The impact of climate change risk on long-term asset allocation

Jean-Charles Bertrand* Guillaume Coqueret†
Nicholas McLoughlin‡ Stéphane Mesnard§

December 2023

Abstract

We propose a framework for long-term cross-asset portfolio choice in which the estimation of the covariance matrix is subject to climate risk. We model the future volatility and correlation of assets as a linear function of three types of forward-looking variables: their long-term future average, climate-aware projections of economic indicators, and scenarios for the temperature anomaly. We analyze the shifts from a baseline 60/40 equity/bond allocation when taking climate risk into account. We find that these changes are small and mostly favorable to bonds if the focus is on the estimation of risk components. Including climate-driven expected returns in the optimization substantially alters the compositions, but this time to the benefit of equities. In all cases, the risk-adjusted returns decrease, often significantly, when taking climate impact into account.

1 Introduction

Climate change is now a ubiquitous concern in the asset management industry (Dietz et al. (2016), Bolton and Kacperczyk (2021)). Yet, there is no unambiguous way to take it into account in investment decisions and we refer to Chapter 5 of Coqueret (2022) for a review of sustainable portfolio compositions.

Most often, the question is tackled at a very granular level, with estimates and exposures for individual stocks and bonds, see, e.g., Bressan et al. (2022), Le Guenedal and Roncalli (2022) and Faiella et al. (2023)), and especially when it comes to portfolios that target net-zero emissions based on decarbonization paths (Barahhou et al. (2022)). Aggregate strategies at the asset class level are scarce in the literature[1] even though they are essential to asset managers. The present paper seeks to partly fill this void.

*HSBC Asset Management and HEC Paris, E-mail: jean-charles.bertrand@hsbc.fr
†EMLYON Business School, 23 avenue Guy de Collongue, 69130 Ecully, FRANCE. E-mail: coqueret@em-lyon.com
‡HSBC Asset Management, E-mail: nicholas.mcloughlin@hsbc.com
§HSBC Asset Management, E-mail: stephane.mesnard@hsbc.fr

The views expressed in this article are those of the authors alone and not of HSBC Asset Management.

[1] Hodges et al. (2022) being a notable exception.

Our approach differs from most contributions in the way it seeks to estimate the future components of the covariance matrix of the two asset classes. In the existing literature, which we review below, academics and practitioners aim to reveal *predictive* relationships between the dependent variables (risk, returns and correlations) and broad sets of predictors, including climate proxies. This is somewhat far-fetched, as current economic and climate conditions may not be indicative of future developments. The markets, the economy and the environment are synchronously linked, which is why we propose a model that binds risk to macro indicators in a *contemporaneous* manner. This is all the more relevant that we consider long horizons and that, until now, climate shocks have had little durable effects on prices and returns (Rebonato (2023)): most of the impacts are expected to occur in the future.

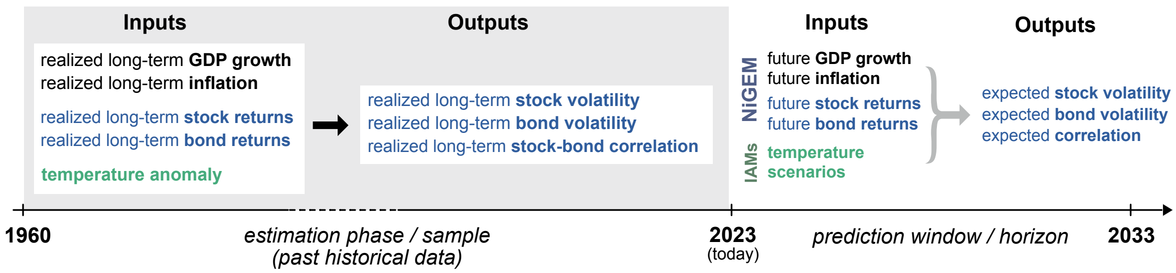

Figure 1: **Diagram of the proposed framework**.

Our main contributions are the following. First, we propose a framework that estimates asset class risks based on future projected and climate-based components such as GDP growth, inflation, and temperature anomaly. In doing so, we break with most approaches that rely on *predictive* specifications in which the risk ten years from now is inferred from *present* economic values. Furthermore, we choose to outsource the special task of economic forecasts to climate experts.

Our empirical results show that long-term volatility forecasts exhibit little dispersion and are concentrated around 14.5% and 7% for US stocks and bonds, respectively. Moreover and crucially, the effect of potential climate-caused losses is fairly similar across the two asset classes. With regards to long-term correlations, our estimates are mostly neutral and between zero and 10% - whereas historical values have shifted between times at -30% and periods at +30%.

As a consequence of this homogeneity in the impact of climate risks across asset classes, asset allocations that are driven by risk considerations only are not *too* impacted by climate shifts. A target allocation of 60% in equities that ignore climate concerns will be shifted by roughly 2-3%, mostly downwards. Indeed, a large majority of scenarios in this case are slightly beneficial to bonds. The marginal changes in optimal allocation come from the similar impact that climate damages have on both classes (moderately higher risk, on average).

If, in addition to risk, we integrate climate-driven expected returns in a mean-variance optimization, then the story is completely different. First, the 60/40 allocation can dramatically change, e.g., either towards 40/60, or towards 80/20, depending on the scenarios. In this case, most of the shifts are favorable to stocks because the climate impacts are more stringent on bond returns compared to equity returns.

Our framework is flexible and allows the shrinkage of the distributions of predicted risks and returns towards climate-agnostic priors. This helps assess the sensitivity of our results to changes in the inputs. We find our results are robust in the sense that large deviations from the 60/40 portfolio require substantial adjustments in the distributions of volatilities or expected returns. Moreover, models with quadratic impacts of temperature anomaly do not change forecasts much, compared to linear impacts.

We also investigate shifts for investors that are much more risk averse and who have an initial target allocation of 20% in equities. Our findings indicate that integrating climate risks into both risks and returns in this case has a milder impact, with most scenarios leading to a range of 15% to 25% for the equity pocket.

Lastly and crucially, a genuine concern is that of the impact of climate not only on optimal portfolio composition, but on performance. We document that climate risks are unambiguously detrimental to risk-adjusted returns. Accounting for potential physical and transition losses can reduce the latter by up to 50%.

The paper is structured as followed. Section 2 outlines the literature review. The details of our methodology as well as the data used in our empirical work are presented in Section 3. Our first empirical results (long-term models and predictions) are gathered in Section 4. Their impact on asset allocation is investigated in Section 5, and robustness check follow in Section 6. Finally, we conclude in Section 7.

2 Literature review

2.1 Variable dynamics

Handling uncertainty in financial and economic conditions is most often tackled by proposing dynamic models of state variables, which we write in vector form $\boldsymbol{x}_t$ at time t. The most general formulation of such models is as follows:

$$f(\boldsymbol{x}_t) = \sum_{k=1}^{K} g_k(\boldsymbol{x}_{t-k}) + \boldsymbol{e}_t. \tag{1}$$

As we will recall below, the most common form of such specifications is the vector autoregression (VAR) which sets $f(\boldsymbol{x}) = \boldsymbol{x}$, $K = 1$ and $g(\boldsymbol{x}) = \boldsymbol{A}\boldsymbol{x}$ for some matrix $\boldsymbol{A}$. In some contributions, the focus is set on only one dependent variable (i.e., one component of $\boldsymbol{x}_t$), for instance, asset return volatility and in this case the function f is simply a selection operator that discards the other variables. In this case, the model becomes a simple predictive regression

$$y_t = \sum_{k=1}^{K} g_k(\boldsymbol{x}_{t-k}) + \boldsymbol{e}_t, \tag{2}$$

which may or may not be linear and/or additive. The error terms $\boldsymbol{e}_t$ are usually assumed to follow the standard assumptions (i.i.d. Gaussian).

The most popular modelling choice is undoubtedly the **VAR model** and it commonly includes both financial and climate-related variables. For instance, Shen et al. (2019) combine 21 variables including temperature change and specify particular restrictions on their autoregressive links. Cosemans et al. (2022) focus on equities and use Bayesian estimations of a VAR model that combines returns, the log price-dividend ratio and the temperature anomaly (this is in fact just an adaptation of the framework of Avramov et al. (2018)).

They propose 4 different priors depending on investor types and they outline portfolio compositions that follow from an expected utility maximization routine. Bayesian VARs are also used in Prosperi and Zanin (2023) and the variables include the market and green (minus brown) factor returns, the carbon price, the short-term and long-term risk-free rate, one inflation swap and the oil price. The above contributions rely on such dynamics for the purpose of asset allocation. Financial returns are included in the list of state variables and their statistical properties (upon estimation of the models) are then exploited in utility-maximizing schemes to derive optimal portfolio compositions. Lastly, VAR models that include disaster instrumental variables are mentioned in Baker et al. (2023).

Examples of **predictive regressions** such as the one defined in Equation (2) are also abundant, especially when it comes to forecasting financial risks. For instance, Bonato et al. (2023a) forecast realized volatility in commodity exchange markets with climate-based predictors (temperature, precipitation and wind).[2] In a follow-up paper (Bonato et al. (2023b)), they specify a bagging predictive model based on 14 climate predictors that outperforms a benchmark prediction that relies on financial realized volatilities only. Temperature data is ubiquitous and also used in Salisu et al. (2022) to explain tail risks in S&P 500 returns (with extended GARCH models). Models that seek to predict future volatility with past volatility only (Cardinale et al. (2021)) or with macro-economic predictors (Engle et al. (2013)) are also well documented.[3] Lastly, there are also contributions that seek to explain the cross-section of returns, based on proxies climate risk and attention, and we refer for instance to the panel models of Bolton and Kacperczyk (2021), to Ardia et al. (2022) and to Le Tran et al. (2022).

Finally, it must also be underlined that VAR models can lead to expressions for conditional variance. For example, it is shown in Avramov et al. (2018) that the variance of future aggregate returns,[4] conditional on current data is the sum of four terms: the variance of i.i.d. innovations for returns, a mean reversion term, an uncertainty component about future mean returns and a proxy for estimation risk.

2.2 Climate scenarios and stress testing

The foundational assumption in climate-based asset pricing is that returns are dependent on some state variables that characterize the future or current climate situation. Large economic models, such as the one developed by the National Institute of Economic and Social Research are built to take into account links between the financial and economic sphere on one side and the environment on the other. Their National Institute Global Econometric Model (NiGEM) is commonly used by many institutions, including the Network for Greening the Financial System (NGFS) and thus many central banks. Climate scenarios within the model are inspired by IPCC reports and we refer to Hantzsche et al. (2018) for an overview of the modules upon which the model is built.

Another early attempt in this direction is the work of Bansal et al. (2017) in which the authors propose a temperature-augmented long-run risks (LRR-T) equilibrium model in

[2] See also Campos-Martins and Hendry (2023) for a focus of climate risks on the oil and gas industries.

[3] The topic of volatility forecasting is vast and has generated an extensive literature. We refer to Poon and Granger (2003) for a survey on mostly linear methods and to Ge et al. (2022) for a recent account of models that rely on neural networks.

[4] We underline that we are interested in aggregate risks. For individual asset risks, we point for instance to Bressan et al. (2023). This article leverages highly granular data on corporations in order to determine risks at the intra-firm asset level (e.g., for individual production plants).

which consumption generates emissions and temperature rises, which triggers random natural disasters.[5] Upon calibration, they find that "*a one degree Celsius increase in temperature leads to about -5% decline in equity valuations*". In a follow-up paper (Bansal et al. (2021)), the authors derive an expression for conditional expected excess returns of assets. It is driven by dividend beta and exposure to endogenous temperature variations.

Relatedly, Venturini (2022) surveys the literature that seeks to reveal links between asset prices or returns and climate variables (temperature and precipitation). This review shows the diversity in results because there are as many papers that conclude that climate risk should be (or is) priced as those that find the opposite. Another important survey is that of Campiglio et al. (2023). In their exhaustive summary, the authors discriminate between two types of studies. First, they compile the contributions that have sought to test if asset prices have been, in the past, affected by climate risks. Second, they focus on forward looking methods and models that propose prospective analyses.

This is also exactly the aim of climate scenarios and stress testing. This latter topic is almost exclusively focused on assessing the risks for the financial system. Acharya et al. (2023) discuss the current approaches that are used by regulators (central banks) to assess systemic risks stemming from climate-driven shocks, and their impact on the balance sheets of banks. Climate stress testing is similarly surveyed by Reinders et al. (2023b), in which the authors list six climate shocks, along with their propagation mechanisms in different modelling frameworks, from IAMs to macroeconomy-focused approaches, disaster models and valuation channels. The authors have another paper dedicated to this latter subject (Reinders et al. (2023a)). Lastly, on this topic, Fuchs et al. (2023) run causality tests (difference-in-difference) for which the treatment date is the launch of the French climate stress test. They report that "*stress tested banks charge higher interest rates for borrowers with high-transition risk*".

