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

Mathijs Cosemans
Xander Hut
Mathijs van Dijk

Abstract

We propose a novel approach for measuring the impact of climate change on long-horizon equity risk and optimal portfolio choice. Our method combines historical data about the impact of climate change on return dynamics with prior beliefs elicited from the temperature long-run risk (LRR-T) model of Bansal, Kiku, and Ochoa (2019). Our Bayesian framework incorporates this prior information to obtain more precise estimates of long-term climate risks than existing methods that solely rely on historical data. We document two key findings. First, disasters lead to persistently lower risk-free rates, whereas the negative impact on market returns is transitory. As a result, the expected equity risk premium increases after disasters. Second, an investor with LRR-T beliefs perceives stock markets to be riskier over longer horizons because disasters induced by climate change reduce mean reversion in returns. Combining these results, we find that for investors with a horizon longer than 25 quarters, the optimal allocation to equity decreases when accounting for climate change because the increase in perceived market risk outweighs the increase in the market risk premium. In contrast, for short-term investors the allocation to equity increases relative to investors with uninformative prior beliefs about the effect of climate change on future returns.

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

JEL classification: C11, G11, G12

*All at Rotterdam School of Management, Erasmus University. Our e-mail adresses are mcosemans@rsm.nl, hut@rsm.nl, and madijk@rsm.nl, respectively. For helpful comments and suggestions, we thank Doron Avramov, Jules Tinang, Dirk Schoenmaker, Marta Szymanowska, and participants of the 2019 GRASFI PhD workshop, 2020 Erasmus PhD seminar, and 2020 Netspar Taskforce meeting. We thank Netspar and INQUIRE Europe for financial support.

1 Introduction

A growing body of research examines how climate change affects socio-economic outcomes, ranging from agricultural production and labor productivity to health and civil conflict.[1] Colacito, Hoffmann, and Phan (2019) show that rising temperatures can reduce U.S. economic growth by up to one-third over the next century. Climate change also affects financial markets. Hong, Li, and Xu (2019) find lower earnings growth and stock returns for food companies in countries that experience prolonged periods of drought. Bansal, Kiku, and Ochoa (2019) argue that temperature change is a source of long-run economic risk that carries a significantly positive premium in equity markets.

Investors are increasingly aware that climate change can pose significant investment risks that should be integrated in their portfolio strategies. 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 conducted by Krueger, Sautner, and Starks (2020)).[2]

Existing work studies the impact of climate change on *current* asset prices. However, given the long-term nature of climate change, long-horizon investors such as pension funds are more concerned about its effects over longer holding periods. These long-term effects are difficult to measure due to two key problems:data unavailability and peso problems. Estimating the climate risk exposure of assets is challenging because data availability for climate risk factors is limited and because currently available data samples may not include sufficient realizations of severe climate change effects. Robert Stambaugh summarizes the uncertainty about the long-horizon impact of climate change on asset prices by noting that when we "expand the horizon to the next several decades, the possible effects of global warming range from negligible to catastrophic."[3]

We address these issues by proposing a novel approach for measuring the impact of climate

[1] Dell, Jones, and Olken (2014) provide an overview of work on the economic impact of climate change.

[2] 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."

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

change on long-horizon equity risk and return that supplements historical data with prior information derived from economic theory. In particular, we set up a Vector Autoregression model (VAR) to analyze the long-run dynamics of equity returns and include temperature change as a predictor in the VAR model to proxy for climate risk. Since the frequency and economic impact of past climate-induced disasters may not be representative of their future impact in a scenario of prolonged climate change, historical data may not be very informative about the impact of climate change on long-horizon returns. We therefore estimate the VAR parameters using a Bayesian approach that combines historical data with theoretical prior beliefs about return dynamics.[4]

We elicit these beliefs from the temperature long-run risk (LRR-T) model of Bansal, Kiku, and Ochoa (2019), in which rising temperatures influence asset prices by increasing the likelihood of future climate disasters that lower economic growth. Because of investors' concerns about the implications of temperature increases for future growth, climate risk can be reflected in current asset prices even though the impact of climate change in historical data is limited. By imposing a structure on the relation between temperature change and returns, the LRR-T model provides prior information about the impact of climate change. Incorporating this information yields more precise estimates of long-term climate risks than existing methods that solely rely on past data.

Our main results are as follows. First, we show that an investor with LRR-T beliefs perceives stock markets to be riskier over longer horizons because disasters induced by climate change reduce mean reversion in returns. Intuitively, because in the LRR-T model temperature change is persistent, climate-induced disasters tend to occur in clusters during high-temperature periods, i.e., disasters are more likely to be followed by further negative shocks to consumption and dividends than by positive shocks. As a result, the long-horizon predictive variance of returns is higher for investors who form beliefs based on the LRR-T model than for investors with an uninformative prior.

Second, we find that the investor with LRR-T beliefs expects the market risk premium to increase after a climate disaster occurs. In particular, whereas disasters cause a persistent negative shock to risk-free rates, the negative impact on expected market returns is transitory because prices

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

rapidly account for lower expected future dividends after a disaster strikes.

Third, we quantify the impact of climate change on portfolio choice for investors who can choose between allocating their wealth to the market portfolio and the risk-free asset. We show that for long-term investors with a horizon beyond 25 quarters, the optimal allocation to equity decreases when accounting for climate change because the increase in perceived market risk outweighs the increase in the market risk premium. In contrast, for short-term investors the increase in the expected market risk premium is sufficiently large to increase the optimal allocation to equity. The impact of climate change on the predictive return variance increases with the investment horizon because of the positive trend in temperature that increases the occurrence of clustered climate disasters. Assuming more severe climate change scenarios amplifies these effects.

Finally, we document larger (smaller) expected returns for investors in a portfolio that is (non-) vulnerable to temperature change, as compared to each other and the market. This is visible in both the historical data and our theoretical model. Investors turn out to be fairly (perhaps even more than fairly) compensated for additional risk in vulnerable portfolios. Therefore, long-only investors who can combine the risk-free asset with vulnerable and non-vulnerable portfolios tend to invest in vulnerable portfolios in a similar fashion as they would otherwise do in the market, and leave non-vulnerable portfolios out of their optimal allocation entirely.

Our work contributes to two strands of literature. First, we add to the growing body of literature that studies how climate change affects asset prices. Existing work shows that carbon emissions (Ilhan, Sautner, and Vilkov (2020) and Bolton and Kacperczyk (2019)), temperature increases (Balvers, Du, and Zhao (2017), Barnett (2020), and Kumar, Xin, and Zhang (2019)), and drought trends (Hong, Li, and Xu (2019)) impact asset prices. Whereas these papers focus on the short-term consequences of climate change, we study the implications for long-horizon investors.

Second, we extend the literature on long-horizon equity risk. Climate change adds to known determinants of long-run risk such as mean reversion (Barberis (2000)) and uncertainty about future expected returns (Pastor and Stambaugh (2012)). Our paper builds on the approach proposed by Avramov, Cederburg, and Lucivjanska (2018) to formulate prior beliefs about return dynamics based on economic theory. We contribute to their work by incorporating climate change as a new source of long-horizon risk and by studying its impact on optimal portfolio choice for long-term investors. We show that climate change has a significant impact on the predictive distribution of equity returns, increasing both the market risk premium and the long-horizon return variance.

The paper proceeds as follows. In Section 2 we introduce our approach for modeling long-horizon equity returns. In Section 3 we discuss how we incorporating prior beliefs implied by the LRR-T model in our model. Section 4 presents the predictive regression estimates, the long-horizon implications from these estimates on variance ratios and correlations, and the impact of climate change on optimal portfolio choice of long- and short-term investors. Section 5 concludes.

2 Long-Horizon Portfolio Choice

Section 2.1 introduces the predictive VAR model that we use to characterize the long-horizon dynamics of equity returns. Section 2.2 explains the Bayesian approach used to estimate the VAR parameters model and the methodology used to construct optimal portfolios.

2.1 Long-horizon forecasts

We estimate a Vector Autoregression (VAR) model on quarterly data from 1947Q1 to 2019Q4 to analyze the long-horizon dynamics of equity risk and returns, as is common in the strategic asset allocation literature (e.g., Barberis (2000) and Pastor and Stambaugh (2012)). To capture climate risk, we augment the model with a return predictor related to climate change through the inclusion of the long-term trend in temperature change in the model. Specifically, the VAR is given as

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

where $r_{m,t+1}$ is the log market return, $p_{t+1} - d_{t+1}$ is the log price-dividend ratio of the market portfolio, $r_{f,t+1}$ is the ex ante risk-free yield to maturity and ΔT_{t+1} is the long-term temperature innovation, proxied by the first difference of the five-year moving average of the temperature anomaly.[5] The set of VAR parameters (a, B, Σ) can be split into parts as

$$a = \begin{bmatrix} a_r \ a_x' \end{bmatrix}', \quad a_x = \begin{bmatrix} a_{pd} \ a_{rf} \ a_T \end{bmatrix}',$$
$$B = \begin{bmatrix} b_r \\ B_x \end{bmatrix}, \quad b_r = \begin{bmatrix} b_{r,pd} \ b_{r,rf} \ b_{r,T} \end{bmatrix}, \quad B_x = \begin{bmatrix} b_{pd,pd} \ b_{pd,rf} \ b_{pd,T} \\ b_{rf,pd} \ b_{rf,rf} \ b_{rf,T} \\ b_{T,pd} \ b_{T,rf} \ b_{T,T} \end{bmatrix}, \tag{2}$$

and

$$\Sigma = \begin{bmatrix} \sigma_r^2 \ \Sigma_{xr}' \\ \Sigma_{xr} \ \Sigma_x \end{bmatrix}, \quad \Sigma_{xr} = \begin{bmatrix} \sigma_{rpd} \\ \sigma_{rrf} \\ \sigma_{rT} \end{bmatrix}, \quad \Sigma_x = \begin{bmatrix} \sigma_{pd}^2 \ \sigma_{pdrf} \ \sigma_{pdT} \\ \sigma_{rfpd} \ \sigma_{rf}^2 \ \sigma_{rfT}' \\ \sigma_{Tpd} \ \sigma_{Trf} \ \sigma_T^2 \end{bmatrix}. \tag{3}$$

