Climate Change and Long-Horizon Portfolio Choice: Combining Insights from Theory and Empirics*

Mathijs Cosemans†

Xander Hut‡

Mathijs van Dijk§

May 13, 2022

Abstract

We propose a novel approach for incorporating climate risk in long-horizon portfolio choice. Our method combines information from historical data about the impact of climate change on return dynamics with prior beliefs elicited from a temperature long-run risks (LRR-T) model. Relative to investors with no prior views about climate change effects, investors with LRR-T beliefs perceive stocks to be much riskier over longer horizons, because climate change increases estimation risk and uncertainty about future expected returns and weakens their beliefs in mean reversion. Because the perceived increase in long-horizon risk is not fully offset by higher equity premiums, investors' optimal allocation to equities decreases sharply with the horizon.

JEL classifications: C11, G11, G12, G23

Keywords: climate change, long-run risks, prior beliefs, optimal portfolio choice, uncertainty

*For helpful comments and suggestions, we thank Raphael Abiry (discussant), Kevin Andrew (discussant), Doron Avramov, Servaas van Bilsen (discussant), Sebastian Gehricke (discussant), Felix Miebs (discussant), Marta Szymanowska, Jules Tinang, Bas Werker, and seminar participants at the University of Edinburgh, the 2022 INQUIRE UK Spring Seminar, the 2021 Australasian Finance and Banking Conference, 2021 GRASFI Conference, 2021 CEF-Group Climate Finance Symposium, 2021 FMA Annual Meeting, 2021 Netspar Pension Day, 2021 IWFSAS Conference on Financial System Stability, 2021 $(EC)^2$ Conference on Econometrics of Climate, 2021 New Zealand Finance Meeting, and the 2021 CefES Conference on European Studies. This work made use of the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-2954, which is partly financed by the Dutch Research Council (NWO). We thank Netspar and INQUIRE Europe for financial support.

†Rotterdam School of Management, Erasmus University, E-mail: mcosemans@rsm.nl.

‡Rotterdam School of Management, Erasmus University, E-mail: hut@rsm.nl.

§Rotterdam School of Management, Erasmus University, E-mail: madijk@rsm.nl.

1 Introduction

Investors are increasingly aware that climate change can pose significant investment risks that should be integrated in their portfolio strategies.[1] Regulators also recognize that financial institutions are exposed to environmental risk factors and encourage them to quantify these risks. For instance, new European regulations (IORP II) require pension funds to include climate risk in their own-risk assessment. However, despite their efforts, investors struggle to fully grasp the potential impact of climate change on the value of their portfolio and are searching for approaches to quantify these risks (see, e.g., the survey evidence provided by Krueger, Sautner, and Starks (2020)).

Given the long-term nature of climate change, long-horizon investors such as pension funds are concerned about its effects on asset returns and risk over longer holding periods. Estimating the long-horizon climate risk exposure of assets is challenging because currently available historical data samples may not include sufficient realizations of severe climate change effects and because of substantial uncertainty about the long-horizon impact of climate change on asset prices. Stambaugh summarizes this uncertainty by noting that when we "expand the horizon to the next several decades, the possible effects of global warming range from negligible to catastrophic."[2]

We address these issues by proposing a novel approach for measuring the impact of climate change on long-horizon equity risk and portfolio choice that supplements historical data with prior information derived from economic theory. We set up a Vector Autoregression (VAR) model to analyze the long-run dynamics of equity returns and include temperature as a predictor in the VAR to proxy for climate risk. Since the frequency and impact of past climate-induced disasters may not be representative of their future impact in a scenario of prolonged climate change, we estimate the VAR parameters using a Bayesian approach that combines historical data with theoretical prior beliefs about the impact of temperature change on return dynamics. We construct the model-based prior following the method proposed by Avramov, Cederburg, and Lucivjanska (2018).[3]

[1] In his 2020 letter to CEOs available at www.blackrock.com/uk/individual/larry-fink-ceo-letter, Larry Fink notes that "climate change is almost invariably the top issue that clients around the world raise with BlackRock."

[2] This quote is from https://www.nytimes.com/2009/03/29/your-money/stocks-and-bonds/29stra.html.

[3] Avramov, Cederburg, and Lucivjanska (2018) point out that "asset pricing theory could provide additional guidance about important aspects of the return process for which the sample evidence is not particularly informative."

We elicit these beliefs from the temperature long-run risks (LRR-T) model of Bansal, Kiku, and Ochoa (2019), in which rising temperatures influence asset prices by increasing the likelihood of future climate-driven natural disasters that lower economic growth.[4] Because of investors' concerns about the consequences of temperature increases for future growth, climate risk is reflected in current asset prices even though the impact of climate change in historical data is limited.[5] By imposing a structure on the relation between temperature and returns, the LRR-T model provides prior information about the impact of climate change on future return dynamics.

Our Bayesian approach to incorporating climate risk into portfolio choice has several advantages. First, it allows for the accommodation of various views about the impact of climate change on return dynamics. We consider four different investor types. The *dogmatic* investor is convinced that climate change affects returns in the way implied by the temperature long-run risks model and assigns full weight to the model-based prior belief. The *agnostic* investor has no prior views about the effect of temperature change on future returns and lets the data speak by giving full weight to the return dynamics implied by the data. The *Bayesian* investor has faith in her prior beliefs derived from economic theory, but is aware that these beliefs may be inaccurate and updates them based on observed data. In our implementation, she assigns equal weights to the LRR-T prior and the data. For comparison, we also consider a Bayesian investor with prior beliefs derived from the original LRR model of Bansal and Yaron (2004) that does not feature climate change. We quantify the effect of these prior beliefs on the perceived riskiness of stock markets over different holding periods and on the optimal portfolio choice for long-term investors.

Second, the Bayesian framework allows us to assess the impact of uncertainty about the financial impact of climate change on optimal portfolio choice. We explicitly incorporate parameter uncer-

[4]Colacito, Hoffmann, and Phan (2019) show that rising temperatures can reduce U.S. economic growth by up to one-third over the next century. Dell, Jones, and Olken (2012) find that a temperature increase of 1 degree Celsius in a given year reduces economic growth in poor countries by 1.3 percentage points on average. Dell, Jones, and Olken (2014) provide an overview of work on the economic impact of temperature change.

[5]The LRR-T model captures the effect of physical climate risk on economic growth and asset prices, i.e., the systematic risk resulting from more frequent extreme weather events due to climate change. Investment portfolios can also be exposed to transition risks that arise when moving towards a low-carbon economy. Temperature change can affect these transition risks by increasing the likelihood of changes in government policies that increase a company's cost of doing business or decrease its asset values, such as the introduction of a carbon tax or the risk of stranded assets when known fossil fuel reserves cannot be burned. Although these transition risks are not explicitly captured by the LRR-T model, they are incorporated in our analysis to the extent that they are reflected in past return data.

tainty into the portfolio optimization by integrating over the posterior distribution of parameters in the return-forecasting model to form a predictive distribution for returns. In contrast, traditional frequentist methods treat the point estimates of the model parameters as known, thereby completely ignoring estimation risk. To quantify the effect of parameter uncertainty, we compare the portfolios formed based on the predictive distributions with and without parameter uncertainty.

Our main results are as follows. First, we show that Bayesian investors with beliefs formed based on the temperature long-run risks model perceive stocks to be much riskier over longer horizons than investors with uninformative beliefs or investors with prior beliefs based on the original LRR model that does not incorporate climate change. At a horizon of 10 years (25 years), the LRR-T investors perceive the annualized variance to be 1.1 (1.4) times the one-year variance. For these investors stocks therefore appear more volatile in the long run than in the short run. In contrast, investors with uninformative priors or LRR priors perceive the annualized variance to be 0.50 or 0.75 times the one-year variance at a 25-year horizon, respectively, so for these investors stocks appear safer in the long run. Investors with LRR-T priors consider stocks to be riskier over long holding periods because climate change increases estimation risk and uncertainty about future returns and weakens their beliefs in mean reversion. Intuitively, because shocks to temperature are persistent in the LRR-T model, climate-induced disasters tend to occur in clusters during high-temperature periods. The model therefore predicts that disasters are likely followed by other disasters, which implies that negative shocks to consumption and dividends are more likely to be followed by further negative shocks than by positive shocks, thereby reducing mean reversion in returns.

Second, we find that investors with LRR-T beliefs predict higher equity risk premia at all horizons than investors with LRR beliefs or investors with uninformative priors about the VAR parameters. The higher equity risk premium is mainly due to lower risk-free returns. In particular, because in the LRR-T model disasters have a long-lasting effect on future economic growth, agents increase their precautionary savings and have higher demand for the risk-free asset, which lowers its return. Because the risk-free rate is persistent and temperature is expected to keep increasing in the future, LRR-T investors predict decreasing risk-free returns over the horizon. In contrast, investors who base their forecasts purely on historical data expect risk-free rates to increase over time from their current low levels to their historical average. LRR-T investors further predict that market returns decrease with the horizon because rising temperatures increase the likelihood of disasters that depress stock markets. However, because the decrease in market returns is more than offset by the decrease in risk-free returns, the market risk premium increases with the horizon.

Third, we show that incorporating prior information about the impact of climate change on asset prices has a significant impact on the optimal portfolio choice of long-horizon buy-and-hold investors who allocate wealth between the market portfolio of stocks and the risk-free asset. Because the higher equity premium does not fully offset the perceived increase in long-horizon equity risk, investors with LRR-T beliefs allocate much less capital to equities at long horizons than investors with no prior views on the impact of climate change. For a horizon of 25 years, the equity allocation of LRR-T investors is 10 percentage points lower than that of the other investors. We further find that for all investor types, the optimal allocation to stocks starts to decrease for horizons longer than 10 years due to increased parameter uncertainty, corroborating the conclusion of previous work that it is important to account for parameter uncertainty when forming optimal portfolios.[6]

Our work adds to the growing literature that studies how climate change affects asset prices.[7] Existing work shows that asset prices are affected by carbon emissions (Bolton and Kacperczyk (2021, 2022), Ilhan, Sautner, and Vilkov (2021)), temperature increases (Balvers, Du, and Zhao (2017), Barnett (2021), Addoum, Ng, and Ortiz-Bobea (2020), Kumar, Xin, and Zhang (2019)), drought periods (Hong, Li, and Xu (2019)), and broader climate-related disasters (Correa, He, Herpfer, and Lel (2021), Huynh and Xia (2021)). Whereas these papers focus on the consequences of climate change for short-term asset returns, we study its implications for long-term returns.

Our paper also contributes to the literature on the term structure of equity risk and its implications for optimal portfolio choice. Barberis (2000) finds that equity risk decreases with the horizon in the presence of mean reversion in stock returns. As a result, long-term investors should have a higher allocation to stocks than short-term investors (Campbell, Chan, and Viceira (2003), Campbell and Viceira (2005)). However, Pastor and Stambaugh (2012) and Johannes, Korteweg,

[6]See, e.g., Kandel and Stambaugh (1996), Barberis (2000), Pastor and Stambaugh (2012), Johannes, Korteweg, and Polson (2014).

[7]Giglio, Kelly, and Stroebel (2021) provide a recent overview of the climate finance literature.

and Polson (2014) demonstrate that stocks become riskier over longer horizons because the effect of mean reversion is more than offset by estimation risk and uncertainty about current and future expected returns. Avramov, Cederburg, and Lucivjanska (2018) and Carvalho, Lopes, and McCulloch (2018) show that the term structure of predictive variance crucially depends on the economic model used to form prior beliefs about return dynamics. We contribute to this literature by incorporating climate change as a new source of long-horizon risk and by studying its impact on long-horizon portfolio choice.

The paper proceeds as follows. Section 2 introduces our approach for modeling the impact of climate change on stock return dynamics. Section 3 explains how we elicit prior beliefs about return dynamics from the LRR-T model. Section 4 presents the long-horizon risk and return forecasts and the implications of climate change for optimal portfolio choice. Section 5 concludes.

2 Long-horizon return dynamics and portfolio choice

2.1 Predictive vector autoregression framework

We estimate a first-order Vector Autoregression (VAR) model to analyze the long-horizon dynamics of equity risk and returns, as is common in the literature (e.g., Barberis (2000) and Campbell, Chan, and Viceira (2003)). To capture climate risk, we augment the model with the temperature anomaly, defined as the deviation of the temperature in a specific period from the long-term average temperature calculated over a base period. The VAR structure therefore incorporates information about the impact of temperature on future stock returns. Specifically, the VAR model is given by

$$\begin{bmatrix} r_{i,t+1} \\ p_{t+1} - d_{t+1} \\ r_{f,t+1} \\ T_{t+1} \end{bmatrix} = a + B \begin{bmatrix} p_t - d_t \\ r_{f,t} \\ T_t \end{bmatrix} + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, \Sigma), \tag{1}$$

where $r_{i,t+1}$ is the log return on the market portfolio of stocks, $p_{t+1} - d_{t+1}$ is the log price-dividend ratio of the market portfolio, $r_{f,t+1}$ is the risk-free rate, and T_{t+1} is the temperature anomaly.

The first equation in the VAR system models expected stock market returns as a function of the predictor or state variables. The other equations specify the dynamics of the predictor variables. The price-dividend ratio is commonly used as a return predictor in prior work (e.g., Fama and French (1988), Campbell and Shiller (1988)). The inclusion of the price-dividend ratio and the risk-free rate in the VAR is also motivated by the temperature long-run risks (LRR-T) model discussed in Section 3, which implies that both variables are directly related to expected stock market returns. The LRR-T model also predicts that the temperature anomaly is a useful state variable for predicting future returns.

2.2 Predictive VAR estimation

We estimate the predictive VAR in Equation (1) using Bayesian methods that incorporate prior information about the VAR parameters. Most important for our purposes, the Bayesian approach allows for the accommodation of different views about the impact of climate change on the dynamics of asset returns and state variables. In the paper, we consider four types of investors with different prior beliefs about the VAR parameters. Table 1 summarizes the different investor types.

The first investor is the *agnostic* investor who has no prior views about the relation between the predictors and future asset returns and assigns full weight to the return dynamics implied by the historical data. For this investor we specify an uninformative prior for the VAR parameters and estimate the VAR following the standard approach of Barberis (2000).

The second investor type is the *dogmatic* investor who forms beliefs about the VAR parameters based on the implications of an asset pricing model. In our implementation, investors elicit prior beliefs about the joint dynamics of returns and predictor variables from the temperature long-run risks (LRR-T) model. The dogmatic investor does not use any information embedded in the historical data and therefore gives full weight to the dynamics implied by the model-based prior.

Because the LRR-T model does not yield a closed-form expression for the VAR parameters, we implement the procedure of Del Negro and Schorfheide (2004) and Avramov, Cederburg, and Lucivjanska (2018) to construct the model-based prior distributions for the VAR parameters. We first simulate quarterly stock returns, price-dividend ratios, risk-free rates, and temperature anomalies from the LRR-T model. Conceptually, we form the model-based prior by estimating the VAR using the simulated data. Specifically, for each variable we simulate 250,000 time series of 292 quarters to match the length of our historical sample (1947Q1-2019Q4). We then estimate the population moments of the variables in the VAR model by calculating the averages of the sample moments across the 250,000 simulations. We use these population moments to construct the model-based prior distribution for the VAR parameters. Intuitively, the dogmatic investor forms the model-based prior based on a pseudosample of data from the LRR-T model that has the same size as the historical data. Following Avramov, Cederburg, and Lucivjanska (2018), we control the precision of the prior belief by scaling the population moments to match the length of our sample period. The investor with the model-based prior therefore allows for potential misspecification of the LRR-T model, such that ex ante she does not know the parameters of the VAR model with certainty.

The third investor we consider is the *Bayesian LRR-T* investor who combines information from the LRR-T based prior and historical data to make inferences about long-horizon risk and return. This investor essentially shrinks the VAR parameters estimated using historical data to the VAR parameters implied by the LRR-T model. By scaling the population moments in the prior distribution, we let the investor assign equal weight to the prior and the historical data. Further details about the VAR estimation procedure are provided in Section A.3 of the Online Appendix.

The fourth investor type is the *Bayesian LRR* investor who derives prior beliefs about the VAR parameters from the LRR model of Bansal and Yaron (2004) that does not incorporate climate risk and updates these beliefs upon observing the historical data. This investor completely ignores the impact of climate change on return dynamics and therefore does not include the temperature anomaly as a predictor in the VAR in Equation (1). The Bayesian investor with LRR beliefs forms the benchmark for the Bayesian investor with LRR-T beliefs. By comparing the results for these two investor types we can examine the effect of accounting for temperature change on inferences about long-horizon return dynamics and on the optimal portfolio choice of long-term investors.

Apart from having different beliefs about the parameters of the VAR model, the four investors also differ in terms of the predictor values they use for making out-of-sample forecasts at the end of the sample period. We allow for these differences to ensure that the predictor values used by each investor are consistent with her views about the relative informativeness of the sample data and the model-based prior. The *agnostic* investor constructs forecasts based on the mean values of the price-dividend ratio (p/d) and risk-free rate (r_f) observed over the last five years of the historical sample. We use the five-year average instead of the end-of-sample value to smooth out short-term fluctuations in the predictors that are uninformative for long-horizon dynamics. The *dogmatic* investor makes forecasts based on the mean value of p/d and r_f computed over the last five years of data simulated from the LRR-T model.[8] The *Bayesian LRR(-T)* investor forms forecasts based on the simple average of the five-year mean of p/d and r_f implied by the LRR(-T) model and their five-year sample mean. All investors except for the *Bayesian LRR* investor generate forecasts based on the average temperature anomaly observed over the last five years in the sample and use the exogenous process described in the next section to project future temperature changes.

[Table 1 about here.]

2.3 Climate change scenario

The standard VAR framework is useful for modeling the dynamics of asset returns and financial state variables. However, it is less suitable for modeling climate change because the state variables included in the VAR have no direct theoretical link to future temperature. Figure 1 illustrates this point by plotting the temperature anomaly forecasts implied by the VAR models corresponding to each investor type. These projected temperature paths are not in line with climate change scenarios constructed by the IPCC (2021). For example, the intermediate climate change scenario of the IPCC (SSP2-4.5) predicts global warming of 2 degrees Celsius for the mid-term period 2041-2060 that corresponds to the end of our 25-year forecast horizon. In fact, the two VAR models estimated using historical data predict the temperature anomaly to *decrease* over the next 25 years. The LRR-T model does project rising temperatures but the magnitude of the increase is smaller than in most climate change scenarios deemed plausible by the IPCC and agencies such as NASA.[9]

We therefore disregard the temperature forecasts implied by the VARs and specify an exogenous

[8] As an alternative, we also considered forecasts based on the *population* means of p/d and r_f computed from the 250,000 simulations of the LRR-T model. This yields similar conclusions.

[9] See, e.g., the climate projections of the NASA Ames Research Center at https://www.nasa.gov/nex/gddp.

temperature process that is consistent with the intermediate climate change scenario of the IPCC. Using the same temperature projection for all investor types also ensures that differences in results across investors are purely driven by their different views about the impact of climate change on future returns and not by differences in temperature forecasts. We thus assume that all investors know the true parameters of the process that governs temperature change and do not disagree about the existence of climate change, but have different beliefs about its economic impact.[10]

We model the future temperature anomaly using the following autoregressive process:

$$T_{t+1} = \gamma + \rho T_t + \eta_{t+1}, \qquad \eta_{t+1} \sim N(0, \sigma_\eta^2), \tag{2}$$

with ρ set equal to the persistence of the temperature process in the LRR-T model (0.9971) and γ chosen such that the temperature anomaly at the end of our forecast horizon (k=100 quarters) is equal to 2 degrees.[11] We set σ_η^2 equal to 0.001 to match the *very likely* (5 - 95%) range of the temperature anomaly in the intermediate scenario defined by the IPCC for the mid-term period 2041-2060.[12] We simulate 250,000 future temperature anomaly paths based on Equation (2). Figure 1 plots the mean and the 5% and 95% percentile values of these simulated trajectories.

[Figure 1 about here.]