Beyond asset pricing, similar models are also proposed for macroeconomic quantities, such as output growth rates. For instance, Kahn et al. (2021) estimate an equation of the form (1) where the dependent variable is the GDP growth of countries. The model comes from an economic equilibrium and the results suggest a negative relationship between output growth and persistent extreme changes (positive or negative) in temperature. Similarly, Kalkuhl and Wenz (2020) run panel models of the form[6]

$$g_{t,i} = \alpha \Delta T_{t,i} + \beta T_{t,i} \Delta T_{t,i} + \gamma_1 T_{t,i} + \gamma_2 T_{t,i}^2 + X + e_{t,i}, \tag{3}$$

where $T_{t,i}$ includes temperature and precipitation levels of year t in region i, $g_{t,i}$ per capita growth rate of output and X includes control variables and fixed effects. They report that "*1°C temperature increase in a region with an annual mean temperature of 25°C, reduces gross regional product by about 3.5% in that region*". More recent estimates in Winter and Kiehl

[5] Their work is inspired from the literature on disaster risk. We for instance refer to Tsai and Wachter (2015) who summarize the contributions on asset pricing under disaster risk. The idea is anticipative here because the models are built on the premise that the shock comes from consumption and not from climate directly (though both could be connected of course). Disaster risk is included via a compound Poisson process within the diffusion of consumption and dividends. This allows the use of theoretical results on jump diffusion process to derive an analytical form for the equity premium in such models. In fact, the volatility in non-turbulent times can also be computed, as is shown in Wachter (2013). It links the instantaneous volatility with the intensity of occurrence of a disaster, but it is synchronous and not predictive.

[6] Zhao et al. (2018) propose a similar specification and discriminate between rich and poor areas in the i index of the equation. In a similar fashion, Newell et al. (2021) lay out a model with higher orders, i.e., that go beyond quadratic terms in temperature and precipitation.

(2023) indicate contractions of global output of 30% or more by 2050, while Casey et al. (2023) confirm a strong geographical heterogeneity of climate risks.

3 Method and data

3.1 Model specification

If we assume the position of an investor with mean-variance utility, we require long-term estimates of average returns, volatilities and correlations for portfolio allocation. There are at least three ways to compute these values:

- First, it is possible to estimate a VAR and derive long-term forecasts from current values, as in Hoevenaars et al. (2008) or Shen et al. (2019). One drawback from this approach is that long-term quantities are typically obtained by adding (compounding) short-term ones, which is not necessarily accurate or stable.
- A second approach is to rely on a theoretical model for asset prices (and their returns) that yields closed-form values for averages and variances, akin to Tsai and Wachter (2015) for instance. The main issue is then that parameters are estimated on past data, whereas the impact of climate change is expected to worsen in the future in ways that is not yet reflected in historical data. Another problem is that these types would need to be extended to several assets, in order to extract correlations.
- Finally, the third option, which we follow in this paper, is to determine simple relationships between the long-term unknowns and other state variables, which are themselves influenced by climate change. For instance, Engle et al. (2013) model volatility as an augmented AR(1) model in which macroeconomic predictors are added to the autoregressive component. In a very different spirit, Gostlow (2022) models asset returns as a linear combination of risk factors plus an additional climate risk component. This approach has the added benefit of transparency; it is tractable and simple to augment with further extensions to increase sophistication.

In the present article we seek to leverage expert-based forward-looking estimates, with a particular focus on producing accurate forecasts for long-term volatilities and correlations.

3.2 Data sources and statistics

The crux of our empirical work will be to estimate long-term variances and correlations based on long-term averages, as well as on long-term GDP growth and inflation.[7] To estimate the model, we will need past values for these quantities. To craft forward-looking forecasts, we shall rely on climate-driven predictions.

First, our primary source for forward-looking climate-based are the Phase 3 outputs of the NiGEM model (see Hantzsche et al. (2018)). These outputs are determined by three dimensions:

[7] Growth and inflation have been shown to drive rates in developed countries (Davis et al. (3121)). For stock returns, we point to Boudoukh and Richardson (1993) (inflation) and Baker et al. (2005) and Faugère and Van Erlach (2006) (growth). For the correlation between equities and bonds, we point to Campbell et al. (2020) - though Duffee (2023) arrives at more nuanced conclusions.

- the **climate model**, which can be either MESSAGEix-GLOBIOM (van Ruijven and Min (2020)), GCAM (Wise et al. (2014)) or REMIND-MAgPIE (Hilaire and Bertram (2020))
- the **climate scenario** which belongs to one of the following 6 categories: Below 2°C, Current Policies, Delayed Transition, Divergent Net Zero, Nationally Determined Contributions (NDCs) and Net Zero 2050. All scenarios are relative to a **baseline** model that ignores climate risks entirely.
- the **severity** that is taken into account: none (baseline), physical risks only, transition risks only, and both physical and transition risks.

In Figure 2, we depict the corrections that climate models imply for long-term GDP growth, compared to a baseline scenario. We observe clear heterogeneity across models, whereas differences are less pronounced between transition and combined risks in terms of severity. This is due to the fact that physical risk, according to the models, has a milder impact on economic output in the long run. The equivalent plot for equity returns is postponed to Figure 16 in Appendix A.4.

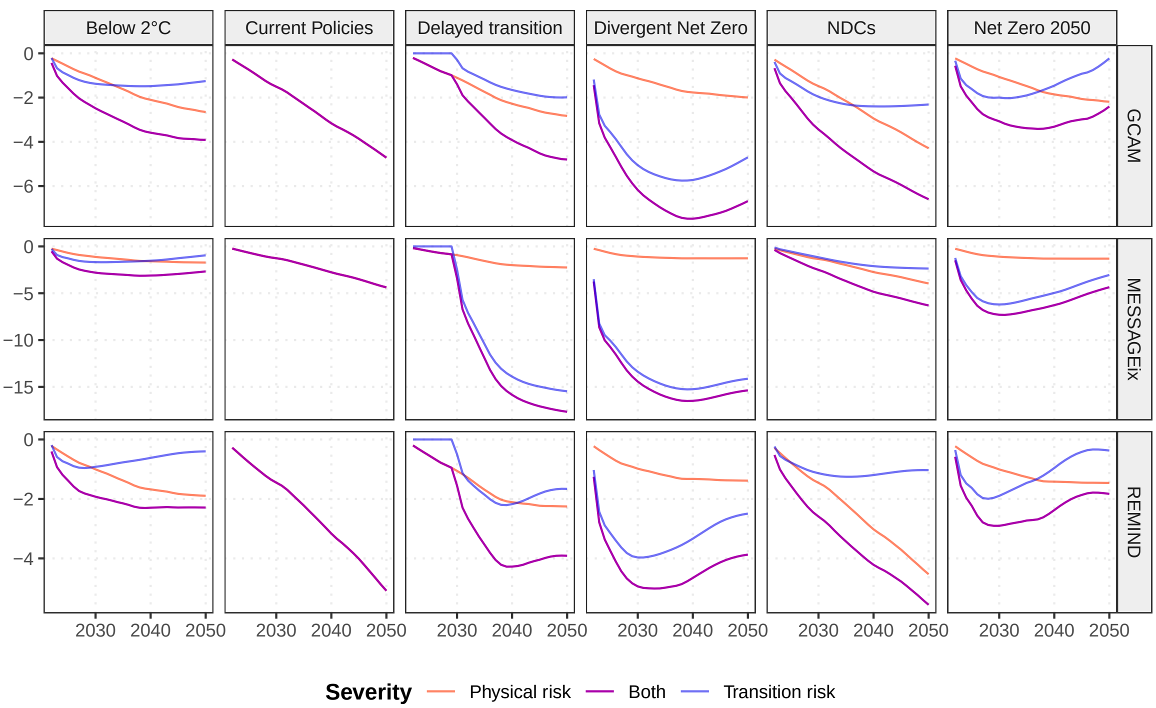

Figure 2: **Climate adjustment for US GDP growth**. We plot the time-series of correction for GDP growth implied by climate models, scenarios and severities compared to the baseline situation (in %). The values are in percentages of adjustment to the baseline growth value. For the period 2023-2033, it is equal to +2.2%. Hence a -10% adjustment implies a growth of +2%.

For past (realized) series, there are three components: economic (GDP and inflation), climate (temperature anomaly) and financial (asset returns). For the data on temperature anomaly, we refer to the National Centers for Environmental Information from which we extract monthly series for all 48 states in mainland US. The series are then simply averaged to obtain an aggregate value at the national level. The temperature anomaly

(ΔT henceforth) is then averaged over a 10 year horizon and we provide an illustration thereof in Figure 13 in the Appendix.

With regard to financial time-series, we work with equity and bond data from the Stocks, Bonds, Bills, and Inflation (SBBI) Data provided by Morningstar via the CFA Research Foundation Investment Data Alliance. The time-series of the 10 year average and standard deviation of returns is plotted in Figure 14, along with the corresponding long-term correlation. For the latter, we observe two regimes in which the correlation is either around +30% or in the vicinity of -30%. Similar patterns are obtained in Brixton et al. (2023), Molenaar et al. (2023) and Duffee (2023), but for the latter the effect is reversed because the correlation is with Treasury yields. While the plot shows points slightly before 1960, all estimations of coefficients will henceforth rely on samples that start in January 1960 and end in April 2023.

Finally, the last inputs we require are of economic type, namely GDP growth and inflation. They are obtained from the portal of the Federal Reserve of Saint Louis with the *GDP* and *CPIAUCSL* identifiers. We plot the corresponding time-series in Figure 15 in the Appendix. Henceforth, it is naturally the relative variations of these indices that are used as independent variables in our models. For example, we will use the notation ΔGDP_t for the rate of increase $GDP_t/GDP_{t-1} - 1$, and the same will apply to inflation.

4 Results

4.1 Linking volatility and correlation with the economy

The present section focuses volatilities and correlations. The integration of estimates for long-term returns will be treated as an extension of our baseline results (in Section 5.3). While many contributions seek to unveil *predictive* relationships between the volatility and state variables, we propose to work with *synchronous* regressions of the form

$$\sigma_t = a + b_\mu \mu_t + b_{gdp} \Delta GDP_t + b_{cpi} \Delta CPI_t + b_T \Delta T + e_t, \tag{4}$$

where μ_t is the long-term asset return, ΔGDP is GDP (relative) growth, ΔCPI is inflation and ΔT is temperature anomaly. All terms are synchronous, i.e., computed over the same rolling periods of 10 years. Naturally, by construction, these variables are all highly autocorrelated. From an inference standpoint, this is a major issue because of biases in OLS estimators, as revealed in the seminal work of Stambaugh (1999) (see also the more recent results of Boudoukh et al. (2022)). However, from a predictive perspective, this is much less an issue. For the sake of completeness, we provide four types of results in Table 1.

First, we propose two kinds of models: those that do not include long-term return averages μ_t as predictors (Model A), and those that do (Model B). This is because we want to be able to propose models that are entire free from average return estimations, which is the case of Model A. More fundamentally, we are agnostic with regard to the possible link between risk and returns and wish to explore both options.[8]

Next, for each model, in Panel 1, we document the estimates from standard OLS regressions. In addition, in Panel 2, to correct for the autoregressive bias in estimates,

[8] The risk-return relationship is an old debate in financial economics with abundant and often contradictory results. We refer for instance to the contributions of Gehr (1979), Campbell (1996), Malkiel and Xu (1997), Harrison and Zhang (1999), Whitelaw (2000) and Linton and Perron (2003) - to name but a few.