We use the predictive VAR from Equation (1) to make long-horizon forecasts for the market return[6] (cumulative log returns $r_{m,t,t+k} = r_{m,t+1} + \cdots + r_{m,t+k}$) and the state variables in the model conditional on the estimated set of VAR parameters (a, B, Σ) and the last observed value of the state variables x_t, given as

$$\mathbb{E}[r_{m,t,t+k} \mid a, B, \Sigma, x_t] = ka_r + b_r(I - B_x)^{-1} \Big((kI - (I - B_x)^{-1}(I - B_x^k))a_x + (I - B_x^k)x_t \Big), \tag{4}$$

[5] Since we assume that the Campbell and Shiller (1988) decomposition holds in our holds in our model, we do not include dividend growth in this VAR specification. The VAR coefficients related to dividend growth can be computed directly from the coefficient estimates for the market return and the price-dividend ratio, as is done in Avramov, Cederburg, and Lucivjanska (2018). However, in our analysis, we do not study the VAR dynamics of dividend growth.

[6] We use the same predictive regression for the returns of portfolios that are vulnerable and non-vulnerable to temperature change in additional analyses.

where $x_t = \left[p_t - d_t \; r_{f,t} \; \Delta T_t\right]'$. Avramov, Cederburg, and Lucivjanska (2018) show that the conditional variance of the long-horizon return forecast equals

$$\begin{aligned}\mathrm{Var}(r_{m,t,t+k} \mid a, B, \Sigma) = & \underbrace{k\sigma_r^2}_{\text{i.i.d. uncertainty}} + \underbrace{\sum_{i=1}^{k-1} 2b_r(1-B_x)^{-1}(I-B_x^i)\Sigma_{xr}}_{\text{mean reversion}} \\ & + \underbrace{\sum_{i=1}^{k-1}(b_r(1-B_x)^{-1}(I-B_x^i))\Sigma_x(b_r(1-B_x)^{-1}(I-B_x^i))'}_{\text{uncertainty about future expected returns}}\end{aligned} \tag{5}$$

The i.i.d. uncertainty stays constant per-period. Mean reversion causes long-horizon variance to decrease, which is why Siegel (2008) argues that equities are less risky for long-horizon investors. The third component is uncertainty about future expected returns. The idea of this component is that predictors in the VAR model are time-varying, and uncertainty about the level of future predictors will increase the predictive variance of the dependent variable as well. We study the changes to the variance over the horizon by analyzing the variance ratio for market returns $VR_k = \frac{\mathrm{Var}(r_{m,t,t+k}|D_t)}{k\mathrm{Var}(r_{m,t,t+1}|D_t)}$.

2.2 Bayesian estimation

We estimate the parameters of the VAR using a Bayesian approach that complements historical data with prior information derived from economic theory. Specifically, we impose a structure on the relation between the temperature trend and economic and dividend growth by specifying prior beliefs based on the LRR-T model from Bansal, Kiku, and Ochoa (2019). We derive the implications of the LRR-T model for return dynamics by applying the framework of Avramov, Cederburg, and Lucivjanska (2018). The model and the construction of the model-based VAR parameters are discussed in Sections 3.1 to 3.4.

Our Bayesian approach to incorporating climate risk into portfolio choice allows for the accommodation of different investor beliefs about the pricing of climate risk in equity and risk-free bond markets. We quantify the effect of these prior beliefs on the perceived riskiness of stock markets over different holding periods and on optimal portfolio choice for long-term investors. We consider three prior investor beliefs. First, the agnostic investor has no prior views about the impact of climate change on equity return dynamics and, therefore, lets the data speak. This is equivalent to a frequentist approach for estimating the VAR model that ignores prior information and gives full weight to the return dynamics implied by historical data.[7] Second, the climate risk believer is convinced that climate change affects equity returns in the way implied by the temperature long-run risk model. This dogmatic investor therefore assigns full weight to the model-based prior beliefs. Historical data are not taken into account by this investor. Third, the Bayesian investor has faith in her prior beliefs derived from theory, but is aware that these beliefs may be inaccurate and updates them based on observed data. She assigns equal weights to the prior and the data. The posterior specification for the VAR model from this investor is discussed in Appendix A.3. Our approach easily extends to different beliefs. For example, one could analyze a climate risk denier that fixes the coefficient from market returns on temperature change in the VAR model to a value of zero.

2.2.1 Sampling and portfolio choice

We draw posterior observations from the predictive VAR from Equation (1) from the posterior VAR distribution in Appendix A.3 with direct sampling until we have 5,000 accepted draws. In the VAR model based on the data with uninformative prior (agnostic investor) and the model-based prior (climate risk believer), all draws are accepted. In the VAR model that combines historical data with the model-based prior (Bayesian investor), we apply the non-negativity constraint proposed in Campbell and Thompson (2008) by only accepting draws with positive equity premiums. The methodology to apply this constraint in a Bayesian setting is also discussed in Pettenuzzo, Timmermann, and Valkanov (2014) and Avramov, Cederburg, and Lucivjanska (2018). The posterior mean of these draws is the set of VAR parameters on which we base our portfolio choice.

We compute the optimal allocation to the risky portfolios and the risk-free asset for a buy-and-hold long-only investor. Optimal portfolios for various investment horizons k are constructed by maximizing expected power utility with respect to the predictive distribution of future stock

[7] Technically, we estimate the posterior VAR for this investor on historical data with an uninformative multivariate Jeffreys prior.

returns.[8]. Formally, the investor maximizes expected utility at time $t + k$ conditional on the estimated VAR parameters (a, B, Σ)

$$\max_{w_{t,t+k}} \quad E_t[U(W_{t+k})|(a, B, \Sigma)], \tag{6}$$

where end-of-period wealth of our buy-and-hold investor is $W_{t+k} = w_{t,t+k}R_{t,t+k}$, with $w_{t,t+k}$ the vector of optimal portfolio weights to the available assets and $R_{t,t+k}$ the cumulative return on these assets. Power utility is given as $U(W_{t+k}) = \frac{W_{t+k}^{1-A}}{1-A}$, in which $A = 5$ is the risk aversion level. We compare two investors, one who invests in the market portfolio and the risk-free asset, and one who combines portfolios that are vulnerable or non-vulnerable to temperature innovations with the risk-free asset. We use the numerical methodology described in Appendix A.4 to solve the optimal weights for investment horizon k from 1 to 100 quarters. In a nutshell, we use a grid search over possible investment weights to compute the weights that give the highest average utility over 250,000 draws of future returns forecasted from the estimated VAR parameters until horizon k. As Hoevenaars, Molenaar, Schotman, and Steenkamp (2014), we explicitly account for reinvestment risk in the risk-free asset and, therefore, allow for correlated returns between the risk-free asset and risky assets.

3 Incorporating Theoretical Prior Information

Section 3.1 outlines the LRR-T model that is used to form prior beliefs about the impact of climate change on financial market variables. Section 3.2 describes the data including the construction of vulnerable and non-vulnerable portfolios and Section 3.3 discusses the calibration of the LRR-T model. Section 3.4 illustrates the implications from the LRR-T model and the simulation approach used to derive the model-implied prior beliefs for the VAR parameters.

[8]Power utility is often assumed in related work, e.g. Pastor and Stambaugh (2012), Diris (2014), Hoevenaars, Molenaar, Schotman, and Steenkamp (2014), and Johannes, Korteweg, and Polson (2014)

3.1 The temperature long-run risk 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 temperature change 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

$$U_t = \left[(1-\delta)C_t^{\frac{1-\gamma}{\theta}} + \delta(\mathbb{E}_t[U_{t+1}^{1-\gamma}])^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}, \tag{7}$$

with C_t aggregate consumption at time t, $\delta \in (0,1)$ the investor time preference, γ the coefficient of risk aversion, $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$, and ψ the elasticity of intertemporal substitution (IES). The log of aggregate consumption growth ($\Delta c_{t+1} = \log(C_{t+1}/C_t)$) and log dividend growth of the portfolio i[9] ($\Delta d_{t+1,i} = \log(D_{t+1,i}/D_{t,i})$) follow

$$\begin{aligned}
\Delta c_{t+1} &= \mu_c + \sigma_t\eta_{t+1} + X_{t+1},\\
\Delta d_{t+1,i} &= \mu_d + \pi_d\sigma_t\eta_{t+1} + \phi_i X_{t+1} + \varphi_d\sigma_t u_{t+1},\\
\sigma_{t+1}^2 &= \bar{\sigma}^2 + \nu(\sigma_t^2 - \bar{\sigma}^2) + \sigma_w w_{t+1},\\
&\eta_{t+1}, u_{t+1}, w_{t+1} \sim N i.i.d.(0,1),
\end{aligned} \tag{8}$$

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 N with Poisson distributed increments as

$$\begin{aligned}
X_{t+1} &= \rho X_t + d\Delta N_{t+1},\\
\Delta N_{t+1} &\sim \text{Poisson}(\lambda_t = \Delta t(\lambda_0 + \lambda_1 T_t)),
\end{aligned} \tag{9}$$

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 long-term temperature level, in turn

[9] Portfolio i is the market or the portfolio (non-)vulnerable to temperature innovations, $i = m, v, nv$

based on the atmospheric carbon concentration ε_t as

$$\begin{aligned} T_{t+1} &= \chi \varepsilon_{t+1}, \\ \varepsilon_{t+1} &= \nu_\varepsilon \varepsilon_t + \mu_\varepsilon + \Theta(\mu_c + \sigma_t \eta_{t+1}) + \sigma_\zeta \zeta_{t+1}, \\ \zeta_{t+1} &\sim N i.i.d.(0,1). \end{aligned} \tag{10}$$

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

$$z_{t(,i)} = A_{0(,i)} + A_{1(,i)} T_t + A_{2(,i)} X_t + A_{3(,i)} \sigma_t^2. \tag{11}$$

Portfolios differ in the exposure of future dividend growth on disaster impact X_{t+1}, through the portfolio-specific parameter ϕ_i in Equation (8), which is higher for portfolios more vulnerable to temperature change. The analytical solution of 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.[10] 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 market portfolio. The log price-dividend ratio is the log difference between the real S&P 500 price and the corresponding monthly real dividend.