2.4 Long-horizon risk and return

We draw the parameters of the predictive VAR in Equation (1) from their posterior VAR distribution in Appendix A.3. We compute the predictive variance of returns conditional on each of these draws of the VAR parameters and the observations of the state variables at time t. The predictive variance for the k-period return from time N to time $N+k$ can be decomposed as:

$$\text{Var}(r_{i,N,N+k} \mid D_N) = \mathbb{E}[\text{Var}(r_{i,N,N+k} \mid \Theta, D_N|D_N] + \text{Var}[\mathbb{E}(r_{i,N,N+k} \mid \Theta, D_N|D_N], \tag{3}$$

[10]Our framework can be generalized to account for uncertainty about the parameters of the temperature process. We do not incorporate this type of parameter uncertainty because our goal is to study the long-horizon implications of different beliefs about the economic *impact* of climate change rather than the *magnitude* of this change.

[11]In particular, we specify $\gamma = (T_{N+100} - \rho^{100}\bar{T}_{N-19,N}) \times (1-\rho)/(1-\rho^{100})$, with T_{N+100} set equal to 2 and $\bar{T}_{N-19,N}$ equal to 1.18, which is the average temperature anomaly over the last five years (20 quarters) in our sample.

[12]The IPCC (2021) reports a best estimate of 2 degrees Celsius with a *very likely* range of 1.6 to 2.5 degrees.

where $r_{i,N,N+k} = r_{i,N+1} + \cdots + r_{i,N+k}$ is the cumulative log return on asset i over horizon k, Θ denotes the set of parameters in the VAR model in Equation (1) and D_N is the data observable at time N. The first term is the expectation of the conditional variance of k-period returns. The second term is the variance of the conditional expected k-period return implied by the VAR, which accounts for uncertainty about the VAR parameters. This term captures the notion that investors perceive returns to be riskier because they do not know the true values of the VAR parameters.

Avramov, Cederburg, and Lucivjanska (2018) demonstrate that the conditional variance of long-horizon asset returns in the predictive VAR framework can be written as:

$$\begin{aligned}
&\mathrm{Var}(r_{i,N,N+k} \mid D_N) \\
&= \underbrace{\mathbb{E}[k\sigma_{r_i}^2 \mid D_N]}_{\text{i.i.d. uncertainty}} + \underbrace{\mathbb{E}\left[\sum_{j=1}^{k-1} 2b_{r_i}(I - B_x)^{-1}(I - B_x^j)\Sigma_{xr_i} \mid D_N\right]}_{\text{mean reversion}} \\
&+ \underbrace{\mathbb{E}\left[\sum_{j=1}^{k-1}(b_{r_i}(I - B_x)^{-1}(I - B_x^j))\Sigma_x(b_{r_i}(I - B_x)^{-1}(I - B_x^j))' \mid D_N\right]}_{\text{uncertainty about future expected returns}} \\
&+ \underbrace{\mathrm{Var}\left[ka_{r_i} + b_{r_i}(I - B_x)^{-1}\left((kI - (I - B_x)^{-1}(I - B_x^k))a_x + (I - B_x^k)x_N^{(l)}\right) \mid D_N\right]}_{\text{estimation risk}},
\end{aligned} \tag{4}$$

where $x_N^{(l)}$ corresponds to the predictor values used by investor type l for constructing out-of-sample forecasts at time N, as described in Section 2.2. The elements of the set of VAR parameters a, B, and Σ in Equation (1) are defined as follows:

$$\begin{aligned}
&a = \begin{bmatrix} a_{r_i} & a_x' \end{bmatrix}', \quad a_x = \begin{bmatrix} a_{pd} & a_{r_f} & a_T \end{bmatrix}', \\
&B = \begin{bmatrix} b_{r_i} \\ B_x \end{bmatrix}, \quad b_{r_i} = \begin{bmatrix} b_{r_i,pd} & b_{r_i,r_f} & b_{r_i,T} \end{bmatrix}, \quad B_x = \begin{bmatrix} b_{pd,pd} & b_{pd,r_f} & b_{pd,T} \\ b_{r_f,pd} & b_{r_f,r_f} & b_{r_f,T} \\ b_{T,pd} & b_{T,r_f} & b_{T,T} \end{bmatrix},
\end{aligned} \tag{5}$$

and

$$\Sigma = \begin{bmatrix} \sigma^2_{r_i} & \Sigma'_{xr_i} \\ \Sigma_{xr_i} & \Sigma_x \end{bmatrix}, \quad \Sigma_{xr_i} = \begin{bmatrix} \sigma_{r_i,pd} \\ \sigma_{r_i,r_f} \\ \sigma_{r_i,T} \end{bmatrix}, \quad \Sigma_x = \begin{bmatrix} \sigma^2_{pd} & \sigma_{pd,r_f} & \sigma_{pd,T} \\ \sigma_{r_f,pd} & \sigma^2_{r_f} & \sigma'_{r_f,T} \\ \sigma_{T,pd} & \sigma_{T,r_f} & \sigma^2_T \end{bmatrix}. \tag{6}$$

The conditional variance in Equation (4) consists of four components.[13] The first term captures the uncertainty from i.i.d. shocks and contributes the same variance per period at all investment horizons, i.e., the variance of the cumulative k-period return increase linearly with k. The second term reflects mean reversion in returns that arises when a negative shock to returns is offset by positive shocks to expected returns and vice versa, thereby reducing long-horizon variance. The third term reflects the uncertainty about future expected returns that contributes positively to predictive variance and increases with the horizon k. This component is due to time variation in conditional expected returns and exists even if investors would know the model parameters.[14]

The fourth and last term can be interpreted as estimation risk and reflects the fact that investors do not know the values of the parameters governing the joint dynamics of returns and predictors. Estimation risk contributes positively to the predictive variance and its impact increases with the horizon. The Bayesian framework explicitly accounts for the effect of parameter uncertainty on predictive variance by integrating over the posterior distribution of parameters in the VAR model to form a predictive distribution for returns. In contrast, in traditional frequentist methods, the point estimates of the model parameters are treated as known, thereby ignoring estimation risk. As pointed out by Pastor and Stambaugh (2012), parameter uncertainty plays a role in all four components in Equation (4). The first three terms are expectations of quantities that are random due to uncertainty about the values of the underlying VAR parameters.[15] The last term only

[13] The sum of the first three terms on the right-hand side of Equation (4) corresponds to $\mathbb{E}[\text{Var}(r_{i,N,N+k} \mid \Theta, D_N|D_N]$ in Equation (3) and the fourth term corresponds to $\text{Var}[\mathbb{E}(r_{i,N,N+k} \mid \Theta, D_N|D_N]$ in that equation.

[14] As noted by Pastor and Stambaugh (2012) and Avramov, Cederburg, and Lucivjanska (2018), mean reversion requires time variation in expected returns. The negative effect of mean reversion on long-horizon predictive variance is therefore at least partially offset by the positive effect of uncertainty about future expected returns.

[15] In particular, each term is the expectation of a function of the parameters. When the posterior distribution of these functions is skewed, the posterior mean can be significantly affected by less likely extreme parameter values.

depends on parameter uncertainty and would be zero if investors knew the true parameters.[16] We study the changes to the variance over the horizon by analyzing the variance ratio $\mathrm{VR}_{k,i} = \frac{\mathrm{Var}(r_{i,N,N+k}|D_N)}{k\mathrm{Var}(r_{i,N,N+1}|D_N)}$ and its underlying components.[17]

We also analyze the long-horizon returns implied by the VAR by sampling returns on the risky asset and the risk-free asset from the predictive distribution. We simulate 250,000 future return paths of length k conditional on the draws of the VAR parameters and the state variables observable at time t and calculate predictive returns as the simple average of the simulated return paths.

2.5 Optimal portfolio choice

We construct optimal portfolios for long-only buy-and-hold investors. The investment universe consists of the market portfolio of stocks and the risk-free asset. In contrast to Barberis (2000) and Pastor and Stambaugh (2012) who assume that the real return on the riskless asset is constant and known to the investor, we explicitly account for reinvestment risk and uncertainty about the future risk-free rate. Because of this roll-over risk, T-bills become increasingly risky when the investment horizon grows. By accounting for time variation in the real risk-free rate we also capture horizon effects in correlations between the risk-free asset and the risky asset that can have a substantial impact on the optimal asset allocation of long-term investors (see, e.g., Campbell and Viceira (2005) and Hoevenaars, Molenaar, Schotman, and Steenkamp (2014)).

We construct optimal portfolios for various investment horizons k by maximizing expected utility with respect to the predictive distribution of future returns. Formally, at time N the investor maximizes expected utility over terminal wealth k periods in the future conditional on the data observable up to time N:

$$\max_{w_N} \quad E_N[U(W_{N+k})|D_N], \tag{7}$$

where W_{N+k} is the end-of-period wealth given by $W_{N+k} = w'_N R_{N,N+k}$, with w_N the vector of optimal portfolio weights and $R_{N,N+k}$ the vector of cumulative gross returns.

[16] Pastor and Stambaugh (2012) consider an additional source of long-horizon variance that captures uncertainty about the current value of expected returns when predictors are imperfect. We adopt the Bayesian framework of Avramov, Cederburg, and Lucivjanska (2018) that does not incorporate this additional uncertainty component.

[17] We scale each of these components by $k\mathrm{Var}(r_{i,N,N+1} \mid D_N)$ to determine their contribution to the variance ratio.

We assume that investors have power utility of the form[18]

$$U(W_{N+k}) = \frac{W_{N+k}^{1-A}}{1-A}. \tag{8}$$

We set the risk aversion parameter A equal to 5. We use the numerical approach described in Online Appendix B.1 to compute the optimal weights for investment horizon k, which ranges from 1 to 100 quarters. In short, we use a grid search over possible portfolio weights to find the weights that yield the highest average utility over 250,000 return paths of length k sampled from the predictive distribution of returns.

3 Incorporating model-based prior information

3.1 Temperature long-run risks model

Theoretical beliefs about our VAR parameters are based on an adjusted version of the LRR-T model of Bansal, Kiku, and Ochoa (2019). The LRR-T model imposes a structure on the relation between the temperature anomaly and financial market variables such as stock market returns and price-dividend ratios. The representative investor with Epstein and Zin (1989) recursive preferences optimizes lifetime utility

$$U_t = \left[(1-\delta)C_t^{\frac{1-\gamma}{\theta}} + \delta(\mathbb{E}_t[U_{t+1}^{1-\gamma}])^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}, \tag{9}$$

with C_t aggregate consumption at time t, $\delta \in (0,1)$ the investor time preference, γ the coefficient of relative risk aversion, $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$, and ψ the elasticity of intertemporal substitution (EIS). The log of aggregate consumption growth ($\Delta c_{t+1} = \log(C_{t+1}/C_t)$) and the log dividend growth of portfolio

[18] Power (CRRA) utility is commonly assumed in previous work, e.g., Balduzzi and Lynch (1999), Barberis (2000), Jurek and Viceira (2011), Pastor and Stambaugh (2012), Diris, Palm, and Schotman (2014), Hoevenaars, Molenaar, Schotman, and Steenkamp (2014), Johannes, Korteweg, and Polson (2014), Carvalho, Lopes, and McCulloch (2018).

i ($\Delta d_{t+1,i} = \log(D_{t+1,i}/D_{t,i})$) follow

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + \sigma_t \eta_{t+1} + X_{t+1}, \\ \Delta d_{t+1,i} &= \mu_d + \pi_d \sigma_t \eta_{t+1} + \phi_i X_{t+1} + \varphi_d \sigma_t u_{t+1}, \\ \sigma^2_{t+1} &= \bar{\sigma}^2 + \nu(\sigma^2_t - \bar{\sigma}^2) + \sigma_w w_{t+1}, \\ &\eta_{t+1}, u_{t+1}, w_{t+1} \sim i.i.d.\ N(0,1), \end{aligned} \tag{10}$$

with mutually independent shocks η_{t+1}, u_{t+1}, and w_{t+1} scaled by time-varying volatility σ_t. X_{t+1} is the adverse economic impact of temperature-driven disasters on consumption and dividend growth, based on a disaster process $\tilde{N}_t$ with Poisson distributed increments as

$$\begin{aligned} X_{t+1} &= \rho X_t + d\Delta \tilde{N}_{t+1}, \\ \Delta \tilde{N}_{t+1} &\sim \text{Poisson}(\lambda_t = \Delta t(\lambda_0 + \lambda_1 T_t)), \end{aligned} \tag{11}$$

where $\rho < 1$ is the persistence of economic disaster impact, $d < 0$ is the initial disaster-related growth shock to consumption and dividends, and T_t is the temperature anomaly, which is driven by the atmospheric carbon concentration ε_t:

$$\begin{aligned} T_{t+1} &= \chi \varepsilon_{t+1}, \\ \varepsilon_{t+1} &= \nu_\varepsilon \varepsilon_t + \mu_\varepsilon + \Theta(\mu_c + \sigma_t \eta_{t+1}) + \sigma_\zeta \zeta_{t+1}, \\ \zeta_{t+1} &\sim i.i.d.\ N(0,1). \end{aligned} \tag{12}$$

This structure allows for a feedback loop between consumption growth and the atmospheric carbon concentration, by including μ_c and η_{t+1} in the process for the latter. We assume that the log of the wealth-consumption ratio z_t and the log of the price-dividend ratio of portfolio i, $z_{t,i}$, are given as

$$z_{t(,i)} = A_{0(,i)} + A_{1(,i)} T_t + A_{2(,i)} X_t + A_{3(,i)} \sigma^2_t. \tag{13}$$

The analytical solution for the price-dividend and wealth-consumption ratios, using the Campbell and Shiller (1988) decomposition, is presented in Appendix A.1. We discuss the adjustments made to the LRR-T model of Bansal, Kiku, and Ochoa (2019) in Appendix A.2.

3.2 Data

Market returns, dividend growth, and price-dividend ratios are from the Irrational Exuberance dataset available on Robert Shiller's website.[19] As a proxy for market returns we use the monthly real log returns including dividends on the S&P 500 index. Dividend growth is the log difference of the monthly real dividends on the index. The price-dividend ratio is the log difference between the S&P 500's price and the monthly dividend payment on the index.

Monthly ex-ante real risk-free returns are constructed following Beeler and Campbell (2012). We use the seasonally unadjusted consumer price index (CPI) from the Bureau of Labor Statistics to construct quarterly and yearly inflation as the log difference between the CPI levels at the end of the current period and the end of the previous period. To construct ex post real risk-free yields we subtract the quarterly log inflation from the log CRSP Treasuries three month risk-free yields. Ex ante risk-free rates are the predicted value from the regression of ex post real risk-free yields on an intercept, nominal risk-free yields and the annual log inflation divided by four.

We obtain the U.S. temperature anomaly from the nClimDiv dataset of the National Oceanic and Atmospheric Administration (NOAA).[20] The monthly temperature anomaly is defined as the temperature in a given month minus the average temperature in that same month over the base period 1901-1946. This base period is chosen to end just before our sample to avoid any look-ahead bias. We convert the temperature anomaly to degrees Celsius as in Bansal, Kiku, and Ochoa (2019) and compute the quarterly anomaly as the average of the monthly anomalies in each quarter. Consumption data is obtained from the National Income and Product Accounts (NIPA). Quarterly per capita real consumption growth is the seasonally-adjusted aggregate nominal consumption expenditures on nondurables and services (NIPA Table 2.3.5), adjusted with the price deflator series (NIPA Table 2.3.4) and divided by population size (NIPA Table 2.1). Consumption growth is quarterly because monthly consumption data is unavailable for the early years of our sample. All other variables are monthly and time-aggregated to a quarterly frequency using the approach

[19] http://www.econ.yale.edu/~shiller/data.htm.
[20] https://www.ncdc.noaa.gov/cag/national/time-series.

of Bansal, Kiku, and Yaron (2016).

The sample period is 1947-2019. This starting date follows Beeler and Campbell (2012) and Barnett (2021) and balances the need for a longer sample to make more accurate long-term return forecasts with the fact that data from periods preceding the general awareness of climate change is likely uninformative about the impact of climate change on long-horizon return dynamics.

3.3 Model calibrations

Table 2 reports parameter values from our calibration of the LRR-T model. This calibration is chosen to match the first and second moments of consumption growth, dividend growth, the price-dividend ratio, the market return, and the risk-free rate over our sample period 1947-2019. In this calibration, the temperature anomaly is zero in expectation at the start of the simulated sample, with a long-term temperature expectation that is one degree Celsius higher. This one-degree temperature increase is in line with the smoothed temperature increase reported by the NOAA, which rises from 0.1 over the period 1943-1947 to 1.18 over the period 2015-2019.

[Table 2 about here.]

The first and second moments of observed consumption growth, dividend growth, price-dividend ratios, and returns in the data are compared with the implications from the LRR-T model calibrations in Table 3. As a benchmark we also report the moments implied by the LRR model of Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2012) that does not feature climate risk.[21] The moments from the LRR(-T) models are the medians from 250,000 simulations, each consisting of 936 monthly observations to match the length of our historical sample 1947-2019 after a burn-in period of 60 months. The monthly data is time-aggregated to quarterly values. The calibrated LRR-T model fits the data well in most aspects, similar to the original LRR model. There are two moments for which the model does not match the data.

First, as is commonly observed with LRR models (e.g., Beeler and Campbell (2012)), the implied price-dividend ratio level and variance is too low. For example, Bansal, Kiku, and Yaron

[21] We recalibrate some parameters in the Bansal, Kiku, and Yaron (2012) configuration to better match the observed data moments in our sample. Specifically, we set $\rho = 0.981$, $\mu_d = 0.0021$, $\phi = 2.0$, $\pi = 2.0$, and $\varphi = 5.0$.

(2012) report expected price-dividend ratios 9% below the data estimate in their LRR population moments, comparable to our results. Second, the expected temperature anomaly is higher in the LRR-T simulations than observed in our sample. This is driven by the concave increase in the LRR-T autoregressive structure for the temperature level, whereas historical temperature increases are more linear. Therefore, the temperature anomaly in the early years of our LRR-T simulations is higher than in the corresponding decades in the data.

[Table 3 about here.]

3.4 Model implied risk premia

With the calibrations for the LRR(-T) models, the theoretical market risk premium for the LRR and LRR-T models can be decomposed. The market risk premium in the LRR-T model equals

$$\begin{aligned} \ln \mathbb{E}_t[R_{m,t+1}] - r_{f,t} &= -\text{Cov}_t(m_{t+1} - \mathbb{E}_t[m_{t+1}], r_{m,t+1} - \mathbb{E}_t[r_{m,t+1}]) \\ &= \underbrace{\lambda_\eta \beta_{m,\eta} \sigma_t^2}_{\text{growth premium}} + \underbrace{\lambda_\varsigma \beta_{m,\varsigma} \sigma_\varsigma^2}_{\text{temp premium}} + \underbrace{\lambda_w \beta_{m,w} \sigma_w^2}_{\text{volatility premium}} + \underbrace{\lambda_X \beta_{m,X} (\lambda_0 \Delta t + \lambda_1 \Delta t T_t)}_{\text{disaster premium}}, \end{aligned} \tag{14}$$

with the derivation and analytical solutions for lambdas and betas given in Appendix A.1.

There are four components of the market risk premium. The first term, related to shocks η_t, is the growth premium, as it is directly related to shock to consumption growth. The second component is the temperature premium, related to the uncertainty in the temperature process that impacts the future expected disaster-occurrence. Note that this term is related to the uncertainty in the temperature process, not to the actual temperature level. Therefore, even when the temperature anomaly is forecasted to decrease, this premium will remain stable as a reward for the risk of future changes in the number of expected disasters. Third, we have a volatility risk premium that affects the time-varying volatility of shocks to consumption growth. Finally, the last component is the disaster premium, which is based on the variance of the current disaster process, that is in turn a function of the temperature anomaly. Because of this final component, higher temperature levels increase the market risk premium directly.

Comparing the LRR and LRR-T risk premia, we see a similar structure, but the focus clearly changes to a climate related risk premium in the latter model. In the LRR model, the risk premium has three components: short-run risk, long-term growth risk, and volatility risk. The first component has the same structure in the LRR and LRR-T models, and corresponds to the growth premium discussed above. The long-term growth risk is related to the exposure to shocks in the persistent component of the LRR model, and as such comparable with the disaster premium in the LRR-T model. Finally, volatility risk impacts the risk premium in two ways in the normal LRR model: by introducing (short-term) noise in the direct shock to consumption growth η, and by affecting the size of the shock to the persistent component. The first aspects of the volatility risk corresponds to the same component of the LRR-T risk premium. The second aspect is comparable to the LRR-T temperature premium, which is based on the noise in the temperature anomaly, a measure for the future expected disaster shock impact.