Model A: without long-term averages										
Panel 1: Standard OLS										
	Equities					**Bonds**				
Variable	Estim.	Sd. Err.	t-stat.	p-val	R^2	Estim.	Sd. Err.	t-stat.	p-val	R^2
Intercept	0.145	0.000	407.93	0.000	0.616	0.071	0.000	185.76	0.000	0.710
ΔCPI_t	0.965	0.054	18.02	0.000	0.616	1.373	0.057	23.90	0.000	0.710
ΔGDP_t	-0.456	0.075	-6.05	0.000	0.616	-0.756	0.081	-9.35	0.000	0.710
ΔT	0.008	0.001	6.70	0.000	0.616	0.008	0.001	6.36	0.000	0.710
Panel 2: Amihud and Hurvich (2004) regressions										
Intercept	0.145	0.000	407.91	0.000	0.616	0.071	0.000	193.12	0.000	0.731
ΔCPI_t	0.977	0.055	17.84	0.000	0.616	1.400	0.057	24.78	0.000	0.731
ΔGDP_t	-0.473	0.077	-6.13	0.000	0.616	-0.792	0.080	-9.94	0.000	0.731
ΔT	0.008	0.001	6.36	0.000	0.616	0.007	0.001	6.08	0.000	0.731
Model B: with long-term averages										
Panel 1: Standard OLS										
	Equities					**Bonds**				
Variable	Estim.	Sd. Err.	t-stat.	p-val	R^2	Estim.	Sd. Err.	t-stat.	p-val	R^2
Intercept	0.145	0.000	421.85	0.000	0.641	0.071	0.000	224.87	0.000	0.802
μ_t	-0.051	0.007	-7.31	0.000	0.641	0.206	0.011	18.79	0.000	0.802
ΔCPI_t	0.906	0.052	17.30	0.000	0.641	1.103	0.050	22.24	0.000	0.802
ΔGDP_t	-0.396	0.073	-5.40	0.000	0.641	-0.608	0.067	-9.04	0.000	0.802
ΔT	0.008	0.001	7.39	0.000	0.641	0.006	0.001	5.61	0.000	0.802
Panel 2: Amihud and Hurvich (2004) regressions										
Intercept	0.145	0.000	425.56	0.000	0.647	0.071	0.000	227.36	0.000	0.806
μ_t	-0.052	0.007	-7.42	0.000	0.647	0.205	0.011	18.74	0.000	0.806
ΔCPI_t	0.925	0.053	17.43	0.000	0.647	1.127	0.050	22.50	0.000	0.806
ΔGDP_t	-0.422	0.074	-5.67	0.000	0.647	-0.640	0.068	-9.39	0.000	0.806
ΔT	0.008	0.001	6.99	0.000	0.647	0.005	0.001	5.17	0.000	0.806

Table 1: **Synchronous links between long-term volatility and macroeconomic indicators**. We report the estimates from regression (4) for equities (left) and bonds (right). We present models without long-term averages first (Model A) and with them next (Model B). The upper sub-panels pertain to simple OLS estimators, and the second ones to the augmented method from Amihud and Hurvich (2004). The sample runs from January 1960 to April 2023. Independent variables are demeaned before estimation and the corresponding means are 10.70% for stock returns, 6.41% for bond returns, 3.62% for inflation, 6.39% for economic growth and 71.79% for ΔT. The reported R^2 are the *adjusted* R^2.

we run augmented regressions as advised in Amihud and Hurvich (2004). The idea is to start by estimating a first order vector-autoregression (VAR(1)) for the predictors: $\boldsymbol{X}_t = \boldsymbol{a} + \boldsymbol{X}_{t-1}\boldsymbol{A} + \boldsymbol{e}_t$ and to add the fitted residuals $\hat{\boldsymbol{e}}_t$ in the original model (4). Note that there is one vector of residual values for each predictor.

Fortunately, as we can see in the table, the various specifications do not move the numbers much: a token of stability for our results. We note that all coefficients are statistically significant at all the usual levels of confidence. Consequently, the simple OLS model seems robust enough and we will henceforth stick with it for our predictions in the next section.

In terms of raw coefficients, we see that for equities, volatility is negatively linked to long-term returns and GDP growth but positively linked to inflation and rising temperature. The only difference for bonds is with long-term average (μ_t): in this case the link is positive.[9] We underline that working with demeaned series allows to interpret the intercept in the models as the historical average of the dependent variable. A last comment on pure fit: the reported adjusted R^2 are quite high, which comes at least partly from the autocorrelation in the variables due to overlapping samples. Persistence is problematic for inference but beneficial to plain forecasting. In addition, the fit is higher for bonds than for equities and adding long-term return averages (μ_t in Model B) is also more beneficial to bonds than to stocks.

The next step is to perform a similar estimation for long-term *correlation* between the two asset classes. In this case, we need to make some changes compared to the specification of Equation (4). First, we acknowledge that correlation is a bounded and oscillating indicator (see Figure 14), whereas the temperature anomaly is bound to be increasing in the foreseeable future. Hence, we should not include ΔT in the model. Just as in the case for volatility, it is not obvious whether adding the long-term return averages is a good idea or not, hence we propose specifications without (Model A) and with them (Model B). Lastly, we propose an intermediate choice that is inspired from Molenaar et al. (2023) and includes inflation and the long-term bond return as independent variables (Model C). The results are compiled in Table 2 and just as before in Table 1, there is no major shift when we take into account the Amihud and Hurvich (2004) correction in the augmented model.

Variable	Estim.	Sd. Err.	t-stat.	p-val	R^2	Estim.	Sd. Err.	t-stat.	p-val	R^2
	Model A: model *without* long-term averages									
	Standard OLS					**Amihud and Hurvich (2004) model**				
Intercept	0.056	0.007	8.56	0.000	0.601	0.056	0.006	8.68	0.000	0.609
ΔCPI_t	2.215	0.684	3.24	0.001	0.601	2.232	0.683	3.27	0.001	0.609
ΔGDP_t	8.514	0.670	12.71	0.000	0.601	8.525	0.668	12.76	0.000	0.609
	Model B: model *with* long-term averages									
	Standard OLS					**Amihud and Hurvich (2004) model**				
Intercept	0.056	0.005	12.02	0.000	0.798	0.056	0.005	12.11	0.000	0.800
μ_t^{equities}	0.507	0.107	4.71	0.000	0.798	0.486	0.109	4.47	0.000	0.800
μ_t^{bonds}	3.899	0.184	21.16	0.000	0.798	3.954	0.186	21.26	0.000	0.800
ΔCPI_t	-3.816	0.602	-6.34	0.000	0.798	-4.013	0.606	-6.63	0.000	0.800
ΔGDP_t	13.293	0.551	24.12	0.000	0.798	13.454	0.554	24.28	0.000	0.800
	Model C: Molenaar et al. (2023) - inspired									
	Standard OLS					**Amihud and Hurvich (2004) model**				
Intercept	0.056	0.007	8.50	0.000	0.596	0.056	0.007	8.49	0.000	0.596
ΔCPI_t	8.729	0.333	26.19	0.000	0.596	8.729	0.333	26.18	0.000	0.596
μ_t^{bonds}	2.554	0.209	12.23	0.000	0.596	2.555	0.209	12.21	0.000	0.596

Table 2: **Synchronous links between long-term correlation and macroeconomic indicators**. We report the estimates from regression (4) with or without two regressors for long-term averages depending on asset class (μ_t^{equities} for equities and μ_t^{bonds} for bonds), as well as with or without intercepts. The left panel pertains to simple OLS estimators, while the right one to the augmented method from Amihud and Hurvich (2004). Dependent variables were de-meaned before estimation. The R^2 are adjusted.

[9]The negative coefficients for equities is consistent with the aggregate low volatility anomaly (see, e.g., Glosten et al. (1993)), though it has mostly been documented recently in the cross-section of stocks: Ang et al. (2006), Ang et al. (2009), Baker et al. (2011), Li et al. (2016) and Beveratos et al. (2017).

Several important comments are in order. First, GDP growth and long-term average returns of both asset classes are positively linked to correlation. However, for inflation, the link depends on the set of variables included in the model. In the full model, inflation is negatively related to the stock-bond correlation, but in the other two, the relationship reverses.

Given the links revealed in Tables 1 and 2, we are able to plug both the long-term returns and the climate-driven forecasts for the economic indicators (raw values in Figure 15). This will produce the missing parts of the puzzle, namely long-term volatilities and correlations. Given the non-negligible disagreement across the correlation models, we will include all three for the sake of completeness.

4.2 Long-term predictions

We leverage the models of Table 1 and 2 for the sake of forecasting. To this purpose, we plug long-term estimates of the independent variables in the models. For inflation, GDP growth and stock returns, we resort to climate-based values stemming from NiGEM (Phase 3) and for temperature, to IAM outputs.[10] For bonds, we take NiGEM estimates for long term yields and compute returns to be the average yield over the investment horizon plus price return, i.e.,

$$\mu_t^{\text{bonds}} = \frac{1}{H}\sum_{s=0}^{H-1}\hat{y}_{t+s} + \frac{D}{H}(\hat{y}_t - \hat{y}_{t+H}), \tag{5}$$

where $\hat{y}_t$ is the predicted bond yield in a given climate-driven scenario. For simplicity, we assume that the duration of the bond is equal to the investment horizon i.e. $D = H$ (which matches the case of a zero coupon bond with a constant maturity H).

In addition, to add more flexibility in estimates and generate more granularity, we resort to shrinkage in the following form

$$v_{i,n} = a_i v^* + (1 - a_i)v_n, \tag{6}$$

where $a_i \in (0, 1)$ is the shrinkage intensity, v_n is the value (volatility or correlation) of the n^{th} model-scenario and v^* is the anchor value.[11] This enables a Bayesian-type of mixture in which the data-driven estimates can be blended with a prior value, thereby allowing more flexibility. In our empirical application, we choose the a_i to be all percentages between 0% and 80% to avoid blends that are too concentrated towards the target. This makes 81 intensities, which, combined with the 162 scenarii makes 13,122 estimates in total for one given anchor value v^*.

In Figure 3, we produce the distribution of 10 year volatilities (2023-2033) under the models from the standard OLS coefficients of Model A (panel 1) in Table 1. In the upper panels, we illustrate the difference between past realized and future predicted volatility. The contrast in support is striking: for equities, historical volatility has ranged from 11% to 17% with an average of 14.5%, whereas our forecasts lie mostly between 14% and 15%, with thin tails around 18%. Similar conclusions hold for bonds as well, with a support mostly between 6% and 8% for predicted volatility. The forecasted risks from Model B which includes long-term returns are postponed to Figure 17 in Appendix A.5. The results are essentially identical with a minor shift of volatilities which are slightly less conservative in this latter figure.

[10] The latter are retrieved on the AR6 scenario explorer.

[11] We will also sometimes call it the *prior target* or *prior value*.

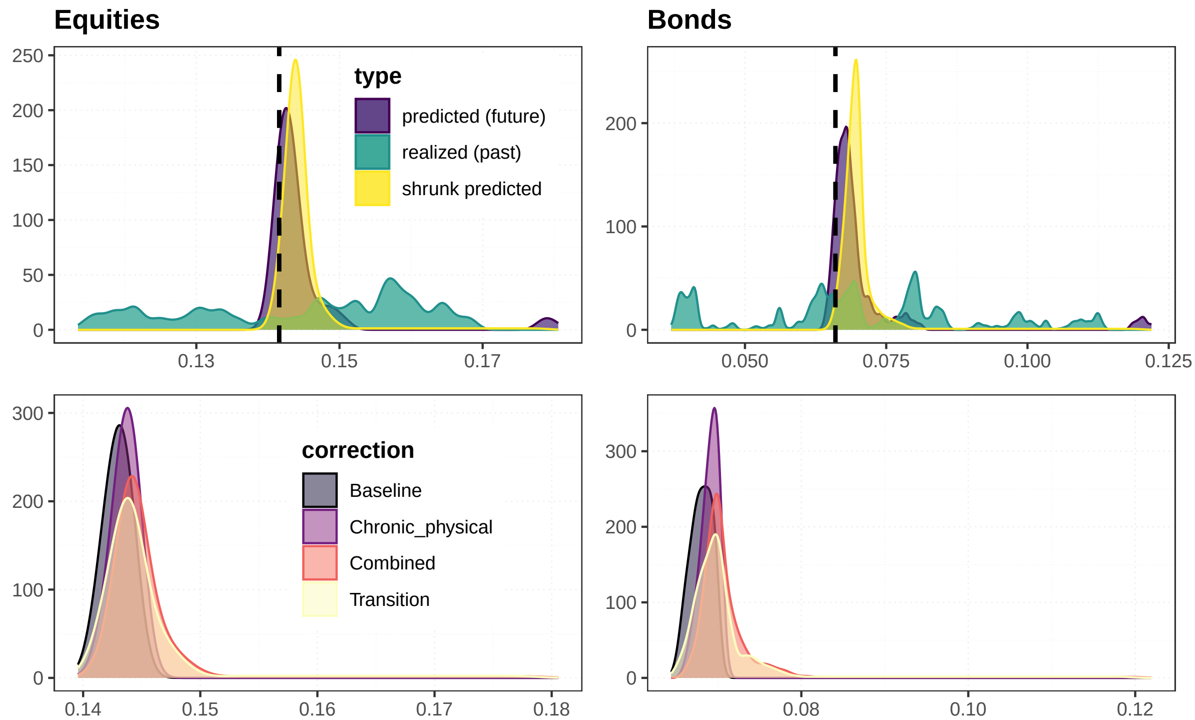

Figure 3: **Long-term volatility prediction**. We plot the distribution of long-term forecasts for volatility (2023-2033) across the two asset classes with Model A panel 1 from Table 1. The upper panels compare historical values with the predicted ones while the lower ones focus on the latter but show the shift that occurs when considering various levels of climate severity. The priors used when post-shrinking in (6) are the historical means (1960-2023): 0.145 for equity volatility and 0.071 for bond volatility. The intensities α_i are between 0.0 and 0.8 with 0.01 increments. The vertical lines in the upper panels mark the average predictions in the *baseline* scenarios (agnostic to climate change): 14.2% and 6.9% for equities and bonds, respectively.