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

[10] `http://www.econ.yale.edu/~shiller/data.htm`.

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 average monthly land-based U.S. temperature anomalies from the nClimDiv dataset of the National Oceanic and Atmospheric Administration (NOAA).[11] We transform these anomalies to degrees Celsius, as in Bansal, Kiku, and Ochoa (2019). We use the first difference of the five-year moving average (MA) of the temperature anomaly as a proxy for the long-term temperature change.

Consumption data is from National Income and Product Accounts (NIPA). Annual 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 from NIPA Table 2.3.4 and divided by population (from NIPA Table 2.1).

We construct portfolios that are vulnerable or non-vulnerable to temperature change based a sort of industry portfolios on contemporaneous exposure to temperature change, building on the work of Balvers, Du, and Zhao (2017). We obtain value weighted monthly industry returns for 49 industries from Kenneth French's website.[12] We run a five-year rolling window regression (with at least 2.5 years of available data) of the log real monthly industry returns on temperature change and log real market returns. We then sort the industries on their exposure to temperature change and take the bottom (top) 5 industries as (non-)vulnerable industries. We take equally weighted returns of these industries in the month after our regression as the return of our vulnerable and non-vulnerable portfolios.

The sample period is 1947-2019.[13] All variables are monthly, except for consumption growth, which is annual because monthly consumption data is unavailable in the early years of our sample. Whenever we report quarterly or annual results, these are time-aggregated from monthly data using

[11] Temperature anomalies are the monthly average temperatures minus the average temperature in that same month for our base period 1901-2000. This base period is arbitrary and has no impact on our results, since we take a first difference that drops the base level.

[12] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[13] Our starting date follows Barnett (2020) 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 equity risk.

the methods from Bansal, Kiku, and Yaron (2016).

3.3 Model calibrations

We calibrate the LRR-T model with the calibration parameters given in Table 1. The main calibration is a scenario in which the temperature anomaly is zero in expectation at the start of our sample, with a long-term temperature expectation that is one degree Celsius higher. We call this the baseline climate change scenario and it is chosen because it closely matches our historical data, where we observe an increase in the temperature anomaly roughly from zero to one degree Celsius[14]. We calibrate the LRR-T model for the baseline scenario to match economic growth and financial moments.

We adjust the baseline LRR-T calibration in two ways for additional analyses. First, we include two scenarios with increased climate change for which we use the same LRR-T calibration, adjusting only the expected temperature change by adjusting the starting atmospheric carbon concentration (ε_0) and the separate trend in carbon emissions (μ_e). The alternative scenarios have expected increases from one to two (moderate scenario) and one to four (severe scenario) degrees Celsius over our sample period. For these scenarios we start at a higher temperature level to match expected climate change going forward, instead of matching observed data. Second, we allow for portfolio-specific impact of climate change by adjusting the dividend growth loading on climate disasters ϕ_i to increase or decrease vulnerability towards climate change. Specifically, we include the market portfolio in the baseline calibration and use ϕ_v and ϕ_{nv} for the vulnerable and non-vulnerable portfolios, respectively.

The first and second moments of observed economic growth, price-dividend ratios and returns in the data are compared with the implications from the LRR-T model calibrations in Table 2. The baseline LRR-T model fits the data well in most aspects. The main moments the model does not match the data are related to the price-dividend ratio, as is commonly observed with LRR calibrations. For example, Bansal, Kiku, and Yaron (2012) report expected price-dividend ratios at 3.07 and standard deviation of the price-dividend ratio at 26% in their LRR population moments.

[14] In our sample, the temperature anomaly is 0.1 in 1948 and 1.2 in 2019. Alternative sources of temperature anomalies show more smoothed versions that are closer to the 0-1 we use in our calibration.

Table 1: **Calibration parameters for the temperature long-run risk model**

LRR-T Baseline								
Preferences	δ	γ	ψ					
	0.998	6	1.5					
Consumption	μ_c	ρ	d	$\bar{\sigma}$	ν	σ_w		
	0.0052	0.99	−0.004166	0.0072	0.999	0.0000028		
Dividend	μ_d	π_d	ϕ_m	ϕ_v	ϕ_{nv}	φ_d		
	0.0062	2.0	1.1	1.35	0.85	5.0		
Climate	ν_e	ε_0	μ_e	Θ	σ_ζ	χ	λ_0	λ_1
	0.9971	−1	0.0095	1	0.4	0.2	0.075	0.075
Scenarios	ε_0	μ_e						
LRR-T Moderate	4	0.0238						
LRR-T Severe	2.1	0.0528						

This table shows the calibrated parameters for the temperature long-run risk model from Equations (8) to (11), based on a monthly decision interval. The baseline scenario matches historically observed climate change, with the temperature anomaly increasing from 0 to 1 degrees Celsius in expectation. The alternative scenarios forecast increased climate change, with future expected increases of the temperature anomaly from 1 to 2 and 1 to 4 degrees Celsius, respectively, for the moderate and severe scenarios. For these alternative scenairos, we only adjust the climate process in the baseline calibration by using different values for ε_0 and μ_e. ϕ_m, ϕ_v and ϕ_{nv} are the parameters for the different growth processes of the market, and the portfolios vulnerable and non-vulnerable to temperature innovations, respectively.

In our calibration, we have a stronger focus on matching equity returns well, for these moments our LRR-T calibration performs better than previous LRR calibrations.

Increased temperature change in the moderate and severe LRR-T scenarios change the model implications in several ways. In general, we find that higher temperatures imply more disasters, resulting in lower growth rates, price-dividend ratios and average returns. The annual standard deviation of the market returns does not increase with more severe climate scenarios, because negative monthly market returns around disasters are followed by a quick recovery. In other words, each climate disaster has a persistent impact on prices and price-dividend ratios, but its effect on market returns is transitory. Expected market returns are, however, significantly lower with more severe climate scenarios because of the increased occurrence of disasters that each have a transitory effect on market returns. In expectation, the market risk premium remains similar with different scenarios for climate change, because expected market returns and expected risk-free

Table 2: Calibration moments for temperature long-run risk models

	Data	LRR-T Baseline	LRR-T Moderate	LRR-T Severe
$E(r_m)$	0.072	0.070	0.049	0.025
$\sigma(r_m)$	0.164	0.160	0.160	0.161
$E(r_f)$	0.006	0.006	-0.013	-0.035
$\sigma(r_f)$	0.019	0.024	0.029	0.040
$E(p-d)$	3.515	3.011	2.801	2.530
$\sigma(p-d)$	0.446	0.226	0.214	0.258
$E(\Delta c)$	0.018	0.018	-0.011	-0.044
$\sigma(\Delta c)$	0.013	0.039	0.046	0.059
$E(\Delta d)$	0.025	0.025	-0.007	-0.043
$\sigma(\Delta d)$	0.061	0.118	0.123	0.130
$E(\Delta T)$	0.015	0.014	0.014	0.041
$\sigma(\Delta T)$	0.112	0.112	0.112	0.115

This table reports the first and second moments of market returns, ex-ante risk-free yields, log price-dividend ratios, consumption and dividend growth, and the temperature innovation. The first column reports the historical moments from monthly data time-aggregated to annual values (1947-2019), the columns on the right show the population moments from 10,000 simulations from the three different calibrations of the LRR-T models from Table 1. The baseline model matches historically observed climate change and the moderate and severe models forecast scenarios with increased climate change.

returns decrease by roughly the same amount. Finally, the standard deviations of consumption growth increase relatively more than those of dividend growth, because consumption growth is more affected by the increased occurrence of disasters than dividend growth[15].

3.4 Model simulations

We derive the implications of the LRR-T model for return dynamics by applying the framework of Avramov, Cederburg, and Lucivjanska (2018). The LRR-T model does not present an analytical solution to the VAR from Equation (1). Therefore, we simulate data from our asset pricing model and estimate the VAR on the simulated data. For each version of the LRR-T model (i.e. for each

[15]Consumption growth is affected more in our calibration because we prioritize matching the observed standard deviation of market returns to matching the observed standard deviation of dividend growth. Effectively, we include a relatively volatile shock u_{t+1} in the dividend growth process of Equation (8) that is not affected by the occurrence of disasters. Therefore, the standard deviation of dividend growth is not affected as much by increased disaster occurrence.

