

Estimating macro-fiscal effects of climate shocks from billions of geospatial weather observations

Berkay Akyapı, Matthieu Bellon, and Emanuele Massetti*

December 19, 2023

Abstract

The literature studying the macroeconomics of weather has focused on temperature and precipitation annual averages, ignoring their distributions over space and time. We construct hundreds of variables from high frequency, high spatial resolution weather measurements. Using the LASSO, we identify the parsimonious subset of variables that can best explain GDP and key macro-fiscal variables. We find that an increase in the occurrence of high temperatures and droughts, and scarcer mild temperatures reduce GDP. These variables substantially improve the share of GDP variations explained by weather. Additional evidence suggests that fiscal policy mitigates these shocks.

Keywords: climate, extreme weather, GDP, fiscal policy, big data

JEL: C33, C55, E62, O40, O44, Q54

*Akyapı: berkayakyapi@ufl.edu; Bellon: m.bellon@esm.europa.eu; Massetti: emassetti@imf.org; We thank anonymous reviewers, seminar participants at the IMF, the World Bank, the ESM, the EAERE 27th annual conference, and the 2023 NGFS conference on Green Finance Research Advances for helpful comments and discussions. The views expressed herein are those of the authors and should not be attributed to the IMF, the ESM, their respective executive boards, or their management teams.

1. Introduction

Climate is changing and will continue to change in the forthcoming decades even if sharp reductions in greenhouse gas emissions will succeed in stabilizing temperature in the second half of the century (IPCC, 2021b). In this context, economists and policy-makers are striving to better understand all the effects of climate on the economy.

Studies of the macroeconomic impacts of weather and climate can be divided in two groups.¹ One strand of the literature estimates the effect of climate on the macro economy by aggregating impacts from sectoral studies using a reduced-form damage function that links global mean temperature to total output losses (e.g. Hope et al., 1993; Nordhaus and Yang, 1996; Tol, 1997; Christensen et al., 2012). Another strand of the literature uses econometric analysis to directly estimate the impact of random changes in temperature, and sometimes precipitation, on GDP per capita or Total Factor Productivity (Dell et al., 2012; Deryugina and Hsiang, 2014; Burke et al., 2015; Abatzoglou et al., 2018; Letta and Tol, 2019; Kalkuhl and Wenz, 2020; Tol, 2021; Newell et al., 2021; Kahn et al., 2021). Both strands of the literature have mostly focused on the effect of annual average weather on GDP and this contrasts with the widespread concerns and evidence from sectoral studies that extreme weather may cause the largest losses.

We fill this gap and contribute to the econometric analysis of the effect of climate on the macro economy along three dimensions. First, we leverage a global dataset of daily measurements of temperature and precipitation with high spatial resolution to construct a large array of weather variables which can capture all sorts of potentially relevant extreme events. However, the large number of potentially relevant weather variables creates a challenge for standard estimation techniques. Therefore, our second contribution is the use of the Least Absolute Shrinkage and Selection Operator (LASSO) to select the weather variables that contribute the most to explaining macroeconomic outcomes. Third, we look beyond the effect of weather on GDP and examine important fiscal aggregates. This extension relative to prior work is motivated by the fact that fiscal policy, if counter-cyclical, potentially absorbs and masks some of the macro effects of weather shocks.

In short, we find that focusing on GDP and weather averages misses most of the macroeconomic impacts of climate shocks. The introduction of a small number of well-selected alternative weather variables goes a long way in improving our understanding of macro-fiscal variations.

Our study starts with the construction of a rich database of weather variables that can be used to conduct macro-fiscal analysis. We rely on dozens of billions of daily temperature and precipitation measurements on a global grid with a 30-Km resolution from the ERA5 dataset produced by the Copernicus Climate

¹Climate is the long-run distribution of weather over several decades (Auffhammer et al., 2013). Weather varies continuously, but it is bounded by its long-term distribution. This distribution can be characterized using averages, but also higher-order moments. We use “weather shocks” or “climate shocks” to indicate short-term changes in a weather variable. Climate change is instead the long-run, slow-moving change in the distribution of weather over several decades.

Change Service (C3S) at the European Centre for Medium-Range Weather Forecasts (ECMWF). There is an intractable number of ways to combine these daily geospatial measurements into country annual variables to match the annual frequency and country-level coverage of GDP per capita. Therefore, we rely on the climate literature to guide the construction of potentially relevant weather variables. We obtain 160 variables that include, for example, the center and tails of the distribution of temperature and precipitation, heat and cold waves, droughts, and intense precipitation. Once merged with macro-fiscal outcomes and after accounting for the introduction of lags, our dataset covers about 500 variables and 203 countries annually over the 1979-2019 period.

Our weather dataset exploits the richness of daily geospatial measurements to capture local and infra-annual shocks. These variables can reflect weather events that are likely missed when averaging over space and time. For example, total annual precipitation fails to capture a local drought if it concurs with high precipitation later in the year or in other parts of the country. These variables also allow us to differentiate between the effects of duration and intensity. For example, the effect of extreme heat (short-lived extreme temperature) can differ from the effect of heat waves (prolonged periods with unusually high temperature). We can also measure shocks that are only relevant because they are deviations from local and seasonal norms (temperatures that can be normal in a country like India could be devastating in a country with a different climate like Mongolia). For each variable, we consider aggregation over space using area weights and population weights.

We rely on a flexible empirical specification to relate weather shocks to macro-fiscal outcomes. In our baseline specification, we regress the first difference of the macroeconomic variable of interest on the first difference of our selected weather variables. Country fixed effects remove time-invariant country characteristics that affect growth, including the linear trends in some climate variables that we observe in the estimation sample. Year fixed effects remove global shocks. We also add lags of all variables to allow for rich dynamic effects and control for auto-correlation. Hence, the effect of weather shocks is estimated using random interannual variation of countries' de-trended weather after removing global shocks.

With approximately 500 independent variables, standard macroeconomic regressions would quickly run into over-fitting and near multi-collinearity issues. To select the variables that can best explain macroeconomic outcomes, we use an algorithm based on the LASSO (Tibshirani, 1996). This algorithm maintains a balance between underfitting and overfitting by imposing a penalty on the inclusion of non-zero coefficients in the regression. Non-zero coefficients are those for which the reduction in standard error must outweigh the penalty incurred by their inclusion in the regression, ensuring that only the most influential variables are retained. In the LASSO, the number of retained variables depends on a hyperparameter whose value can be set to maximize the fit of the model. Because our main concern is the sparsity and interpretability of results, we focus on the fit as measured by the Bayesian Information Criterion but we also discuss the alternative results obtained with different fit criteria.

We find that a handful of weather variables have a significant impact on GDP per capita. Some of these

variables capture harsh droughts and very high temperatures. We estimate that an increase in the occurrence of such weather shocks has a detrimental effect on GDP. Conversely, we find that an increase in mild temperatures have beneficial effects. We find that a one standard deviation in the selected variables leads to impacts of around 0.2 percentage points of GDP.

Our results demonstrate that the effects of specific temperature shocks identified in many micro or sectoral studies have sizeable macro-level effects. The beneficial effect of moderate temperatures and the harmful effect of very high temperature on agricultural output (Schlenker and Roberts, 2009; Blanc and Schlenker, 2017), mortality (Deschênes and Greenstone, 2011), energy consumption (Deschênes and Greenstone, 2011), time allocated to labor (Graff Zivin and Neidell, 2014), and labor productivity (Somanathan et al., 2021) have been widely documented. The impact of droughts is similar to what other studies found (Cantelmo et al., 2019; IMF, 2020).