Based on our calibrations, we find that our LRR model implies that 17% of our risk premium is driven by short-term risk, 32% by long-term growth risk, and 51% by long-run volatility risk. In the LRR-T model, 8% of the risk premium is driven by the growth premium, 63% by the temperature premium and 29% by the disaster premium. The volatility risk premium has negligible impact on the risk premium, since it only adds noise to the short-term consumption growth shocks. However, it is useful to have time-varying volatility in the model to better match the second moment of the market return in our LRR-T calibration. There are two main takeaways from this decomposition. First, the LRR-T risk premia are driven by components that are, in their structure, similar and comparable to what the LRR model implies. Second, temperature levels and climate-related disasters are a very important part of the financial implications of the LRR-T model.

3.5 Model simulations

To illustrate how temperatures impact financial performance in the LRR-T model we show a single simulation from our calibrated model in Figure 2. In this simulation we see that temperatures increase over the sample (in line with the calibrated trend), with quite some variance in the temperature process as is observed in the historical data. In expectation, temperature anomalies start at 0 and increase to 1, but we observe both much higher and lower temperature anomalies in the simulation, with a coincidental low temperature anomaly at the end of our simulated sample.

With increases in temperatures, the expected occurrence of future disasters that affect future consumption and dividend growth rates also goes up. Price-dividend ratios and risk-free rates respond to these expected future adverse events with immediately decreases that are mostly visible in the price-dividend ratio. Climate disasters start occurring after roughly 250 months in this simulation, most clearly visible in the risk-free rate.[22] These disasters have an immediate and persistent impact on risk-free rates and price-dividend ratios. For risk-free rates, persistent impact is caused by the slow response of the wealth-consumption ratio to disasters.[23] In market returns, we do not observe persistent impact of disasters, as decreasing prices keep future market returns relatively stable. There is, however, significant transitory impact from climate disasters on market returns, as the most negative market returns observed in our simulation *are* caused by disasters. Overall, the market risk premium increases after disaster occurrence, because risk-free returns are persistently lower while the market rebounds quickly after the initial transitory shock.

[Figure 2 about here.]

The population moments from 250,000 simulations of the baseline LRR-T model imply the VAR structure shown in Table B1 in Online Appendix B.2. Note that in this estimation, all coefficients are significant by construction, since we base the VAR in this table on the population moments from the model without allowing for misspecification. First, relations between market returns and financial state variables are as expected. Increases in price-dividend ratios decrease expected market returns and increases in risk-free rates increase expected market returns. This VAR structure differs from previous literature by the inclusion of the temperature anomaly as a predictive variable. Temperature levels have a strong negative impact on next period expected market returns. The coefficient from market returns on the temperature anomaly is economically

[22]In this particular simulation path, we observe disasters relatively late by coincidence. In expectation we would observe roughly one disaster each every 100 months. Overall, there are 10 disasters in this simulation, which is what we would expect.

[23]The stability of the wealth-consumption ratio, especially in a long-run risks setting, is also documented by Lustig, Van Nieuwerburgh, and Verdelhan (2013).

large: a temperature anomaly of 1 degree Celsius implies that quarterly expected market returns are decreased by 0.61%, or 2.4% annually. A one standard deviation shock to the temperature anomaly decreases quarterly expected market returns by 0.51%, or 2% annually. This initial impact on market returns is partly offset in long-horizon forecasts, because the temperature increase also decreases price-dividend ratios, but that effect is smaller. The temperature anomaly is highly persistent, which is implied by the LRR-T model because we calibrate a positive trend in the temperature levels. Overall, this VAR shows that the LRR-T model indeed imposes a structure on the impact of temperature levels on equity risk and return.

4 Empirical results

Thus far we have examined the effect of climate change on financial markets through the lens of the LRR-T model. We now discuss how these theoretical implications affect the empirical estimates of a Bayesian investor who combines the market portfolio with a risk-free asset.

4.1 Predictive regressions

We estimate the quarterly predictive VAR from Equation (1) on the sample from 1947Q1 to 2019Q4 for four investor types. The first three investor types are the agnostic (data-based) investor, the dogmatist (climate risk believer) investor, and the Bayesian LRR-T investor discussed in Section 2.1. Finally, we analyze a benchmark Bayesian LRR investor that forms prior beliefs based on the LRR model without climate change. Table 4 shows the posterior mean of 250,000 draws of the VAR coefficients for these investor types, along with their posterior t-statistics.

In panel A of Table 4, we show the posterior VAR for the agnostic investor, based on historical data with the temperature anomaly as a state variable. From the coefficient estimates we observe that price-dividend ratios have negative forecasting power on market returns, while increases in risk-free yields imply higher market returns, as expected. We observe statistically insignificant coefficient estimates on the temperature anomaly. This could be a result of large uncertainty about the impact of climate change on financial markets, because of parameter uncertainty or unavailability of sufficient disasters in historical data. This uncertainty is one of the key reasons to take prior beliefs into account in a Bayesian setting, which helps estimation precision. On the other hand, the coefficient estimates on the temperature anomaly are economically very meaningful. A one standard deviation shock to the temperature anomaly implies a 0.74% change in the forecasted quarterly market return, or almost 3% annually. Based on forecasts for future temperature anomalies of up to two degrees Celsius, impact of temperatures on market returns would be even larger, especially in long-horizon forecasts. Other than market returns, financials are not strongly affected by temperatures in the data. The temperature process itself is included for completeness but not used for forecasts, where an exogenous temperature process is used, as discussed in Section 2.3.

The opposite perspective is a dogmatist investor that bases her beliefs about the VAR model completely on the LRR-T prior. Panel B of Table 4 reports the posterior VAR for this climate risk believer. The VAR coefficients are almost identical to the population estimates for the VAR parameters reported in Table B1. As before, we observe a negative loading from market returns on the temperature anomaly. In the LRR-T model, increasing temperatures decrease the price-dividend ratio, increasing market returns. However, this effect is offset by the realisation of large negative returns around climate disasters, which occur more often with high temperature levels. In our simulations, the second effect is stronger. The difference with the population estimates are in the decreased statistical significance of the results. We introduce model misspecification in these estimates to match the information in the prior beliefs to the historical data sample from 1947Q1 to 2019Q4, as described in Section 2.1. Because of the noise in the agnostic VAR based on historical data, this results in large uncertainty in the VAR of the dogmatic investor as well. However, we still observe economically meaningful negative impact from increased temperatures on financial outcomes.

Comparing panels A and B, it is clear that historical data and economic theory have opposing implications about the impact of climate change on financial performance, which makes it useful to include both views in a Bayesian analysis. A key difference between the agnostic and dogmatist investors is that the former does not show a relation between both the price-dividend ratio and risk-free returns and the temperature anomaly, while this relation is reasonably strong for the latter. This is driven by the impact of climate disasters on risk-free returns in the LRR-T model. In that model, disasters result in additional savings, decreasing risk-free returns.

Panel C of Table 4 reports the posterior VAR for the Bayesian investor that gives equal weights to implications from historical data (panel A) and the baseline LRR-T prior (panel B). As expected, the posterior coefficients and correlations of this VAR are generally in between the reported values in panel A and panel B. Since panels A and B present opposite results, the Bayesian investor has a relatively small loading on the temperature anomaly, which is visible in the small absolute posterior coefficients. This small coefficient *does* have significant impact on predictive distributions of market returns, especially for long-horizons. Small changes in the predictive VAR parameters become increasingly important after several quarterly forecasts.

Finally, panel D of Table 4 is the benchmark Bayesian LRR investor who combines return dynamics implied by historical data with those implied by the LRR prior. As expected, the general structure of the VAR is very similar to Panel C, the only real difference is the fact that the temperature anomaly is not included in the model. Therefore, comparing the forecasts from this investor with the implications from the Bayesian LRR-T investor highlights the impact of the temperature anomaly on our models, instead of other differences in the LRR and LRR-T models.

[Table 4 about here.]

4.2 Long-horizon risk

We now use the 250,000 draws of the VAR coefficients around their posterior mean in Table 4 to forecast long-horizon variance of market returns. For this, we combine draws of the coefficients with the exogenous temperature process discussed in Section 2.3 in the long-horizon variance formula from Equation (4) to forecast variance ratios for horizons up to 100 quarters. In Figure 3, the per period variance by horizon is given for the agnostic (data-based), dogmatic (LRR-T based), and Bayesian LRR and LRR-T investors (combining data with model based priors).

Within a few quarters, the variance ratios quickly diverge. In the data, we find strong mean reversion, which is also documented by, among others, Barberis (2000) and Siegel (2014). Intuitively, current low (high) returns are offset by future high (low) returns, because of predictability in returns in the data combined with negative correlation between current and future returns. Therefore, variance ratios quickly decrease when the horizon gets longer. As opposed to earlier work, the large coefficient from market returns on the temperature anomaly increases the uncertainty about future expected returns for the agnostic investor, as the future temperature level is highly uncertain. At the longest horizon of our forecasts, the decrease in the additional mean reversion combined with the continuously increasing uncertainty about future expected returns increases the variance ratio again, as compared to horizons between 50 and 75 quarters. Overall, the variance ratios over the horizon of our agnostic investor is similar to the results in Avramov, Cederburg, and Lucivjanska (2018), but increased predictability in our setting from including the temperature anomaly in our model increases both mean reversion and uncertainty about future expected returns.

At first sight on the overall variance ratios, the benchmark LRR Bayesian investor seems to be close to the data-based agnostic investor. However, the underlying components driving long-horizon variance for these two investors are very different. The key difference is that the data implies stronger predictability from the price-dividend ratio, which increases both mean reversion and uncertainty about future expected return compared to the Bayesian LRR investor.

The dogmatic investor with its LRR-T prior shows very high long-horizon variance ratios, with per period market volatility predicted up to 2.75 times higher than the one-quarter ahead forecast. There are two key drivers of the difference with the agnostic investor. First, mean reversion is weaker. Intuitively, mean reversion decreases because current low returns caused by climate-induced disasters are followed by future low returns caused by even more disasters. This effect is driven by the fact that disasters in the LRR-T simulations occur in clusters, in periods with high temperature levels. Therefore, the negative correlation between current and future returns that is needed for mean reversion is much weaker than in the data. Second, estimation risk has increased significantly, because the disasters in the LRR-T simulations increase the uncertainty around the parameter estimates.

Influenced by the LRR-T prior, the LRR-T Bayesian investor has relatively high variance ratios, monotonically increasing with the investment horizon. At the longest horizon of 100 quarters, the variance ratio of the LRR-T Bayesian investor is twice as large as that of the benchmark LRR Bayesian and the agnostic investors. The underlying components of the variance ratio are between what is implied by the data and the prior beliefs, except for the uncertainty about future expected returns. Lower uncertainty is driven by the relatively small coefficient from market returns on the temperature anomaly for the Bayesian investor, a logical result from the opposite signs of corresponding coefficient for the agnostic and dogmatic investors. Even though the coefficient on the temperature anomaly is smaller for the LRR-T Bayesian investor, the inclusion of temperatures in the model is sufficient to increase uncertainty about future expected returns compared to the benchmark LRR investor.

[Figure 3 about here.]

For our asset allocation decision, the riskiness of the short-term bond returns is also relevant to consider. We assume that our short-term bond is risk-free in the sense that there is no credit risk. However, over longer horizons, reinvestment (roll-over) risk affects investments that are risk-free in the short run. The variance of the risk-free return relative to the risk of the market return is presented in Figure B1 in Online Appendix B.2. For all models, risk-free returns have more variance with longer horizons, as reinvestment risk increases with the horizon. Both the data and the LRR-T based models imply similar relative riskiness between risk-free and market returns. The Bayesian LRR investor forecasts less reinvestment risk for the risk-free asset. This observation aligns with the calibration moments in Table 3, where LRR based risk-free return simulations are less volatile than in the data and the LRR-T model.

Overall, including climate change through LRR-T prior beliefs results in higher predictive long horizon risk for the LRR-T Bayesian investor investing the market portfolio. This effect is large, with long-run variances up to twice as high as as the benchmark LRR Bayesian investor. Next to market returns, the LRR-T Bayesian investor also predicts more reinvestment risk in the risk-free asset.

4.3 Long-horizon returns

Predictive returns also differ significantly for the four different investors and by horizon. Figure 4 shows the average forecast of risk-free returns, market returns and market risk premium from

250,000 simulated return paths. Each return path is forecasted from a draw of the VAR posterior distribution for each investor type to account for parameter uncertainty, combined with the exogenous temperature process discussed in Section 2.3.

The risk-free return forecasts are initially all very similar and close to zero, in line with the current observed short-term T-bill rate. With forecasts over the horizon, we see opposite trends for the agnostic and dogmatic investors. Forecasts from the data push long-term risk-free returns towards the historical average, which is higher than the current risk-free rate levels. In the LRR-T model, however, future temperature increases decrease the risk-free rate. This reflects the negative impact that climate disasters have on the risk-free returns in the LRR-T simulations. Disasters persistently affect future consumption and dividend growth and agents save to offset this, increasing the demand for the risk-free asset, decreasing its return. Since the risk-free rate is very persistent and temperatures keep increasing in our forecasts, the dogmatic investor predicts decreasing risk-free returns over our horizon. We note that this decrease is not set to continue indefinitely if one extends the horizon, since risk-free rates are not 100% persistent and climate models generally do not forecast unbounded temperature increases. The LRR-T Bayesian investor forecasts risk-free returns in between the data and the LRR-T model. Without temperature increases in the model, the benchmark LRR Bayesian investor forecasts risk-free returns similar to the agnostic investor.

The predictive market returns reported in Figure 4 are very similar for the agnostic and LRR Bayesian investors, both increasing by horizon. The agnostic investor has positive VAR coefficients from market returns on the temperature anomaly and the risk-free returns, as visible in panel A of Table 4. The increase in predictive market returns over the horizon is, therefore, driven by the increase in predictive risk-free returns and the increasing temperature levels in the exogenous temperature process. For the LRR Bayesian investor, temperatures are not included in the model. However, a similar horizon pattern for predictive market returns is driven by a larger loading on the increasing risk-free returns, as visible in panel D of the same table.

The most surprising of the return forecasts is the quick increase in the short-term market returns expected by the dogmatic investor. This is a direct result of the implementation of an exogenous temperature process, where the starting value of our out-of-sample temperature anomaly is the recently observed level in the data. This temperature level is higher than in the LRR-T simulated data used by the dogmatic investor, which is why there is a jump in the temperature levels from the in-sample estimation to the out-of-sample forecasts. This initial jump in our forecasts quickly decreases the price-dividend ratio, resulting in a sharp increase in expected market returns. We do not feel that this is a realistic forecast and, therefore, do not conclude anything from the first ten quarters of the LRR-T prior forecasts.[24] At longer horizons, increasing temperatures combined with decreasing risk-free returns drive expected market returns down. This corresponds well to the observed negative disaster-related returns in the LRR-T simulations, as presented in Figure 2 above.

The market returns predicted by the LRR-T Bayesian investor are stable over the horizon, as increased temperatures balance the negative impact from decreases in the risk-free returns. In other words, the direct impact from the forecasted temperature increase on market returns is offset by its indirect impact on market return through the risk-free asset.

All four models imply that the market risk premium increases with the investment horizon, although the underlying mechanism differs. The agnostic investor forecasts increasing market returns with the horizon, mainly driven by decreasing price-dividend ratios towards the historical average and increasing temperatures. The LRR Bayesian benchmark investor shows similar risk premium forecasts, at a lower level because of slightly higher risk-free returns. For the dogmatic and LRR-T Bayesian investors, decreasing risk-free rates are an important aspect of the increasing risk premium over the horizon. While the horizon effects for the market risk premium are similar for all models, the risk premia are at different levels.

[Figure 4 about here.]

Finally, we discuss the correlations of the returns on the risk-free asset and the market portfolio. Figure 5 documents that for all our investors these returns start relatively uncorrelated and increase over the horizon. This increased correlation is mainly driven by the positive coefficient from market

[24]One could argue that this result implies that the exogenous temperature process should not be the same for all our models. However, we believe that it is more important to keep future temperature processes similar to compare the differences in the implications of climate change between the models than to make these short-term market return forecasts more stable.

returns on lagged risk-free returns in the VAR models of Table 4.

[Figure 5 about here.]

Overall, the LRR-T Bayesian investor predicts higher market risk premia over all investment horizons than the agnostic and LRR Bayesian investors. This effect is driven by predictive risk-free returns that decrease over the horizon, combined with relatively high short-term market return predictions.

4.4 Portfolio choice

We have now documented two key results. First, the LRR-T Bayesian investor predicts significantly higher long-horizon riskiness of the market portfolio than our benchmark LRR Bayesian investor, because of the uncertainty in the temperature anomaly forecast and increased estimation risk. Second, including beliefs from the LRR-T model results in higher predictive market risk premia, because forecasted climate-related decreases in consumption growth negatively affect long-horizon risk-free returns. Next to these key results, the LRR-T model implies relatively more reinvestment risk and higher correlations between market and risk-free returns than in the data. These results each have implications for portfolio choice, as we discuss in this section.

We analyze a long-only investor that can invest in the market portfolio and a risk-free asset with reinvestment risk. The top panel of Figure 6 shows the optimal allocation to equities for the agnostic, dogmatic, and Bayesian investors. The optimal weight in equities increases quickly with the investment horizon for the agnostic investor, because there is very strong mean reversion in the related predictive VAR. At the longest horizon market returns become riskier again and predictive risk-free returns become higher, which brings the optimal weight to the market portfolios down. The dogmatic investor initially sees a quick increase in the optimal allocation to the market portfolio driven by the sharp increase in the market risk premium, a consequence of our exogenous temperature process. Over the horizon, the dogmatic investor keeps allocating less and less to the market because the perceived long-horizon risk of the market portfolio increases quickly.

As expected, the allocation to the market portfolio of the LRR-T Bayesian investor is in between the allocations of the agnostic and dogmatic investors. Comparing the LRR-T and LRR Bayesian investors, we find that the inclusion of the temperature anomaly in the model has economically meaningful impact on optimal portfolios, which changes over the horizon. The investor with LRR-T prior beliefs has higher weights to equities than the LRR Bayesian investor for horizons of up to roughly 25 quarters, with the opposite result for longer investments horizons. This is mainly driven by the two key results discussed above: the inclusion of climate change in the model implies higher market risk premia for all horizons, but risk increases considerably over the horizon. Therefore, the allocation is initially higher, but becomes lower when the increased risk outweighs the higher market risk premium. Especially in the long run, these differences are economically meaningful: the optimal weights to equities are roughly 20% lower for the LRR-T Bayesian investor that invests for 100 quarters, compared to the benchmark LRR Bayesian investor.

The riskiness of the market portfolio and the risk-free asset are affected very differently by parameter uncertainty. In Hoevenaars, Molenaar, Schotman, and Steenkamp (2014), parameter uncertainty does not have significant impact on portfolio choice, but with the inclusion of the temperature anomaly in our model we find higher estimation risk for equities. Reinvestment risk of risk-free returns is, on the other hand, mainly driven by uncertainty about future expected returns, because of the loadings from risk-free returns on the temperature anomaly. Therefore, when parameter uncertainty is accounted for, equity returns become considerably more risky, relative to risk-free returns. The longer the investment horizon, the more important the inclusion of parameter uncertainty becomes for our portfolio choice. In the bottom panel of Figure 6, we therefore observe similar weights to the market portfolio for short horizons with and without parameter uncertainty, but significantly lower allocations to equities in the long run. The differences between the allocations of the different investors are similar to before, with the same key results driving the allocation over the horizon.

[Figure 6 about here.]

4.5 Alternative climate change scenarios

In this section, we examine the impact of two alternative climate scenarios on the long-horizon risk-return tradeoff and optimal portfolio choice. The first alternative scenario assumes that the temperature anomaly increases to 4 degrees Celsius at the end of the forecast horizon (100 quarters from now), rather than to 2 degrees as assumed in the base case scenario described in Section 2.3. In the second scenario the temperature anomaly still increases to 2 degrees, but the uncertainty about the future temperature path is increased. In particular, we set σ_η^2 in Equation (2) equal to 0.06 to match the historical volatility of the temperature anomaly. Recall that in the default scenario, we set σ_η^2 equal to 0.001 to match the uncertainty around the expected temperature anomaly in the IPCC forecasts. In other words, in the baseline scenario the variance is chosen to match the uncertainty around the expectation (standard error of the mean), which is much lower than the historical volatility of the temperature anomaly that we match in the alternative scenario.