One important takeaway from these figures is that both asset classes see a very similar pattern in the shift of their risk: a minority of scenarios predict lower volatility and a majority of them indicate marginally larger risks. A very small proportion of cases imply much larger risks. The fact that the impact of climate is relatively similar for both asset classes (a consequence of the similarity of signs and magnitudes from Model A of Table 1) will explain much of our results below.

In the lower panels of Figure 3, we show the distribution of predicted volatility, but conditional on the severity of climate scenarios. As the severity increases from baseline to combined (physical plus transition risk), the distribution slowly shifts to the right and the average risk increases, as one would expect.

In Figure 4, we depict the distribution of predicted correlation at the 10 year horizon. In line with our analysis above, we provide three different specifications, the last one being inspired by a recent publication. Notably, we remark that the historical correlation has a much wider support, with two modes around ± 0.3 (see the bottom panel of Figure 14). Hence, the predicted correlations seem to represent some middle ground and are more concentrated around zero.

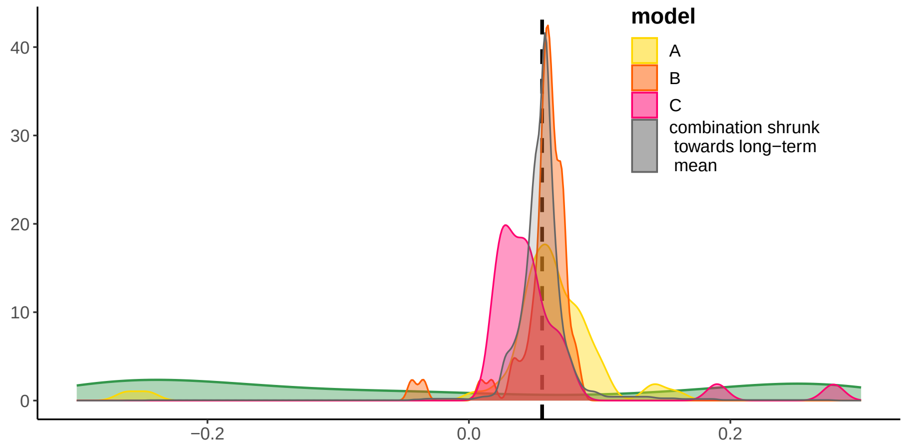

Figure 4: **Distribution of long-term correlation predictions**. We plot the distribution of long-term forecasts for correlation between the two asset classes. The models are those of Table 2. The grey distribution corresponds to all models, shrunk towards an agnostic value equal to the realised (past) average of 5.6% - a relatively agnostic prior. The shrinkage intensities are percents between 0.0 and 0.8 (81 values in total). The vertical dashed line highlights the historical average correlation and the green line at the bottom is the historical density. Predictors were demeaned before the forecast.

5 Implications for long-term asset allocation

5.1 Optimal allocation without expected returns

In the present paper, we are mostly interested in the risk component of portfolio optimization. Below, we present a method that allows to bypass the specification of expected returns in optimal mean-variance allocation. Let us assume an investor with a quadratic utility and budget constraint who seeks to solve the following

$$\max_w \left\{ w'\mu - \frac{\gamma}{2} w'\Sigma w \right\}, \quad s.t. \quad w'e = 1 \tag{7}$$

The first order condition (FOC) is $\mu - \gamma\Sigma w + \lambda e = 0$, where λ is the Lagrange multiplier set to satisfy some constraint and e is a vector of ones, with e' being the transpose of e. The optimal weights have the form

$$w^* = \gamma^{-1}\Sigma^{-1}(\mu + \lambda e). \tag{8}$$

Our goal in this section is to identify the changes to a benchmark portfolio, which we take to be 60%/40% in equity/bonds, when there is a climate shock to the estimation of the covariance matrix Σ. But we do not know the two values of the expected return vector μ. The only thing we need is an arbitrary long-term average return, which we fix at some unspecified constant $\bar{m}$. From (8), we have that $\mu = \gamma\Sigma w - \lambda e$, and we assume we do not know the level of λ, e.g., if we overlook the constraint. This level can be fixed to match the target return $\bar{m}$. In this case, we obtain $\tilde{\lambda} = \gamma w \Sigma w - \bar{m}$ (multiplying the last expression by w', with $\bar{m} = w'\mu$). Hence, the expected returns that are required to reach the target $\bar{m}$ are given by

$$\tilde{\mu} = \gamma\Sigma w - \tilde{\lambda} e. \tag{9}$$

Now, suppose there is a climate shock and the new (updated) covariance matrix shifts to $\tilde{\Sigma}$ - but the risk aversion stays the same. The new optimal weights will be:

$$\gamma^{-1}\tilde{\Sigma}^{-1}(\tilde{\mu}+\xi e)=\gamma^{-1}\tilde{\Sigma}^{-1}(\gamma\Sigma w-\tilde{\lambda}e+\xi e)=\gamma^{-1}\tilde{\Sigma}^{-1}(\gamma\Sigma w+\hat{\lambda}e), \tag{10}$$

where the adjusted Lagrange multiplier $\hat{\lambda}=\xi-\tilde{\lambda}$ should be chosen to saturate the budget constraint. Hence, the new portfolio composition is not driven by μ, nor by the arbitrary $\bar{m}$. Moreover, it is easy to show that $\hat{\lambda}$ is in fact proportional to γ,[12] hence the optimal allocation is also unrelated to γ. It depends only on the two covariance matrices (Σ before the climate adjustment and $\tilde{\Sigma}$ after) and on the original portfolio composition $w=(0.6,0.4)$.

In Figure 5, we show the distribution of optimal portfolio weights when starting from a 60%/40% allocation and then switching the baseline covariance matrix with all those that incorporate climate-driven shocks. A large majority of equity allocation shifts lie in a range of 57%-61%, which is a relatively narrow interval. Nevertheless, a few scenarios lead to more pronounced changes, reducing the equity exposure to 53% of the portfolio. We underline that the largest shifts are those that include transition risks, which are considered more financially material by the climate models.

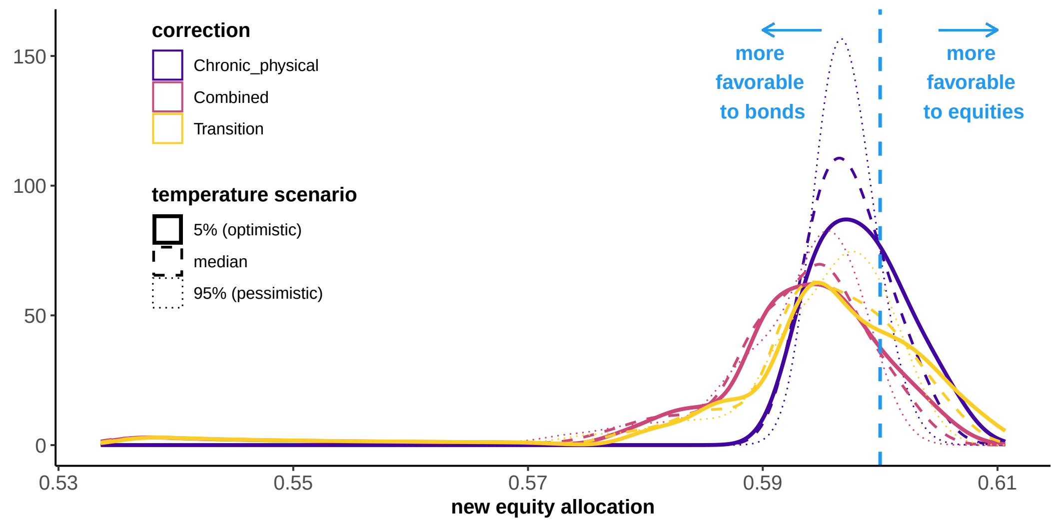

Figure 5: **Climate-based adjustment to the 60%/40% allocation - without long-term returns**. We plot the shift in optimal portfolio composition when taking into account climate-based values for volatilities and cross-asset correlations. The forecasts come from Model A in Table 1 (volatility) and Model A in Table 2 (correlation). We report the distribution (across scenarios) of the **equity** share of the portfolio weights given in Equation (10). Σ represents the *baseline* covariance matrix and $\tilde{\Sigma}$ the one that stems from predictions with climate-driven corrections. The density lines are depicted for three levels of correction and three temperature scenarios from the IAM output. The vertical line represents the 60% starting point of the target allocation.

These outliers to the left of the distributions, which are all tied to transition and combined corrections, come from scenarios from the MESSAGEix-GLOBIOM model and the *Delayed transition* scenario. The latter imply a volatility for equities of 18%, which are the most pessimistic, as shown in Figure 3. The correlation is at 8% for these scenarios, which are also the least diversifying for model B in Figure 4.

[12] The budget constraint imposes $1=\gamma^{-1}e'\tilde{\Sigma}^{-1}(\gamma\Sigma w+\hat{\lambda}e)$, i.e., $\hat{\lambda}=\gamma\frac{(1-e'\tilde{\Sigma}^{-1}\Sigma w)}{e'\tilde{\Sigma}^{-1}e}$.

5.2 Covariance matrix based on long-term returns

In the previous subsection, we focused on risks that were driven by economic and temperature indicators solely. In tables 1 and 2, we proposed models that also rely on long-term return predictions to explain volatilities and correlations. Below in Figure 6, we plot the shift in allocation that correspond to the following models: model B in Table 1 and models A and C in Table 2.

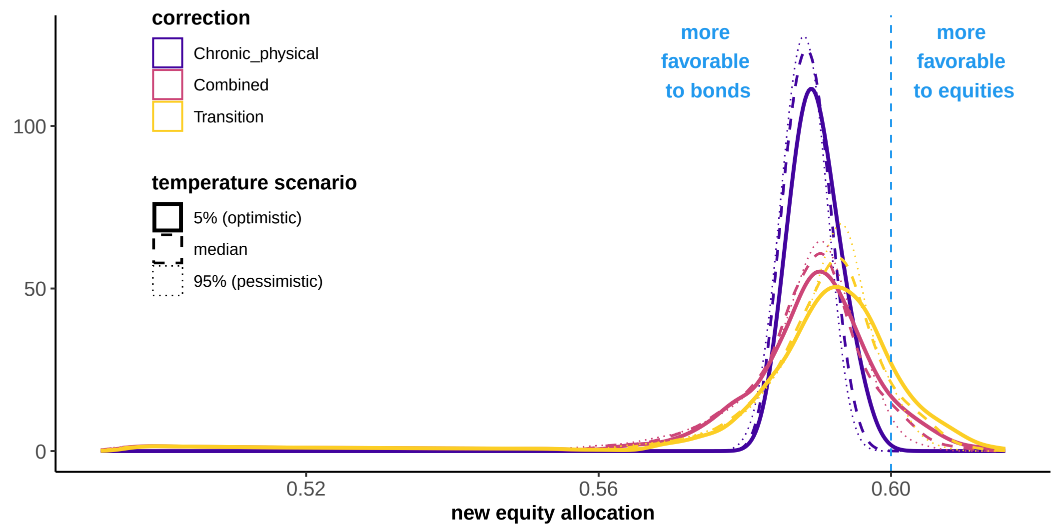

Figure 6: **Climate-based adjustment to the 60%/40% allocation - with long-term returns only impacting the covariance matrix**. We plot the shift in optimal portfolio composition when taking into account climate-based values for volatilities and cross-asset correlations. The forecasts come from Model B in Table 1 (volatility) and Models B and C in Table 2 (correlation). We report the distribution (across scenarios) of the **equity** share of the portfolio weights given in Equation (10). Σ represents the *baseline* covariance matrix and $\tilde{\Sigma}$ the one that stems from predictions with climate-driven corrections. The density lines are depicted for three levels of correction and three temperature scenarios from the IAM output. The vertical line represents the 60% starting point of the target allocation.