We empirically test the persistence of weather impacts using impulse response functions estimated with the local projection method proposed in Jordà (2005). We find evidence that each instance of a weather shock has a persistent and stable impact on GDP levels over a 7-year horizon. However, we define weather shocks as random inter-annual variations from mean levels. Therefore, “good” and “bad” shocks are symmetric by construction in all countries. Their average value is equal to zero over our estimation period and our specification implies that their average impact on average growth (and GDP level) is also equal to zero. Our results do not imply that GDP is unaffected by climate change. On the contrary, trends observed in some of our climate variables may have already negatively affected average GDP growth and would be expected to remain detrimental in the future.

However, our empirical setup is not ideal to quantify the effect of climate change on long-run growth rates. We remove the average levels and linear trends of weather variables by using first differences and country fixed effects. This strengthens our strategy to identify the effect of weather shocks which are arguably unpredictable in this setup. However, it removes changes to the long-run distribution of weather (both changes in average conditions and in the frequency of tail events) which would capture climate change and its effects. Our setup is also not ideal to quantify the extent of future adaptation capable of reducing climate change impacts.

We confirm the robustness of our results with a battery of alternative specifications and heterogeneity analysis. The selection of relevant climate variables is similar when we use the Elastic Net (an alternative machine-learning operator addressing some of the LASSO’s limitations) and different sets of fixed effects. The estimated magnitude of the selected weather shocks is robust to the introduction of a range of controls capturing the presence of violent conflicts, inflation dynamics, and external shocks, and to alternative estimators that can correct biases arising from time dependence in panel data.

Heterogeneity analysis confirms our main results and highlights meaningful differences across country groups. Overall, we find that the effect of weather shocks is larger in countries that are more oriented

towards agriculture, in countries that are poorer, in the first half of the panel than in the second half, and in countries with less democratic institutions.

One of our key results is that our selected climate variables perform much better in explaining GDP variations than temperature and precipitation averages used in the literature. We confirm this result for a wide range of metrics by establishing comparisons with two central papers in the literature, [Burke et al. \(2015\)](#) and [Kahn et al. \(2021\)](#). For example, we measure how the within R-squared calculated on out-of-sample data improves after we introduce climate variables in a specification without climate variables. We find that adding our selected climate variables in the GDP regressions they consider can double or triple the percentage increase in the out-of-sample within R-squared. This result emphasizes that changes in weather extremes are more important than changes in average conditions to explain GDP variations and that the overall macroeconomic importance of climate may be larger than previously thought.

Nevertheless, we find that the total amount of variation in GDP per capita attributable to weather is small. Our selection of climate variables can at most increase the within R-square by a few percents. This is an indication that weather is not the main driver of GDP variations globally on average.

We additionally conduct a systematic analysis of the composition and the mitigating and amplifying effects of fiscal responses by examining government revenue, expenditure and debt, together with GDP. To keep our analysis compact, we use the LASSO to select the most relevant climate variable for each of these three fiscal outcomes. The procedure selects three new variables: the length of the longest day cold wave, mean precipitation in wet days, and mean precipitation in the driest months.

We find that the response of fiscal variables to weather shocks tends to mitigate the effects of a weather shock by implying larger fiscal deficits when the shock has negative consequences for GDP. Specifically, the expenditure-to-GDP ratio increases with harsh droughts while the revenue ratio falls when temperatures become less favorable. We also find that other weather shocks have significant and rich impacts on fiscal aggregates but are hard to interpret. The rich patterns we uncover suggest that the characteristics of the fiscal responses to weather shocks are complex but significant, and deserve more granular analyses.

The rest of the paper is organized as follows. The next section describes our empirical specification and the algorithm to select relevant climate variables. The third section explains how we construct the weather variables and summarizes the main characteristics of our dataset. The fourth and fifth sections present results, first for GDP per capita, and then for fiscal variables. The last section concludes.

2. Methods

2.1. Empirical model specification

Weather shocks can have potentially complex dynamic effects on GDP. We start by relating GDP per capita in country i at time t ($y_{i,t}$) to a vector of weather variables ($\mathbf{X}_{i,t}$) with a very flexible specification:

$$\ln y_{i,t} = \sum_{k=0}^K a_{k,i} t^k + \sum_{l=1}^L \theta_l \ln y_{i,t-l} + \sum_{p=0}^P \beta'_p \mathbf{X}_{i,t-p} + \mathbf{c}' \mathbf{Z}_t + \varepsilon_{i,t} \quad (1)$$

where $\sum_{k=0}^K a_{k,i} t^k$ are country-specific polynomial trends in weather patterns or economic activity, \mathbf{Z}_t is a vector of variables capturing global shocks, and $\varepsilon_{i,t}$ is the error term. This specification encompasses various models estimated in the literature (Hsiang, 2010; Dell et al., 2012; Deryugina and Hsiang, 2014; Burke et al., 2015; Kalkuhl and Wenz, 2020; Kahn et al., 2021), potentially allowing weather variables to have persistent dynamic effects on GDP.

To address serial-correlation and the fact that country GDP levels are non-stationary, we estimate equation (1) in first difference. It becomes a standard ARDL equation for GDP per capita growth:

$$\Delta \ln y_{i,t} = \sum_{k=0}^{K-1} \alpha_{k,i} t^k + \sum_{l=1}^L \theta_l \Delta \ln y_{i,t-l} + \sum_{p=0}^P \beta'_p \Delta \mathbf{X}_{i,t-p} + \mathbf{c}' \Delta \mathbf{Z}_t + \epsilon_{i,t}. \quad (2)$$

This equation continues to have polynomial trends but of order $K - 1$ instead of K .² We test alternative restrictions on the order of the polynomial and on the vector $\Delta \mathbf{Z}$. Note that persistent effects of weather shocks would be revealed by a long trail of significant β'_p .

We don't allow for a relationship between GDP growth and levels of the weather variables because GDP growth is stationary whereas many weather variables exhibit trends and are not stationary. Table B.1 in the online appendix presents evidence that average temperature and most variables built using temperature data are trended in most countries, as noted in the context of this literature by Kahn et al. (2021).³ This implies that GDP growth and these level variables cannot be related without additional manipulation (Tol, 2019; Kahn et al., 2021).

In our specification, such trends imply that the average change in weather variables takes significant and different values across countries. We systematically include country fixed effects to control for these

²The new coefficients of the polynomial trends are defined based on those in equation (1) by the equation $\sum_{k=0}^K a_{k,i} (t^k - (t-1)^k) = \sum_{k=0}^{K-1} \alpha_{k,i} t^k$, where the left-hand side is equal to $\sum_{k=0}^K a_{k,i} \left(t^k - \sum_{q=0}^k \binom{k}{q} t^q (-1)^q \right)$, and also equal to $\sum_{q=0}^{K-1} t^q \left(\sum_{r=0}^{K-q} \binom{r+1}{q} a_{r,i} (-1)^{r+1} \right)$ after arranging terms. This means $\alpha_{k,i} \equiv \sum_{r=0}^{K-k} \binom{r+1}{k} a_{r,i} (-1)^{r+1}$.