For each climate scenario, we construct the predictive variances, predictive returns, and optimal portfolio weights for different horizons. We focus on the Bayesian LRR-T investor. Figure 7 shows that in the scenario with the higher temperature increase, we observe a sharp increase in the perceived riskiness of equity at long horizons. For an investment horizon of 100 quarters, the variance ratio goes up from 1.39 in the base case scenario to 1.95 in this alternative scenario. The variance decomposition shows that this increase is driven by higher estimation risk. As shown in Equation (4), the impact of parameter uncertainty increases for larger values of the predictor variables, including the temperature anomaly. We also find higher variance ratios in the scenario with higher uncertainty about the future temperature anomaly. As expected, the variance decomposition shows that this increase is due to greater uncertainty about future expected returns. In both alternative scenarios stocks thus appear riskier to long-term investors than in the default temperature scenario, but the increase in predictive variance is attributable to different factors.

[Figure 7 about here.]

Figure 8 shows that in the high-temperature scenario, the predictive equity risk premium is higher than in the baseline scenario, particularly at long horizons. This pattern is consistent with the implication of the temperature long-run risks model that higher temperatures raise the equity premium by increasing the disaster premium (see Equation (14)). In contrast, we do not observe higher equity premiums in the scenario with more uncertainty about future temperature paths.

Although temperature volatility carries a positive risk premium in the LRR-T model under Epstein-Zin preferences, Bansal, Kiku, and Ochoa (2019) point out that this premium vanishes under CRRA utility. Because we construct optimal portfolios for investors with CRRA preferences, we do not include temperature volatility as a return predictor in the VAR model.

[Figure 8 about here.]

Figure 9 illustrates the consequences of the changes in the term structure of predictive risk and return for portfolio choice.[25] Because in the high temperature scenario the increase in the perceived riskiness of equity at long horizons is offset by the increase in the anticipated equity risk premium, the optimal allocation to equity is very similar to the allocation in the base case scenario. The optimal allocation to equity is somewhat lower in the high temperature uncertainty scenario due to the increase in risk without a corresponding increase in risk premiums. In all three temperature scenarios, the optimal allocation to equities decreases sharply with the investment horizon.

[Figure 9 about here.]

5 Conclusion

We propose a new approach for measuring the impact of climate change on long-term equity returns and portfolio choice. Since past data is not very informative about the impact of climate change on future stock returns, we estimate the parameters of the return-forecasting model using a Bayesian approach that complements historical data with prior information about return dynamics implied by a temperature long-run risks (LRR-T) model. The model predicts that rising temperatures affect asset prices by increasing the likelihood of climate-driven disasters that lower economic growth. For comparison, we also consider investors with uninformative priors who base their forecasts purely on historical data and investors with beliefs derived from the long-run risks (LRR) model that does not feature climate change. We study the effect of these different beliefs on the perceived riskiness of equities over various holding periods and on the optimal portfolio choice for long-term investors.

[25] Figures B2 and B3 in the Online Appendix show that the alternative specifications of the temperature process also affect the relative riskiness of the market portfolio and the risk-free asset and the predictive correlations between the returns on both assets.

We document three key findings. First, we show that investors with beliefs formed based on the temperature long-run risks model perceive stocks to be much riskier over longer horizons than investors with uninformative beliefs or investors with prior beliefs based on the original LRR model that does not incorporate climate change. For investors with LRR-T beliefs, stocks appear more volatility in the long run than in the short run because climate change increases estimation risk and uncertainty about future returns and weakens their beliefs in mean reversion. In contrast, investors with uninformative priors or LRR priors view stocks as safer in the long run.

Second, LRR-T investors predicts higher equity risk premia at all horizons than investors with LRR beliefs or investors with uninformative priors. The higher equity risk premium is mainly due to lower risk-free returns. In particular, because in the LRR-T model disasters have a long-lasting effect on future economic growth, agents increase their precautionary savings and have higher demand for the risk-free asset, which lowers its return. LRR-T investors further predict that market returns decrease with the horizon because rising temperatures increase the likelihood of disasters that depress stock markets. However, because the decrease in market returns is more than offset by the decrease in risk-free returns, the market risk premium increases with the horizon.

Third, we show that including prior information about the effect of climate change on asset prices has a significant impact on the optimal portfolio choice of long-horizon investors who allocate wealth between the market portfolio and the risk-free asset. Because the higher equity premium does not fully offset the perceived increase in long-horizon equity risk, investors with LRR-T beliefs allocate much less capital to equities at long horizons than investors with no prior views on the impact of climate change. We further find that for all investor types, the optimal allocation to stocks starts to decrease for horizons longer than 10 years due to increased parameter uncertainty.

Overall, our results demonstrate that incorporating prior information about the impact of climate change on return dynamics can have a significant impact on the perceived riskiness of stocks and the asset allocation of long-horizon investors. Our framework can be readily extended to examine the portfolio implications of other asset pricing models that include climate risk. Extending the analysis to a dynamic portfolio strategy provides another fruitful avenue for future work.

A Appendix

A.1 Solution to temperature long-run risks model

The intertemporal marginal rate of substitution (IMRS) equals

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{t+1,c}, \tag{A.1}$$

where $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$.

A.1.1 The wealth-consumption ratio

We use the Euler equation,

$$\mathbb{E}_t[\exp(m_{t+1} + r_{t+1,c})] = 1, \tag{A.2}$$

to solve the wealth-consumption ratio z_t from Equation (13), filling in the IMRS from Equation (A.1) and the Campbell-Shiller decomposition for the return on the consumption asset, $r_{t+1,c} = \kappa_0 + \Delta c_{t+1} + \kappa_1 z_{t+1} - z_t$. We find

$$\begin{aligned} m_{t+1} + r_{t+1,c} = {} & \theta \log \delta + \theta \kappa_0 + \theta(\kappa_1 - 1) A_0 + (1 - \gamma + \theta \kappa_1 A_1 \chi \Theta) \mu_c + \theta \kappa_1 A_1 \chi \mu_\varepsilon \\ & + (1 - \gamma + \theta \kappa_1 A_1 \chi \Theta) \sigma_t \eta_{t+1} + \theta \kappa_1 A_1 \chi \sigma_\zeta \zeta_{t+1} + \theta(\kappa_1 \nu - 1) A_3 \sigma_t^2 \\ & + \theta \kappa_1 (1 - \nu) A_3 \bar{\sigma}^2 + \theta \kappa_1 A_3 \sigma_w w_{t+1} + ((1 - \gamma)\rho + \theta(\kappa_1 \rho - 1) A_2) X_t \\ & + (1 - \gamma + \theta \kappa_1 A_2) d \Delta N_{t+1} + \theta(\kappa_1 \nu_\varepsilon - 1) A_1 T_t. \end{aligned} \tag{A.3}$$

With the expectation

$$\begin{aligned} \mathbb{E}_t[m_{t+1} + r_{t+1,c}] = {} & \theta \log \delta + \theta \kappa_0 + \theta(\kappa_1 - 1) A_0 + (1 - \gamma + \theta \kappa_1 A_1 \chi \Theta) \mu_c + \theta \kappa_1 A_1 \chi \mu_\varepsilon \\ & + \theta(\kappa_1 \nu - 1) A_3 \sigma_t^2 + \theta \kappa_1 (1 - \nu) A_3 \bar{\sigma}^2 + ((1 - \gamma)\rho + \theta(\kappa_1 \rho - 1) A_2) X_t \\ & + (1 - \gamma + \theta \kappa_1 A_2) d \lambda_0 \Delta t + ((1 - \gamma + \theta \kappa_1 A_2) d \lambda_1 \Delta t + \theta(\kappa_1 \nu_\varepsilon - 1) A_1) T_t \end{aligned} \tag{A.4}$$

and variance

$$\begin{aligned}\mathrm{Var}_t(m_{t+1}+r_{t+1,c}) = {} & (1-\gamma+\theta\kappa_1 A_1\chi\Theta)^2\sigma_t^2 + (\theta\kappa_1 A_1\chi)^2\sigma_\zeta^2 + (\theta\kappa_1 A_3)^2\sigma_w^2 \\ & + (1-\gamma+\theta\kappa_1 A_2)^2 d^2\lambda_0\Delta t + (1-\gamma+\theta\kappa_1 A_2)^2 d^2\lambda_1\Delta t T_t.\end{aligned} \tag{A.5}$$

We now go back to the Euler equation from (A.2) and take logs and use Jensen's inequality to find the solution to the price-consumption ratio from Equation (13):

$$\begin{aligned}\theta(1-\kappa_1\nu)A_3 &= 0.5(1-\gamma+\theta\kappa_1 A_1\chi\Theta)^2, \\ \theta(1-\kappa_1\rho)A_2 &= (1-\gamma)\rho, \\ \theta(1-\kappa_1\nu_\varepsilon)A_1 &= (1-\gamma+\theta\kappa_1 A_2)d\left(1+\frac{(1-\gamma+\theta\kappa_1 A_2)d}{2}\right)\lambda_1\Delta t, \\ (1-\kappa_1)A_0 &= \log\delta+\kappa_0+\left(1-\frac{1}{\psi}+\kappa_1 A_1\chi\Theta\right)\mu_c+\kappa_1 A_1\chi\mu_\varepsilon \\ & \quad + \kappa_1(1-\nu)A_3\bar{\sigma}^2 + 0.5\theta((\kappa_1 A_1\chi)^2\sigma_\zeta^2+(\kappa_1 A_3)^2\sigma_w^2) \\ & \quad + \left(1-\frac{1}{\psi}+\kappa_1 A_2\right)d\left(1+\frac{(1-\gamma+\theta\kappa_1 A_2)d}{2}\right)\Delta t\lambda_0.\end{aligned} \tag{A.6}$$

A.1.2 The risk-free rate

Again, we use the Euler equation to solve the risk-free rate. We start from the IMRS of Equation (A.1) and fill in

$$\begin{aligned}m_{t+1} = {} & \theta\log\delta+(\theta-1)\kappa_0+(\theta-1)(\kappa_1-1)A_0+((\theta-1)\kappa_1 A_1\chi\Theta-\gamma)\mu_c+(\theta-1)\kappa_1 A_1\chi\mu_\varepsilon \\ & + (\theta-1)\kappa_1 A_3(1-\nu)\bar{\sigma}^2+(\theta-1)\kappa_1 A_1\chi\sigma_\zeta\zeta_{t+1}+(\theta-1)\kappa_1 A_3\sigma_w w_{t+1} \\ & + ((\theta-1)\kappa_1 A_1\chi\Theta-\gamma)\sigma_t\eta_{t+1}+(\theta-1)(\kappa_1\nu-1)A_3\sigma_t^2 \\ & + ((\theta-1)(\kappa_1\rho-1)A_2-\gamma\rho)X_t+((\theta-1)\kappa_1 A_2-\gamma)d\Delta N_{t+1}+(\theta-1)(\kappa_1\nu_\varepsilon-1)A_1 T_t.\end{aligned} \tag{A.7}$$

We then have expectation

$$\begin{aligned}\mathbb{E}_t[m_{t+1}+r_{f,t}] = r_{f,t} &+ \theta\log\delta + (\theta-1)\kappa_0 + (\theta-1)(\kappa_1-1)A_0 + ((\theta-1)\kappa_1 A_1\chi\Theta - \gamma)\mu_c \\ &+ (\theta-1)\kappa_1 A_1\chi\mu_\varepsilon + (\theta-1)\kappa_1 A_3(1-\nu)\bar{\sigma}^2 + (\theta-1)(\kappa_1\nu-1)A_3\sigma_t^2 \\ &+ ((\theta-1)(\kappa_1\rho-1)A_2 - \gamma\rho)X_t + ((\theta-1)\kappa_1 A_2 - \gamma)d\lambda_0\Delta t \\ &+ (((\theta-1)\kappa_1 A_2-\gamma)d\lambda_1\Delta t + (\theta-1)(\kappa_1\nu_\varepsilon - 1)A_1)T_t,\end{aligned} \tag{A.8}$$

and variance

$$\begin{aligned}\mathrm{Var}_t(m_{t+1}+r_{f,t}) = ((\theta-1)\kappa_1 A_1\chi)^2\sigma_\zeta^2 + ((\theta-1)\kappa_1 A_3)^2\sigma_w^2 + ((\theta-1)\kappa_1 A_1\chi\Theta-\gamma)^2\sigma_t^2 \\ + ((\theta-1)\kappa_1 A_2-\gamma)^2 d^2\lambda_0\Delta t + ((\theta-1)\kappa_1 A_2-\gamma)^2 d^2\lambda_1\Delta t T_t.\end{aligned} \tag{A.9}$$

With the Euler equation, logs, and Jensen's inequality we find

$$\begin{aligned}0 &= \mathbb{E}_t[m_{t+1}+r_{f,t}] + 0.5\mathrm{Var}_t(m_{t+1}+r_{f,t}) \\ \Leftrightarrow r_{f,t} &= r_f - ((\theta-1)(\kappa_1\rho-1)A_2-\gamma\rho)X_t \\ &\quad - \left(((\theta-1)\kappa_1 A_2-\gamma)d\left(1+\frac{((\theta-1)\kappa_1 A_2-\gamma)d}{2}\right)\lambda_1\Delta t + (\theta-1)(\kappa_1\nu_\varepsilon-1)A_1\right)T_t \\ &\quad - ((\theta-1)(\kappa_1\nu-1)A_3 + 0.5((\theta-1)\kappa_1 A_1\chi\Theta-\gamma)^2)\sigma_t^2, \\ r_f &= -\theta\log\delta + \gamma\mu_c - (\theta-1)(\kappa_0 + (\kappa_1-1)A_0 + \kappa_1 A_1\chi\Theta\mu_c) \\ &\quad - (\theta-1)(\kappa_1 A_1\chi\mu_\varepsilon + \kappa_1 A_3(1-\nu)\bar{\sigma}^2) \\ &\quad - 0.5[((\theta-1)\kappa_1 A_1\chi)^2\sigma_\zeta^2 + ((\theta-1)\kappa_1 A_3)^2\sigma_w^2] \\ &\quad - ((\theta-1)\kappa_1 A_2-\gamma)d\left(1+\frac{((\theta-1)\kappa_1 A_2-\gamma)d}{2}\right)\Delta t\lambda_0.\end{aligned} \tag{A.10}$$

A.1.3 The price-dividend ratio

We similarly solve the price-dividend ratio of portfolio i, $z_{t,i}$ from Equation (13). The Campbell-Shiller decomposition for the returns of portfolio i gives us

$$\begin{aligned} r_{t+1,i} &= \kappa_{0,i} + \Delta d_{t+1,i} + \kappa_{1,i} z_{t+1,i} - z_{t,i} \\ &= \kappa_{0,i} + \mu_d + (\kappa_{1,i} - 1) A_{0,i} + \kappa_{1,i} A_{1,i} \chi \mu_\varepsilon + \kappa_{1,i} A_{1,i} \chi \Theta \mu_c \\ &\quad + \kappa_{1,i} A_{3,i} (1-\nu) \bar{\sigma}^2 + \varphi_d \sigma_t u_{t+1} + (\pi_d + \kappa_{1,i} A_{1,i} \chi \Theta) \sigma_t \eta_{t+1} + \kappa_{1,i} A_{1,i} \chi \sigma_\zeta \zeta_{t+1} \\ &\quad + \kappa_{1,i} A_{3,i} \sigma_w w_{t+1} + (\kappa_{1,i} \nu_\varepsilon - 1) A_{1,i} T_t + (\phi_i \rho + (\kappa_{1,i} \rho - 1) A_{2,i}) X_t \\ &\quad + (\phi_i + \kappa_{1,i} A_{2,i}) d \Delta N_{t+1} + A_{3,i} (\kappa_{1,i} \nu - 1) \sigma_t^2 \end{aligned} \tag{A.11}$$

Combining this with the IMRS from Equation (A.7) we find the expectation

$$\begin{aligned} \mathbb{E}_t[m_{t+1} + r_{t+1,i}] = \theta \log \delta &+ (\theta - 1)\kappa_0 + \kappa_{0,i} + (\theta - 1)(\kappa_1 - 1) A_0 + (\kappa_{1,i} - 1) A_{0,i} + \mu_d \\ &+ (((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i}) \chi \Theta - \gamma) \mu_c + ((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i}) \chi \mu_\varepsilon \\ &+ ((\theta - 1)\kappa_1 A_3 + \kappa_{1,i} A_{3,i})(1 - \nu) \bar{\sigma}^2 \\ &+ ((\theta - 1)(\kappa_1 \nu - 1) A_3 + (\kappa_{1,i} \nu - 1) A_{3,i}) \sigma_t^2 \\ &+ ((\theta - 1)(\kappa_1 \rho - 1) A_2 + (\kappa_{1,i} \rho - 1) A_{2,i} + \phi_i \rho - \gamma \rho) X_t \\ &+ [(\theta - 1)(\kappa_1 \nu_\varepsilon - 1) A_1 + (\kappa_{1,i} \nu_\varepsilon - 1) A_{1,i} \\ &+ ((\theta - 1)\kappa_1 A_2 + \kappa_{1,i} A_{2,i} + \phi_i - \gamma) d \lambda_1 \Delta t] T_t \\ &+ ((\theta - 1)\kappa_1 A_2 + \kappa_{1,i} A_{2,i} + \phi_i - \gamma) d \lambda_0 \Delta t, \end{aligned} \tag{A.12}$$

and variance

$$\begin{aligned} \mathrm{Var}_t(m_{t+1} + r_{t+1,i}) = ((\theta - 1)\kappa_1 A_3 &+ \kappa_{1,i} A_{3,i})^2 \sigma_w^2 \\ &+ ((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})^2 \chi^2 \sigma_\zeta^2 \\ &+ \varphi_d^2 \sigma_t^2 + (((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i}) \chi \Theta + \pi_d - \gamma)^2 \sigma_t^2 \\ &+ ((\theta - 1)\kappa_1 A_2 + \kappa_{1,i} A_{2,i} + \phi_i - \gamma)^2 d^2 \lambda_0 \Delta t \\ &+ ((\theta - 1)\kappa_1 A_2 + \kappa_{1,i} A_{2,i} + \phi_i - \gamma)^2 d^2 \lambda_1 \Delta t T_t. \end{aligned} \tag{A.13}$$

Which we fill in in the Euler equation and take logs to find

$$
\begin{aligned}
0 &= \log \mathbb{E}_t[\exp(m_{t+1} + r_{t+1,i})] = \mathbb{E}_t[(m_{t+1} + r_{t+1,i})] + 0.5\mathrm{Var}_t(m_{t+1} + r_{t+1,i}) \\
&= \theta \log \delta + (\theta - 1)\kappa_0 + \kappa_{0,i} + (\theta - 1)(\kappa_1 - 1)A_0 + (\kappa_{1,i} - 1)A_{0,i} + \mu_d \\
&\quad + (((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})\chi\Theta - \gamma)\mu_c + ((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})\chi\mu_\varepsilon \\
&\quad + C_N(1 + 0.5C_N)\lambda_0 \Delta t + ((\theta - 1)\kappa_1 A_3 + \kappa_{1,i} A_{3,i})(1 - \nu)\bar{\sigma}^2 \\
&\quad + 0.5(((\theta - 1)\kappa_1 A_3 + \kappa_{1,i} A_{3,i})^2 \sigma_w^2 + ((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})^2 \chi^2 \sigma_\zeta^2) \\
&\quad + [(\theta - 1)(\kappa_1 \nu - 1)A_3 + (\kappa_{1,i}\nu - 1)A_{3,i} \\
&\quad + 0.5\varphi_d^2 + 0.5(((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})\chi\Theta + \pi_d - \gamma)^2]\sigma_t^2 \\
&\quad + ((\theta - 1)(\kappa_1 \rho - 1)A_2 + (\kappa_{1,i}\rho - 1)A_{2,i} + \phi_i \rho - \gamma\rho)X_t \\
&\quad + ((\theta - 1)(\kappa_1 \nu_\varepsilon - 1)A_1 + (\kappa_{1,i}\nu_\varepsilon - 1)A_{1,i} + C_N(1 + 0.5C_N)\lambda_1 \Delta t)T_t, \\
C_N &= ((\theta - 1)\kappa_1 A_2 + \kappa_{1,i} A_{2,i} + \phi_i - \gamma)d.
\end{aligned}
\tag{A.14}
$$