In this case, the shifts are slightly more pronounced, with most of the support between 56% and 62%. The long left tail persists and is associated to the same scenario as in Subsection 5.1 and even reaches values of 50% for the equity allocation. One common feature with Figure 5 is that the most extreme shifts are associated with transition risks, not with physical ones. This is in line with previous results in the literature (Faccini et al. (2023)).

5.3 Allocation with NiGEM expected returns

The present paper is focused on the impact of climate change on the **risk** of asset classes, not on their expected returns. Nevertheless, the latter are naturally paramount to investors. Hence, in this section, we document the change to mean-variance optimization (MVO) schemes that take expected returns as inputs. The latter are notably hard to estimate and we point to Ma et al. (2023) for a recent overview on this matter.

More specifically, because MVOs are notoriously leveraged, we introduce an additional constraint to reduce this effect[13]

$$\max_{w} \left\{ w'\mu - \frac{\gamma}{2} w'\Sigma w \right\}, \quad s.t. \quad w'e = 1, \quad w'w \leq \delta, \tag{11}$$

which is solved by

$$w^* = (\gamma\Sigma + 2\eta I)^{-1}(\mu + \lambda e), \tag{12}$$

where the two constants λ and η are chosen so as to saturate the constraints (λ for the budget and η for diversification through δ in (11)). Simply put, the covariance matrix and the vector of expected returns are shrunk. In the asymptotic case when $\lambda, \eta \to \infty$, we obtain a uniform allocation, $w = (1/2, 1/2)$. In order to obtain a 60/40 allocation for the baseline estimates from the model, we fix $\gamma = 0.5$ and $\eta = 0.028$. This makes a rather low risk aversion, but larger values will be considered in Section 6.3 below.

There are an infinite number of choices for expected returns. To remain consistent with climate-driven estimates, we pursue below an exercise that take as inputs the long-term returns generated by the NiGEM model (Phase 3). This is an arbitrary choice and serves only as illustration of the flexbility of our approach. We do not attest to their predictive power or accuracy, which is outside the scope of this paper.

First, in the left panel of Figure 7, we show the distribution of the expected returns generated by the NiGEM model, which we have shrunk according to Equation (6) with 0.7% as prior for bond returns and 2.7% for equities.[14] These values are the averages across all climate scenarios (relatively agnostic priors) and correspond to the μ vector in Equations (11) and (12). With the vertical dashed lines, we also provide the baseline estimates, which are those that do not take into account physical and transition risk (1.94% for bonds and 3.29% for equities).

In the right panel, we depict the shift in allocation between the baseline scenario and all the climate scenarios and adjustments. The parameters in the optimization were calibrated to yield a 60/40 split. Again, we observe different patterns between corrections. The corrections that are based on physical damages only are much more favorable to stocks and almost all changes imply *higher* equity exposure. In contrast, the other two corrections can be much more favorable to bonds, with some extreme scenarios for which equities make up less than half of the portfolio. Clearly, the results are much more sensitive because the range of the shifts is much larger than in Figures 5 and 6. This was predictable, given the sensitivity of optimal weights to expected returns (see, e.g., Chopra and Ziemba (1993)).

It is crucial to understand the major difference between Figure 7 and Figures 5 and 6, which is that the former predicts much more favorable outcomes for stocks compared to the other two. This can be rationalized with the help of the left panel of Figure 7 (which is inferred from the values of Figure 16). Indeed, we see that the shifts, mostly to the left of the distributions compared to the baseline values (dotted vertical lines) is more marked for bonds. Therefore, the climate models predict that physical and transition risks are slightly more detrimental for fixed-income securities than for stocks. This is why, in this case, climate-driven portfolios will favor stocks over bonds.

[13] We refer to Jagannathan and Ma (2003) and Coqueret (2015) for a detailed account on why such regularized optimization can be beneficial. We follow the latter reference here.

[14] Note that this value is based on price returns only and thus omits dividends, which may explain its low magnitude compared to historical levels.

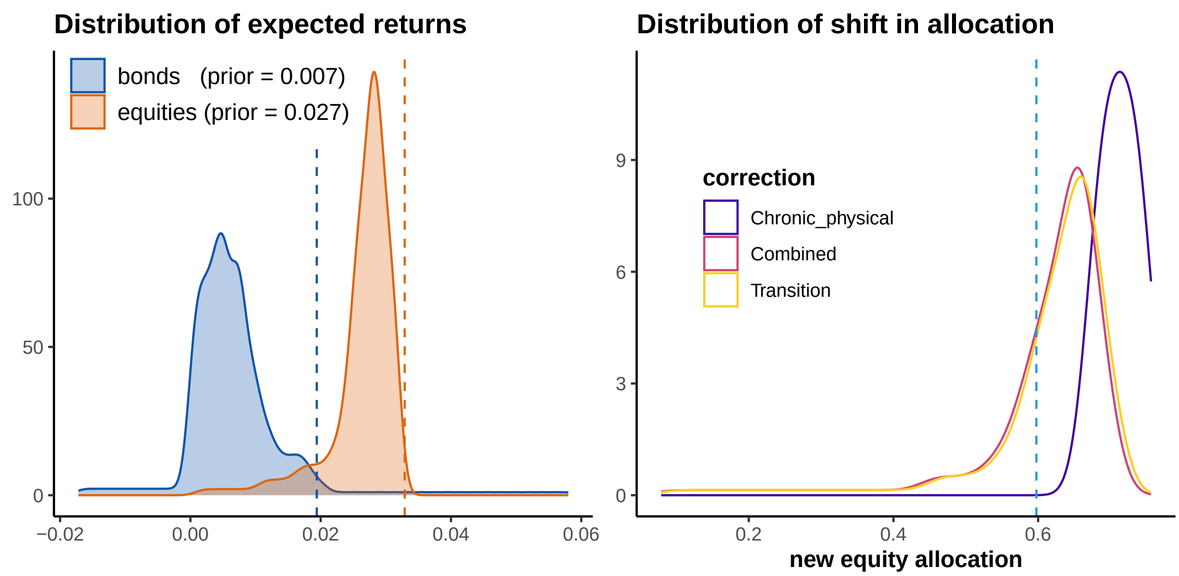

Figure 7: **Adjustments with expected returns**. In the left plot, we display the distribution of shrunk expected returns. The shrinkage intensities are between 0.0 and 0.8 (21 values in total). The vertical dashed lines depict the baseline values, i.e., those that correspond to the scenario without transition nor physical risk. In the right plot, we feature the shift in allocation when plugging NiGEM expected returns (μ) in (12). The initial parameters $\eta = 0.028$ and $\gamma = 0.5$ are calibrated so as to obtain a 60/40 mix for the baseline climate scenario.

5.4 Projected performance

The right panel of Figure 7 seems to indicate that climate-based portfolios should be more tilted towards equities, hence more risky. In this subsection, we document the risk-adjusted performance of portfolios under the assumption that returns and covariance matrices are the ones from the previous subsection.

In Figure 8, we show the projected risk-adjusted return of portfolios, depending on the proportion invested in stocks. The lines represent the median values across all scenarios and forecasts used in Section 5.3. The shaded areas mark the 10%-90% quantiles. For the baseline scenarios, there is a small interval because we consider 3 temperature levels (5%, median, and 95% quantiles in IAM outputs) in our long-term forecasts for volatilities and correlations (which come from the models in Table 1 and 2).

Plainly, the best performance when accounting for climate impact (a ratio of 0.25) is reached at around 30% equity for the top 90% green scenarios. Reversely, the worst one if we don't consider physical and transition risk is also 0.25, obtained at 80%+ equity levels on the orange curve. Hence, as expected, climate risk is clearly detrimental to performance. It is noteworthy to underline that the shaded area is the thinnest around 60% of equity: this is where the dispersion across scenarios is the smallest.

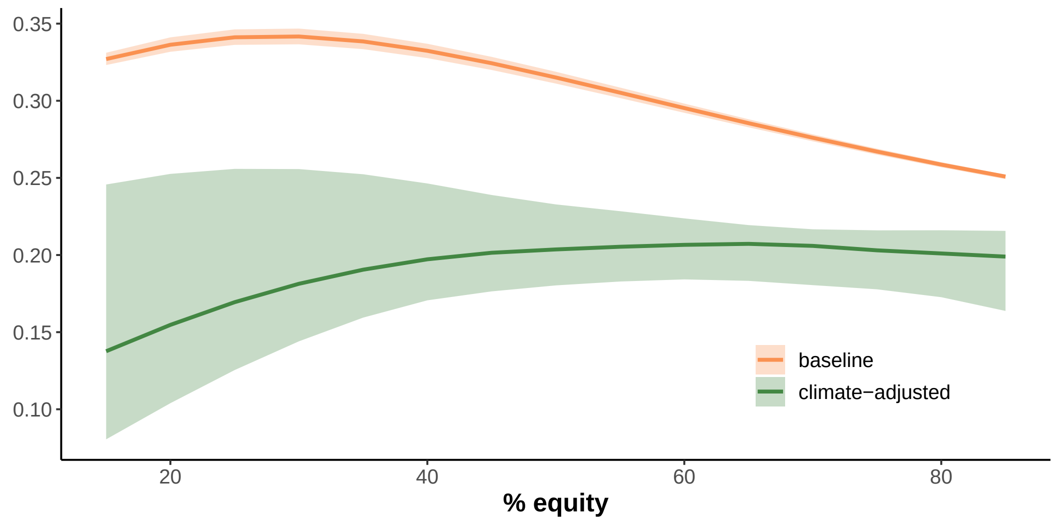

Figure 8: **Projected risk-adjusted performance**. We plot the ratio of expected returns to volatility for various levels of equity allocation (x-axis). The straight lines show the median across all scenarios and values from Section 5.3. The shaded areas mark the $[0.1, 0.9]$ quantile zone.

6 Robustness checks

6.1 Quadratic impact of temperatures

In many models including the DICE of Nordhaus (2014, 2017), the economic damages incorporate a quadratic term of the temperature anomaly. Below, in Table 3, we propose estimates for an extension of Equation (4) that integrates such an effect. The quadratic term is associated with a negative term that mimics that of ΔT, which makes sense because the two variables are positively correlated. For the sake of completeness, we also provide in Panel B results when removing ΔT.

The results are comparable with those of Table 1. Notably, the coefficients for inflation remain positive (inflation increases risk) and those of growth negative (economic expansion reduces risk). The R^2 have the same magnitude as before, and models for bonds seem to provide a better fit, compared to those for stocks.

Furthermore, in Figure 9, we show the corresponding predicted long-term volatilities for both asset classes. The differences with Figures 3 and 17 are subtle, thereby highlighting the stability of our initial results to a small change in model and forecasting protocol.

6.2 Shrinkage priors

Our framework allows for some flexibility in a Bayesian-like fashion. Indeed, while the estimates produced by the expert forecasts rely on both extensive data and sophisticated models, an investor might want to shift them based on her priors. For instance, in Section 5.3, we are able to shift the distribution of expected returns via shrinkage. We can also do this for volatility estimates, as is shown in Figure 10. Therein, we fix the prior for bond volatility to 7% and we test four values for equity priors from 13% to 16%, shown with colors.

Variable	Equities Estim.	Sd. Err.	t-stat.	p-val	R^2	Bonds Estim.	Sd. Err.	t-stat.	p-val	R^2
Panel A: with ΔT										
Intercept	0.145	0.000	443.014	0.000	0.675	0.071	0.000	190.185	0.000	0.723
ΔCPI_t	0.788	0.052	15.279	0.000	0.675	1.267	0.059	21.594	0.000	0.723
ΔGDP_t	-0.199	0.073	-2.734	0.006	0.675	-0.603	0.083	-7.281	0.000	0.723
ΔT	0.013	0.001	11.414	0.000	0.675	0.011	0.001	8.439	0.000	0.723
$(\Delta T)^2$	-0.010	0.001	-11.689	0.000	0.675	-0.006	0.001	-6.122	0.000	0.723
Panel B: without ΔT										
Intercept	0.145	0.000	409.391	0.000	0.619	0.071	0.000	181.925	0.000	0.698
ΔCPI_t	1.228	0.037	33.152	0.000	0.619	1.638	0.041	40.210	0.000	0.698
ΔGDP_t	-0.934	0.037	-25.518	0.000	0.619	-1.221	0.040	-30.352	0.000	0.698
$(\Delta T)^2$	-0.006	0.001	-7.118	0.000	0.619	-0.003	0.001	-2.807	0.005	0.698

Table 3: **Including squared temperature anomalies in volatility models**. We report the estimates from regression (4) (with added $(\Delta T)^2$ term) for equities (left) and bonds (right). The sample runs from January 1960 to April 2023. Independent variables are demeaned before estimation and the corresponding means are 10.70% for equity returns, 6.41% for bond returns, 3.62% for inflation, 6.39% for economic growth, 71.79% for ΔT and 53.86% for $(\Delta T)^2$. The reported R^2 are the *adjusted* R^2.