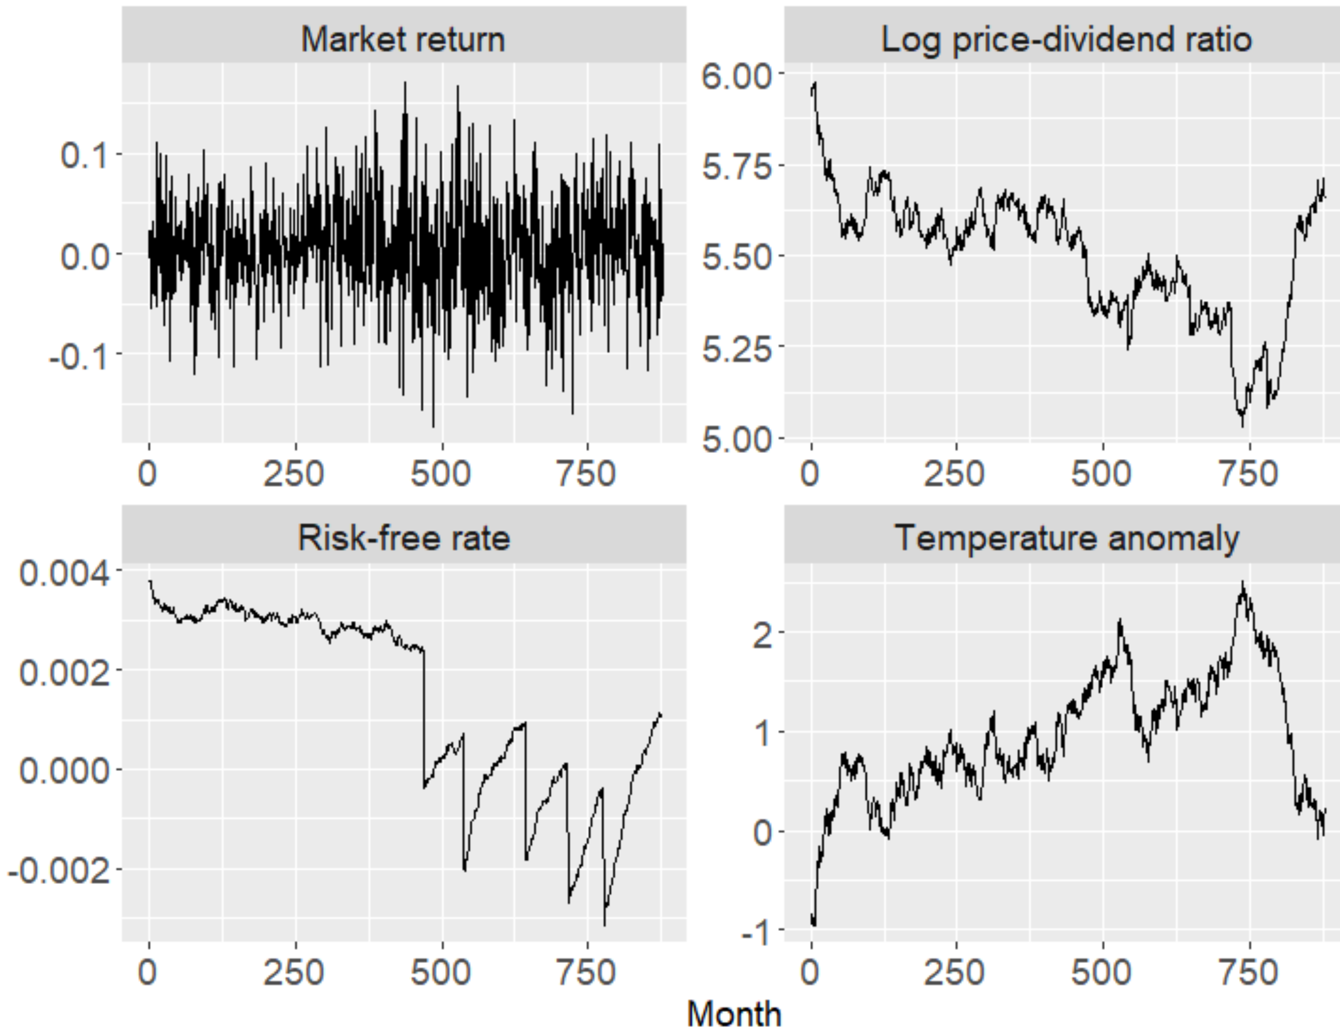

Figure 1: A single simulation from the baseline LRR-T model.
This figure shows the simulated market return, price-dividend ratio, risk-free rate and temperature anomaly from the baseline calibration of the LRR-T model in Table 1. We simulate 876 months, to match our 1947-2019 historical data sample.

calibration), we simulate 10,000 samples based on the processes in Equations (8)-(11), combined with the calibration from Table 1 and the solution to the fixed point problem discussed in Appendix A.1. Each simulated sample matches our historical data sample of 292 quarters from 1947Q1 to 2019Q4. We compute the levels, variances and correlations of the variables in the VAR model as implied by the LRR-T model as the mean of the VAR estimations for these 10,000 samples.

Our Bayesian investors combines the parameter estimates from the VAR model estimated using historical data on stock market returns, price-dividend ratios, risk-free rates, and temperature innovations, with the VAR parameters implied by the simulations from the LRR-T model. We give the same weight to the historical data and prior information. To achieve this, we set the misspecification of the model-implied prior of the VAR parameters to match the misspecification in the historical data, by scaling the prior density to the number of observations in our sample (N). This is visible in the posterior distributions presented in Appendix A.3.

To illustrate how temperature change impacts financial performance we show a single simulation from the baseline LRR-T calibration in Figure 1[16]. 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 shown in Figure 1. With increases in temperatures, the expected occurrence of future disasters that affect future consumption and dividend growth rates is increased. Price-dividend ratios and risk-free rates respond to these expected future adverse events with immediately decreases. Climate disasters start occurring in the second half of the simulation, most clearly visible in the risk-free rate[17]. These disasters have an immediate and persistent impact on risk-free rates and price-dividend ratios. 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.

The population moments from 10,000 simulations of the baseline LRR-T model imply the VAR structure shown in Table 3. For now, all coefficients are significant by construction, since we base the VAR in this table on the population moments from the model without allowing for

[16]This is a single simulation that we have chosen semi-randomly - this is not the general pattern from a large group of simulations, but a random outcome.

[17]We note that, in this simulation, we observe disasters relatively late by coincidence - in expectation we would observe roughly one disaster each every 100 months.

Table 3: Population estimates for VAR parameter and covariances from the temperature long-run risk model

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{m,t+1}$	0.232	−0.050	2.589	−0.058
$p_{t+1} - d_{t+1}$	0.309	0.929	2.170	−0.161
$r_{f,t+1}$	−0.001	0.000	0.981	−0.001
ΔT_{t+1}	−0.005	0.001	−0.007	0.963

This table shows the model coefficients of the VAR from Equation (1) as implied by the population moments of the baseline LRR-T model. The model is estimated on a quarterly time interval matching 1947Q1-2019Q4.

misspecification. First, relations between market returns and financial predictive 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 a temperature innovation as a predictive variable. Temperature increases have a strong negative impact on next period expected market returns. The coefficient from market returns on the temperature trend implies that a one standard deviation shock to temperature change decreases quarterly expected market returns by 0.65%, or 2.6% annually. This initial impact on market returns is partly offset in long-horizon forecasts, because the temperature increase also decreases price-dividend ratios. The temperature innovation is highly persistent, which is implied by the LRR-T model because we calibrate a positive trend in the temperature anomaly. Overall, this VAR shows that the LRR-T model indeed imposes a structure on the impact of temperature changes 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. Section 4.1 presents predictive regression estimates for an investor who forms prior beliefs based on the LRR-T model and for an investor who fully relies on historical data. Section 4.2 shows the implications of these estimates for long-horizon return variances and correlations. Section 4.3 combined these results in portfolio choice of investors with different beliefs over the horizon.

4.1 Predictive regressions

We estimate the quarterly predictive VAR from Equation (1) on the sample from 1947Q1 to 2019Q4 for three investor types: an *agnostic* investor who uses a data-based VAR with uninformative prior, a *dogmatic* investor, referred to as climate risk believer, who assumes that the VAR parameters follow the structure implied by the baseline LRR-T model, and a Bayesian investor who estimates the VAR parameters by combining historical data with the theoretical insights from the LRR-T model. Table 4 shows the posterior mean of the VAR coefficients for these investor types, along with their posterior standard deviations.

In panel A of Table 4, we show the posterior VAR based on historical data. 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. The inclusion of temperature change in the model does not make the VAR much more informative for long-horizon return forecasts. We observe statistically insignificant, but economically meaningful, positive coefficient estimates on the temperature trend. The positive relation between the historical temperature trend and financial outcomes is expected, since we have observed both increasing temperatures and financial growth in our sample.

