicking portfolios and exact arbitrage pricing. Journal of Finance 42, 1–9.

Jurczenko, E., Teiletche, J., 2022. Macro factor-mimicking portfolios. Working paper. doi:10.2139/ssrn.3363598.

Lamont, O., 2001. Economic tracking portfolios. Journal of Econometrics 105, 161–184.

Li, Q., Shan, H., Tang, Y., Yao, V., 2023. Corporate climate risk: Measurements and responses. Review of Financial Studies 37, 1778–1830.

Pedersen, L.H., Fitzgibbons, S., Pomorski, L., 2021. Responsible investing: The ESG-efficient frontier. Journal of Financial Economics 142, 572–597.

Pukthuanthong, K., Roll, R.W., Wang, J.L., Zhang, T., 2022. A new method for factor-mimicking portfolio construction. Working paper. doi:10.2139/ssrn.3341604.

Pàstor, L., Stambaugh, R.F., Taylor, L.A., 2021. Sustainable investing in equilibrium. Journal of Financial Economics 142, 550–571.

Sautner, Z., van Lent, L., Vilkov, G., Zhang, R., 2023a. Firm-level climate change exposure. Journal of Finance 78, 1449–1498.

Sautner, Z., van Lent, L., Vilkov, G., Zhang, R., 2023b. Pricing climate change exposure. Management Science 69, 7540–7561.

Table 1: Full Sample Estimation Results

This table shows full sample results for my two-step framework for four different universes and seven specifications depending on the number of climate betas included as regressors (columns). In all cases, the hedging target corresponds to AR(1) innovations of the MCCC index. In step (1) the climate betas are estimated with the three Fama and French (1993) factors as controls and a 36 months estimation window. For step (2), all seven models include a constant and the control variables *size* and *book-to-market* (not reported), and the estimation period is from January 2013 to December 2023 (T = 132 months). Bold values indicate statistical significance at the 10% level or less.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Individual Stocks (STOCK)							
$Z_{t-1}^{\beta^{d}\,\prime} r_t$	**−1.57**			−0.57	−1.38		−0.22
$Z_{t-1}^{\beta^{w}\,\prime} r_t$		**−2.67**		**−2.44**		**−2.78**	**−2.66**
$Z_{t-1}^{\beta^{m}\,\prime} r_t$			1.08		0.89	**1.21**	**1.18**
Adj.R^2	.03	.07	.03	.08	.05	.10	.10
GICS Level 1 Portfolios (GICS1)							
$Z_{t-1}^{\beta^{d}\,\prime} r_t$	−0.66			−0.24	−0.71		−0.30
$Z_{t-1}^{\beta^{w}\,\prime} r_t$		**−1.02**		**−0.92**		**−1.01**	**−0.89**
$Z_{t-1}^{\beta^{m}\,\prime} r_t$			0.44		0.51	0.41	0.44
Adj.R^2	.02	.05	.02	.05	.03	.05	.05
GICS Level 2 Portfolios (GICS2)							
$Z_{t-1}^{\beta^{d}\,\prime} r_t$	**−1.14**			**−1.13**	−0.72		−0.74
$Z_{t-1}^{\beta^{w}\,\prime} r_t$		−0.67		−0.01		−0.25	0.10
$Z_{t-1}^{\beta^{m}\,\prime} r_t$			**1.19**		**0.91**	**1.16**	**0.91**
Adj.R^2	.05	.02	.07	.05	.08	.07	.08
GICS Level 3 Portfolios (GICS3)							
$Z_{t-1}^{\beta^{d}\,\prime} r_t$	**−1.27**			−0.83	**−1.04**		−0.73
$Z_{t-1}^{\beta^{w}\,\prime} r_t$		**−1.96**		−1.49		**−1.51**	−1.11
$Z_{t-1}^{\beta^{m}\,\prime} r_t$			**1.86**		**1.70**	**1.60**	**1.56**
Adj.R^2	.04	.05	.07	.06	.09	.09	.10

Table 2: Out-of-Sample Hedging Performance
This table presents out-of-sample Pearson correlations between the monthly excess returns of various hedge portfolios and various hedging targets. Each column within panels represents a specification of regression (2) that incorporates one, two, or three frequency-specific climate betas obtained from regressions (1). To facilitate comparisons, in addition to the main evaluation period (2016-2023), I report results for the 2016-2019 period, corresponding approximately to the period in Alekseev et al. (2022).