We set all terms in front of T_t, X_t and σ_t^2 to zero to solve for $z_{t,i}$:

$$
\begin{aligned}
(1 - \kappa_{1,i}\nu)A_{3,i} &= 0.5(((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})\chi\Theta + \pi_d - \gamma)^2 \\
&\quad + 0.5\varphi_d^2 + (\theta - 1)(\kappa_1 \nu - 1)A_3, \\
(1 - \kappa_{1,i}\rho)A_{2,i} &= (\theta - 1)(\kappa_1 \rho - 1)A_2 + \phi_i \rho - \gamma\rho, \\
(1 - \kappa_{1,i}\nu_\varepsilon)A_{1,i} &= (\theta - 1)(\kappa_1 \nu_\varepsilon - 1)A_1 + C_N(1 + 0.5C_N)\lambda_1 \Delta t, \\
(1 - \kappa_{1,i})A_{0,i} &= \theta \log \delta + (\theta - 1)\kappa_0 + \kappa_{0,i} + (\theta - 1)(\kappa_1 - 1)A_0 + \mu_d \\
&\quad + (((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})\chi\Theta - \gamma)\mu_c + ((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})\chi\mu_\varepsilon \\
&\quad + C_N(1 + 0.5C_N)\lambda_0 \Delta t + ((\theta - 1)\kappa_1 A_3 + \kappa_{1,i} A_{3,i})(1 - \nu)\bar{\sigma}^2 \\
&\quad + 0.5(((\theta - 1)\kappa_1 A_3 + \kappa_{1,i} A_{3,i})^2 \sigma_w^2 + ((\theta - 1)\kappa_1 A_1 + \kappa_{1,i} A_{1,i})^2 \chi^2 \sigma_\zeta^2), \\
C_N &= ((\theta - 1)\kappa_1 A_2 + \kappa_{1,i} A_{2,i} + \phi_i - \gamma)d.
\end{aligned}
\tag{A.15}
$$

A.1.4 The equity risk premium

As discussed by Bansal and Yaron (2004), the risk premium of any asset i is based on the covariance between the unexpected returns and innovations in the SDF m_{t+1} as

$$\ln\mathbb{E}_t[R_{i,t+1}] - r_{f,t} = \mathbb{E}_t[r_{i,t+1} - r_{f,t}] + 0.5\text{Var}_t(r_{i,t+1}) = -\text{Cov}_t(m_{t+1} - \mathbb{E}_t[m_{t+1}], r_{i,t+1} - \mathbb{E}_t[r_{i,t+1}]). \tag{A.16}$$

The innovation in the SDF based on Equation (A.1) equals

$$m_{t+1} - \mathbb{E}_t[m_{t+1}] = -\lambda_\eta \sigma_t \eta_{t+1} - \lambda_\zeta \sigma_\zeta \zeta_{t+1} - \lambda_w \sigma_w w_{t+1} - \lambda_X(\Delta N_{t+1} - (\lambda_0 \Delta t + \lambda_1 \Delta t T_t)), \tag{A.17}$$

where

$$\begin{aligned} \lambda_\eta &= \gamma + (1-\theta)\kappa_1 A_1 \chi \Theta, \\ \lambda_\zeta &= (1-\theta)\kappa_1 A_1 \chi, \\ \lambda_w &= (1-\theta)\kappa_1 A_3, \\ \lambda_X &= (\gamma + (1-\theta)\kappa_1 A_2)d. \end{aligned} \tag{A.18}$$

Based on Equation (A.11), the innovation in the return of portfolio i is given by

$$\begin{aligned} r_{i,t+1} - \mathbb{E}_t[r_{i,t+1}] = \varphi_d \sigma_t u_{t+1} &+ (\pi_d + \kappa_{1,i} A_{1,i} \chi \Theta)\sigma_t \eta_{t+1} + \kappa_{1,i} A_{1,i} \chi \sigma_\zeta \zeta_{t+1} \\ &+ \kappa_{1,i} A_{3,i} \sigma_w w_{t+1} + (\phi_i + \kappa_{1,i} A_{2,i}) d(\Delta N_{t+1} - (\lambda_0 \Delta t + \lambda_1 \Delta t T_t)) \end{aligned} \tag{A.19}$$

This brings us to the conditional risk premium of asset i:

$$\begin{aligned} \ln\mathbb{E}_t[R_{i,t+1}] - r_{f,t} &= -\text{Cov}_t(m_{t+1} - \mathbb{E}_t[m_{t+1}], r_{i,t+1} - \mathbb{E}_t[r_{i,t+1}]) \\ &= \lambda_\eta \beta_{i,\eta} \sigma_t^2 + \lambda_\zeta \beta_{i,\zeta} \sigma_\zeta^2 + \lambda_w \beta_{i,w} \sigma_w^2 + \lambda_X \beta_{i,X}(\lambda_0 \Delta t + \lambda_1 \Delta t T_t), \end{aligned} \tag{A.20}$$

where lambdas are given in Equation (A.18) and betas are

$$\begin{aligned}
\beta_{i,\eta} &= \pi_d + \kappa_{1,i} A_{1,i} \chi \Theta, \\
\beta_{i,\zeta} &= \kappa_{1,i} A_{1,i} \chi, \\
\beta_{i,w} &= \kappa_{1,i} A_{3,i}, \\
\beta_{i,X} &= (\phi_i + \kappa_{1,i} A_{2,i}) d.
\end{aligned} \tag{A.21}$$

A.1.5 Solving the fixed-point problem

The approximation constants $\kappa_{0(,i)}$ and $\kappa_{1(,i)}$ are defined as

$$\begin{aligned}
\kappa_{0(,i)} &= \log(1 + \exp(\bar{z}_{(i)})) - \kappa_{1(,i)} \bar{z}_{(i)}, \\
\kappa_{1(,i)} &= \frac{\exp(\bar{z}_{(i)})}{1 + \exp(\bar{z}_{(i)})},
\end{aligned} \tag{A.22}$$

where $\bar{z}_{(i)}$ is the mean of the wealth-consumption ratio z_t or price-dividend ratio of portfolio i, $z_{t,i}$. We solve the mean of these ratios by numerically (through iteration) solving the fixed point problem

$$\begin{aligned}
\bar{z}_{(i)} &= A_{0(,i)} + A_{1(,i)} \bar{T}_t + A_{2(,i)} \bar{X}_t + A_{3(,i)} \bar{\sigma}_t^2 \\
&= A_{0(,i)} + A_{1(,i)} \left(\frac{1}{n} \mathbb{E}_0 \left[\sum_{t=1}^{n} T_t \right] \right) + A_{2(,i)} \left(\frac{1}{n} \mathbb{E}_0 \left[\sum_{t=1}^{n} X_t \right] \right) + A_{3(,i)} \bar{\sigma_t}^2,
\end{aligned} \tag{A.23}$$

where the expected averages $\bar{T}_t$ and $\bar{X}_t$ over the sample from period $t = 1$ to $t = n$ are given as

$$\begin{aligned}
\frac{1}{n} \mathbb{E}_0 \left[\sum_{t=1}^{n} T_t \right] &= \left(\frac{T_0}{n} - \frac{\chi(\mu_\varepsilon + \Theta \mu_c)}{n(1 - \nu_\varepsilon)} \right) \left(\frac{\nu_\varepsilon - \nu_\varepsilon^{n+1}}{1 - \nu_\varepsilon} \right) + \frac{\chi(\mu_\varepsilon + \Theta \mu_c)}{1 - \nu_\varepsilon}, \\
\frac{1}{n} \mathbb{E}_0 \left[\sum_{t=1}^{n} X_t \right] &= \frac{1}{n} \left(-X_0 + \sum_{t=0}^{n} \mathbb{E}_0[X_t] \right) \\
&= \frac{X_0}{n} \left(\frac{\rho - \rho^{n+1}}{1 - \rho} \right) + \frac{1}{n} \left(\frac{d \Delta t \lambda_0}{1 - \rho} + \frac{d \Delta t \lambda_1}{1 - \rho} \frac{\chi(\mu_\varepsilon + \Theta \mu_c)}{1 - \nu_\varepsilon} \right) \left(n - \frac{\rho - \rho^{n+1}}{1 - \rho} \right) \\
&\quad + \frac{1}{n \nu_\varepsilon} \left(\frac{d \Delta t \lambda_1}{1 - \left(\frac{\rho}{\nu_\varepsilon} \right)} \right) \left(T_0 - \frac{\chi(\mu_\varepsilon + \Theta \mu_c)}{1 - \nu_\varepsilon} \right) \left(\frac{1 - \nu_\varepsilon^{n+1}}{1 - \nu_\varepsilon} - \frac{1 - \rho^{n+1}}{1 - \rho} \right).
\end{aligned} \tag{A.24}$$

The sample used in our regressions starts in 1947Q1 and runs until 2019Q4, 292 observations of quarterly data. In our simulations, we construct a similar sample. We simulate $n = 939$ months of data, resulting in a sample of 292 quarters after we drop the first five years of our simulation as burn-in and the last three months because of a lead in the risk-free yields. We calibrate the model to set the expected starting temperature in our simulations to zero, i.e. $\mathbb{E}_0[T_{60}] = 0$, based on our burn-in of 60 months. The risk-free yields are led by one period, since the assumed absence of credit risk implies that the risk-free return realized at the end of quarter Q_i is already observed in Q_{i-1}.

A.2 Adjustments to the LRR-T model

The LRR-T model presented in Section 3.1 extends the LRR-T model of Bansal, Kiku, and Ochoa (2019) in several ways. First, in Equation (10) we allow for more flexibility in the dividend process by letting it have its own mean μ_d and own loading π_d on the shock η_{t+1}, rather than tying these parameters to the corresponding parameters in the consumption process. We further allow for time-varying volatility of consumption and dividends as in Bansal and Yaron (2004).

Second, in Equation (11) we assume a persistence parameter $\rho < 1$ for the economic impact of disasters X, instead of $\rho = 1$. We believe that it is reasonable that climate-related disasters have a persistent impact on consumption and dividend growth, but that the impact is unlikely to be permanent. In the same equation, we let the *increments* of the disaster process ($\Delta\tilde{N}$) be Poisson distributed, instead of the overall process $\tilde{N}$. This adjustment makes the current disaster intensity dependent on recent temperature levels instead of on the historical path of temperature growth.

Third, in Equation (12), we allow for a separate trend μ_ε in the atmospheric carbon concentration to capture potential trends in emissions that are unrelated to consumption growth.

Fourth, in Equation (13), we allow the log wealth-consumption ratio z_t and the log price-dividend ratio $z_{t,m}$ to also depend on the time-varying variance σ_t^2 and on the economic disaster impact X_t. The dependence on σ_t^2 directly follows from the inclusion of time-varying volatility in Equation (10). We include X_t because disasters have a persistent impact on future dividend growth, thereby lowering the wealth-consumption ratio and price-dividend ratio.

A.3 Posterior VAR distributions

We adopt a Bayesian approach to estimate the predictive VAR model given by

$$\begin{bmatrix} r_{i,t+1} \\ p_{t+1} - d_{t+1} \\ r_{f,t} \\ T_{t+1} \end{bmatrix} = C' \begin{bmatrix} 1 \\ p_t - p_d \\ r_{f,t-1} \\ T_t \end{bmatrix} + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, \Sigma). \tag{A.25}$$

where C is the VAR coefficients matrix and Σ is the variance-covariance matrix of the residuals. As shown in the Online Appendix of Avramov, Cederburg, and Lucivjanska (2018), the posterior distribution of the VAR parameters with the model-based prior conditional on the set of parameters of the asset pricing model Θ_M and on the data observable up to time N, D_N, is given by

$$\begin{aligned} \Sigma \mid D_N, \Theta_M &\sim IW((\omega_M + \omega_D)N\hat{\Sigma}(\Theta_M), (\omega_M + \omega_D)N - 4), \\ C \mid \Sigma, D_N, \Theta_M &\sim N(\hat{C}(\Theta_M), \Sigma \otimes (\omega_M N \Gamma^*_{xx} + \omega_D X'X)^{-1}), \end{aligned} \tag{A.26}$$

in which

$$\begin{aligned} \hat{\Sigma}(\Theta_M) &= \frac{1}{(\omega_M + \omega_D)N}[(\omega_M N \Gamma^*_{yy} + \omega_D Y'Y) \\ &\quad - (\omega_M N \Gamma^*_{xy} + \omega_D Y'X)(\omega_M N \Gamma^*_{xx} + \omega_D X'X)^{-1}(\omega_M N \Gamma^*_{xy} + \omega_D X'Y)], \\ \hat{C}(\Theta_M) &= (\omega_M N \Gamma^*_{xx} + \omega_D X'X)^{-1}(\omega_M N \Gamma^*_{xy} + \omega_D X'Y), \end{aligned} \tag{A.27}$$

where N is the number of observations in our sample (292 quarters) and Γ^*_{xx}, Γ^*_{xy}, and Γ^*_{yy} are the population moments implied by the asset pricing models in Section 3.1, as defined in the Online Appendix of Avramov, Cederburg, and Lucivjanska (2018). Finally, $\frac{\omega_M}{\omega_M+\omega_D}$ and $\frac{\omega_D}{\omega_M+\omega_D}$ are the weights given to the model-based prior and the historical data, respectively. For the Bayesian investors with prior beliefs elicited from the LRR or LRR-T model, we assign equal weights to the informative prior and the sample data by setting $\omega_M = \omega_D = 1$. For the agnostic investor we set $\omega_M = 0$ and $\omega_D = 1$. In this case, Equation (A.26) reduces to the posterior distribution corresponding to the VAR estimated using historical data and an uninformative multivariate Jeffreys prior. For the dogmatic investor we set $\omega_M = 1$ and $\omega_D = 0$. Equation (A.26) then reduces to the prior distribution implied by the long-run risks model.

References

Addoum, J., D. Ng, and A. Ortiz-Bobea, 2020, Temperature shocks and establishment sales, *Review of Financial Studies* 33, 1331–1366.

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

Balduzzi, P., and A. Lynch, 1999, Transaction costs and predictability: some utility cost calculations, *Journal of Financial Economics* 52, 47–78.

Balvers, R., D. Du, and X. Zhao, 2017, Temperature shocks and the cost of equity capital: Implications for climate change perceptions, *Journal of Banking and Finance* 77, 18–34.

Bansal, R., D. Kiku, and M. Ochoa, 2019, Climate change risk, Working Paper, Duke University.

Bansal, R., D. Kiku, and A Yaron, 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183–221.

——— , 2016, Risks for the long run: Estimation with time aggregation, *Journal of Monetary Economics* 82, 52–69.

Bansal, R., and A. Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Barberis, N., 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.

Barnett, M., 2021, Climate change and uncertainty: An asset pricing perspective, Working Paper, Arizona State University.

Beeler, J., and J. Campbell, 2012, The long-run risks model and aggregate asset prices: An empirical assessment, *Critical Finance Review* 1, 141–182.

Bolton, P., and M. Kacperczyk, 2021, Do investors care about carbon risk?, *Journal of Financial Economics, Forthcoming.*

——— , 2022, Global pricing of carbon-transition risk, Working Paper, Columbia Business School.

Campbell, J., Y. Chan, and L. Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.

Campbell, J., and R. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Campbell, J., and L. Viceira, 2005, The term structure of the risk-return tradeoff, *Financial Analysts Journal* 61, 34–44.

Carvalho, C., H. Lopes, and R. McCulloch, 2018, On the long-run volatility of stocks, *Journal of the American Statistical Association* 113, 1050–1069.

Colacito, R., B. Hoffmann, and T. Phan, 2019, Temperature and growth: A panel analysis of the united states, *Journal of Money, Credit, and Banking* 51, 313–368.

Correa, R., A. He, C. Herpfer, and U. Lel, 2021, The rising tide lifts some interest rates: Climate change, natural disasters, and loan pricing, Working Paper.

Del Negro, M., and F. Schorfheide, 2004, Priors from general equilibrium models for VARs, *International Economic Review* 45, 643–673.

Dell, M., B. Jones, and B. Olken, 2012, Temperature shocks and economic growth: Evidence from the last half century, *American Economic Journal: Macroeconomics* 4, 66–95.

———, 2014, What do we learn from the weather? The new climate-economy literature, *Journal of Economic Literature* 52, 740–798.

Diris, B., F. Palm, and P. Schotman, 2014, Long-term strategic asset allocation: An out-of-sample evaluation, *Management Science* 61, 2185–2202.

Epstein, L., and S. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Fama, E., and K. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.

Giglio, S., B. Kelly, and J. Stroebel, 2021, Climate finance, *Annual Review of Financial Economics* 13, 15–36.

Hoevenaars, R., R. Molenaar, P. Schotman, and T. Steenkamp, 2014, Strategic asset allocation for long-term investors: Parameter uncertainty and prior information, *Journal of Applied Econometrics* 29, 353–376.

Hong, H., F. Li, and J. Xu, 2019, Climate risks and market efficiency, *Journal of Econometrics* 208, 265–281.

Huynh, T., and Y. Xia, 2021, Panic selling when disaster strikes: Evidence in the bond and stock markets, *Management Science, Forthcoming.*

Ilhan, E., Z. Sautner, and G. Vilkov, 2021, Carbon tail risk, *Review of Financial Studies* 34, 1540–1571.

IPCC, 2021, Climate change 2021: The physical science basis. Working group I contribution to the sixth assessment report of the intergovernmental panel on climate change, Geneva.

Johannes, M., A. Korteweg, and N. Polson, 2014, Sequential learning, predictability, and optimal portfolio returns, *Journal of Finance* 69, 611–644.

Jurek, J., and L. Viceira, 2011, Optimal value and growth tilts in long-horizon portfolios, *Review of Finance* 15, 29–74.

Kandel, S., and R. Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance* 51, 385–424.

Krueger, P., Z. Sautner, and L. Starks, 2020, The importance of climate risks for institutional investors, *Review of Financial Studies* 33, 1067–1111.

Kumar, A., W. Xin, and C. Zhang, 2019, Climate sensitivity and predictable returns, Working Paper, University of Miami.

Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan, 2013, The wealth-consumption ratio, *The Review of Asset Pricing Studies* 3, 38–94.

Pastor, L., and R. Stambaugh, 2012, Are stocks really less volatile in the long run?, *Journal of Finance* 67, 431–478.

Siegel, J., 2014, *Stocks for the Long Run* (McGraw Hill, New York, NY).

Figure 1: Temperature anomaly forecasts

This figure shows the temperature anomaly forecasts implied by the predictive VAR estimates in panels A-C of Table 4 and by the exogenous climate change process described in Section 2.3. The temperature anomaly is measured in degrees Celsius and the horizon ranges from 1 to 100 quarters. The VARs are estimated using the historical data and an uninformative prior (`Data+uninf. prior`), the historical data and an informative prior derived from the temperature long-run risks model (`Data + LRR-T prior`), or the informative LRR-T prior alone (`LRR-T prior`). The starting point for the forecasts is the average temperature anomaly observed over the last five years in the historical sample (for the VAR model `Data+uninf. prior` and the `exogenous scenario`), the average temperature anomaly observed over the last five years in the data simulated from the LRR-T model (for the VAR model `LRR-T prior`), and the average of the five-year mean implied by the LRR-T model and the five-year sample mean (for the VAR model `Data + LRR-T prior`). We simulate 250,000 future temperature anomaly paths based on each model and plot the mean of these simulated trajectories. For the exogenous climate scenario we also show the 5% and 95% percentiles.