On the other hand, we have a dogmatic 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, based on the baseline LRR-T model. The VAR coefficients are almost identical to the population estimated for the VAR parameters reported in Table 3. The difference with the population estimated 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 2.2. We still observe economically meaningful negative impact from increased temperatures on financial outcomes, but our sample is

Table 4: VAR parameter estimates for the agnostic, dogmatic and Bayesian investor

Panel A: Data (uninformative prior) - Agnostic

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{m,t+1}$	0.116	−0.020	0.713	0.029
	0.049	0.010	0.925	0.084
$p_{t+1} - d_{t+1}$	0.074	0.985	1.302	0.014
	0.051	0.010	0.955	0.087
$r_{f,t+1}$	0.002	0.000	0.883	0.001
	0.001	0.000	0.022	0.002
ΔT_{t+1}	−0.035	0.008	0.832	0.072
	0.034	0.007	0.656	0.059

Panel B: LRR-T prior - Dogmatic

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{m,t+1}$	0.234	−0.050	2.602	−0.061
	0.140	0.032	1.132	0.182
$p_{t+1} - d_{t+1}$	0.309	0.929	2.169	−0.160
	0.097	0.022	0.784	0.126
$r_{f,t+1}$	−0.001	0.000	0.982	−0.001
	0.002	0.000	0.017	0.003
ΔT_{t+1}	−0.005	0.001	−0.011	0.963
	0.013	0.003	0.111	0.018

Panel C: Data with LRR-T prior - Bayesian

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{m,t+1}$	0.049	−0.007	1.083	0.011
	0.024	0.005	0.403	0.046
$p_{t+1} - d_{t+1}$	0.032	0.993	0.439	−0.001
	0.025	0.005	0.386	0.051
$r_{f,t+1}$	0.000	0.000	0.962	0.000
	0.001	0.000	0.010	0.002
ΔT_{t+1}	0.008	−0.001	0.177	0.299
	0.018	0.004	0.263	0.039

This table shows the posterior means of the parameter estimates of VAR from Equation (1). Posterior standard deviations are reported below the coefficients. Panel A reports the VAR for the agnostic investor, based on uninformative priors. Panel B reports the VAR for the climate risk believer, based on the baseline LRR-T prior. Panel C reports the VAR for the Bayesian investor that combines data with the baseline LRR-T prior. The VARs are based on a quarterly sample from 1947Q1 to 2019Q4.

relatively short for long-horizon forecasts[18], which may be why posterior coefficients on temperature

[18]Generally, samples used for long-horizon forecasts are longer than 100 years. For example, Avramov, Cederburg, and Lucivjanska (2018) use 141 years of historical data. As discussed above, we are forced to analyze a shorter sample because we believe that historical data from longer ago is not informative about the impact of climate change on market returns.

change are now statistically insignificant. 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.

Panel C of Table 4 reports the posterior VAR for the Bayesian investor that combines implications from historical data (panel A) with 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 generally does not seem to give a lot of weight to temperature change, which is visible in the small absolute posterior coefficients on temperature change. These small coefficients *do* have significant impact on long-horizon predictive distributions of market returns, however. Small changes in the predictive VAR parameters become increasingly important after several quarterly forecasts.

Table 5 shows similar predictive regressions for the vulnerable and non-vulnerable portfolios. In these portfolios, we see stronger predictive power from temperature change on returns, especially for the agnostic investor in panel A and the Bayesian investor in panel C. The agnostic investor finds negative (positive) predictive power from temperature change on the returns of the (non-)vulnerable portfolio. Increasing temperatures decrease vulnerable portfolios, which seems to suggest lower expected future dividends. Non-vulnerable assets are positively affected by temperature increases. Both these coefficients are statistically insignificant, but economically meaningful. However, the difference between these two portfolios does have a statistically significant loading on temperature change. In the LRR-T prior, we also find a negative loading from vulnerable portfolio returns on temperature change. However, in contrast to the data, the non-vulnerable portfolio is also negatively affected by temperature increases. This is in line with the model set-up, where all portfolios are always negatively affected by disasters. The Bayesian investor shows coefficients between those of the agnostic and the dogmatic investors.

4.2 Long-horizon implications

We now use the VARs from Table 4 and the long-horizon variance from Equation (5) to forecast variance ratios for horizons up to 100 quarters. In Figure 2, the per period variance by horizon is

Table 5: Predictive regression for vulnerable and non-vulnerable returns

Panel A: Data (uninformative prior) - Agnostic

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{v,t+1}$	0.130	-0.023	-0.136	-0.042
	0.068	0.014	1.306	0.117
$r_{nv,t+1}$	0.140	-0.027	0.300	0.136
	0.076	0.015	1.443	0.131

Panel B: LRR-T prior - Dogmatic

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{v,t+1}$	0.307	-0.066	3.108	-0.076
	0.144	0.033	1.188	0.186
$r_{nv,t+1}$	0.137	-0.029	1.926	-0.034
	0.132	0.030	1.092	0.171

Panel C: Data with LRR-T prior - Bayesian

	Intercept	$p_t - d_t$	$r_{f,t}$	ΔT_t
$r_{v,t+1}$	0.071	-0.011	0.981	-0.046
	0.033	0.007	0.523	0.066
$r_{nv,t+1}$	0.063	-0.012	1.062	0.080
	0.035	0.007	0.552	0.070

This table shows the posterior means of the parameter estimates of the predictive regression for (non-) vulnerable portfolio returns on the market log price-dividend ratio, the risk-free return and temperature change. Posterior standard deviations are reported below the coefficients. Panel A reports the coefficients for the agnostic investor, based on uninformative priors. Panel B reports the coefficients for the climate risk believer, based on the baseline LRR-T prior. Panel C reports the coefficients for the Bayesian investor that combines data with the baseline LRR-T prior. The regressions are based on a quarterly sample from 1947Q1 to 2019Q4.

given for the agnostic (data-based), dogmatic (LRR-T based), and Bayesian (LRR-T combined with data) investors. Within a few quarters, the variance ratios quickly diverge. This result is mainly driven by the strong difference in mean reversion in the data and the LRR-T prior. Uncertainty about future expected returns is similar for both investors.

In the data, we find extremely strong mean reversion, which is also visible in, among others, Barberis (2000) and Siegel (2008). Therefore, long-horizon variance ratios are only a fraction of the short-term variance ratios in the data. 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. In the LRR-T model, however, there is significantly less mean reversion. 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 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.

Combining historical data with the LRR-T prior, we find that the variance ratio for the Bayesian investor decreases slightly with the horizon. We find that the predictive mean reversion of the Bayesian LRR-T investor is smaller than that of both the LRR-T prior and the historical data. This is an observation Avramov, Cederburg, and Lucivjanska (2018) also report for the Bayesian LRR investor, for which we do not have an explanation.

Another important aspect of portfolio choice over the horizon is the correlation between the assets in the optimization. When the risk-free return is highly correlated with risky portfolio returns, this pushes the investor towards the risky asset because there are smaller diversification benefits from buying the risk-free asset. Similarly, when the vulnerable and non-vulnerable returns are highly correlated, investors are pushed towards the portfolio with the best combination of risk and return.

Figure 3 shows these correlations for different investment horizons. In the short run, risk-free returns are hardly correlated with returns of risky portfolios (neither with the market, nor with vulnerable or non-vulnerable portfolios). In the long run, however, the correlation between the risk-free returns and the market returns becomes positive for every investor, with higher returns implied by the LRR-T model. For the vulnerable and non-vulnerable portfolios, the risk-free return is hardly correlated with the risky returns in the data. These observed correlations all correspond nicely with the estimated coefficients from risky return on the risk-free return in the predictive regressions from Tables 4 and 5. Positive (negative) loadings in those regressions imply more positive (negative) correlations with the risk-free return over the horizon. The correlation between the returns of the vulnerable and non-vulnerable portfolios is relatively high, but decreasing by horizon for both the data-based and Bayesian investors.

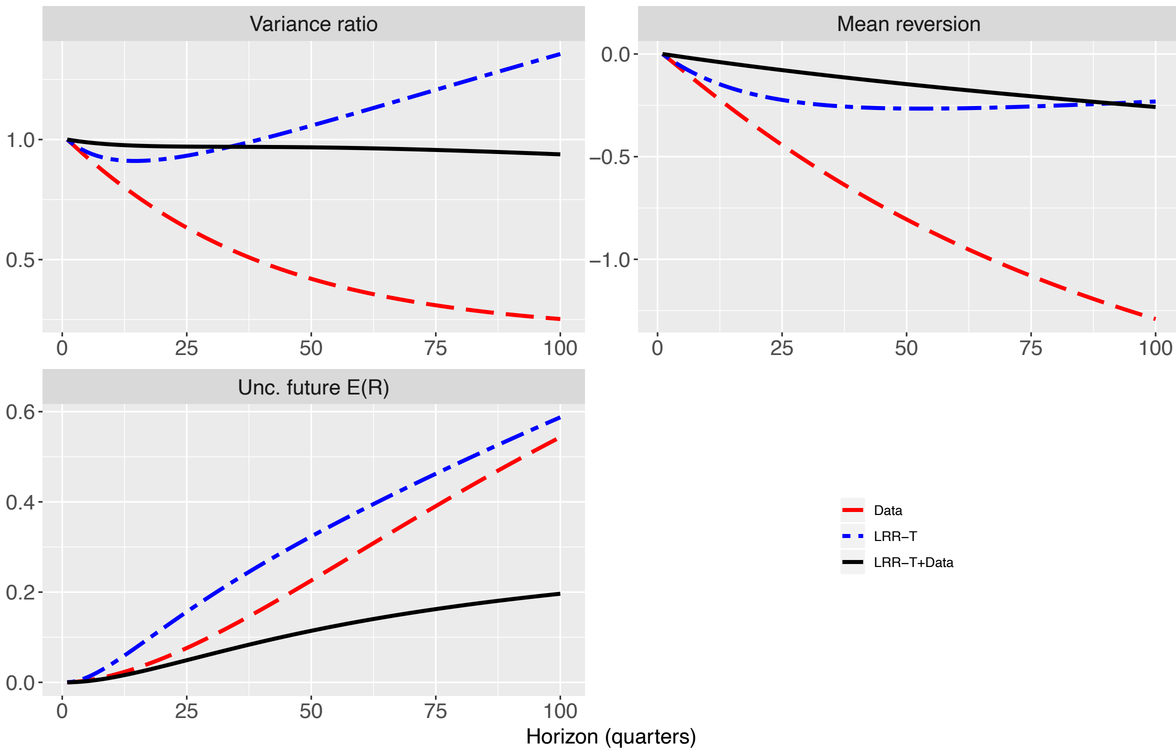

Figure 2: Predictive variance ratios and its components by horizon.
This figure shows the predictive variance ratio (top left) and the underlying components related to mean reversion (top right) and uncertainty about future expected returns (bottom left). The variance ratios of the Bayesian investor (LRR-T+Data) are based on the VAR from panel C of Table 4, combining the views of the agnostic investor (Data) and climate risk believer (LRR-T) from panels A and B of Table 4. Variance ratios are observed for investment horizons from 1 to 100 quarters.