	Panel A: MCCC Index							**Panel B: International Summit Index**						
Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	d	w	m	d,w	d,m	w,m	d,w,m	d	w	m	d,w	d,m	w,m	d,w,m
Evaluation Period: 2016-2019 (≈ Alekseev et al., 2022)														
STOCK	.25	.42	.15	.42	.33	.47	.46	.06	−.30	.12	−.04	.14	.09	.09
GICS1	−.19	−.10	.15	−.13	.14	.04	.03	−.02	.04	.14	.04	.13	.06	.04
GICS2	.04	−.22	.03	−.10	.18	−.14	.09	−.12	−.43	.12	−.40	.02	−.28	−.34
GICS3	−.21	.07	.12	−.13	.05	.15	−.03	−.08	.11	−.12	.10	.00	.02	.04
Evaluation Period: 2016-2023														
STOCK	.01	.24	−.02	.22	.01	.24	.21	.09	−.16	.06	.00	.11	.02	.06
GICS1	−.08	.01	.07	−.06	.06	.05	.02	−.02	.09	.05	.07	.04	.08	.06
GICS2	.06	−.16	.18	−.11	.14	.04	.01	−.11	−.23	.05	−.25	.02	−.08	−.14
GICS3	.01	.12	.24	.05	.19	.21	.15	−.07	.07	−.05	.04	−.03	−.00	00

	Panel C: Global Warming Index							**Panel D: Natural Disaster Index**						
Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	d	w	m	d,w	d,m	w,m	d,w,m	d	w	m	d,w	d,m	w,m	d,w,m
Evaluation Period: 2016-2019 (≈ Alekseev et al., 2022)														
STOCK	−.04	−.41	.24	−.29	.20	.09	.07	.09	.15	−.05	.16	.04	.13	.15
GICS1	.02	−.12	−.24	−.07	−.05	−.22	−.13	−.13	.11	−.09	.03	−.20	.03	−.05
GICS2	.03	−.02	.08	.03	.06	.05	.07	−.20	−.03	−.08	−.10	−.21	−.13	−.20
GICS3	−.05	.07	.01	.00	−.05	.06	−.00	−.12	.05	−.04	−.03	−.20	−.01	−.09
Evaluation Period: 2016-2023														
STOCK	−.05	.00	−.02	.05	−.03	.00	.04	−.05	.11	−.11	.07	−.12	.03	.01
GICS1	−.16	−.01	−.03	−.15	−.10	−.07	−.12	.06	−.08	−.04	−.05	.01	−.07	−.04
GICS2	−.02	−.19	−.01	−.06	−.05	−.14	−.10	−.09	−.10	−.26	−.11	−.19	−.07	−.10
GICS3	−.17	−.05	−.24	−.11	−.29	−.12	−.18	−.11	−.13	−.00	−.19	−.15	−.05	−.15

	Panel E: U.S. Climate Policy Index						
Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	d	w	m	d,w	d,m	w,m	d,w,m
Evaluation Period: 2016-2019 (≈ Alekseev et al., 2022)							
STOCK	.05	−.21	.12	−.04	.12	−.07	.02
GICS1	−.08	−.06	−.05	−.14	−.03	−.10	−.16
GICS2	−.10	−.14	.06	−.15	−.11	.01	−.14
GICS3	.07	.18	.16	.10	.13	.18	.14
Evaluation Period: 2016-2023							
STOCK	−.10	−.03	−.11	−.13	−.16	−.20	−.27
GICS1	−.33	−.17	−.12	−.22	−.33	−.22	−.28
GICS2	−.17	.11	.10	.04	.02	.04	−.04
GICS3	.08	−.12	.00	.02	−.00	−.06	−.07

Figure 1: Hedging Targets Series

This figure displays the various hedging target series used in this paper. All series correspond to AR(1) innovations of a given climate index. In each panel, the left, mid and right plots correspond to daily, weekly and monthly frequency, respectively.