³For example, temperature trends range from 0.07 to 0.6 °C per decade across countries. The positive trend in the prevalence of days with maximum temperature above 35 °C is about seven times larger than average in the country with the fastest trend, and is negative in some countries.

country-specific trends.⁴

Our specification is not totally immune to bias from changes in the slope of trends. To handle time-varying trends, [Kahn et al. \(2021\)](#) subtract the 30-year moving average from each climate variable and take first differences. While effective, this would be very costly for us because our weather data starts in 1979 unlike their data that starts in 1960. In practice, this does not seem to be a major problem because many variables (Table B.1 in the online appendix) and especially those selected for our main specification do not show unambiguous evidence of a significant break in the linear trend over the panel years.

2.2. Dynamic effects of weather shocks: the local projection method

The complex dynamic effect of weather on GDP might not be immediately revealed by the estimation results from equation (2). There might be persistent weather effects because weather shocks themselves are persistent, because of feedback effects if current GDP per capita depends on past GDP levels, or because of a combination of both.

We use a local projection method following [Jorda \(2005\)](#) to estimate impulse response functions from a shock to one or more of our independent weather variables. As shown in Jorda’s seminal paper, this procedure is more robust to misspecification than auto-regressions, easily accommodates flexible specifications, and allows for a simple visualization of the dynamic responses to weather shocks. We estimate variants of equation (2) where the dependent variables are long-differences between GDP per capita between time $t + h$ and time t :

$$\ln y_{i,t+h} - \ln y_{i,t-1} = \sum_{k=0}^{K-1} \alpha_{k,i}^h t^k + \sum_{l=1}^L \theta_l^h \Delta \ln y_{i,t-l} + \sum_{p=0}^P \beta_p^{h'} \Delta \mathbf{X}_{i,t-p} + \gamma^{h'} \Delta \mathbf{Z}_t + \epsilon_{i,t}^h \quad (3)$$

where h indexes the estimation horizon measured in years. Equation (2) corresponds to horizon 0 and coefficient estimates $\beta^{0'}$ capture the contemporaneous effects of weather shocks. We consider dynamics up to horizon 7 as in [Acevedo et al. \(2020\)](#). In case of a persistent effect on GDP growth rates, the coefficient $\beta_p^{h'}$ would be expected to become increasingly large in absolute value as time goes by. If instead the effect of the shock is mean-reverting, $\beta_p^{h'}$ would be expected to converge to 0 as time goes by. Coefficients that remain largely unchanged over time are a sign that, after the initial shock, growth goes back to normal but GDP levels remain persistently affected.

2.3. Selecting relevant weather variables and estimating their effects

Our most important contribution to the literature is to study the effect of a wide set of climate variables. In total, we include 160 weather variables in \mathbf{X} as described in Section 3. Most of these variables have

⁴To see this more clearly, consider a trended variable x evolving as $x_{i,t} = \theta_i t + v_{i,t}$, where θ_i is a time-invariant trend for country i , and $v_{i,t}$ is a random component with zero mean. The first difference is $\Delta x_{i,t} = \theta_i + \Delta v_{i,t}$. The panel average of first differences is $\overline{\Delta x_{i,t}} = 1/(T-1) \sum_t \Delta x_{i,t} = \theta_i + 1/(T-1) \sum_t \Delta v_{i,t}$. The joint use of first differences and fixed effects removes the trend from all weather variables, as $\Delta x_{i,t} - \overline{\Delta x_{i,t}} = \Delta v_{i,t} - \overline{\Delta v_{i,t}}$.

never been studied before. Counting lagged values, we examine approximately 500 variables.

If all the variables that we consider were entered simultaneously in equation (2), the model might still be estimated thanks to our large panel, but estimation would easily run into over-fitting issues. To avoid this problem, we use the Least Absolute Shrinkage and Selection Operator (LASSO) (Tibshirani, 1996) in a process where machine learning (ML) and expert judgement concur in selecting a parsimonious number of relevant variables. As a robustness check, we consider another operator that addresses over-fitting issues, the Elastic Net. It has the advantage of nesting the LASSO under some specific parameterization and can potentially better handle environments where the variables to be selected are highly cross-correlated.

The LASSO selects coefficients to minimize the sum of squared errors in equation (2) plus a weighted penalty term equal to the sum of the absolute value of each coefficient. The weight attributed to the penalty term is a hyper-parameter λ that needs to be selected before minimizing the loss function. Specifically, the LASSO solves the following problem:

$$\min_{\boldsymbol{\theta}, \boldsymbol{\beta}} L(\boldsymbol{\theta}, \boldsymbol{\beta}) + \lambda(\|\boldsymbol{\theta}\|_1 + \|\boldsymbol{\beta}\|_1), \quad (4)$$

where

$$L(\boldsymbol{\theta}, \boldsymbol{\beta}) = \left(\Delta \ln y_{i,t} - \sum_{k=0}^{K-1} \alpha_{k,i} t^k - \sum_{l=1}^L \theta_l \Delta \ln y_{i,t-l} - \sum_{p=0}^P \boldsymbol{\beta}'_p \Delta \mathbf{X}_{i,t-p} - \gamma' \Delta \mathbf{Z}_{i,t} \right)^2$$

$$\lambda(\|\boldsymbol{\theta}\|_1 + \|\boldsymbol{\beta}\|_1) = \lambda \left(\sum_l |\theta_l| + \sum_{j,p} |\beta_{j,p}| \right)$$

Intuitively, the LASSO chooses coefficient estimates by comparing benefits measured by a reduction in the sum of squared errors in equation (2) with costs measured by the size of non-zero coefficients. When a coefficient $\beta_{j,p}$ is shrunk to zero, the variable is effectively omitted from the regression. The penalty term measures the costs associated with having a model with too many variables. The larger is λ , the smaller are coefficient estimates and the smaller the number of variables selected.

Before implementing the LASSO operator, we follow Belloni et al. (2014) and we impose that the model must use country fixed effects and, in some specifications, year fixed effects or country quadratic trends. We do so in two stages, first by regressing all the dependent and independent variables on the selected fixed effects and trends, and second by applying the LASSO to the estimated residuals from the first stage (the “partialed-out” variables).

We choose the value of the hyperparameter λ that optimizes a fit criteria estimated using the variables resulting from implementing the LASSO. There is no universal “optimal” way to choose the fit criteria (and λ), and optimality conditions must be set by the analyst. In the ML literature, this is known as the “no free lunch theorem”: there is no optimization algorithm that is capable of guiding the identification

of a prior for the penalty weight when starting the analysis (Adebayo and Fokoue, 2019). Section B.6 in the online appendix has more details about implementing the optimization process.

We consider a range of fit criteria. We follow the guidance in Zou et al. (2007, 10-11) and focus primarily on results obtained with the Bayesian Information Criterion (BIC) because our main concern is the interpretability, and therefore the sparsity, of the model to be selected. Additionally, we systematically examine the within R-squared that we calculate on out-of-sample data and the Akaike information criterion (AIC).

When using the out-of-sample within R-squared, we compute it in a way to address well-known limitations. We use R-squared measures because they relate to the share of the variance of the dependent variable that can be explained by independent variables and are therefore easy to interpret. However, the R-squared gets mechanically larger as the number of independent variables increases. To address this issue, we first focus on the within R-squared, which is the standard R-squared measure obtained after removing the estimated fixed effects from all other variables.