4.3 Optimal portfolio choice

We have now documented three key results. First, as illustrated in Figure 1, the LRR-T risk-premium for the market increases after the occurrence of disasters, because the risk-free rate is persistently lower while the impact on market returns is transitory. Thus, the future expected risk premium is larger in the LRR-T model. Second, Figure 2 shows that variance over the horizon is affected significantly when climate change is taken into account through the LRR-T model. Specifically, the per period variance of market returns increases increases by horizon for climate risk believers (LRR-T investors), while it decreases for agnostic investor who base their analysis on historical data. Third, correlations between returns on the risk-free asset and the market increase

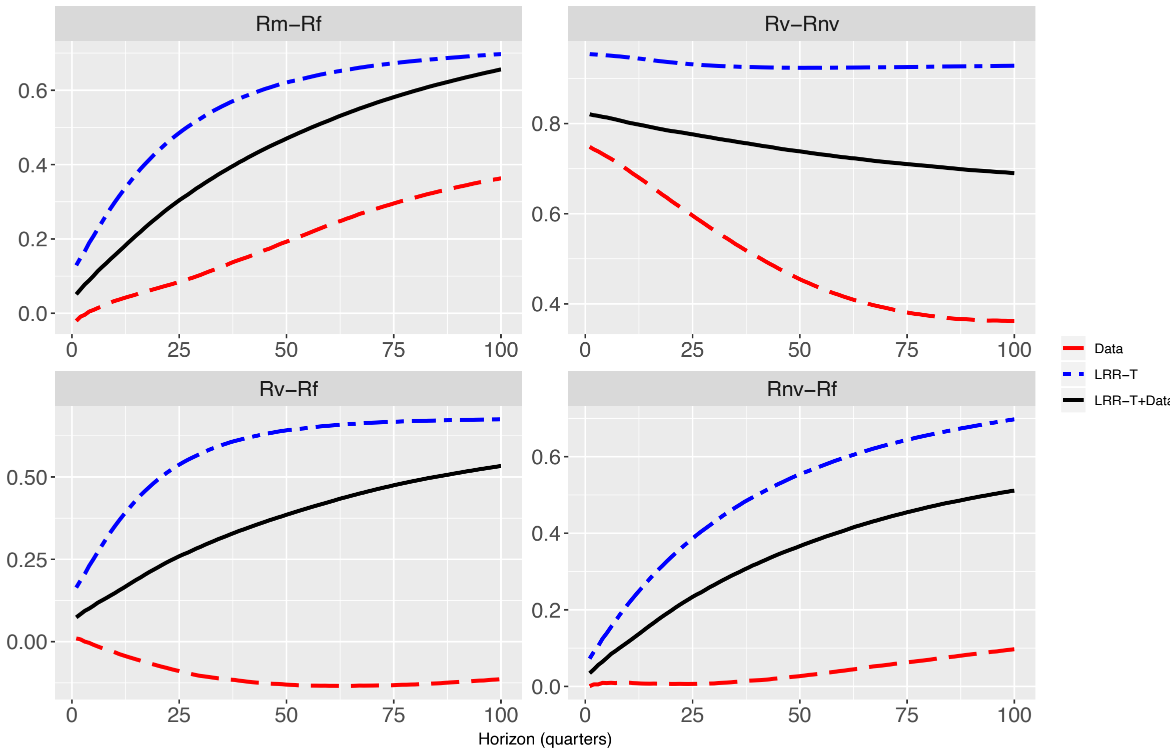

Figure 3: Predictive correlations by horizon.
This figure shows the predictive correlations between the market portfolio and the risk-free asset (top left), the vulnerable and non-vulnerable portfolios (top right) and the (non-)vulnerable portfolio and the risk-free asset (bottom). The correlations of the Bayesian investor (LRR-T+Data) are based on the forecasts from the predictive regressions in panel C of Tables 4 and 5, combining the views of the agnostic investor (Data) and climate risk believer (LRR-T) from panels A and B of Tables 4 and 5. Predictive correlations are observed for investment horizons from 1 to 100 quarters.

by horizon, both in the data and in the LRR-T model. At the same time, correlations between risky assets decrease by horizon. These results each have implications for portfolio choice, as we discuss in this section.

First, we analyze a long-only investor that can invest in the market portfolio and a risk-free asset. The top panel of Figure 4 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 as documented by Barberis (2000) and Siegel (2008). Dogmatic investors that invest based on LRR-T implications have higher weights to equities than the agnostic investor for horizons of

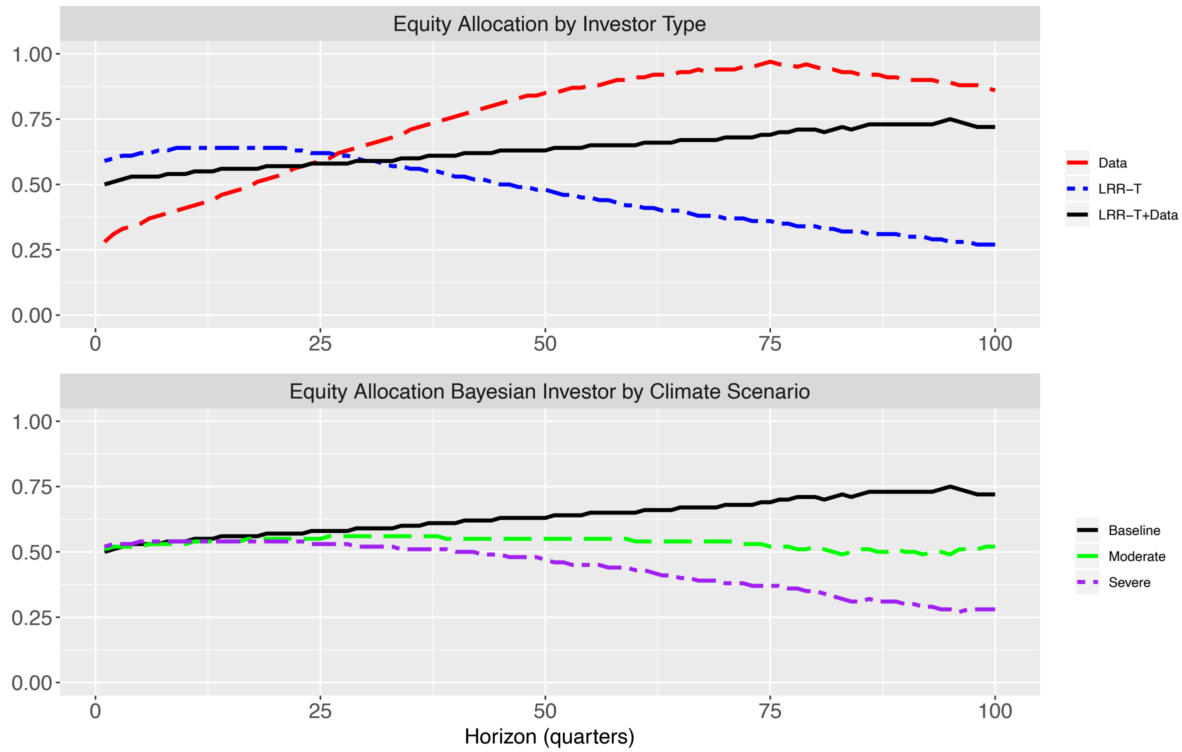

Figure 4: Optimal weight to equities by horizon for different investor types and climate scenarios.
This figure shows the optimal weight to equities for an investor with risk aversion parameter A=5, based on predictive returns and variances from the VAR models in Table 4 for the agnostic investor (Data), climate risk believer (LRR-T) and Bayesian investor (LRR-T+Data). We optimize the weight to equities with the risk-free asset as alternative investment for investment horizons from 1 to 100 quarters. In the top panel, results are from the baseline LRR-T model. The bottom panel shows the results from the Bayesian investor with alternative (baseline, moderate and severe) climate scenarios. We do not allow for short selling.

up to roughly 25 quarters, with the opposite result for longer investments horizons. In the short run, the increase in the future risk premium implied by LRR-T outweighs the increased variance from decreased mean reversion. In the long run, dogmatic investors deal with constantly increasing variance ratios over the investment horizon, resulting in lower allocations to equities.

The Bayesian investor shows continually increasing optimal weights to equities for all horizons. These results correspond to the continues decrease in variance ratios documented in Figure 2. However, the increase in optimal allocation to equities is stronger than would be expected from the relatively minor decrease in variance ratios alone. There are two other effects that push this investor towards the market. First, the risk-free rate has significant reinvestment risk over longer horizons, because disasters result in persistent decreases in risk-free returns. Second, the correlation between the risk-free returns and market returns are strongly positive at longer horizons, as documented in Figure 3.

Now that we have discussed the optimal portfolios for different investor types, we assess the impact of alternative climate change scenarios on our results. The bottom panel of Figure 4 reports the optimal allocation to equities for Bayesian investors taking into account the three different climate scenarios in the LRR-T priors shown in the calibration in Table 1. For long investment horizons, more severe climate change implies higher predictive variances that are not offset by increased risk premia and therefore lead to lower allocations to equities. We find that different climate scenarios do not impact short-term portfolio choice, because in each of these different LRR-T calibrations disasters are not expected in the short run. This highlights that taking climate change into account in strategic asset allocation decisions is particularly important for long-term investors such as pension funds.