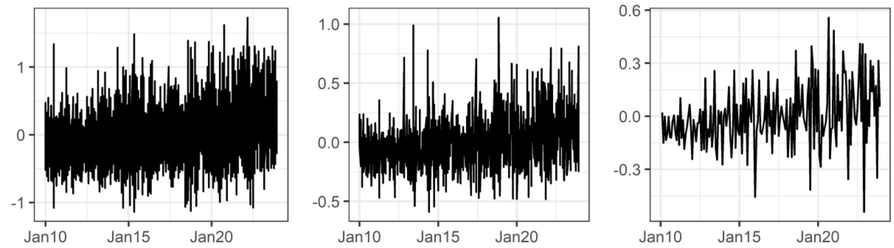

(a) MCCC Index

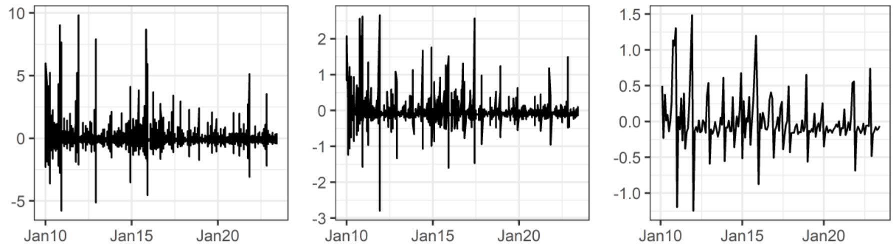

(b) International Summit Index

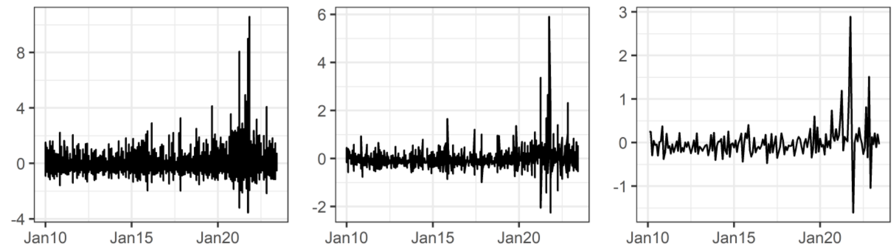

(c) Global Warming Index

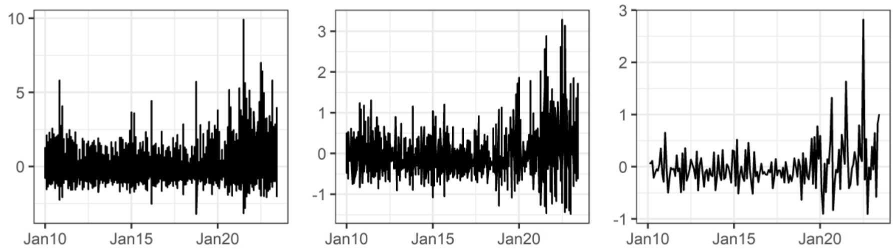

(d) Natural Disaster Index

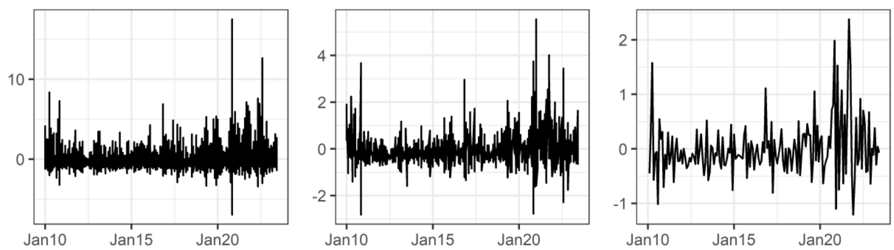

(e) U.S. Climate Policy Index

Figure 2: Climate Betas Distributions

This figure displays the distributions of climate betas for the STOCK, GICS1, GICS2 and GICS3 universes, where the hedging target series in regressions (1) corresponds to AR(1) innovations of the MCCC Index. In each panel, the left, mid and right plots correspond to daily, weekly and monthly climate betas, respectively. Blue dots indicate significance at the 10% level or less.