Additionally, we compute the within R-squared on out-of-sample data. We randomly divide the dataset into 5 equally-sized subsets, where the model is trained on the union of 4 sets (training set) and the remaining set serves as the evaluation set. We compute the within R-squared using the coefficient estimates obtained on the training set and using the data from the evaluation set. This process is iterated 5 times (each iteration corresponds to selecting one of the five sets as the evaluation set) and the final criterion is computed as the average within R-squared across iterations. Furthermore, we consider two different methods for splitting the data into different sets. In the first method, we ignore the panel structure of the data and assign each observation to one of the 5 sets randomly. Because ignoring serial correlation typically leads to a downward bias in the determination of the hyper-parameter (and the selection of too many climate variables), we experiment by randomly assigning each country with all its observations to one of the 5 sets. This way, we preserve the time dimension in each set. In a majority of cases, we confirm that the second assignment method leads to smaller selections. However, the two methods produce similar results and, in what follows, we choose to focus on the results from the second method for conciseness.

The LASSO produces “biased” (or “regularized”) coefficient estimates because the penalty term shrinks them. To estimate the “unbiased” effect of weather shocks, we finally re-estimate the model with the climate variables selected by the LASSO using standard OLS methods. For theoretical justification, see Belloni and Chernozhukov (2013).

We additionally implement the Elastic-Net (EN) as a robustness check. The EN can outperform the LASSO and encourages a grouping effect, where strongly correlated predictors tend to be in or out of the model together. Compared with the LASSO, the EN’s objective is to minimize an equation with an additional term penalizing non-zero coefficients:

$$\min_{\boldsymbol{\theta}, \boldsymbol{\beta}} L(\boldsymbol{\theta}, \boldsymbol{\beta}) + \phi\lambda(\|\boldsymbol{\theta}\|_1 + \|\boldsymbol{\beta}\|_1) + (1 - \phi)\lambda(\|\boldsymbol{\theta}\|_2 + \|\boldsymbol{\beta}\|_2), \quad (5)$$

where

$$(1 - \phi)\lambda(\|\boldsymbol{\theta}\|_2 + \|\boldsymbol{\beta}\|_2) = (1 - \phi)\lambda\left(\sum_l \theta_l^2 + \sum_{j,p} \beta_{j,p}^2\right) \quad \text{and} \quad 0 \leq \phi \leq 1$$

The new third and last term in equation (5) penalizes the sum of the squares of coefficients. Compared to the other penalty term already also present in the LASSO, the sum of squared coefficient penalizes high coefficients more and very small coefficients less. Therefore, the EN is less likely to shrink coefficients all the way to zero and can retain correlated variables in the model. The parameter ϕ in equation (5) governs the balance between the two penalty terms: as the value of ϕ increases, the variable selection performed by both LASSO and EN becomes more similar. As with the LASSO implementation, we choose the hyper-parameters ϕ and λ that maximize the same fit criteria.

3. Data

3.1. Weather data sources and aggregation over time and space

We start from hourly gridded reanalysis temperature and precipitation data from 1979 to 2019 from the ERA5 dataset compiled by the European Centre for Medium-Range Weather Forecasts (Hersbach et al., 2018).⁵ The grid resolution varies with latitude with cells of 30×30 km at most (at the equator).⁶

Weather data is available at a much higher spatial and temporal resolution than typical country macroeconomic data. Our goal is to reduce the millions of weather measurements in every country and year to construct a manageable number of potentially meaningful climate variables. We aggregate raw ERA5 weather data over space and time to construct annual country-level variables using the cloud computing power of Google Earth Engine (GEE) (Gorelick et al., 2017).⁷

We can use the high spatial and temporal granularity of temperature and precipitation data to reveal weather events that would be lost when averaging weather variables over an entire country during a whole year. Country and year averages may bias estimates of weather impacts in at least three important ways. First, averages miss local and infra-year extreme weather events if events of opposite nature cancel out each other. For example, droughts in a specific region in summer can coincide with intense precipitation

⁵Reanalysis data is generated using models that combine a variety of weather observations and past short-term weather forecasts from different datasets (e.g., weather stations, satellites, ocean gauges, weather balloons) to remove biases in measurement and to create a coherent, long-term record of past weather into one regularly spaced grid. For more details, see <https://www.ecmwf.int/en/about/media-centre/focus/2020/fact-sheet-reanalysis> and <https://cds.climate.copernicus.eu/cdsapp#!/home>.

⁶See Section B.1 in the online appendix for more details.

⁷See https://developers.google.com/earth-engine/datasets/catalog/ECMWF_ERA5_DAILY.

in another part of the country or in a different season. Country annual averages would then be unable to reflect these extreme events. Second, the relationship between a weather shock and the economy can be non-linear and dependent on duration, spatial coverage, and intensity. For example, both prolonged high temperatures (a heatwave) and short-lived but very high temperature could impact the economy in different ways. Even if average temperature might approximately reflect the occurrence of various hot weather events, it would fail to capture the different effects associated with the different characteristics of these events. Third, with high-frequency data it is possible to calculate deviations from normal seasonal weather that may not be reflected in averages or as outliers in the full annual distribution. For example, unusually high precipitation in central Europe in summer could have an impact on the economy even if the same level of precipitation would be totally normal and irrelevant in another time of the year or in other regions that are more accustomed to heavy precipitation.

We build variables that describe the distribution of temperature and precipitation as well as notable extreme events following the climate literature (Kim et al., 2020; Perkins and Alexander, 2013). When the literature uses similar alternatives, we include all of them and let the LASSO or the EN select the most relevant option. While we could have used ML techniques to reduce the full matrix of weather measurements into country-year variables, we choose to start from definitions of weather events that are frequently used in the climate literature to obtain results that are easier to interpret and can be linked to other work.

We construct both unweighted and population-weighted country-level variables.⁸ Both weighted and unweighted variables have been used in the literature with convincing arguments, but it is ultimately an empirical question which one explains better GDP. For this reason we include both sets of variables in the LASSO exercise, an opportunity to illustrate how ML methods can be used to address practical questions as the choice of weighting schemes.

3.2. Main variable definitions

Defining extreme weather. A practical problem for empirical research is the lack of unambiguous definitions of extreme weather. For this reason we use many different ways to characterize the full distribution of temperature and precipitations. Table A.1 and the online appendix B.2 provide a complete description of all variables, exact formulas, and summary statistics for the selected variables. In general, a weather event is extreme if some of its characteristics exceed some thresholds. Specifically, definitions are ambiguous about the thresholds to use with respect to major characteristics, like intensity, frequency, or duration.

The literature has used both “relative” and “absolute” thresholds. In some cases, extreme weather is defined as weather “that is rare at a particular place and/or time of year” (Cubasch et al., 2013, 134).

⁸We obtain grid cell level population data by using the Socioeconomic Data and Application Center’s UN WPP-Adjusted Population count dataset. See <https://sedac.ciesin.columbia.edu/data/set/gpw-v4-population-count-adjusted-to-2015-unwpp-country-totals-rev11> for a detailed description. We use fixed population weights – population in year 2000 – to avoid introducing trends.

This suggests the use of thresholds that are specific to locations and seasons (“relative thresholds”). For example, definitions can rely on the 90th percentile of the local distribution of temperature at a certain time of the year. Relative thresholds account for the importance of adaptation to average conditions and emphasize the effects of deviations from local averages. In other cases, extreme weather is defined using thresholds that are constant across space and time (“absolute thresholds”). For example, maximum daily temperature greater than 40 °C are generally considered harmful everywhere. Absolute thresholds are better suited to capture physical limits beyond which weather causes damages, no matter when or where it occurs (e.g., [IPCC, 2021a](#)). We consider both “relative” and “absolute” weather extremes as they can both be relevant for the economy in different ways.