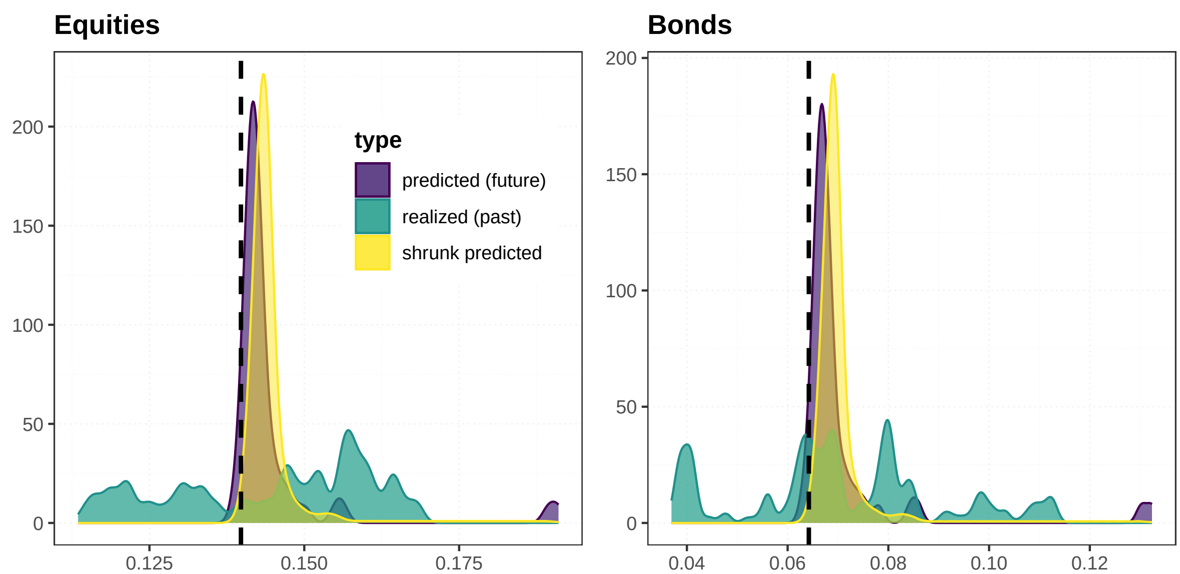

Figure 9: **Long-term volatility prediction with quadratic temperature effect**. We plot the distribution of long-term forecasts for volatility across the two asset classes with Model A panel 1 from Table 3. The priors used when post-shrinking in (6) are the historical means (1960-2023): 0.145 for equity volatility and 0.071 for bond volatility. The intensities α_i are between 0.0 and 0.8 with 0.01 increments. The vertical lines in the upper panels mark the average predictions in the *baseline* scenarios (no adjustment).

Figure 10 shows the level of prior (which shifts the distribution of volatility) that is required to significantly favor or penalize bonds (or equities). With a 13% volatility for equities, we are clearly in a territory that is favorable for equities. With a 16% volatility, it

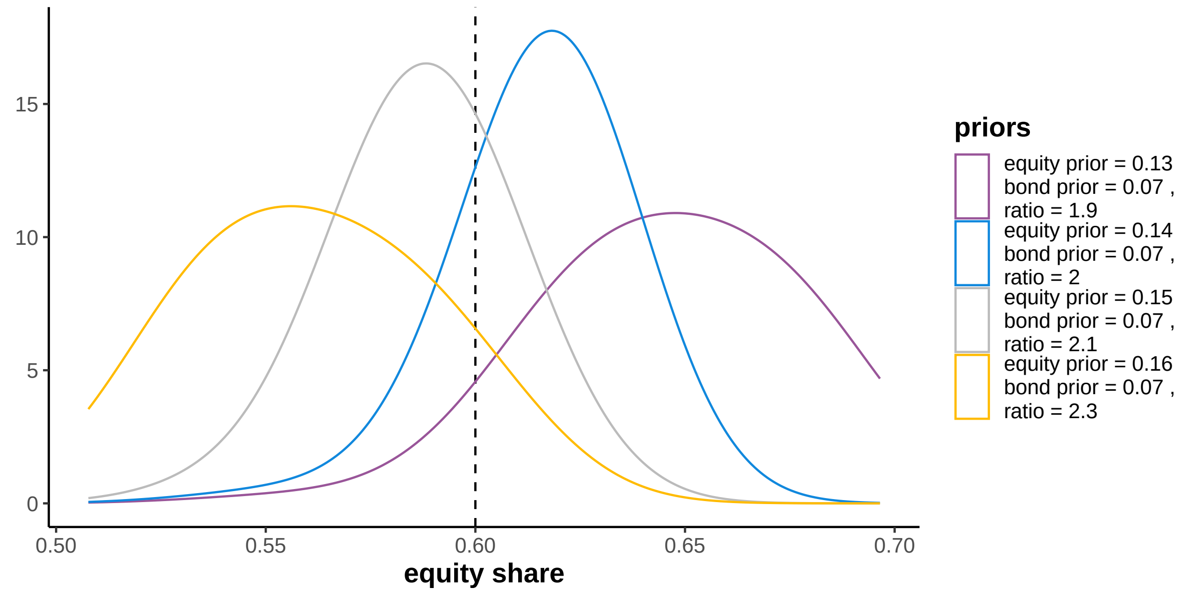

Figure 10: **Alternative priors for volatilities**. We plot the distribution of the equity share when shifting from the baseline scenario to a climate adjusted one. The calibration is such that the baseline equity weight is 60%. Each color is associated with a different value of prior for equities, while the one for bonds is fixed to 7%.

is the opposite. Between these two extreme values, there is a continuum of distributions and we reach an equilibrium (the allocation stays the same at 60/40) at 14.5%.

In Figure 11, we show the impact of alternative priors for **expected returns** on the shift of allocation when taking into account climate scenarios both for the estimation of the covariance matrix *and* the (shrunk) expected returns. It shows that a distribution of changes that would be balanced around the initial 60/40 target requires a strong tilt: bonds need to have expected returns that are barely below those of equities (3.2% versus 3.5%).

6.3 More conservative targets

In this section, we are interested in the change in allocation when starting from less risky initial positions. To ease the exposition, we will fix η to its previous value ($\eta = 0.028$) and link the riskiness of the portfolio only to the risk aversion. When $\gamma = 0.5$, we have that the initial share of equities is equal to 60%. We then increase γ in order to reach shares of 40% and 20%. This requires risk aversions of $\gamma = 2.1$ and $\gamma = 7.6$, respectively. Results are depicted in Figure 12 and show that for the most conservative target (20% of stocks), all corrections types align and are less favorable for equities.

7 Conclusion

We propose to model long-term estimates of risk based on expert climate forecasts. Importantly, the links between variables are synchronous and not predictive. The rationale for this is that realized risk is a consequence of realized (current) economic conditions - which we augment with environmental considerations.

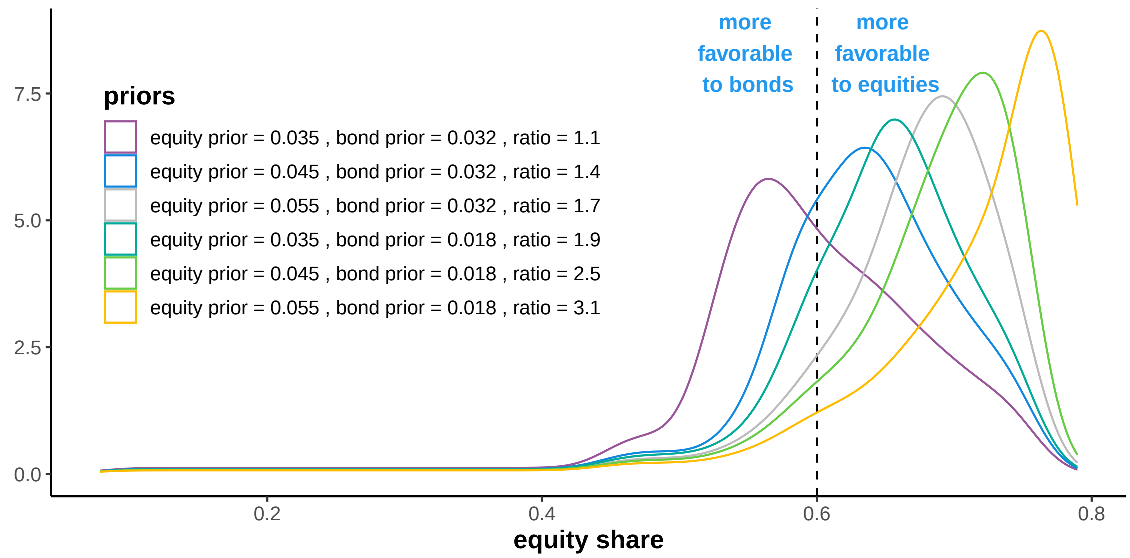

Figure 11: **Alternative priors for expected returns**. We plot the distribution of the equity share when shifting from the baseline scenario to a climate adjusted one. The calibration is such that the baseline equity weight is 60%. Each color is associated with a different couple of priors for expected returns.

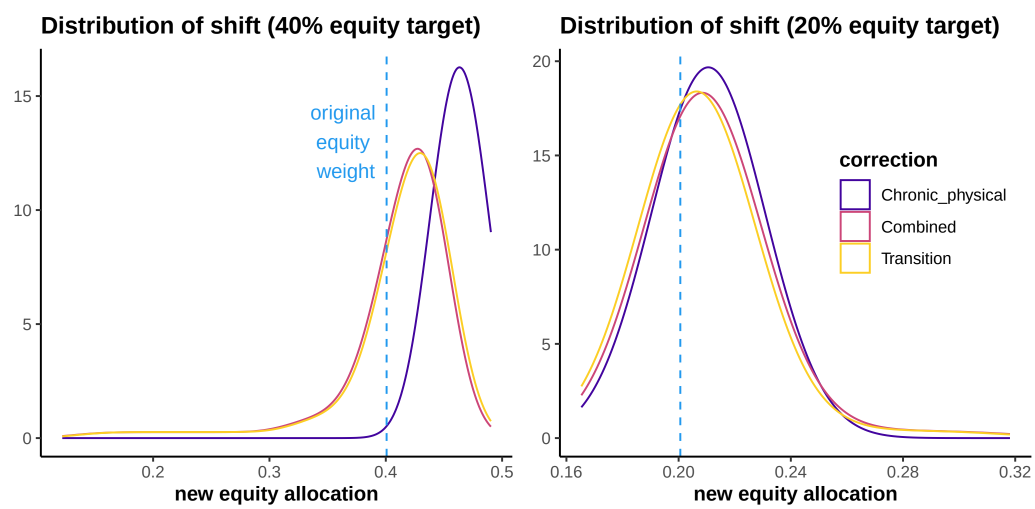

Figure 12: **Adjustments with expected returns - conservative allocations**. In the left plot, we display the distribution of shrunk expected returns The shrinkage intensities are between 0.0 and 0.8 (21 values in total). In the right plot, we feature the shift in allocation when plugging NiGEM expected returns (μ) in (12). The parameters $\eta = 0.028$ and $\gamma = \{2.1, 7.6\}$ are calibrated so as to obtain the 40/60 and 20/80 mixes for the baseline climate scenario.

Our results show that allocations that are driven by risk considerations only are not too impacted by potential climate losses. An initial target of 60% of equities is likely to shift within a range of 57%-63%. This is because the models predict that both asset classes will be harmed by climate impacts in a similar fashion. Thus, the relative risk of one class versus the other is only mildly altered.

However, this is not true for expected returns. The latter are reduced more for bonds than for equities, hence mean-variance portfolios that rely on climate-based estimates of returns are mostly favorable to equities. Our results are robust to several changes in the protocol, including quadratic temperature anomaly and post-hoc changes in distributions of risks and returns.