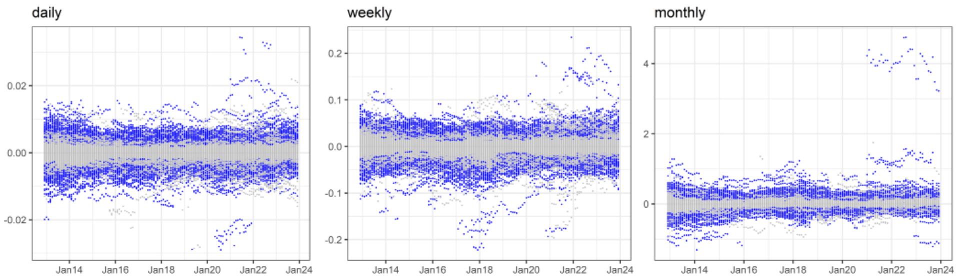

(a) **Individual Stocks (STOCK)**

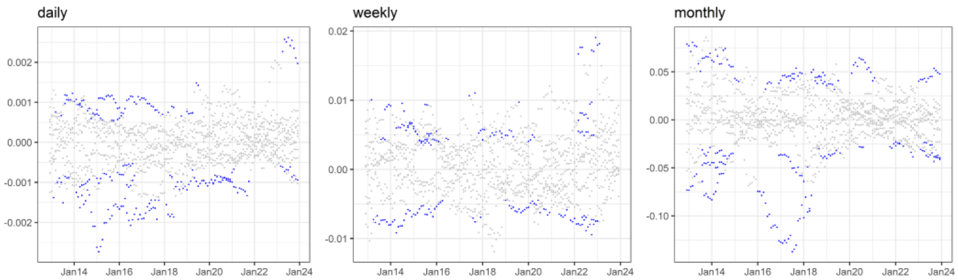

(b) **GICS Level 1 Portfolios (GICS1)**

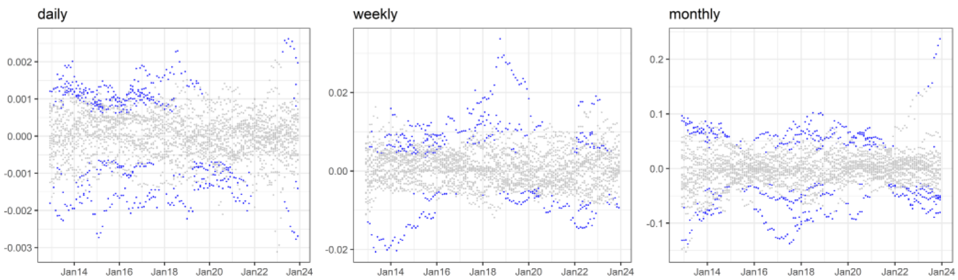

(c) **GICS Level 2 Portfolios (GICS2)**

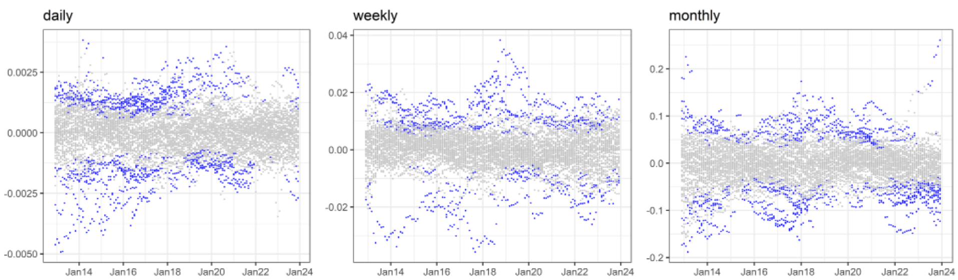

(d) **GICS Level 3 Portfolios (GICS3)**

Figure 3: Climate Betas Frequency Pairwise Correlations
This figure displays the distributions of the frequency-pairwise—Daily-Weekly, Daily-Monthly and Weekly-Monthly—cross-sectional correlations of climate betas for the STOCK, GICS1, GICS2 and GICS3 universes and various hedging targets. For the MCCC index (all other indices) the correlations are based on $T = 133$ months from January 2016 to December 2023 ($T = 127$ months from January 2016 to June 2023). The red points indicate the means.