Finally, we are interested in investors that allow for cross-sectional differences in long-horizon portfolio choice by including portfolios that are either vulnerable or non-vulnerable to temperature change in their investment universe. Figure 5 shows the optimal allocations for an investor that combines these investments with the risk-free asset. The most striking result in this figure is the fact that it is never optimal for a long-only investor to invest in the non-vulnerable portfolio, independent of the views of the investor. Both in the data and the LRR-T prior, the expected returns for the vulnerable assets are significantly higher than those for the non-vulnerable assets. Combined with the high correlation between the vulnerable and non-vulnerable portfolios documented in Figure 3, investors are pushed towards the vulnerable portfolio. Next to this, the implied allocation to the risky portfolio is similar as to the market in the dogmatic LRR-T prior. In the data, the vulnerable portfolio has less mean reversion, which results in smaller risky allocations in the long run.

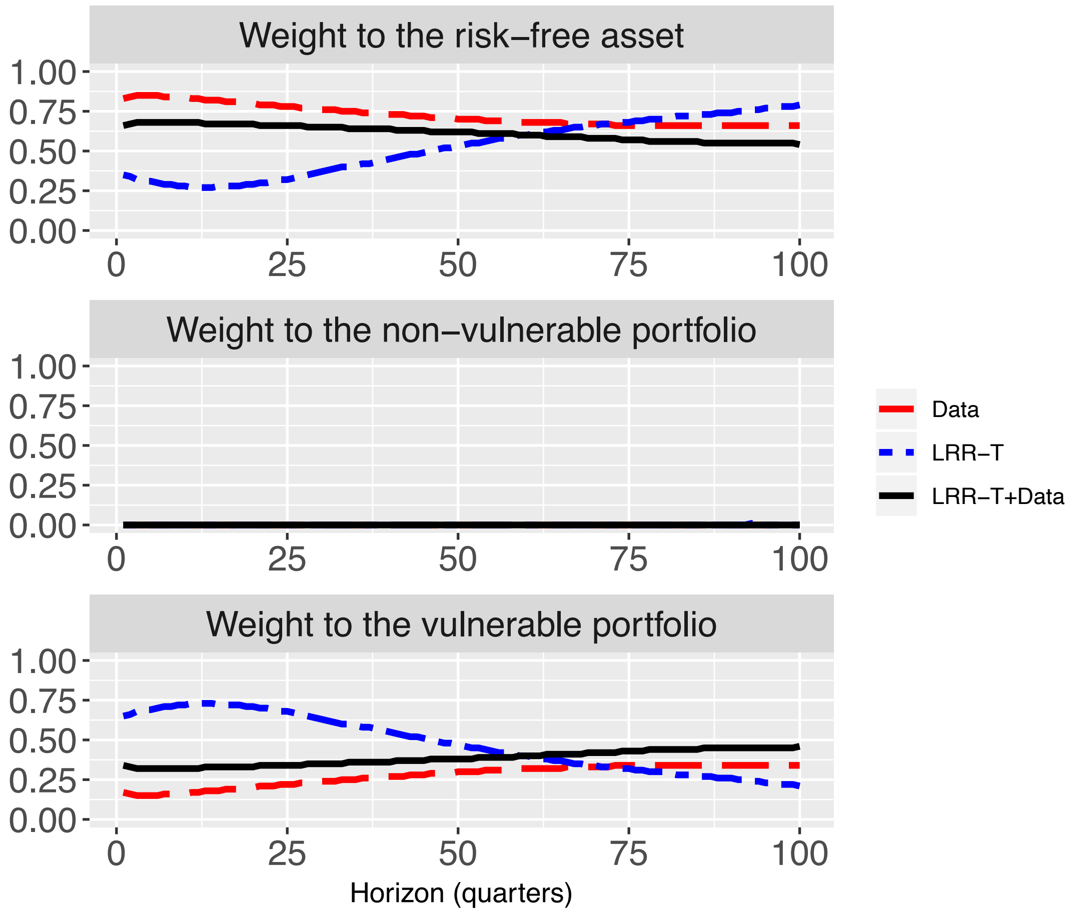

Figure 5: Portfolio choice with vulnerable and non-vulnerable portfolios by horizon. This figure shows the optimal weight to the risk-free asset, non-vulnerable, and vulnerable portfolios for an investor with risk aversion parameter A=5, based on predictive returns and variances from the regression models in Tables 4 and 5. Results are shown for the agnostic investor (Data), climate risk believer (LRR-T) and Bayesian investor (LRR-T+Data). We optimize portfolios for investment horizons from 1 to 100 quarters. We do not allow for short selling.

5 Conclusion

We propose a novel approach for measuring the impact of climate change on long-term equity risk and optimal portfolio choice. We characterize the long-horizon dynamics of equity returns by specifying a VAR model that includes temperature change as a predictor. Because historical data may not be very informative about the impact of climate change on future stock market returns, we estimate the parameters of the VAR using a Bayesian approach that complements historical data with prior information derived from economic theory. Specifically, we elicit prior beliefs from the temperature long-run risk (LRR-T) model of Bansal, Kiku, and Ochoa (2019).

We document four key findings. First, an investor with LRR-T beliefs perceives stock markets to be riskier over longer horizons because disasters induced by climate change reduce mean reversion in returns. Mean reversion decreases because climate disasters tend to cluster in periods with relatively high temperature levels. In other words, current disasters with an adverse impact on market returns are often quickly followed by new disasters, increasing the correlation between current and future returns. Second, the investor with LRR-T beliefs expects the market risk premium to increase after a climate disaster occurs. In particular, whereas disasters cause a persistent negative shock to risk-free rates, the negative impact on expected market returns is transitory because prices rapidly adjust after a disaster strikes.

Third, we show that for investors with a horizon longer than 25 quarters, the optimal allocation to equity decreases when accounting for climate change because the increase in perceived riskiness of stocks outweighs the increase in the market risk premium. In contrast, for short-term investors the increase in market risk premium is sufficiently large to increase the optimal allocation to equity relative to an investor with uninformative beliefs about the effect of climate change on returns. Finally, we document that the risk premium for portfolios vulnerable to temperature innovations is sufficiently large in both the data and the LRR-T beliefs that the long-only optimal allocation to portfolios that are non-vulnerable to temperature innovations are zero for all horizons.

References

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.

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.

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

Barnett, M., 2020, Climate change and uncertainty: An asset pricing perspective, Working Paper.

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, 2019, Do investors care about carbon risk?, Working Paper.

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 S. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509–1531.

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.

Dell, M., B. Jones, and B. Olken, 2014, What do we learn from the weather? the new climate-economy literature, *Journal of Economic Literature* 52, 740–798.

Diris, B., 2014, Model uncertainty for long-term investors, Working Paper, Erasmus University Rotterdam.

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.

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.

Ilhan, E., Z. Sautner, and G. Vilkov, 2020, Carbon tail risk, *Review of Financial Studies* , Forthcoming.

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

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.

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

Pettenuzzo, D., A. Timmermann, and R. Valkanov, 2014, Forecasting stock returns under economic constraints, *Journal of Financial Economics* 114, 517–553.

Siegel, J. J., 2008, *Stocks for the Long Run* (McGraw Hill, New York, NY).

A Appendix

A.1 Solution to the temperature long-run risk model

The price-consumption ratio follows $z_t = A_0 + A_1 T_t + A_2 X_t + A_3 \sigma_t^2$, where

$$
\begin{aligned}
\theta(1-\kappa_1\nu)A_3 &= 0.5(1-\gamma+\theta\kappa_1 A_1\chi\Theta)^2,\\
\theta(1-\kappa_1\rho)A_2 &= (1-\gamma)\rho,\\
\theta(1-\kappa_1\nu_\varepsilon)A_1 &= (1-\gamma+\theta\kappa_1 A_2)d\left(1+\frac{(1-\gamma+\theta\kappa_1 A_2)d}{2}\right)\lambda_1\Delta t,\\
(1-\kappa_1)A_0 &= \log\delta+\kappa_0+\left(1-\frac{1}{\psi}+\kappa_1 A_1\chi\Theta\right)\mu_c+\kappa_1 A_1\chi\mu_\varepsilon\\
&\quad+\kappa_1(1-\nu)A_3\bar{\sigma}^2+0.5\theta((\kappa_1 A_1\chi)^2\sigma_\zeta^2+(\kappa_1 A_3)^2\sigma_w^2)\\
&\quad+\left(1-\frac{1}{\psi}+\kappa_1 A_2\right)d\left(1+\frac{(1-\gamma+\theta\kappa_1 A_2)d}{2}\right)\Delta t\lambda_0.
\end{aligned}
\tag{12}
$$

The price-dividend ratio for portfolio i (either the market (m), or the (non-)vulnerable portfolio (v or nv)) equals $z_{t,i} = A_{0,i} + A_{1,i}T_t + A_{2,i}X_t + A_{3,i}\sigma_t^2$, where

$$
\begin{aligned}
(1-\kappa_{1,i}\nu)A_{3,i} &= 0.5(((\theta-1)\kappa_1 A_1+\kappa_{1,i}A_{1,i})\chi\Theta+\pi_d-\gamma)^2\\
&\quad+0.5\varphi_d^2+(\theta-1)(\kappa_1\nu-1)A_3,\\
(1-\kappa_{1,i}\rho)A_{2,i} &= (\theta-1)(\kappa_1\rho-1)A_2+\phi_i\rho-\gamma\rho,\\
(1-\kappa_{1,i}\nu_\varepsilon)A_{1,i} &= (\theta-1)(\kappa_1\nu_\varepsilon-1)A_1+C_N(1+0.5C_N)\lambda_1\Delta t,\\
(1-\kappa_{1,i})A_{0,i} &= \theta\log\delta+(\theta-1)\kappa_0+\kappa_{0,i}+(\theta-1)(\kappa_1-1)A_0+\mu_d\\
&\quad+(((\theta-1)\kappa_1 A_1+\kappa_{1,i}A_{1,i})\chi\Theta-\gamma)\mu_c+((\theta-1)\kappa_1 A_1+\kappa_{1,i}A_{1,i})\chi\mu_\varepsilon\\
&\quad+C_N(1+0.5C_N)\lambda_0\Delta t+((\theta-1)\kappa_1 A_3+\kappa_{1,i}A_{3,i})(1-\nu)\bar{\sigma}^2\\
&\quad+0.5(((\theta-1)\kappa_1 A_3+\kappa_{1,i}A_{3,i})^2\sigma_w^2+((\theta-1)\kappa_1 A_1+\kappa_{1,i}A_{1,i})^2\chi^2\sigma_\zeta^2),\\
C_N &= ((\theta-1)\kappa_1 A_2+\kappa_{1,i}A_{2,i}+\phi_i-\gamma)d.
\end{aligned}
\tag{13}
$$