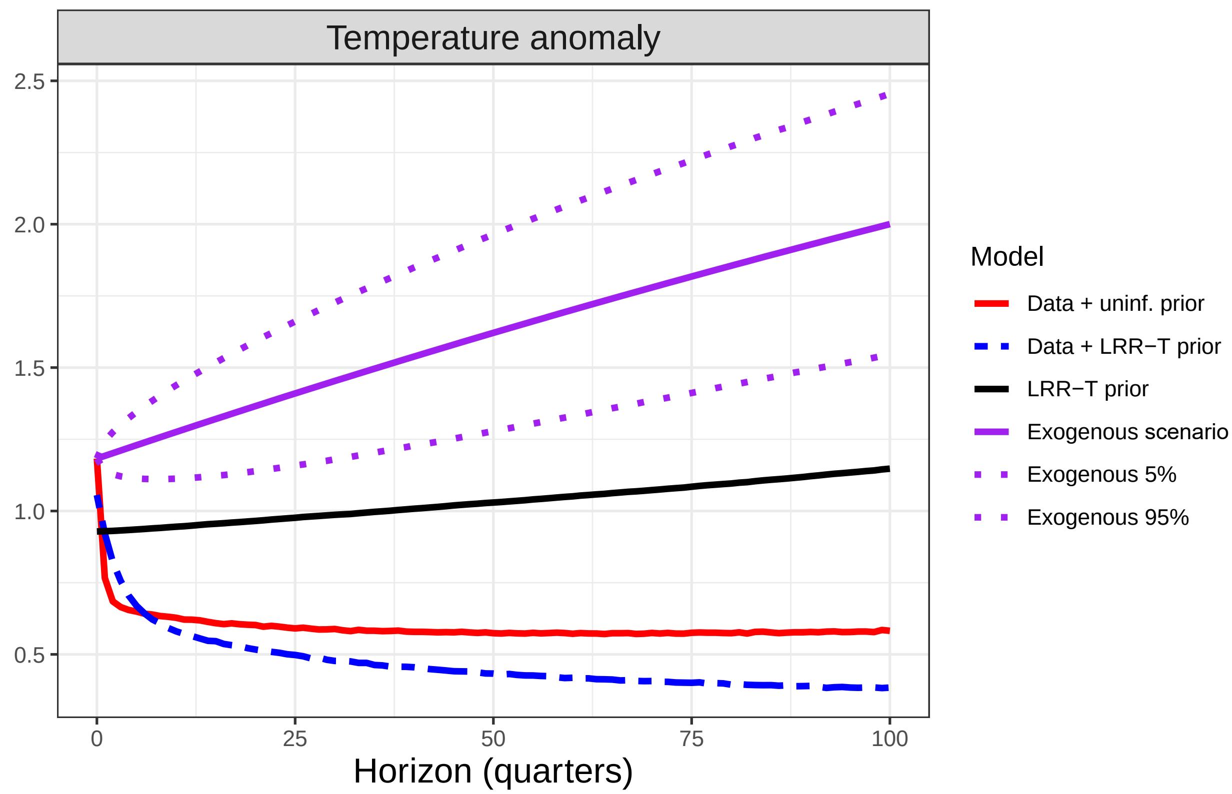

Figure 2: Simulated path from LRR-T model

This figure shows a single simulated path from the LRR-T model using the parameter configuration in Table 2. The stock market return (r_m) is the log real return (including dividends) on the S&P 500 index (in %), the price-dividend ratio (P/D) is the log difference between the S&P 500's price and the monthly dividend payment on the index, the risk-free rate (r_f) is the log ex-ante real risk-free rate, and the temperature anomaly (T) is the temperature (in degrees Celsius) in a given month minus the average temperature in that same month over the base period 1901-1946. We simulate 250,000 samples with a length of 876 months to match the length of our 1947-2019 historical sample. The simulation in the figure is randomly chosen from the 250,000 samples.

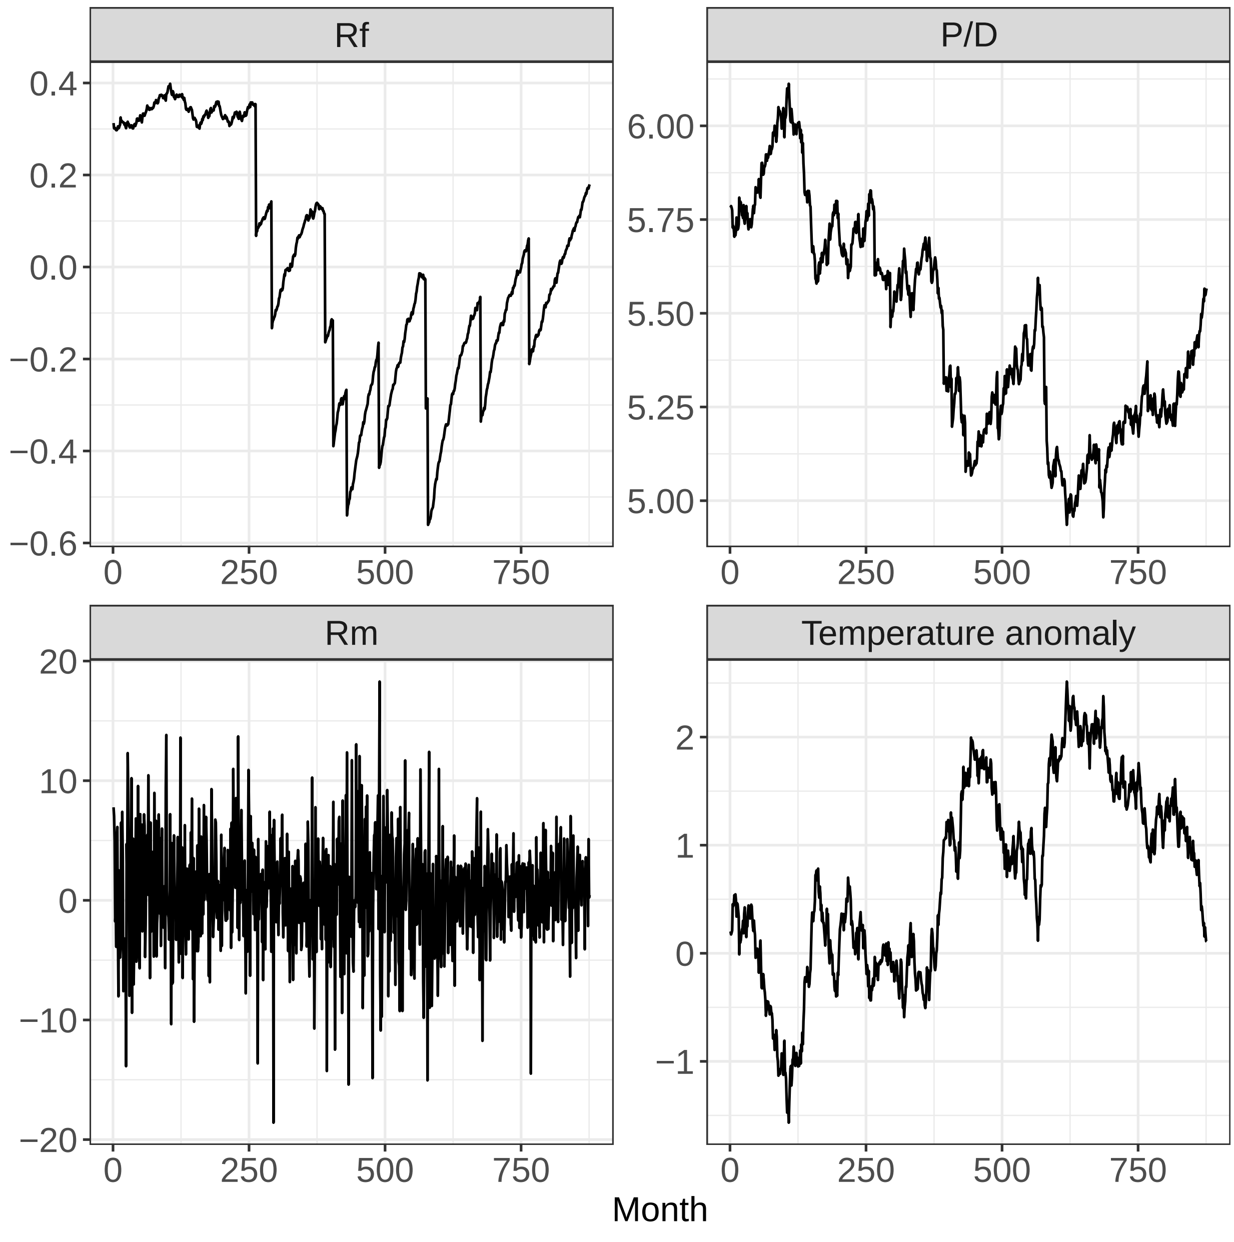

Figure 3: Predictive variance ratios by investment horizon

This figure shows predictive variance ratios for returns on the stock market portfolio (top left) and the underlying components related to mean reversion (top right), uncertainty about future expected returns (bottom left), and estimation risk (bottom right). We plot four different variance ratios that are constructed based on the return dynamics implied by the VAR model in Equation (1). VARs are estimated using the historical sample data and an uninformative prior (`Data+uninf. prior`), the historical data and an informative prior derived from the temperature long-run risks model (`Data + LRR-T prior`), the informative LRR-T prior alone (`LRR-T prior`), or the historical data and an informative prior derived from the standard long-run risks model (`Data + LRR prior`). We decompose each of the variance ratios into three parts according to Equation (4), using 250,000 draws from the posterior distribution of the VAR parameters. The investment horizon ranges from 1 to 100 quarters.

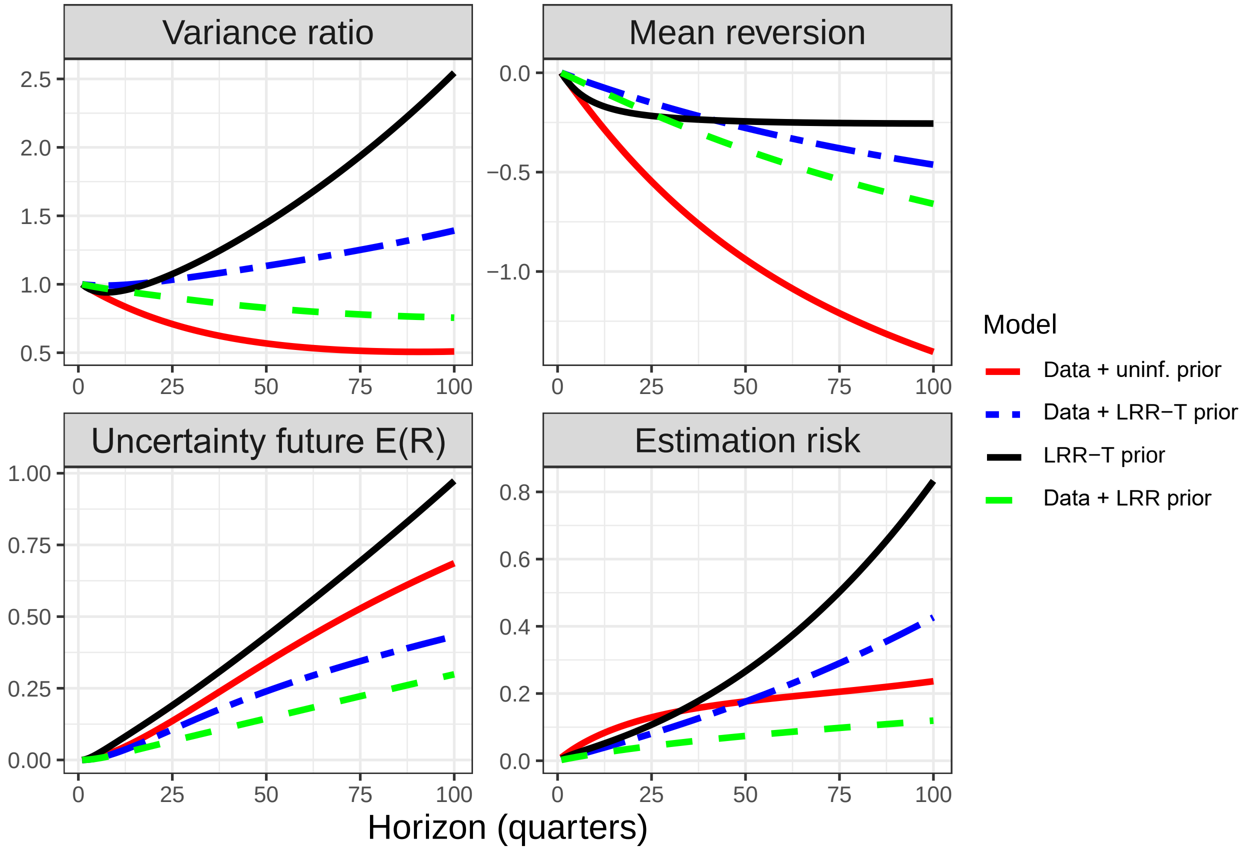

Figure 4: Predictive returns and risk premia by investment horizon

This figure shows forecasts of average risk-free returns, market returns, and market risk premia per quarter (in %), for investment horizons ranging from 1 to 100 quarters. For each variable, we plot four different forecasts implied by the VAR model in Equation (1). VARs are estimated using the historical sample data and an uninformative prior (`Data+uninf. prior`), the historical data and an informative prior derived from the temperature long-run risks model (`Data + LRR-T prior`), the informative LRR-T prior alone (`LRR-T prior`), or the historical data and an informative prior derived from the standard long-run risks model (`Data + LRR prior`). The predictive returns are computed based on 250,000 draws from the predictive distribution corresponding to each model. Specifically, we simulate 250,000 future return paths based on 250,000 draws from the posterior distribution of the VAR parameters. For each of these return paths, we compute the simple average of the quarterly returns over each horizon. The forecasts shown in the figure are the means of the average returns per quarter across the simulated return paths.

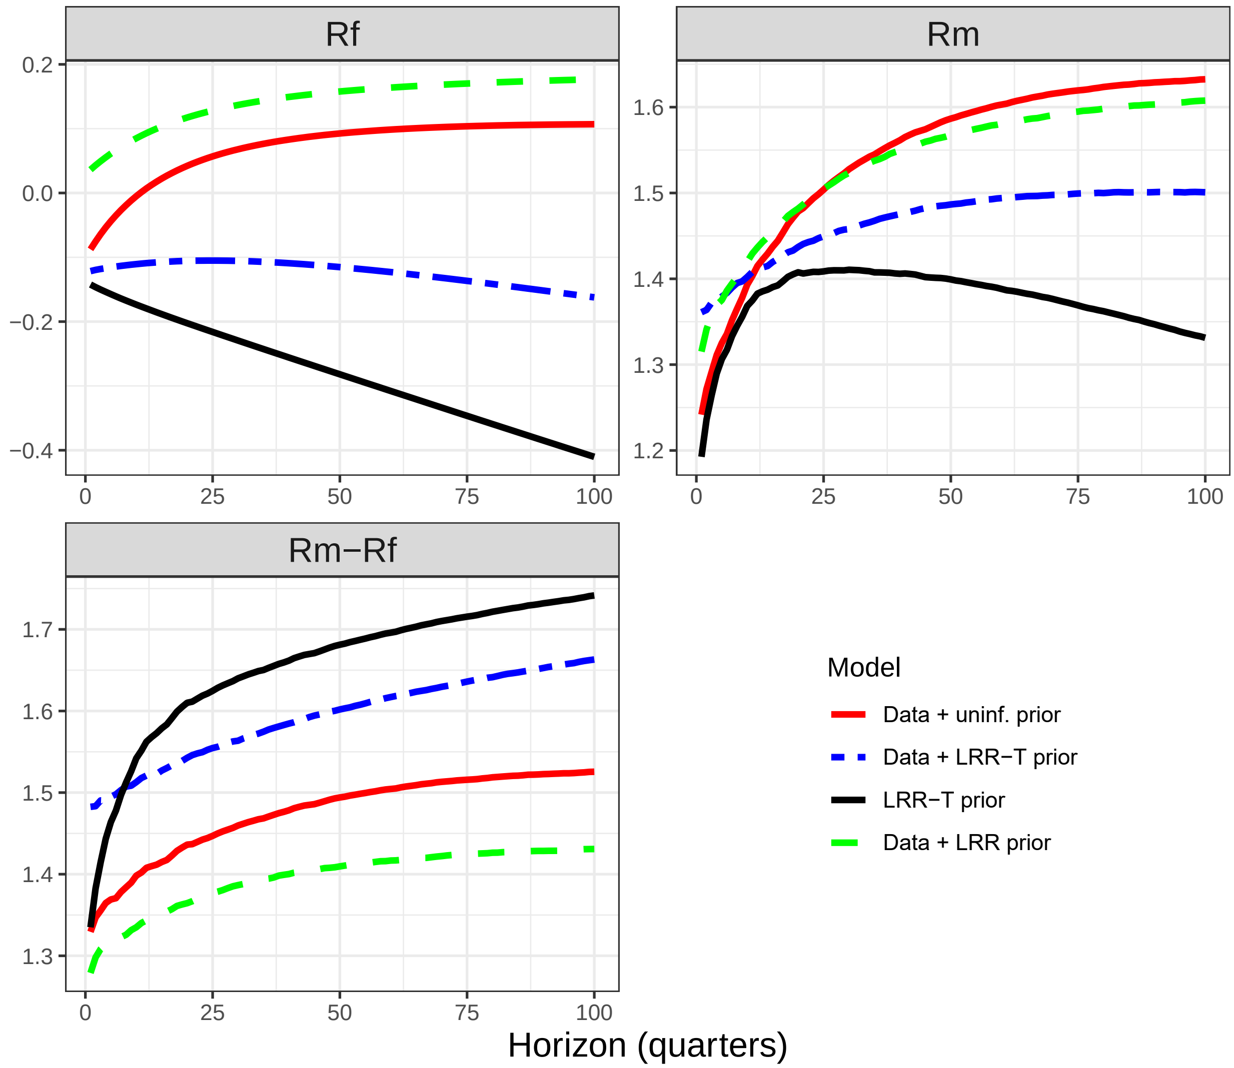

Figure 5: Predictive correlations between risk-free asset and market portfolio

This figure shows predictive correlations between the returns on the risk-free asset and the returns on the market portfolio, for horizons ranging from 1 to 100 quarters. The predictive correlations are based on the return dynamics implied by the VAR in Equation (1). VARs are estimated using the historical sample data and an uninformative prior (`Data+uninf. prior`), the historical data and an informative prior derived from the temperature long-run risks model (`Data + LRR-T prior`), or the informative LRR-T prior alone (`LRR-T prior`). The predictive correlations are computed based on 250,000 draws from the predictive distribution corresponding to each model. We simulate 250,000 future return paths based on 250,000 draws from the posterior distribution of the VAR parameters. For each of these paths, we compute the correlation between the quarterly returns on the risk-free asset and the market portfolio over each horizon. The correlations in the figure are the averages of the correlations across the simulations.