We typically capture extreme events at the country-year level by both counting their occurrences and measuring both their prevalence and intensity. To this end, we count the number of times in which a weather event is observed in each grid-cell and each day of a year. We then divide by the total number of grid-days to calculate an index from 0 to 1 that measures the prevalence of the event over the entire country and the entire year. We also construct a wide range of variables that characterize the intensity of extreme weather over time and space. We do so for temperature, precipitation and wetness/drought as detailed below.

Temperature variables. We consider average temperature, the variance of daily temperature, and the average diurnal temperature range (the difference between the minimum and maximum temperature in a day). We calculate the number of cold nights, cold days, warm nights and warm days using relative thresholds based on the 1979-2019 distributions for every 5-day window centered on each day of the year.

We build various heatwave and coldwave variables based on the climate literature.⁹ We follow [Perkins and Alexander \(2013\)](#) and consider various thresholds to define heat and cold waves in daytime and nighttime. We then count the length of the longest wave, the number waves in a year, the number of days and the average maximum or minimum temperature during such waves. We also follow [Kim et al. \(2020\)](#) and additionally define the duration of cold and warm spells as the number of days exceeding alternative relative temperature thresholds for prolonged periods.

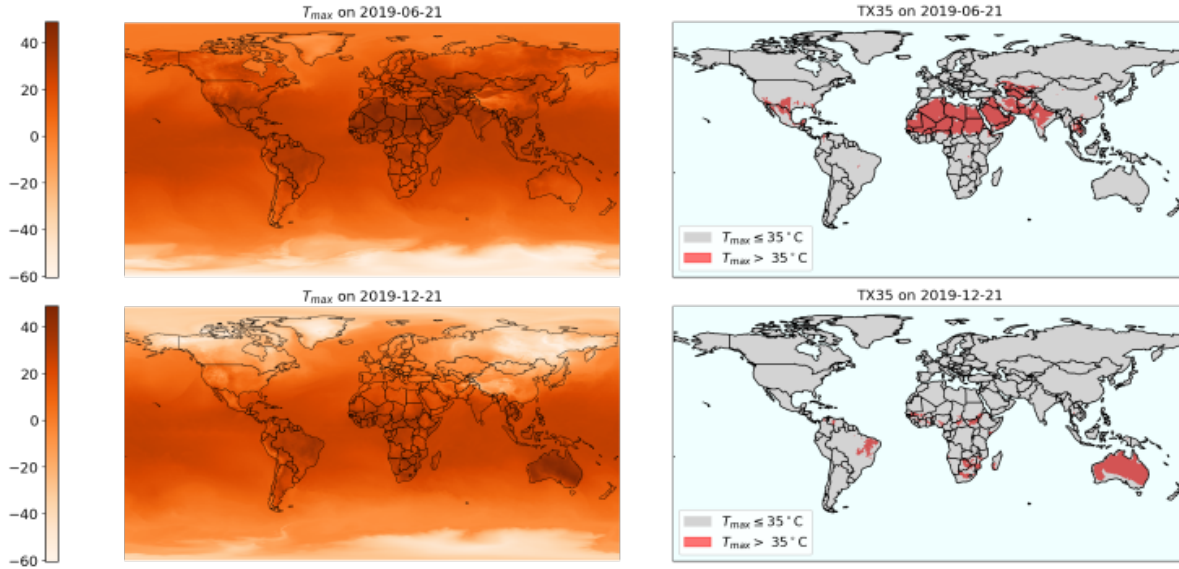
We use the fine spatial resolution of our data to define minimum and maximum variables that are used in the literature to capture local extremes. We compute the annual minimum of night temperatures and the maximum of daytime temperatures for every grid-cell in a country, and average them out over space.

We also define another set of extreme temperature variables using absolute temperature thresholds often used in the climate literature (e.g., [IPCC, 2021a](#)). With absolute thresholds, using the highest possible level of spatial resolution is essential to avoid missing the potentially meaningful events that would otherwise be averaged out. For example, Figure 1 illustrates how country averages would miss many instances

⁹“Heatwaves” and “coldwaves” are loosely defined as prolonged periods with unusually warm or cold temperatures (e.g., [Perkins, 2015](#)).

of days with temperature $35\text{ }^{\circ}\text{C}$ in only parts of a country. To avoid this, we calculate the prevalence of extreme temperature events using different thresholds (e.g., below $0\text{ }^{\circ}\text{C}$, above $35\text{ }^{\circ}\text{C}$ and $40\text{ }^{\circ}\text{C}$) using daily and grid-cell data.

Figure 1: Illustrating the role of high spatial resolution when using absolute thresholds



Notes: This figure illustrates the importance of high spatial resolution when accounting for daily maximum temperatures exceeding $35\text{ }^{\circ}\text{C}$ (TX35). As seen in the top row, at the beginning of summer 2019, only a small share of the US (8%) experienced temperatures higher than $35\text{ }^{\circ}\text{C}$. These temperatures would average out if we were to use country means. Similarly for Brazil in December 2019, the bottom row shows that only 12% of the country crosses the $35\text{ }^{\circ}\text{C}$ threshold. In both cases, country averages would fail to capture these extreme temperatures.

Finally, to capture potential non-linear effects of temperature on macroeconomic variables, we define $3\text{ }^{\circ}\text{C}$ -wide intervals from $-9\text{ }^{\circ}\text{C}$ and below to $30\text{ }^{\circ}\text{C}$ and above and we count how often temperature falls in these intervals over space and time (see for example [Schlenker and Roberts, 2009](#)). This approach allows us to capture the impact of temperature on macro-fiscal variables at different temperature levels imposing minimal restrictions on the temperature response functional form.

Precipitation variables. We use the term *precipitation* throughout this paper because our data measures both rain and snow precipitations (converted into rain equivalents). We sometimes focus on “wet days”, that are days with 1 mm precipitation or more, or on “dry days” with precipitation below 1 mm. Our variable set includes the country-year averages and the variance of daily precipitation, which we construct twice, on all calendar days and on wet days. We also measure precipitation on very wet and extremely wet days, where these days are defined using relative thresholds.

We build several variables to capture extended wet and dry periods. We count the largest number of consecutive dry days, wet days, very wet days, and extremely wet days. We measure precipitation in the longest period of wet, very wet and extremely wet days respectively.

Floods are among the most destructive climate disasters. To capture short but intense precipitation that may cause floods, we use the maximum amount in a year of rainfall in 1-day or 5-day intervals. To capture extreme precipitation at the local level, we also examine total monthly precipitation in each grid cell. We use these to calculate the country average of maximum and minimum monthly precipitation. These intense precipitation indicators are good proxies for tropical cyclones because they capture extreme rainfall often observed during these events. For example, plots of maximum 5-day rainfall over the Caribbean and Southern United States in 2004 reveal the track of category 5 Hurricane Ivan.

As for temperature, we make use of our data high spatial resolution to define precipitation extremes using absolute thresholds. We calculate the number of consecutive days in which a minimum percentage of the country area is experiencing a dry day using different percentage thresholds. Similarly to what we do with temperature, we define four precipitation intervals (divided by 1, 10, and 20 mm thresholds), and measure how often precipitation is in any of these intervals. We define the maximum extent of heavy and very heavy precipitation as the maximum surface of a country where precipitation exceeds 10 mm and 20 mm respectively. We also construct an indicator that measures deviations from a balanced level of precipitation. This indicator measures the absolute deviation from having precipitation between 1 and 10 mm half the time over space and time.