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

$$\begin{aligned}\kappa_{0(,i)} &= \log(1+\exp(\bar{z}_{(i)})) - \kappa_{1(,i)}\bar{z}_{(i)}, \\ \kappa_{1(,i)} &= \frac{\exp(\bar{z}_{(i)})}{1+\exp(\bar{z}_{(i)})},\end{aligned} \tag{14}$$

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

$$\begin{aligned}\bar{z}_{(i)} &= A_{0(,i)} + A_{1(,i)}\bar{T}_t + A_{2(,i)}\bar{X}_t + A_{3(,i)}\bar{\sigma}_t^2 \\ &= A_{0(,i)} + A_{1(,i)}\left(\frac{1}{n}\mathbb{E}_0\left[\sum_{t=1}^{n} T_t\right]\right) + A_{2(,i)}\left(\frac{1}{n}\mathbb{E}_0\left[\sum_{t=1}^{n} X_t\right]\right) + A_{3(,i)}\bar{\sigma_t}^2,\end{aligned} \tag{15}$$

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

$$\begin{aligned}\frac{1}{n}\mathbb{E}_0\left[\sum_{t=1}^{n} T_t\right] &= \left(\frac{T_0}{n} - \frac{\chi(\mu_\varepsilon+\Theta\mu_c)}{n(1-\nu_\varepsilon)}\right)\left(\frac{\nu_\varepsilon-\nu_\varepsilon^{n+1}}{1-\nu_\varepsilon}\right) + \frac{\chi(\mu_\varepsilon+\Theta\mu_c)}{1-\nu_\varepsilon}, \\ \frac{1}{n}\mathbb{E}_0\left[\sum_{t=1}^{n} X_t\right] &= \frac{1}{n}\left(-X_0 + \sum_{t=0}^{n}\mathbb{E}_0[X_t]\right) \\ &= \frac{X_0}{n}\left(\frac{\rho-\rho^{n+1}}{1-\rho}\right) + \frac{1}{n}\left(\frac{d\Delta t\lambda_0}{1-\rho} + \frac{d\Delta t\lambda_1}{1-\rho}\frac{\chi(\mu_\varepsilon+\Theta\mu_c)}{1-\nu_\varepsilon}\right)\left(n - \frac{\rho-\rho^{n+1}}{1-\rho}\right) \\ &\quad + \frac{1}{n\nu_\varepsilon}\left(\frac{d\Delta t\lambda_1}{1-\left(\frac{\rho}{\nu_\varepsilon}\right)}\right)\left(T_0 - \frac{\chi(\mu_\varepsilon+\Theta\mu_c)}{1-\nu_\varepsilon}\right)\left(\frac{1-\nu_\varepsilon^{n+1}}{1-\nu_\varepsilon} - \frac{1-\rho^{n+1}}{1-\rho}\right).\end{aligned} \tag{16}$$

The sample used in our regressions starts in 1947Q1 and runs until 2019Q4, 292 observations of quarterly data. For our simulations, we construct a similar sample. Therefore, we simulate $n = 939$ months of data, resulting in a sample of 292 quarters after we drop the first five years of our simulation to make the first difference of the five-year MA of temperatures and the last three months because of a lead in the risk-free yields.

A.2 Adjustments to the LRR-T model

The LRR-T model presented in Section 3.1 deviates from the LRR-T model of Bansal, Kiku, and Ochoa (2019) in several ways.

First, in Equation (8), the dividend growth process allows for different loadings on the different shocks in the model and we let volatility vary over time. Both adjustments increase the flexibility of the model and are generally implemented in recent versions of LRR model (among others, in Bansal, Kiku, and Yaron (2012)). We allow for portfolio-specific disasters impact through a portfolio-specific ϕ_i.

Second, in Equation (9), we include a persistence $\rho < 1$ of 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 impact should not be indefinite. In the same equation, we let the increments of the disaster process (ΔN) be Poisson distributed, instead of the whole process N. With this adjustment, we make the current disaster intensity dependent on recent temperature levels, instead of on the historical path of temperature growth.

Third, in Equation (10), the atmospheric carbon concentration includes a separate trend μ_ε. This trend is used to generate the different climate scenarios in our analysis.

Finally, in Equation (11), we assume that the log wealth-consumption ratio z_t and the log price-dividend ratio $z_{t,m}$ also depend on the economic disaster impact X_t and time-varying variance σ_t^2. The inclusion of σ_t^2 follows the inclusion of time-varying volatility in Equation (8). We believe that the the inclusion of economic disaster impact X_t in the processes for z_t and $z_{t,m}$ is intuitively reasonable. When temperatures increase, z_t and $z_{t,m}$ should decrease as prices decrease to adjust for expected future disasters. Without sufficient precautionary savings, these ratios should also be affected when disasters occur, because these disasters should have significant cashflow impact.

A.3 Posterior VAR distributions

We follow Avramov, Cederburg, and Lucivjanska (2018) and estimate the VAR model in Equation (1) as

$$\begin{bmatrix} r_{m,t+1} \\ p_{t+1} - d_{t+1} \\ r_{f,t} \\ \Delta T_{t+1} \end{bmatrix} = C' \begin{bmatrix} 1 \\ p_t - p_d \\ r_{f,t-1} \\ \Delta T_t \end{bmatrix} + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, \Sigma). \tag{17}$$

where C is the VAR coefficients matrix and Σ is the variance-covariance matrix of the residuals. As shown in the Appendix of Avramov, Cederburg, and Lucivjanska (2018), the posterior distribution of the VAR parameters with the model-based prior conditional on the observed data for periods $1, \ldots, t$, D_t, and the asset pricing model parameters Θ_M is given by

$$\begin{aligned} \Sigma \mid D_t, \Theta_M &\sim IW((\omega_M + \omega_D)N\hat{\Sigma}(\Theta_M), (\omega_M + \omega_D)N - 4), \\ C \mid \Sigma, D_t, \Theta_M &\sim N(\hat{C}(\Theta_M), \Sigma \otimes (\omega_M N \Gamma^*_{xx} + \omega_D X'X)^{-1}), \end{aligned} \tag{18}$$

in which

$$\begin{aligned} \hat{\Sigma}(\Theta_M) &= \frac{1}{(\omega_M + \omega_D)N}[(\omega_M N \Gamma^*_{yy} + \omega_D Y'Y) \\ &\quad - (\omega_M N \Gamma^*_{xy} + \omega_D Y'X)(\omega_M N \Gamma^*_{xx} + \omega_D X'X)^{-1}(\omega_M N \Gamma^*_{xy} + \omega_D X'Y)], \\ \hat{C}(\Theta_M) &= (\omega_M N \Gamma^*_{xx} + \omega_D X'X)^{-1}(\omega_M N \Gamma^*_{xy} + \omega_D X'Y), \end{aligned} \tag{19}$$

N is the number of observations in our sample (292 quarters), and Γ^*_{xx}, Γ^*_{xy}, and Γ^*_{yy} are the population moments from the asset pricing models in Section 3.1, conditional on the model parameters Θ_M, as introduced in the 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 weight given to the model-based prior and the historical data, respectively. For the Bayesian investor, we set equal weights to the informative prior and historical data by setting $\omega_M = \omega_D = 1$. For the agnostic investor we specify $\omega_M = 0$ and $\omega_D = 1$, and Equation (18) then equals the posterior of the analysis on historical data with a multivariate Jeffreys uninformative prior. For the climate change believer we specify $\omega_M = 1$ and $\omega_D = 0$, and

Equation (18) then equals the prior distribution solely based on an asset pricing model.

A.4 Numerical optimal asset allocation

We numerically solve the optimal asset allocation for our buy-and-hold long-only investor following the methodology from, among others, Barberis (2000). This appendix explains our methodology. We closely follow a similar explanation in the appendix from Diris (2014) in this section.

Sampling from the predictive distribution

We sample $N = 250,000$ paths of length $K = 100$ quarters from the predictive distribution of the asset returns and state variables in our VAR model. We repeat the following two steps N times:

1. For k in $1, \ldots, K$, Sample the returns and state variables in period k conditional on the posterior mean VAR parameters (a, B, Σ) and the predictor variable values in $k-1$. We start forecasting from the last observed value of our predictor variables in the data at period $k = 1$.

2. Re-sample the asset returns and predictor variables in period k when we draw a quarterly return for the risk-free asset below -10%. This step is needed to avoid minus infinity utility, as an investor can never go bankrupt as long as it invests in the risk-free asset with this restriction. In practice, re-sampling is hardly ever required.

Calculation of buy-and-hold portfolio

Based on the N return forecasts from above we now follow the steps below to compute the optimal investment portfolio.

1. Make a grid of portfolio weights. We invest long only, i.e. weights are between zero and one, and use steps of 0.01 for our grid search. Then, for each weight in the grid:

2. Take a set of weights and calculate realized utility for each simulated path.

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

For each horizon $k = 1, \ldots, K$, we choose the optimal vector of weights as the one that results in the largest expected utility from 3.