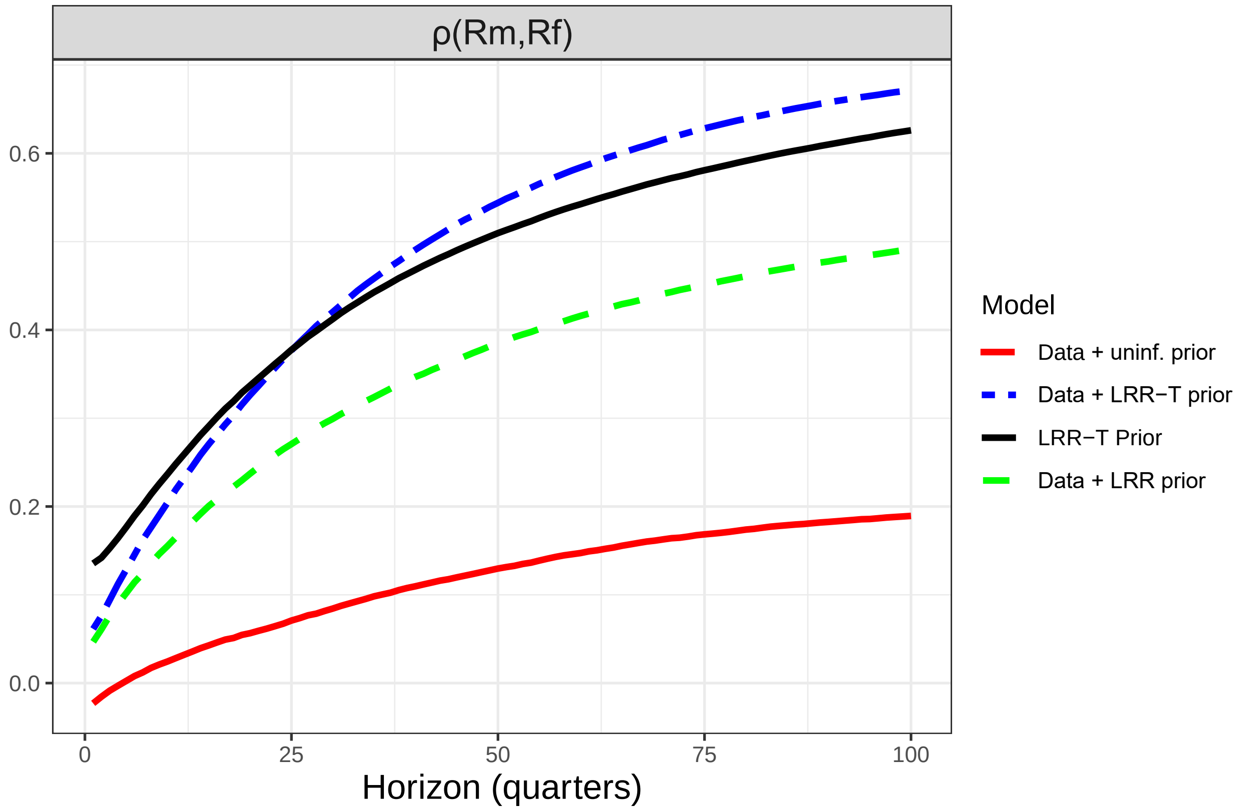

Figure 6: Optimal portfolio weights by investment horizon

This figure shows optimal portfolio weights for an investor who allocates wealth between the market portfolio of stocks and the risk-free asset, with an investment horizon ranging from 1 to 100 quarters. The plots show the optimal allocation to the stock portfolio for a long-only buy-and-hold investor who has power utility with risk aversion $A = 5$. The optimal weights are computed following the approach explained in Section 2.5, using the predictive return distribution implied by the VAR model in Equation (1). VARs are estimated using the historical sample data and an uninformative prior (`Data+uninf. prior`), the historical data and an informative prior derived from the temperature long-run risks model (`Data + LRR-T prior`), the informative LRR-T prior alone (`LRR-T prior`), or the historical data and an informative prior derived from the standard long-run risks model (`Data + LRR prior`). The optimal allocation in the top panel ignores parameter uncertainty and is based on 250,000 samples from the predictive distribution with the parameters of the predictive VAR model set equal to their posterior means. The optimal allocation in the bottom panel accounts for parameter uncertainty and is based on 250,000 samples from the predictive distribution with the parameters of the predictive VAR model drawn from their posterior distribution.

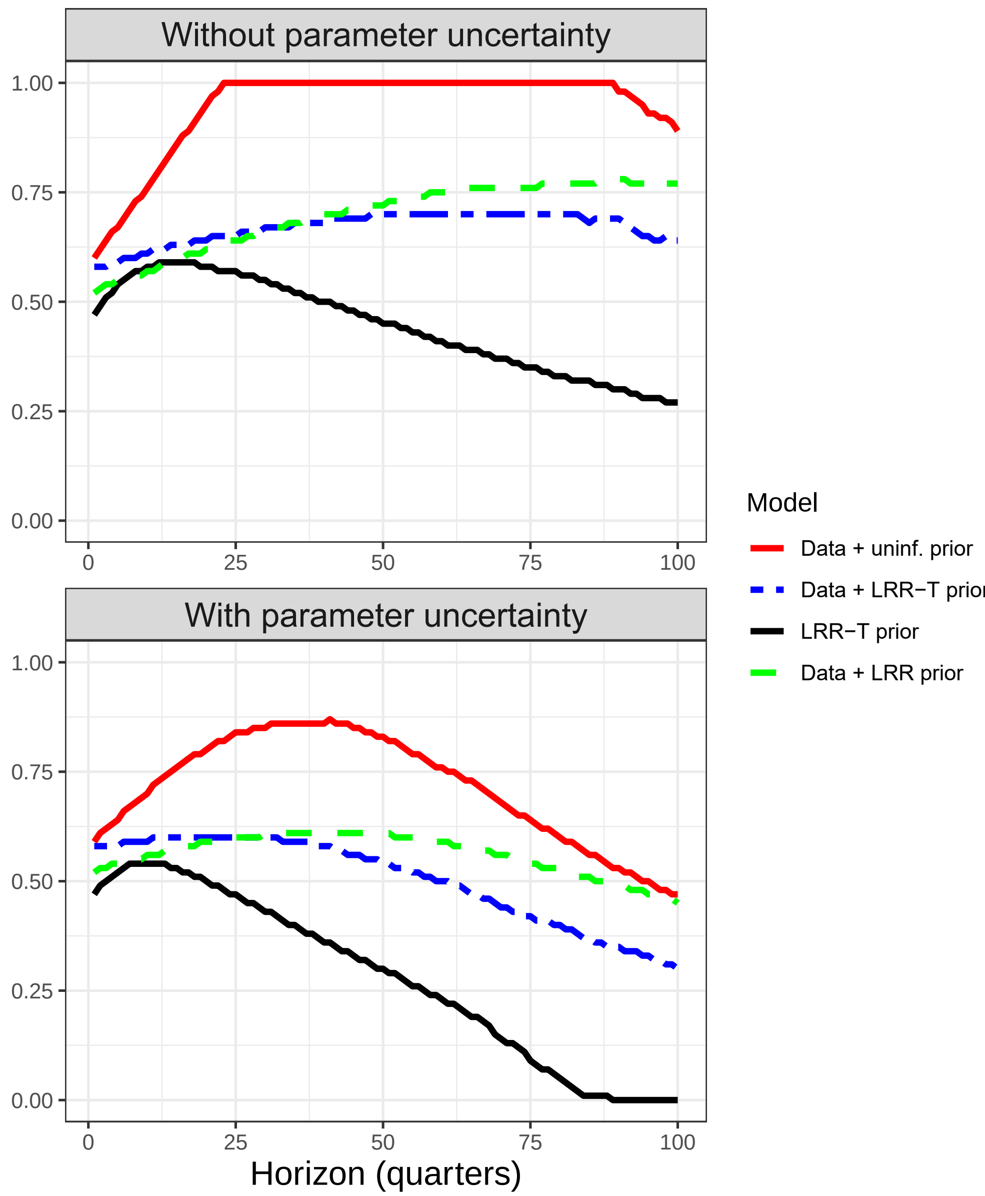

Figure 7: Predictive variance ratios in alternative climate scenarios

This figure shows predictive variance ratios for returns on the stock market portfolio (top left) and the underlying components related to mean reversion (top right), uncertainty about future expected returns (bottom left), and estimation risk (bottom right). We plot three different variance ratios that are constructed based on the return dynamics implied by the VAR model in Equation (1). All VARs are estimated using the historical data and an informative prior derived from the temperature long-run risks model. We consider three different climate scenarios: i) the default scenario described in 2.3 (`Base case`); ii) a scenario in which the temperature anomaly increases to 4 degrees at the end of the forecast horizon (`High increase`); iii) a scenario in which the volatility of the temperature anomaly process in (2) matches the historical volatility (`High uncertainty`). We decompose each of the variance ratios into three parts according to Equation (4), using 250,000 draws from the posterior distribution of the VAR parameters. The investment horizon ranges from 1 to 100 quarters.

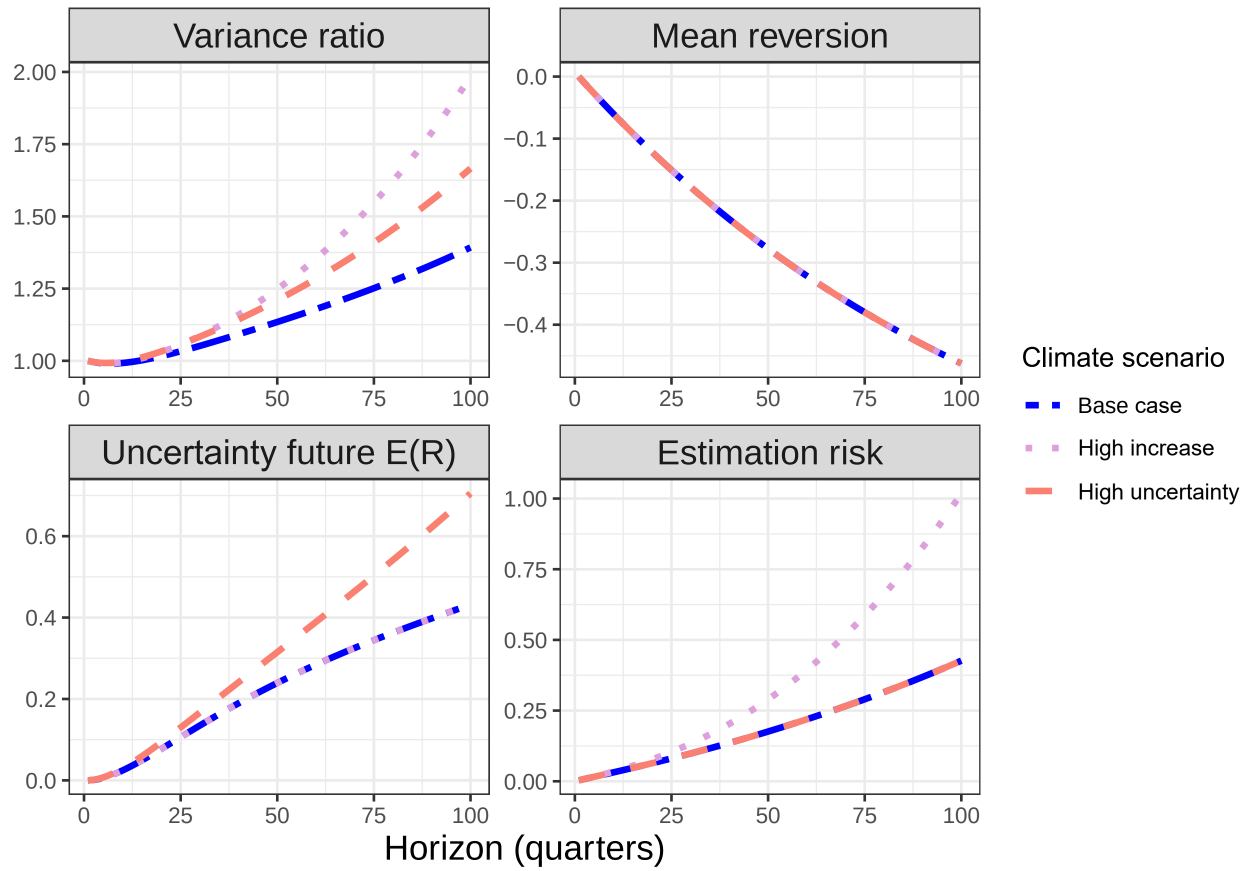

Figure 8: Predictive returns and risk premia in alternative climate scenarios

This figure shows forecasts of average risk-free returns, market returns, and market risk premia per quarter (in %), for investment horizons ranging from 1 to 100 quarters. For each variable, we plot three different forecasts implied by the VAR model in Equation (1). All VARs are estimated using the historical data and an informative prior derived from the temperature long-run risks model. We consider three different climate scenarios: i) the default scenario described in 2.3 (`Base case`); ii) a scenario in which the temperature anomaly increases to 4 degrees at the end of the forecast horizon (`High increase`); iii) a scenario in which the volatility of the temperature anomaly process in (2) matches the historical volatility (`High uncertainty`). The predictive returns are computed based on 250,000 draws from the predictive distribution corresponding to each model. Specifically, we simulate 250,000 future return paths based on 250,000 draws from the posterior distribution of the VAR parameters. For each of these return paths, we compute the simple average of the quarterly returns over each horizon. The forecasts shown in the figure are the means of the average returns per quarter across the simulated return paths.

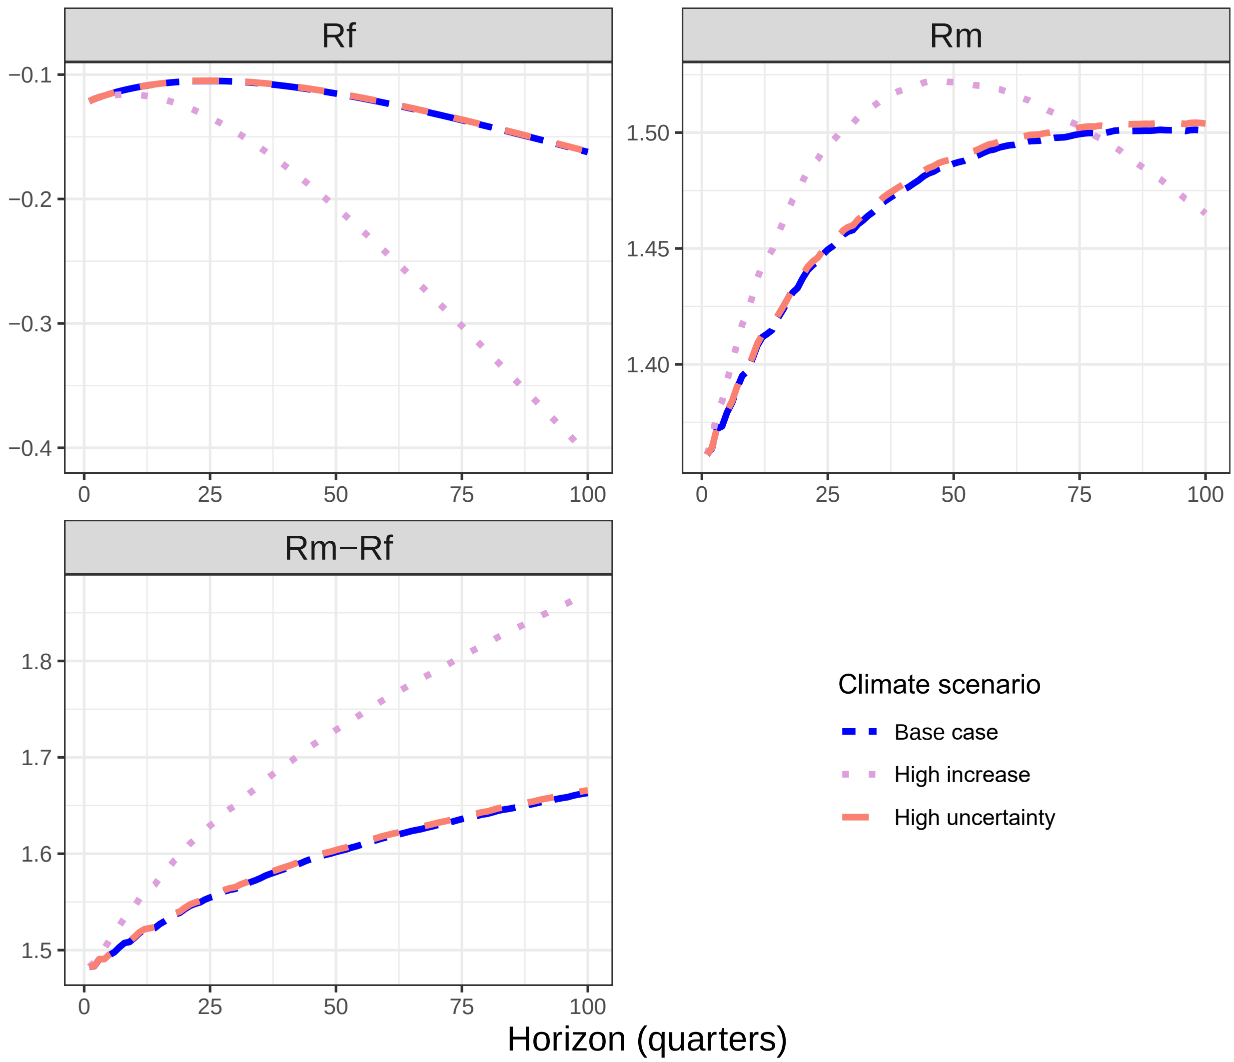

Figure 9: Optimal portfolio weights in alternative climate scenarios

This figure shows optimal portfolio weights for an investor who allocates wealth between the market portfolio of stocks and the risk-free asset, with an investment horizon ranging from 1 to 100 quarters. The plots show the optimal allocation to the stock portfolio for a long-only buy-and-hold investor who has power utility with risk aversion $A = 5$. The optimal weights are computed following the approach explained in Section 2.5, using the predictive return distribution implied by the VAR model in Equation (1). All VARs are estimated using the historical data and an informative prior derived from the temperature long-run risks model. We consider three different climate scenarios: i) the default scenario described in 2.3 (`Base case`); ii) a scenario in which the temperature anomaly increases to 4 degrees at the end of the forecast horizon (`High increase`); iii) a scenario in which the volatility of the temperature anomaly process in (2) matches the historical volatility (`High uncertainty`). The optimal allocation accounts for parameter uncertainty and is based on 250,000 samples from the predictive distribution with the parameters of the predictive VAR model drawn from their posterior distribution.

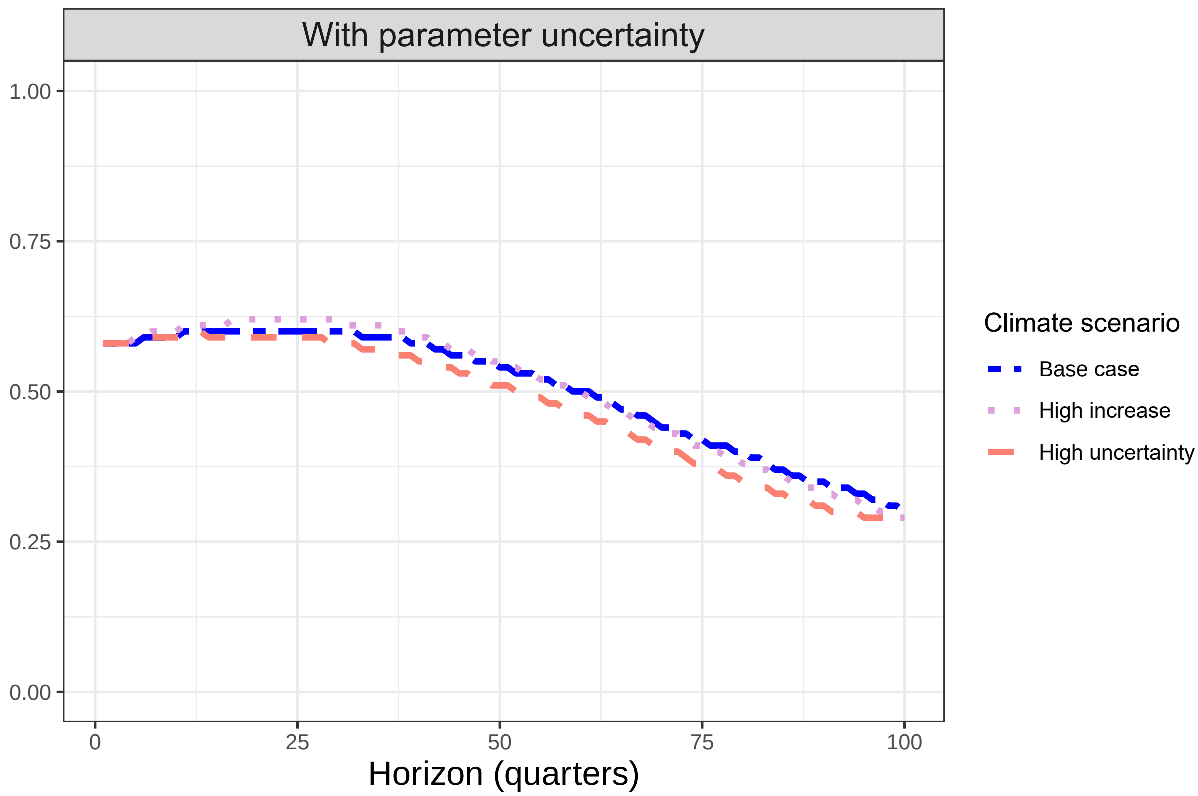

Table 1: Overview of investor types

This table summarizes the investor types considered in the paper. Investors differ in terms of their prior beliefs about the parameters of the predictive VAR model in Equation (1), the use of historical data in the VAR estimation, the predictor variables included in the VAR, and the values of the predictors used for making out-of-sample forecasts. The *agnostic* investor has no prior views about the effect of temperature change on future asset returns and assigns full weight to the return dynamics implied by historical data. The *dogmatic* investor elicits prior beliefs about the impact of climate change on return dynamics from the temperature long-run risks (LRR-T) model outlined in Section 3.1 and assigns full weight to this model-based prior. The *Bayesian LRR-T* investor derives prior beliefs from the LRR-T model and updates these beliefs based on observed data. The *Bayesian LRR* investor derives prior beliefs from the LRR model of Bansal and Yaron (2004) and updates these beliefs based on observed data. Both Bayesian investor types assign equal weight to the prior information and the sample information. The *agnostic* investor constructs forecasts based on the mean value of the price-dividend ratio (p/d) and risk-free rate (r_f) observed over the last five years of the historical sample. The *dogmatic* investor forecasts returns based on the mean value of p/d and r_f observed over the last five years of the data simulated from the LRR-T model. We simulate 250,000 time series of 292 quarters to match the length of our historical sample (1947Q1-2019Q4) and calculate the averages of the five-year means of the predictors across the simulations. The *Bayesian LRR(-T)* investor forms forecasts based on the simple average of the most recent five-year mean of p/d and r_f implied by the LRR(-T) model and their most recent five-year sample mean. All investors except for the *Bayesian LRR* investor include the temperature anomaly (T) as an additional predictor variable in the VAR. Investors generate return forecasts based on the average temperature anomaly observed over the last five years in the sample and use the exogenous process in Equation (2) to project future temperature changes.