Wetness and drought variables. We use the Palmer Drought Severity Index (PDSI) (Palmer, 1965) to introduce a measure of dry and wet periods that combines temperature and precipitation data to estimate cumulative deviations in soil moisture from normal conditions (Dai et al., 2004; Abatzoglou et al., 2018; Lai et al., 2020).¹⁰ The PDSI ranges from -10 to +10, but values below -4 and above +4 are very rare. To capture extreme conditions during a year we build variables measuring the share of total grid-months subject to droughts and harsh droughts (with PDSI respectively below -3 and -4), and to periods with high and very high moisture (with PDSI respectively above 3 and 4). As for precipitation, we also seek to capture the maximum geographical extent of droughts and wet conditions. For each of the four categories, we compute the share of affected grid-cells in the month where the share is at its maximum.

In sum, in our empirical analysis, we consider 45 different temperature variables, 29 precipitation variables, and 8 wetness-drought variables, for a total of 82 unique climate variables. We add the first and second lag of these variables as well as their population-weighted counterparts. In total and after removing perfectly collinear variables, we use a set of 480 climate variables.

Macro-fiscal variables. We use GDP per capita from the World Bank’s World Development Indicators (WDI).¹¹ For fiscal outcomes, we collect variables from the IMF World Economic Outlook (WEO) because it has a wider coverage than the WDI. We use government revenue and expenditure expressed

¹⁰Data downloaded from Google Earth Engine. See <http://www.climatologylab.org/terraclimate.html> and https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE for a detailed description of the datasets.

¹¹Specifically, we use the variable “GDP per Capita constant 2015 US\$” (NY.GDP.PCAP.KD).

in percentage of GDP, also from the WEO database.¹² To examine gross debt, we draw from the Global Debt Database which improves on other databases by constructing long series with a consistent coverage over time (Mbaye et al., 2018).¹³ For some countries (all countries in the case of debt), the fiscal variables only cover the central government, implying that our analysis will miss local governments’ response.

3.3. Summary statistics

We start our analysis by constructing estimation samples that are suited for our analysis. Mainly, we want to avoid that the selection of variables is overly influenced by a few outlier values in the dependent variable. Therefore, we exclude all the observations that have values more than 5 standard deviations away from mean GDP per capita growth. This leads us to remove 0.4 percent of the total initial observations and to use a sample with 203 countries and 6,653 observations.

We use subsamples for robustness checks and further investigations. When we consider the balanced sub-sample of countries with non-missing non-outlier observations from 1984 to 2019, we use a sub-sample with 129 countries and 4,644 observations. When we consider fiscal variables, we exclude observations with values more than 5 standard deviations away from the mean for any of the seven variables of interest (government revenue, expenditure, and debt, both in log and as ratios to GDP, and GDP per capita growth). As a result, we exclude 80 observations (2 percent of the sample with non-missing fiscal values) and use a sub-sample with 165 countries and 3,890 observations.

Our empirical and identification approach relies on inter-annual variation within country. Therefore, we use a standard approach to decompose the variance of variables into *between* and *within* components (see online appendix Section B.3). The *between* standard deviation measures variation of average country weather around the global mean. The *within* standard deviation measures the average deviation from country averages.

GDP per capita grew by 1.6 percent per year on average in our largest sample used for GDP analysis and by 2.0 percent per year on average in the smaller sample used for fiscal policy analysis (Table B.2 in the online appendix). The within standard deviation of GDP per capita growth is large, ranging from 4.7 (larger sample) to 3.7 (smaller sample) percentage points. The ratios of government revenue and expenditure to GDP grew at the same average rate of 0.1 percentage points per year. The inter-annual variation of these variables is substantial, with *within* standard deviations ranging from 3.2 to 4.0 percentage points. Government debt-to-GDP ratios were stable on average but with a large *within* standard deviation of 11.1 percentage points.

Many weather variables exhibit a trend over the sample period. Table B.1 in the online appendix reports

¹²Specifically, we use the variables GGR_NGDP and GGX_NGDP for the general government.

¹³We use “Central government debt, % of GDP” rather than “General government debt, % of GDP” because of wider coverage.

results from a systematic analysis of trends in all weather variables in all countries from 1979 to 2019. We find evidence of positive and statistically significant country-specific trends in a majority of countries for variables related to temperatures. Such trends are more rare for precipitation variables.

In our model specification, we assume that trends are time-invariant. To check the validity of this assumption, we test for a structural break in trends with unknown break date for all variables in every country. For most variables in most countries, we cannot reject the null hypothesis that there are no significant structural breaks. We conclude that the assumption of time-invariant trends is acceptable. We also test all first differences of weather variables for the presence of a unit root and we reject it in all cases with p-values close to zero.

4. GDP Results

We start by describing the selection of the climate variables using the LASSO. We then study the effect of the selected variables on GDP growth by applying alternative restrictions to the general model illustrated in equation (2). Our preferred specification is parsimonious and uses the first two lags of the dependent variable, country fixed effects ($k = 0$ in equation (2)) and controls \mathbf{Z}_t that only include a vector of year dummies. We test alternative models with (1) global GDP instead of year dummies, and (2) an additional quadratic time trend by country.


4.1. Climate variable selection

In our baseline specification, we use country and year fixed effects. After removing perfectly collinear variables, the collinearity test we derived from [Belloni et al. \(2014\)](#) does not suggest dropping any of the remaining 480 weather variables. The random search process based on the BIC selects 5 out of these variables, with $\lambda = 0.033$. The selected variables include the first two lags of GDP per capita growth and three climate variables.


Panel a in Figure 2 shows how the BIC varies when the number of selected variables decrease as we increase λ . We see a rapid improvement (decline) in the BIC as λ increases from low values and as many climate variables are dropped. Once we reach a selection of 5 variables, each incremental increase in λ and the associated exclusion of additional climate variables deteriorates (increase) the BIC substantially. As we increase λ from its optimal value, the smaller selections are strictly nested within the optimal 5-variable selection, meaning that no new climate variables are introduced. This result is not pre-determined because the algorithm does not impose that smaller selections are nested within larger ones.

Implementing the random search to maximize criteria other than the BIC leads to larger selections of variables (Figure 2 panels b and c). With the AIC, we obtain a selection of 36 variables. With the out-of-sample R-squared we obtain selection of 18 variables, which do not change whether we split observations in the sample by country or we split the sample by drawing observations randomly and ignoring the panel


Figure 2: Selection of climate variables impacting GDP (baseline FE specification)
 (a) Selection using the BIC



(b) Selection using the AIC



(c) Selection using the out-of-sample within R2



--- AIC (LHS) - - - BIC (LHS) Out of Sample Within R-Squared (LHS) — # of Selected Variables (RHS)

Note: The figures show the results of implementing the LASSO for different penalty parameters λ . The red lines are similar in every panel and show how the number of selected variables vary with λ . The grey dashed lines in each panel show the variation of different criteria with λ . The within R-squared is calculated on a sub-sample of countries (evaluation set) based on coefficients estimated on the rest of countries (training set) as explained in the main text. The estimated model has GDP per capita growth as the dependent variable and includes country and year effects. The dots indicate the different selection outcomes given by the local optimum for each criteria respectively.