Unfortunately, while the optimal allocations are not necessarily very sensitive to climate risks, portfolio performance is. Risk-adjusted returns are significantly impacted, notably because climate-based returns are overwhelmingly smaller than their climate-agnostic counterparts. Our results indicate that risk-adjusted returns can be between 20% and 50% lower when accounting for climate-driven losses.

Naturally, our findings depend on the choice of a particular model for volatilities and correlations, and we have deployed a parsimonious approach for simplicity. Modern climate models generate many other economic outputs that could be useful as well, such as commodity prices, consumption, or unemployment rate. Depending on beliefs as to which variables drive risk, larger and potentially more complex models (e.g., non-linear) could be envisioned. This is left for future research.

A Additional material

A.1 Temperature anomaly

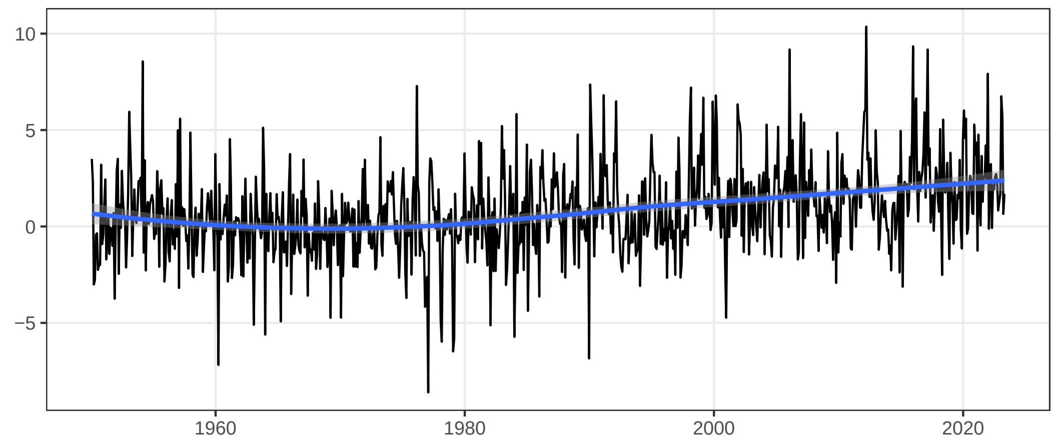

Figure 13: **Temperature anomaly in the US**. We plot the time-series of the simple average of the temperature anomaly of the 48 states in mainland US. The smoothed blue curve is the rolling moving average over 10 years.

A.2 Long-term returns and volatilities

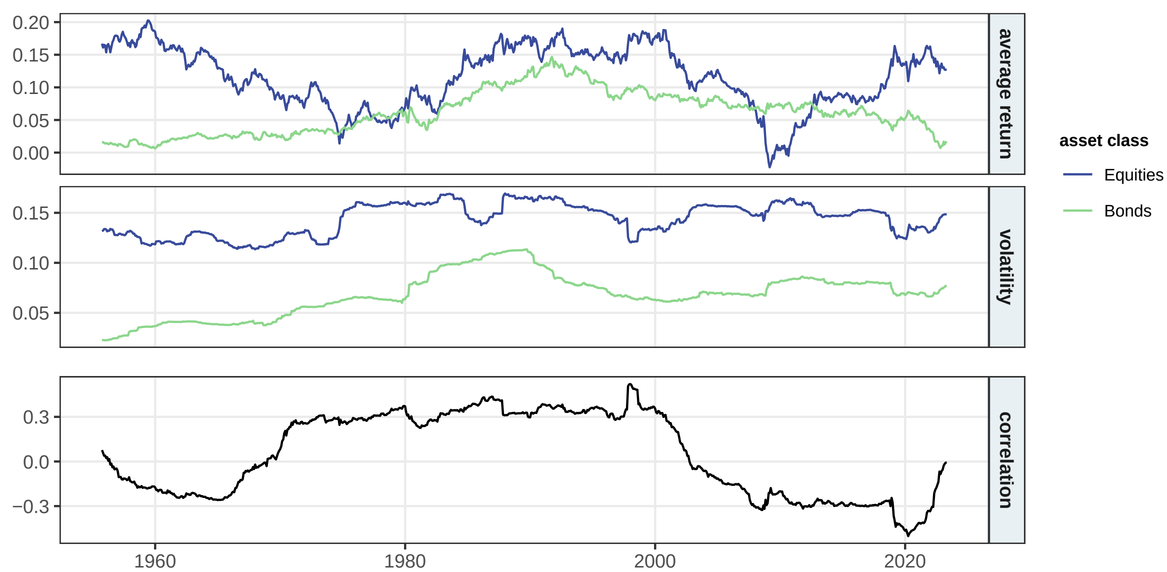

Figure 14: **Financial time-series**. We plot the long term average returns, volatilities and correlation for US equities and sovereign bonds. The first two are averaged over 10 years while the correlation is simply computed over a 10 year sample. Each point is a realized value and hence corresponds to the sample of the *prior* 10 years.

A.3 Macroeconomic indicators

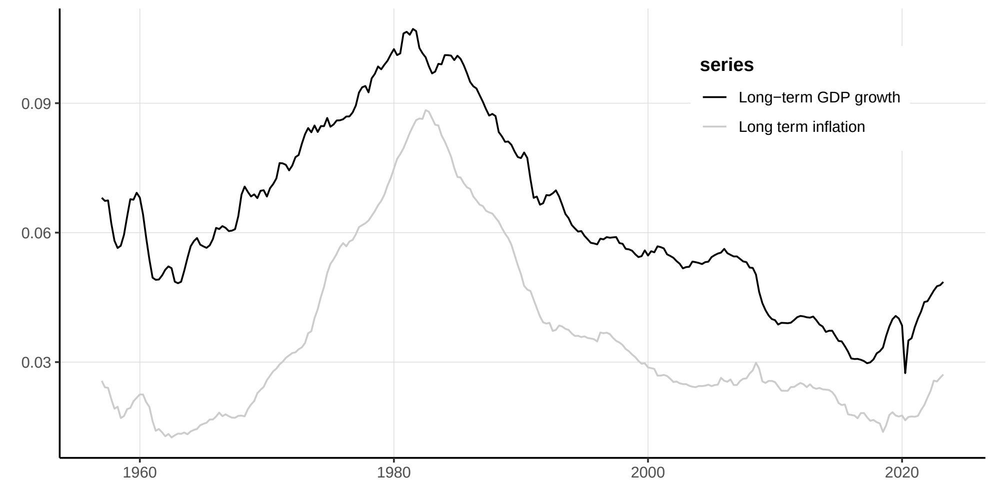

Figure 15: **Macro variables**. We plot the time-series of the two aggregate indicators we use to proxy the economic environment: long-term inflation and GDP growth rates (geometric averages). They are downloaded from the Federal Reserve of Saint Louis data portal (*GDP* and *CPIAUCSL* identifiers).

A.4 Climate adjustments for equity returns

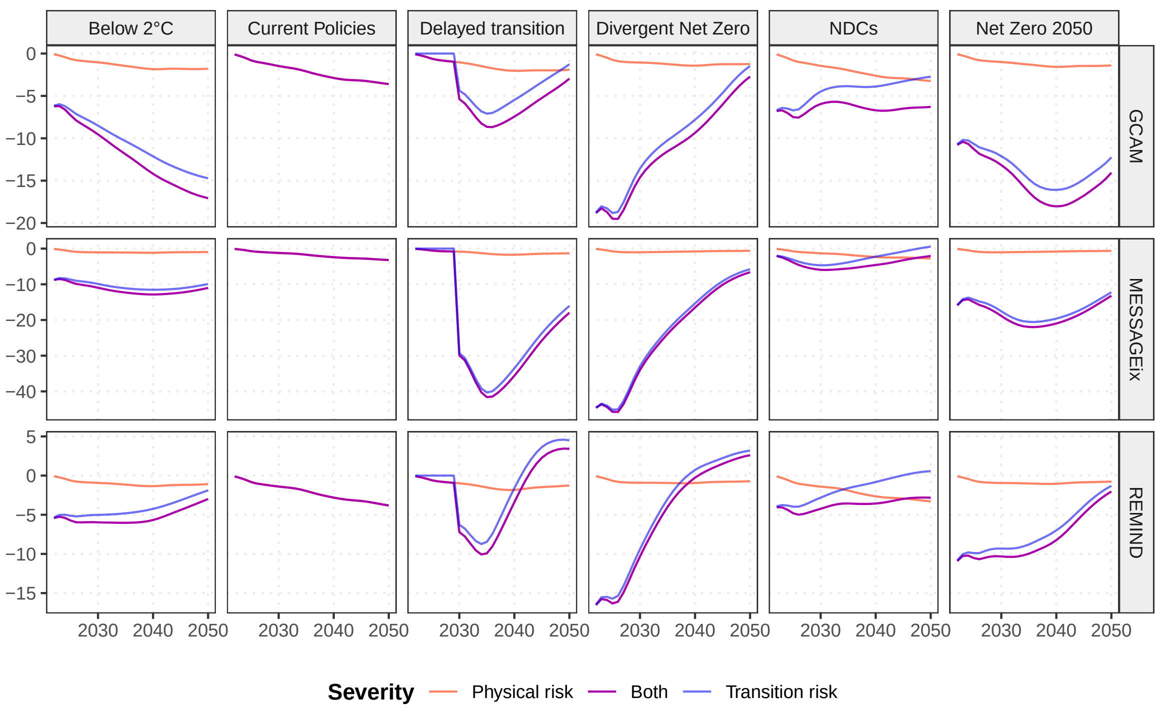

Figure 16: **Climate adjustments for equity returns**. We plot the time-series of correction for equity returns implied by climate models, scenarios and severities compared to the baseline situation (in %). The values are in percentages of adjustment to the baseline growth value. For the period 2023-2033, it is equal to +3.29%.

A.5 Volatility prediction with long-term returns

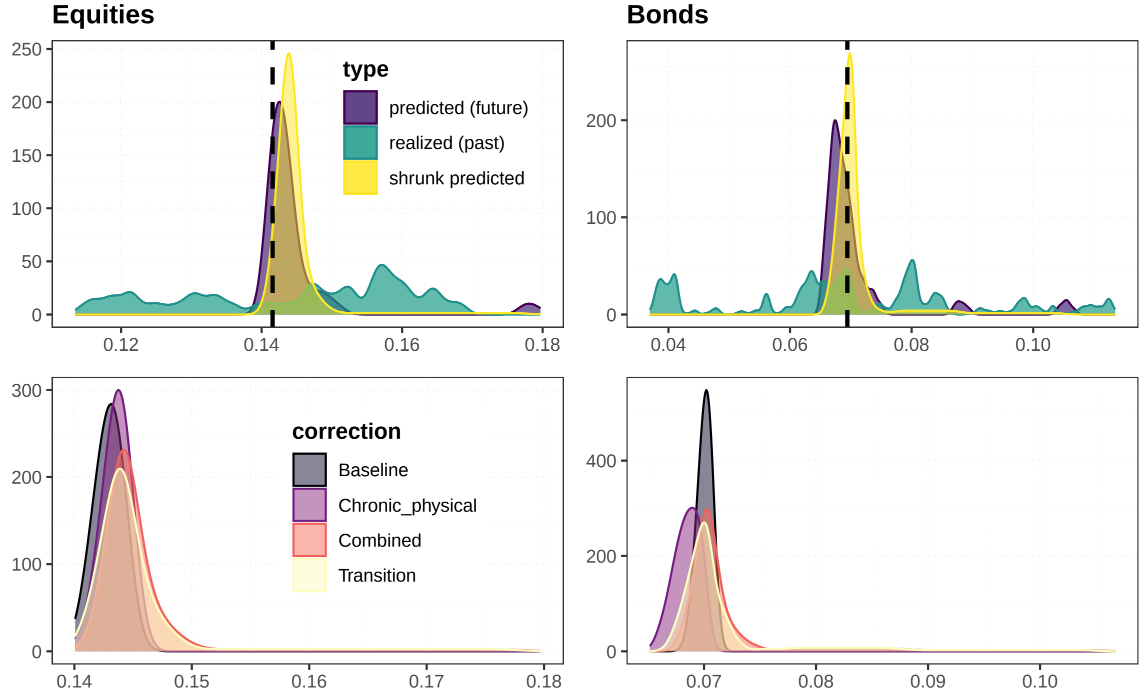

Figure 17: **Long-term volatility prediction - with long-term returns**. We plot the distribution of long-term forecasts for volatility across the two asset classes - from Model B panel 1 in Table 1 The upper panels compare historical values with the predicted ones while the lower ones focus on the latter but show the shift that occurs when considering various levels of climate severity. The priors used when post-shrinking in (6) are the historical means (1960-2023): 0.145 for equity volatility and 0.071 for bond volatility and the intensities α_i are between 0.0 and 0.8 with 0.01 increments. The vertical lines in the upper panels mark the average predictions in the *baseline* scenarios (no adjustment).