Investor type	Model	Prior	Data	Predictors	Predictor value (p/d and r_f)
Agnostic	Data + uninf. prior	Uninformative	Historical	p/d, r_f, T	5-Year sample mean
Dogmatic	LRR-T prior	LRR-T	None	p/d, r_f, T	5-Year LRR-T mean
Bayesian LRR-T	Data + LRR-T prior	LRR-T	Historical	p/d, r_f, T	(1/2)×5-Year LRR-T+(1/2)×5-Year sample
Bayesian LRR	Data + LRR prior	LRR	Historical	p/d, r_f	(1/2)×5-Year LRR+(1/2)×5-Year sample

Table 2: Temperature long-run risks model parameters

This table presents the parameters for the temperature long-run risks (LRR-T) model:

$$\begin{aligned}
\Delta c_{t+1} &= \mu_c + \sigma_t \eta_{t+1} + X_{t+1} \\
\Delta d_{t+1,i} &= \mu_d + \pi_d \sigma_t \eta_{t+1} + \phi_i X_{t+1} + \varphi_d \sigma_t u_{t+1} \\
\sigma_{t+1}^2 &= \bar{\sigma}^2 + \nu(\sigma_t^2 - \bar{\sigma}^2) + \sigma_w w_{t+1} \\
X_{t+1} &= \rho X_t + d\Delta N_{t+1}, \\
\Delta N_{t+1} &\sim \text{Poisson}(\lambda_t = \Delta t(\lambda_0 + \lambda_1 T_t)) \\
T_{t+1} &= \chi \varepsilon_{t+1} \\
\varepsilon_{t+1} &= \nu_\varepsilon \varepsilon_t + \mu_\varepsilon + \Theta(\mu_c + \sigma_t \eta_{t+1}) + \sigma_\zeta \zeta_{t+1},
\end{aligned}$$

where X is the adverse impact of temperature-driven disasters on consumption growth Δc and dividend growth Δd, N is the disaster process, T is the temperature anomaly, and ϵ is the carbon concentration. All parameters are in monthly terms. The preference parameters δ, γ, and ψ denote the investor's risk aversion, elasticity of intertemporal substitution (EIS), and time discount factor, respectively.

Preferences	δ	γ	ψ					
	0.998	5	1.5					
Consumption	μ_c	ρ	d	$\bar{\sigma}$	ν	σ_w		
	0.0052	0.99	-0.004166	0.0072	0.999	0.0000028		
Dividends	μ_d	π_d	ϕ_i	φ_d				
	0.0062	2.0	1.1	5.0				
Temperature	ν_e	ε_0	μ_ε	Θ	σ_ζ	χ	λ_0	λ_1
	0.9971	-1	0.0095	1	0.5	0.2	0.075	0.075

Table 3: Data moments and long-run risks moments

This table reports the first and second moments of market returns, risk-free rates, price-dividend ratios, consumption and dividend growth, and the temperature anomaly. Returns and growth rates are in percentages and all moments except for the temperature anomaly are in logs. The first column reports the historical moments from monthly data time-aggregated to quarterly values from 1947Q1 to 2019Q4. The columns on the right show the population moments implied by the LRR and LRR-T models. For each model, the moment reported is the median from 250,000 simulations, with each simulation matching the length of the 1947Q1-2019Q4 sample.

Moment	Data	LRR	LRR-T
$E(r_m)$	1.81	1.77	1.71
$\sigma(r_m)$	7.25	8.00	8.17
$E(r_f)$	0.13	0.21	0.17
$\sigma(r_f)$	0.47	0.27	0.62
$E(p-d)$	4.89	4.45	4.42
$\sigma(p-d)$	0.44	0.16	0.23
$E(\Delta c)$	0.47	0.45	0.45
$\sigma(\Delta c)$	0.49	1.14	1.43
$E(\Delta d)$	0.66	0.63	0.62
$\sigma(\Delta d)$	1.94	5.83	5.90
$E(T)$	0.33		0.65
$\sigma(T)$	0.83		0.83

Table 4: Predictive VAR estimates

This table reports the posterior means of the parameters for the first-order VAR model in Equation (1). The predictor variables are the log price-dividend ratio ($p-d$), the log ex-ante real risk-free rate (r_f), and the temperature anomaly (T) measured in degrees Celsius. The log real return (including dividends) on the S&P 500 (r_m) is modeled as a function of these predictor variables. Panel A shows results for the VAR estimated using the historical sample data and an uninformative prior. Panel B reports the VAR parameters implied by the temperature long-run risks (LRR-T) prior. Panel C shows the posterior means for the VAR estimated using the sample data and the LRR-T prior. Panel D reports the parameters for the VAR estimated using the sample data and the prior derived from the standard long-run risks (LRR) model. The t-statistics in parentheses are based on the posterior standard deviations. The sample period is 1947Q1-2019Q4.

	Panel A: Data + uninformative prior			
	Intercept	$p_t - d_t$	$r_{f,t}$	T_t
$r_{m,t+1}$	14.28	-2.63	0.81	0.89
	(2.76)	(-2.46)	(0.87)	(1.58)
$p_{t+1} - d_{t+1}$	0.10	0.98	1.37	0.01
	(1.79)	(89.10)	(1.44)	(1.22)
$r_{f,t+1}$	0.19	-0.03	0.89	0.01
	(1.54)	(-1.36)	(40.83)	(0.50)
T_{t+1}	-2.48	0.57	-11.36	0.19
	(-4.65)	(5.13)	(-1.18)	(3.23)

	Panel B: LRR-T prior			
	Intercept	$p_t - d_t$	$r_{f,t}$	T_t
$r_{m,t+1}$	32.05	-6.87	2.74	-0.61
	(1.18)	(-1.14)	(2.29)	(-0.49)
$p_{t+1} - d_{t+1}$	0.93	0.79	3.31	-0.04
	(4.86)	(18.78)	(3.93)	(-4.09)
$r_{f,t+1}$	0.36	-0.08	0.99	-0.03
	(0.90)	(-0.88)	(56.52)	(-1.48)
T_{t+1}	0.42	-0.09	0.97	0.98
	(0.92)	(-0.90)	(0.47)	(46.25)

	Panel C: Data + LRR-T prior			
	Intercept	$p_t - d_t$	$r_{f,t}$	T_t
$r_{m,t+1}$	7.11	-1.23	1.56	0.33
	(1.90)	(-1.54)	(2.75)	(0.90)
$p_{t+1} - d_{t+1}$	0.05	0.99	0.76	0.00
	(1.58)	(141.86)	(1.53)	(0.71)
$r_{f,t+1}$	-0.01	0.01	0.95	-0.02
	(-0.16)	(0.38)	(87.39)	(-2.34)
T_{t+1}	0.74	-0.12	-20.65	0.70
	(2.44)	(-1.87)	(-4.43)	(23.90)

	Panel D: Data + LRR prior		
	Intercept	$p_t - d_t$	$r_{f,t}$
$r_{m,t+1}$	8.22	-1.43	1.48
	(2.26)	(-1.82)	(1.84)
$p_{t+1} - d_{t+1}$	0.05	0.99	1.22
	(1.75)	(146.40)	(1.75)
$r_{f,t+1}$	0.04	0.00	0.91
	(0.57)	(-0.30)	(64.61)

B Online Appendix

B.1 Numerical procedure for optimal asset allocation

In this section we explain the numerical procedure used to solve the optimal asset allocation for the long-only buy-and-hold investor. We first discuss how we sample return paths from the predictive distribution and then explain how we calculate the optimal buy-and-hold portfolios.

Sampling from predictive distribution

We draw $N^* = 250,000$ sample paths of length $K = 100$ quarters from the predictive distribution of asset returns and state variables by repeating the following two steps N times:

1. For each k in $1, \ldots, K$, sample the asset returns and state variables in period k conditional on each draw of the VAR parameters (a, B, Σ) from the posterior distribution and the values of the state variables in period $k - 1$.[26] For the initial period $k = 1$, condition on the average value of the state variables $p - d$ and r_f over the last five years in the sample (for the VAR model `Data+uninf. prior`), the population average of these state variables (for the model `LRR-T prior`), or the mean of the five-year average and the population average (for the model `Data+LRR-T prior`). For all models, the initial value for the state variable T is the average temperature anomaly observed over the last five years in the sample. Future temperature anomaly paths are simulated based on the exogenous process in Equation (2).

2. Re-sample the asset returns and state variables in period k if we draw a quarterly return on the risk-free asset below -10%. This step is needed because the simple return drawn from the predictive distribution can get arbitrarily close to -100% for each asset. In that case, the expected utility of all possible portfolios is minus infinity and the maximization problem in Equation (7) is no longer well defined. Imposing the restriction that the quarterly risk-free rate is at least -10% ensures that expected utility is finite for at least some portfolios, because it implies that the investor's wealth will never go to zero if she invests a positive amount in the risk-free asset. Because -10% is very far below the sample mean of the risk-free rate, this lower bound is almost never hit in our simulations and re-sampling is almost never required.

Calculation of optimal buy-and-hold portfolio

After generating the N^* return paths, we compute the optimal buy-and-hold portfolio as follows:

1. Construct a grid of portfolio weights. We invest long only, i.e., weights are between zero and one, and use weight steps of 1% for our grid search.

2. Pick one set of portfolio weights from the grid and calculate the realized utility for all N^* simulated return paths.

3. Approximate expected utility by computing the mean of these N realized utilities.

4. Repeat steps 2 and 3 for all portfolio weights on the grid. For each horizon $k = 1, \ldots, K$, choose the weights that maximize expected utility.

[26] For the analysis that does not incorporate parameter uncertainty (top panel in Figure 6), returns and state variables are simulated conditional on the posterior mean of the VAR parameters.

B.2 Additional results

In this section we provide additional empirical results omitted from the paper for the sake of brevity.

Table B1: Population estimates for VAR parameters implied by LRR-T model

This table shows the model coefficients of the VAR from Equation (1) estimated on a very large set of 250,000 simulations from the LRR-T model. The model is estimated on a quarterly time interval and each simulation matches the 1947Q1-2019Q4 sample period. Returns are in log percentages, the price-dividend ratio is in logs and the temperature anomaly is in degrees Celsius.

	Intercept	$p_t - d_t$	$r_{f,t}$	T_t
$r_{m,t+1}$	32.07	-6.87	273.55	-0.61
$p_{t+1} - d_{t+1}$	0.93	0.79	3.32	-0.04
$r_{f,t+1}$	0.36	-0.08	99.04	-0.03
T_{t+1}	0.43	-0.09	1.01	0.98

Figure B1: Relative variance of risk-free asset and market portfolio

This figure shows the term structure of predictive variances of returns on the risk-free asset (including roll-over risk) divided by the predictive variances of returns on the market portfolio, for horizons ranging from 1 to 100 quarters. The predictive variances are constructed based on the return dynamics implied by the VAR model in Equation (1). VARs are estimated using the historical sample data and an uninformative prior (`Data+uninf. prior`), the historical data and an informative prior derived from the temperature long-run risks model (`Data + LRR-T prior`), the informative LRR-T prior alone (`LRR-T prior`), or the historical data and an informative prior derived from the standard long-run risks model (`Data + LRR prior`).

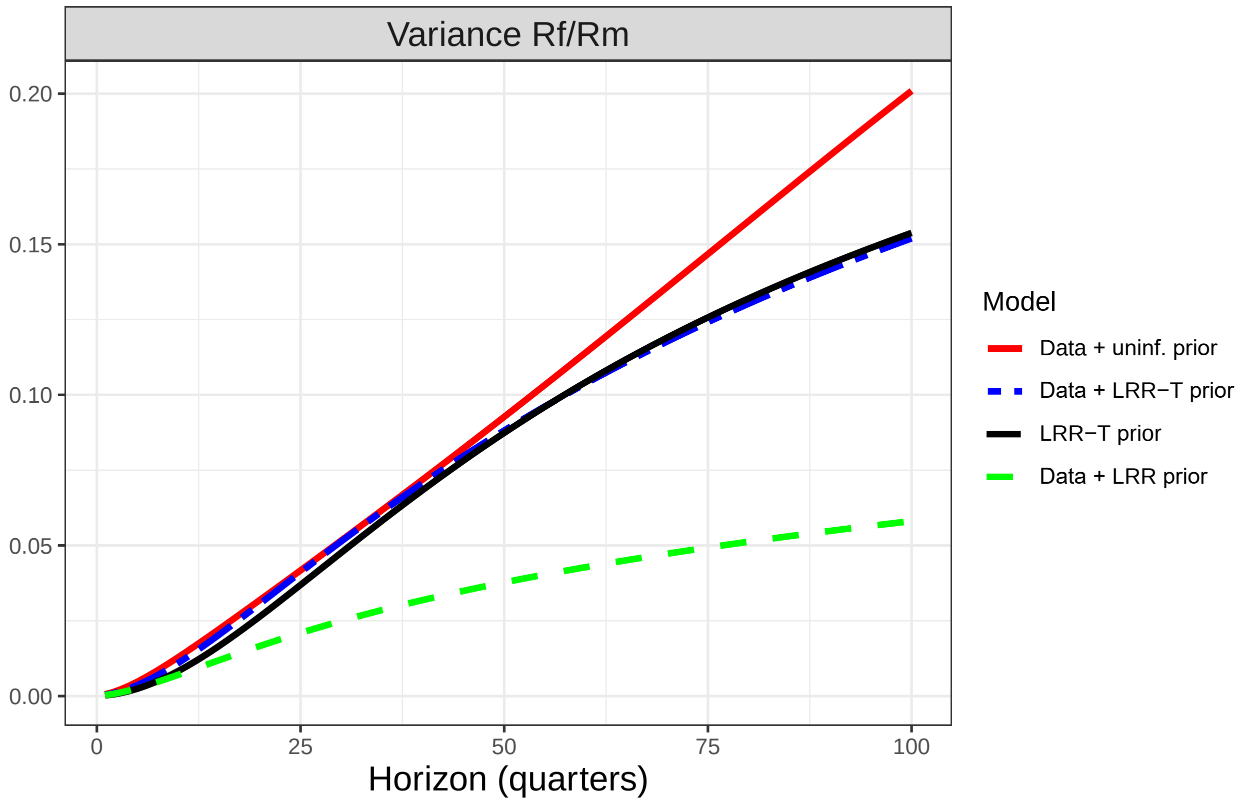

Figure B2: Relative variance in alternative climate scenarios

This figure shows the term structure of predictive variances of returns on the risk-free asset (including roll-over risk) divided by the predictive variances of returns on the market portfolio, for horizons ranging from 1 to 100 quarters. The predictive variances are constructed based on the return dynamics implied by the VAR model in Equation (1). All VARs are estimated using the historical data and an informative prior derived from the temperature long-run risks model. We consider three different climate scenarios: i) the default scenario described in 2.3 (`Base case`); ii) a scenario in which the temperature anomaly increases to 4 degrees at the end of the forecast horizon (`High increase`); iii) a scenario in which the volatility of the temperature anomaly process in (2) matches the historical volatility (`High uncertainty`).

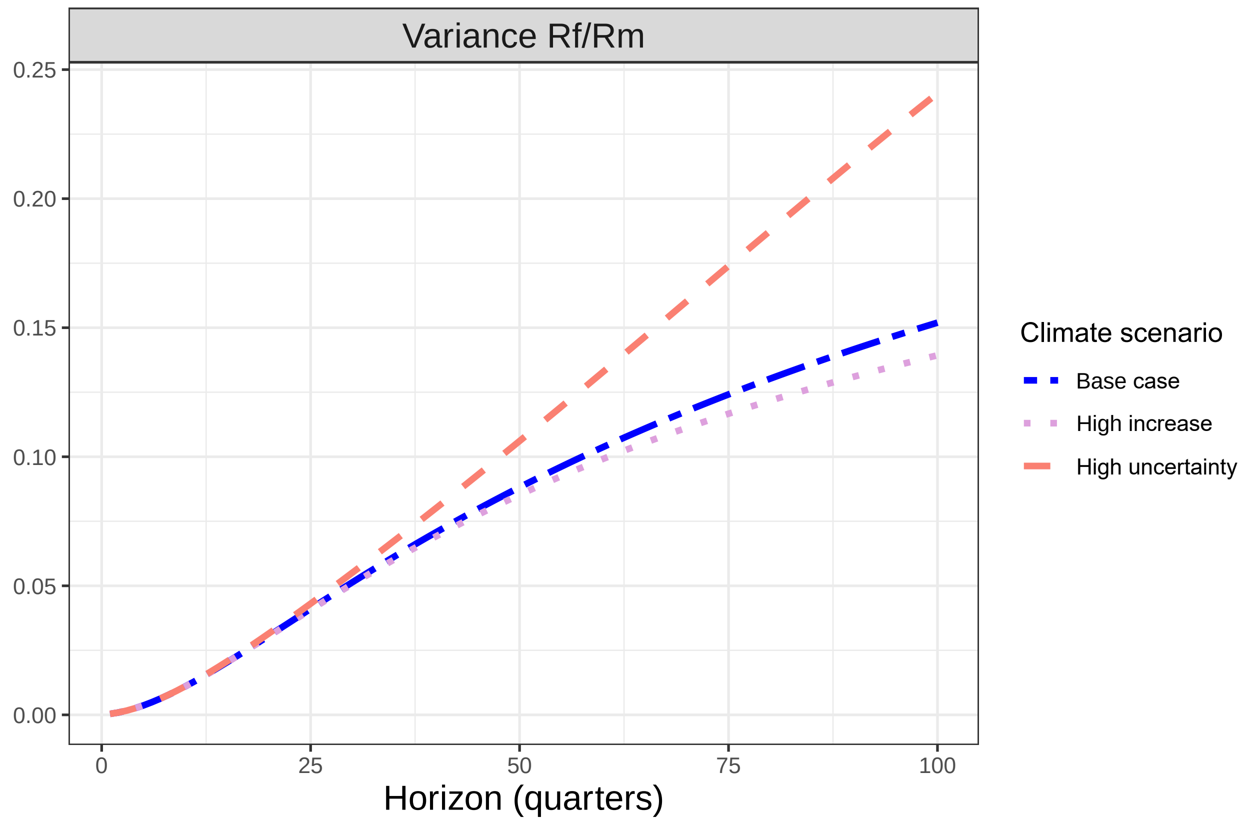

Figure B3: Predictive correlations in alternative climate scenarios

This figure shows predictive correlations between the returns on the risk-free asset and the returns on the market portfolio, for horizons ranging from 1 to 100 quarters. The predictive correlations are based on the return dynamics implied by the VAR in Equation (1). All VARs are estimated using the historical data and an informative prior derived from the temperature long-run risks model. We consider three different climate scenarios: i) the default scenario described in 2.3 (`Base case`); ii) a scenario in which the temperature anomaly increases to 4 degrees at the end of the forecast horizon (`High increase`); iii) a scenario in which the volatility of the temperature anomaly process in (2) matches the historical volatility (`High uncertainty`). The predictive correlations are computed based on 250,000 draws from the predictive distribution corresponding to each model. We simulate 250,000 future return paths based on 250,000 draws from the posterior distribution of the VAR parameters. For each of these paths, we compute the correlation between the quarterly returns on the risk-free asset and the market portfolio over each horizon. The correlations in the figure are the averages of the correlations across the simulations.