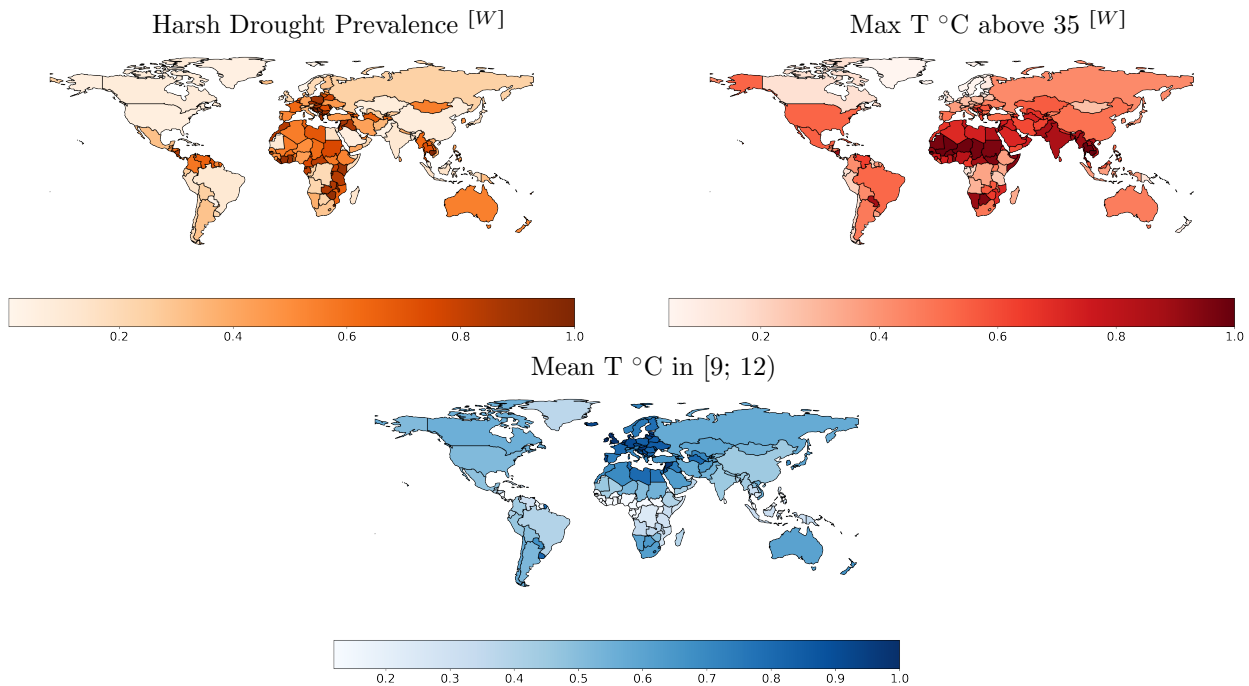
structure. We note that many of the climate variables in these selections are statistically insignificant, indicating that, while they contribute to explaining the variation of the dependent variable, they do not help much to understand the effect of weather on GDP variations (see Tables B.10-B.11 in the online appendix). This lack of significance supports our choice of favoring the selection based on the BIC.

Our preferred selection of climate variables derived using the BIC is robust to the use of alternative cri-

teria. These climate variables are also in the selection based on the AIC and based on the out-of-sample within R-squared. Remarkably, the selections are nested within each other as all variables selected with the BIC are in the set selected with the out-of-sample within R-squared, which is itself in the set selected with the AIC.

We examine the out-of-sample within R-squared to focus on the explanatory power of climate variables. Our preferred specification with 5 variables improves the out-of-sample within R-squared by 29 percent relative to a specification with just the 2 lags of the dependent variable as independent variables, from 0.070 to 0.090. By construction, using out-of-sample within R-squared as selection criterion leads to the largest out-of-sample R-squared, but adding 8 additional weather variables to the 5 variables selected using the BIC only improves the out-of-sample within R-squared by 3 percentage points, from 0.090 to 0.092.

Figure 3: Global distribution of standard deviations of selected weather shocks



Notes: Each panel shows the percentiles of the global distribution of the standard deviation of the first difference of a climate variable over the period 1979-2019. Variable definitions are detailed in appendix in Table A.1. [W] indicates population-weighted variables.

4.2. The effect of weather variables on GDP growth

The three selected weather variables in our preferred specification with country and year effects are the share of grid-days with PDSI less than -4 corresponding to *harsh droughts* (Harsh Drought Prevalence), the share of grid-days with *hot temperature* (Max T °C above 35), and the share of grid-days with *mild temperature* (Mean T °C in [9; 12]). The LASSO selects population-weighted variables except for days with mild temperature. Figure 3 shows the geographic distribution of the country standard deviation of

first differences for three key climate variables by percentile. Summary statistics of first differences of these variables are displayed in Table B.3 and correlation coefficients between GDP growth, the first two lags of GDP growth, and first differences of the selected climate variables are shown in Table B.4 in the online appendix.

The correlation between the first differences and GDP growth is generally very low, ranging from 0.040 to 0.058 in absolute value. However, as a comparison, the correlation between the first difference of average annual temperature and GDP growth is even smaller and equal to only -0.015. Among the selected climate variables, the largest correlation is found between harsh droughts and hot temperatures (0.180). Harsh droughts are also positively correlated with average annual temperature (0.156) because temperature plays a role in the definition of the PDSI drought indicator. Hot temperatures and average annual temperature are highly correlated in comparison with other variables (0.362) but average annual temperature is not retained by the LASSO.

The selected climate variables have an intuitive effect on GDP, as shown in Table 1 where all the variables are standardized using the same standard deviation across all years and countries to facilitate interpretation. We find that harsh droughts and hot temperatures have adverse effects on contemporaneous GDP growth. By contrast, we find that mild temperatures have beneficial effects on GDP. The distributions of temperature by country are typically bell-shaped and the modes of these distributions are above 12 °C in the vast majority of countries. Therefore, a warming shock typically implies fewer mild temperatures and, therefore, GDP losses. All these effects are highly significant.

In our baseline specification (column A) we use the unbalanced panel and year fixed effects. A positive shock equal to one standard deviation of the first difference of harsh droughts in a country reduces GDP growth by 0.23 percentage points. Similarly, an increase in hot temperatures has a large and significant effect on GDP growth. A positive shock equal to one standard deviation in the first difference of hot temperatures reduces GDP growth by 0.17 percentage points. A positive shock equal to one standard deviation in mild temperatures increases the growth rate by 0.15 percentage points.

Our more granular temperature measurements than in other macroeconomic studies reveal the harmful effect of temperatures above 30-33 °C on country GDP that has been widely documented at a more disaggregated level for agricultural output (Schlenker and Roberts, 2009; Blanc and Schlenker, 2017), mortality and energy consumption (Deschênes and Greenstone, 2011), time allocated to labor (Graff Zivin and Neidell, 2014), and labor productivity (Somanathan et al., 2021). As these results remain significant at the macro level they must be large at the sectoral level and/or widespread across sectors. While the micro and macro effects of temperature are not easily comparable due to differences across studies and different coverage of the economy, the impact of harsh droughts is similar to what found by other studies (Cantelmo et al., 2019; IMF, 2020).