References

Acharya, V. V., R. Berner, R. F. Engle, H. Jung, J. Stroebel, X. Zeng, and Y. Zhao (2023). Climate stress testing. *SSRN Working Paper 4413834*.

Amihud, Y. and C. M. Hurvich (2004). Predictive regressions: A reduced-bias estimation method. *Journal of Financial and Quantitative Analysis 39*(4), 813–841.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *Journal of Finance 61*(1), 259–299.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2009). High idiosyncratic volatility and low returns: International and further US evidence. *Journal of Financial Economics 91*(1), 1–23.

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2022). Climate change concerns and the performance of green vs. brown stocks. *Management Science Forthcoming*.

Avramov, D., S. Cederburg, and K. Lučivjanská (2018). Are stocks riskier over the long run? taking cues from economic theory. *Review of Financial Studies 31*(2), 556–594.

Baker, D., J. B. De Long, and P. R. Krugman (2005). Asset returns and economic growth. *Brookings Papers on Economic Activity 2005*(1), 289–330.

Baker, M., B. Bradley, and J. Wurgler (2011). Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal 67*(1), 40–54.

Baker, S. R., N. Bloom, and S. J. Terry (2023). Using disasters to estimate the impact of uncertainty. *Review of Economic Studies Forthcoming*.

Bansal, R., M. Ochoa, and D. Kiku (2017). Climate change and growth risks. *SSRN Working Paper 2892433*.

Bansal, R., M. Ochoa, and D. Kiku (2021). Climate change risk. In C. V.V and B. Litterman (Eds.), *Climate Change Economics: The Role of Uncertainty and Risk*.

Barahhou, I., M. Ben Slimane, T. Roncalli, and N. Oulid Azouz (2022). Net zero investment portfolios-part 1. the comprehensive integrated approach. *SSRN Working Paper 4283998*.

Beveratos, A., J.-P. Bouchaud, S. Ciliberti, L. Laloux, Y. Lempérière, M. Potters, and G. Simon (2017). Deconstructing the low-vol anomaly. *Journal of Portfolio Management 44*(1), 91–103.

Bolton, P. and M. Kacperczyk (2021). Do investors care about carbon risk? *Journal of Financial Economics 142*(2), 517–549.

Bonato, M., O. Cepni, R. Gupta, and C. Pierdzioch (2023a). Climate risks and realized volatility of major commodity currency exchange rates. *Journal of Financial Markets 62*, 100760.

Bonato, M., O. Cepni, R. Gupta, and C. Pierdzioch (2023b). Climate risks and state-level stock market realized volatility. *Journal of Financial Markets*, 100854.

Boudoukh, J., R. Israel, and M. Richardson (2022). Biases in long-horizon predictive regressions. *Journal of Financial Economics 145*(3), 937–969.

Boudoukh, J. and M. Richardson (1993). Stock returns and inflation: A long-horizon perspective. *American Economic Review 83*(5), 1346–1355.

Bressan, G., A. Duranovic, I. Monasterolo, and S. Battiston (2023). Asset-level climate physical risk assessment is key for adaptation finance. *SSRN Working Paper 4062275*.

Bressan, G., I. Monasterolo, and S. Battiston (2022). Sustainable investing and climate transition risk: a portfolio rebalancing approach. *Journal of Portfolio Management 48*(10), 165–192.

Brixton, A., J. Brooks, P. Hecht, A. Ilmanen, T. Maloney, and N. McQuinn (2023). A changing stock–bond correlation: Drivers and implications. *Journal of Portfolio Management 49*(4), 64–80.

Campbell, J. Y. (1996). Understanding risk and return. *Journal of Political Economy 104*(2), 298–345.

Campbell, J. Y., C. Pflueger, and L. M. Viceira (2020). Macroeconomic drivers of bond and equity risks. *Journal of Political Economy 128*(8), 3148–3185.

Campiglio, E., L. Daumas, P. Monnin, and A. von Jagow (2023). Climate-related risks in financial assets. *Journal of Economic Surveys Forthcoming*.

Campos-Martins, S. and D. F. Hendry (2023). Common volatility shocks driven by the global carbon transition. *Journal of Econometrics Forthcoming*.

Cardinale, M., N. Y. Naik, and V. Sharma (2021). Forecasting long-horizon volatility for strategic asset allocation. *Journal of Portfolio Management 47*(4), 83–98.

Casey, G., S. Fried, and E. Goode (2023). Projecting the impact of rising temperatures: The role of macroeconomic dynamics. *IMF Economic Review*, 1–31.

Chopra, V. K. and W. T. Ziemba (1993). The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management 19*(2), 6–11.

Coqueret, G. (2015). Diversified minimum-variance portfolios. *Annals of Finance 11*(2), 221–241.

Coqueret, G. (2022). *Perspectives in Sustainable Equity Investing*. CRC Press.

Cosemans, M., X. Hut, and M. A. Van Dijk (2022). Climate change and long-horizon portfolio choice: Combining insights from theory and empirics. *SSRN Working Paper 3920481*.

Davis, J., C. Fuenzalida, L. Huetsch, B. Mills, and A. M. Taylor (4603121). Global natural rates in the long run: Postwar macro trends and the market-implied r* in 10 advanced economies. *SSRN Working Paper w31787*.

Dietz, S., A. Bowen, C. Dixon, and P. Gradwell (2016). 'climate value at risk'of global financial assets. *Nature Climate Change 6*(7), 676–679.

Duffee, G. R. (2023). Macroeconomic news and stock–bond comovement. *Review of Finance 27*(5), 1859–1882.

Engle, R. F., E. Ghysels, and B. Sohn (2013). Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics 95*(3), 776–797.

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

Faiella, I., E. Bernardini, J. Di Giampaolo, S. Letta, D. Nasti, and M. Fruzzetti (2023). The exposure of investments to climate and environmental risks. In *Financial Risk Management and Climate Change Risk: The Experience in a Central Bank*, pp. 277–312. Springer.

Faugère, C. and J. Van Erlach (2006). The equity premium: consistent with GDP growth and portfolio insurance. *Financial Review 41*(4), 547–564.

Fuchs, L., H. Nguyen, T. Nguyen, and K. Schaeck (2023). Climate stress test, bank lending, and the transition to the carbon-neutral economy. *SSRN Working Paper 4427729*.

Ge, W., P. Lalbakhsh, L. Isai, A. Lenskiy, and H. Suominen (2022). Neural network–based financial volatility forecasting: A systematic review. *ACM Computing Surveys (CSUR) 55*(1), 1–30.

Gehr, A. K. (1979). Risk and return. *Journal of Finance 34*(4), 1027–1030.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance 48*(5), 1779–1801.

Gostlow, G. (2022). Pricing physical climate risk in the cross-section of returns. *SSRN Working Paper 3501013*.

Hantzsche, A., M. Lopresto, and G. Young (2018). Using NiGEM in uncertain times: Introduction and overview of NiGEM. *National Institute Economic Review 244*, R1–R14.

Harrison, P. and H. H. Zhang (1999). An investigation of the risk and return relation at long horizons. *Review of Economics and Statistics 81*(3), 399–408.

Hilaire, J. and C. Bertram (2020). The REMIND-MAgPIE model and scenarios for transition risk analysis.

Hodges, P., H. Ren, K. Schwaiger, and A. Ang (2022). Net-zero investing for multi-asset portfolios seeking to satisfy paris-aligned benchmark requirements with climate alpha signals. *Journal of Portfolio Management 48*(4), 33–58.

Hoevenaars, R. P., R. D. Molenaar, P. C. Schotman, and T. B. Steenkamp (2008). Strategic asset allocation with liabilities: Beyond stocks and bonds. *Journal of Economic Dynamics and Control 32*(9), 2939–2970.

Jagannathan, R. and T. Ma (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance 58*(4), 1651–1683.

Kahn, M. E., K. Mohaddes, R. N. Ng, M. H. Pesaran, M. Raissi, and J.-C. Yang (2021). Long-term macroeconomic effects of climate change: A cross-country analysis. *Energy Economics 104*, 105624.

Kalkuhl, M. and L. Wenz (2020). The impact of climate conditions on economic production. Evidence from a global panel of regions. *Journal of Environmental Economics and Management 103*, 102360.

Le Guenedal, T. and T. Roncalli (2022). Portfolio construction with climate risk measures. *SSRN Working Paper 3999971*.

Le Tran, V., T. Leirvik, M. Parschat, and P. Schive (2022). Climate-change risk and stocks' return. *SSRN Working Paper 4306875*.

Li, X., R. N. Sullivan, and L. Garcia-Feijoo (2016). The low-volatility anomaly: Market evidence on systematic risk vs. mispricing. *Financial Analysts Journal 72*(1), 36–47.

Linton, O. and B. Perron (2003). The shape of the risk premium: evidence from a semi-parametric generalized autoregressive conditional heteroscedasticity model. *Journal of Business & Economic Statistics 21*(3), 354–367.

Ma, R., B. R. Marshall, N. H. Nguyen, and N. Visaltanachoti (2023). Estimating long-term expected returns. *SSRN Working Paper 4493448*.

Malkiel, B. G. and Y. Xu (1997). Risk and return revisited. *Journal of Portfolio Management 23*(3), 9.

Molenaar, R., E. Senechal, L. Swinkels, and Z. Wang (2023). Empirical evidence on the stock-bond correlation. *SSRN Working Paper 4514947*.

Newell, R. G., B. C. Prest, and S. E. Sexton (2021). The gdp-temperature relationship: implications for climate change damages. *Journal of Environmental Economics and Management 108*, 102445.

Nordhaus, W. (2014). Estimates of the social cost of carbon: concepts and results from the dice-2013r model and alternative approaches. *Journal of the Association of Environmental and Resource Economists 1*(1/2), 273–312.

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences 114*(7), 1518–1523.

Poon, S.-H. and C. W. J. Granger (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature 41*(2), 478–539.

Prosperi, L. and L. Zanin (2023). A modelling framework for projections of equity portfolio returns under climate transition scenarios. *SSRN Working Paper 4062090*.

Rebonato, R. (2023). Asleep at the wheel? The risk of sudden price adjustments for climate risk. *Journal of Portfolio Management Forthcoming*.

Reinders, H. J., D. Schoenmaker, and M. van Dijk (2023a). A finance approach to climate stress testing. *Journal of International Money and Finance 131*, 102797.

Reinders, H. J., D. Schoenmaker, and M. A. van Dijk (2023b). Climate risk stress testing: A conceptual review. *SSRN Working Paper 4362342*.

Salisu, A. A., C. Pierdzioch, R. Gupta, and R. van Eyden (2022). Climate risks and us stock-market tail risks: A forecasting experiment using over a century of data. *International Review of Finance*.

Shen, S., A. LaPlante, and A. Rubtsov (2019). Strategic asset allocation with climate change. *SSRN Working Paper 3249211*.

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics 54*(3), 375–421.

Tsai, J. and J. A. Wachter (2015). Disaster risk and its implications for asset pricing. *Annual Review of Financial Economics 7*, 219–252.

van Ruijven, B. and J. Min (2020). The MESSAGEix-GLOBIOM model and scenarios for transition risk analysis.

Venturini, A. (2022). Climate change, risk factors and stock returns: A review of the literature. *International Review of Financial Analysis 79*, 101934.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance 68*(3), 987–1035.

Whitelaw, R. F. (2000). Stock market risk and return: An equilibrium approach. *Review of Financial Studies 13*(3), 521–547.

Winter, D. J. and M. Kiehl (2023). Long-term macroeconomic effects of shifting temperature anomaly distributions. *SSRN Working Paper 4553820*.

Wise, M., K. Calvin, P. Kyle, P. Luckow, and J. Edmonds (2014). Economic and physical modeling of land use in GCAM 3.0 and an application to agricultural productivity, land, and terrestrial carbon. *Climate Change Economics 5*(02), 1450003.

Zhao, X., M. Gerety, and N. V. Kuminoff (2018). Revisiting the temperature-economic growth relationship using global subnational data. *Journal of Environmental Management 223*, 537–544.