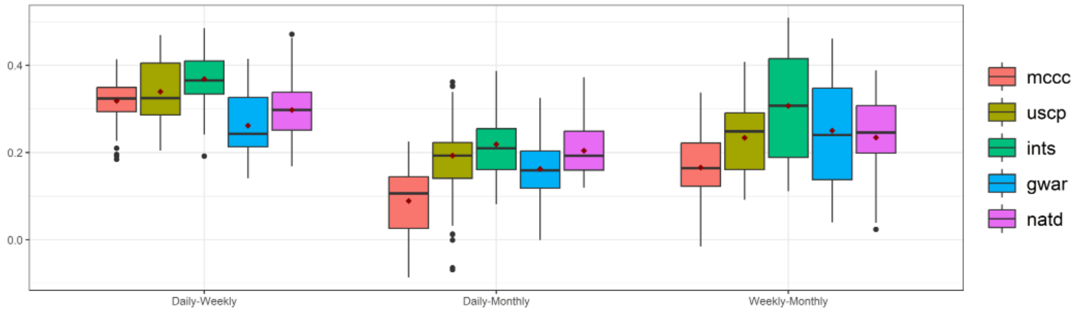

(a) **Individual Stocks (STOCK)**

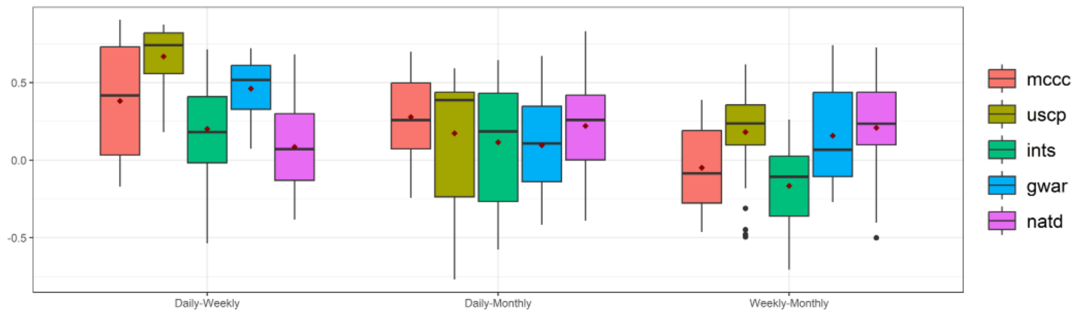

(b) **GICS Level 1 Portfolios (GICS1)**

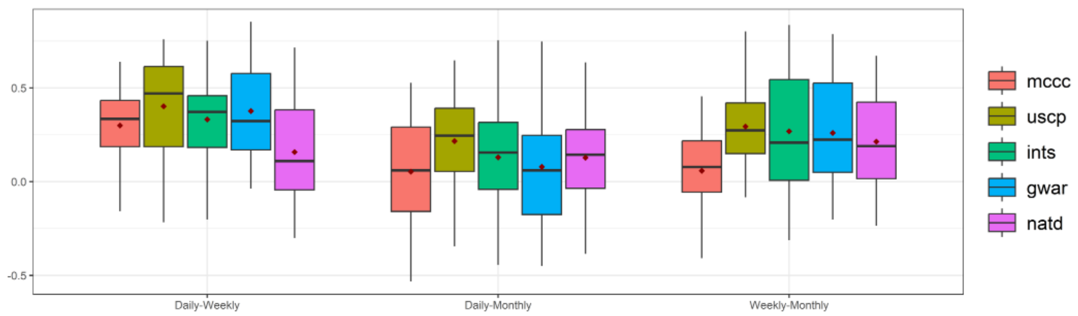

(c) **GICS Level 2 Portfolios (GICS2)**

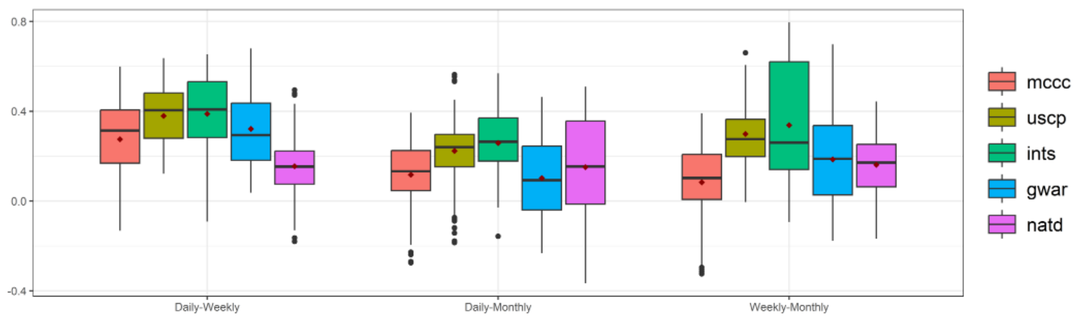

(d) **GICS Level 3 Portfolios (GICS3)**