The selection of significant weather shocks includes variables constructed with both absolute (hot temper-

Table 1: The effect of changes in selected climate variables on GDP per capita growth

	(A)	(B)	(C)	(D)	(E)	(F)	(G)
First difference in							
Harsh Drought Prevalence ^[W]	-0.229*** (0.0541)	-0.262*** (0.0535)			-0.247*** (0.0531)	-0.207*** (0.0538)	-0.310*** (0.0622)
Max T °C above 35 ^[W]	-0.169** (0.0676)		-0.214*** (0.0668)		-0.156** (0.0635)	-0.174** (0.0679)	-0.210** (0.0923)
Mean T °C in [9; 12]	0.145*** (0.0380)			0.162*** (0.0386)			
PPT Minimum ^[W]							-0.193** (0.0777)
Observations	6,653	6,653	6,653	6,653	6,653	6,653	4,644
Year fixed effects	Yes	Yes	Yes	Yes	No	Yes	Yes
World GDP growth	No	No	No	No	Yes	No	No
Country quadratic trends	No	No	No	No	No	Yes	No
Balanced	No	No	No	No	No	No	Yes
R-squared	0.264	0.261	0.261	0.260	0.253	0.349	0.224
Within R-squared	0.0997	0.0967	0.0965	0.0950	0.150	0.0360	0.0790

Note: The table shows the country fixed-effect estimates of the effect of climate variables on the first difference of log real GDP per capita expressed in constant 2015 USD. All climate variables are standardized with standard deviations equal to 100 to ease interpretation and their definitions are detailed in appendix in Table A.1. [W] indicates population-weighted variables. Columns (A)-(F) show results from the full sample of 203 countries for 1979-2019. All regressions include controls such as one or two lags of the dependent variables. Column E additionally includes world growth as a control. Coefficient estimates of controls are reported in the appendix in Table A.2. Column G is estimated on a balanced subsample of 129 countries for 1984-2019. Standard errors are clustered by country and reported in brackets.

atures and mild temperatures) and relative thresholds (harsh droughts). For droughts, we find that drier than average conditions are harmful, no matter what is the average precipitation level in a country. For temperature, we find that the 35 °C threshold is selected over alternative definitions of high temperatures based on the deviations from local and seasonal norms.

Interestingly, the LASSO selects both population-weighted and area-weighted variables, a new result in the literature. This suggests more complex transmission channels than those previously assumed by studies that use only one weighting scheme. The choice of population weights for harsh droughts and hot temperatures indicates that these are particularly important relatively more populated areas, probably indicating harmful effects of extreme temperature on labor productivity and the importance of abundant water supply near to where people live. More mild temperatures have instead beneficial effects also in areas with low population density.

Do our estimates identify the causal effects of weather on GDP? Our econometric strategy removes important potential sources of endogeneity. First differences and country fixed effects capture time-invariant or very slow-moving variables that likely explain average growth rates, like the quality of institutions, sectoral composition, market structures, geography, and climate itself. We also do not have reasons to believe that our specification suffers from reverse causality because annual and infra-annual weather shocks are largely exogenous with respect to contemporaneous economic activity. However, the use of the LASSO or the EN for variable selection does not guarantee causal interpretation of our results.

The LASSO and the Elastic-Net allow us to select a small number of variables that are the best *proxies* for the weather phenomena that impact GDP even as these phenomena preclude easy definition and measurement. Some additional weather variables that would refine the characterization of these phenomena may not be selected, either because their explanatory power is too small in our sample or because they are not part of our initial set of climate variables.

Our strategy works to maximize the fit of the model with a parsimonious number of variables, but the theoretical weather phenomena of interest are measured with error and causal identification is compromised. Our variables capture the effect of the underlying relevant (but unknown and imperfectly measured) weather phenomena on macro-fiscal variables. When discussing results we use “hot temperatures”, “harsh droughts”, and “mild temperatures” instead of the exact definition of the selected variables to emphasize the underlying weather phenomenon and de-emphasize a too strict literal interpretation of our selected variables.

4.3. Robustness of the Selected Weather Shocks

Alternative fit criteria. The estimated effects of our selected weather shocks remain broadly unchanged when including in the OLS regression the variables that are additionally selected when using the AIC and the out-of-sample within R-squared criteria. Some of these variables are correlated with the three climate variables from our main selection. The estimated effects of our three selected variables tend to be smaller but remain significant (Tables B.10 and B.11 in the online appendix).

Separability of weather effects. We start by examining whether our selected climate variables have an effect on their own or whether their estimated effects result from their combination. To this end, we enter each variable separately in our specification with year and country effects as reported in columns B-D of Table 1. Coefficient estimates are similar to those in the baseline column A, suggesting that the main effect of each variable is mostly independent and additional to the effect of the other variables.

Alternative Model Specifications. We confirm the robustness of our variable selection and coefficient estimates by considering alternative model specifications. For each model we repeat the LASSO exercise with the BIC as for our baseline specification. In column E, we drop year fixed effects but we introduce world GDP per capita growth as a candidate variable that can potentially be selected by the LASSO to control for common world-wide developments. Under this specification, the LASSO leads to the selection of five variables (Figure B.2) which indeed includes world GDP per capita growth (Table B.6). The selection includes the two lags of the dependent variable as well as harsh droughts and hot temperatures. It does not include mild temperatures. Column E of Table 1 shows that the coefficients of the two climate variables used in our baseline specification are robust to this new specification.

We additionally consider a specification with country-specific quadratic trends in addition to year effects. The LASSO selects five variables (Figure B.3), including also in this case the same two climate variables

selected in our baseline but excluding mild temperatures and the second lag of the dependent variable (Table B.6). Column F of Table 1 shows that the coefficient estimates from the OLS estimation of the climate variables that are common to the baseline and to this new specification are, once again, very similar.

We also implement our selection method on a balanced subsample with the 129 countries that have non-missing observations continuously from 1984 to 2018. The LASSO selects five variables (Figure B.4), including the same two climate variables consistently selected and a new climate variable, the weighted average over space of precipitation in the driest month, along with two lags of the dependent variable (Table B.6). Results in column G in Table 1 shows that the effects of harsh droughts and of hot temperatures remain robust. The new climate variable is estimated to have a negative GDP effect, implying that an increase in precipitation in the driest months is detrimental.

Elastic-Net instead of LASSO. Our results are strikingly consistent when we use the EN instead of the LASSO for all the model specifications in Table 1 and for all the alternative fit criteria. The optimization process using the richer penalty term in equation (5) results in choosing parameter values that put more weight on the penalty term of the LASSO. As a result, when we select variables to maximize the BIC, we obtain the exact same variable selections for all model specifications, except in the case without fixed effects where the only difference is that hot temperatures are dropped.

Further, Figure A.1 in the appendix shows the results of the optimization and selection process using our baseline specification with country and year effects when we experiment with all the alternative fit criteria. Again, we obtain consistent findings: in addition to the five variables selected with the BIC, about a dozen more climate variables are selected with the AIC and the out-of-sample criteria and our main three climate variables are always included. Additional selection results using the alternative fit criteria are reported in the online appendix in Table B.7.

Inclusion of additional explanatory variables. We confirm the magnitude and relevance of the selected weather shocks under richer specifications with more explanatory variables from the macroeconomic literature (see Table A.3 in the appendix). The magnitude and statistical significance of our three climate variables remain broadly unchanged if we control for additional lags of the dependent variables (the third and fourth lags of GDP per capita growth) and lagged climate variables (first and second lags). Among lagged climate variables, only the lagged value of mild temperatures is significant, and the coefficient estimate suggests that the positive effect of this shock unfolds gradually. We additionally introduce controls that are known to be associated with GDP variations and two of their lags. The estimated effect of violent conflicts (captured by a dummy variable and its lags from the Uppsala Conflict Dataset v23.1 (Gleditsch et al., 2002)) are significant but do not alter our results. In our sample, inflation dynamics (from the World Bank World Development Indicators — WDI) is estimated to be insignificant and has no noticeable effects on our results. Exchange rate dynamics (also from the WDI) and terms of trade dynamics (from Gruss and Kebhaj, 2019) are estimated to have significant effects but leave the estimated effects of our climate variables essentially unchanged.

Alternative estimation choices. Table A.4 in the appendix shows that our results are robust to the use of various alternative estimation choices and estimators. Estimation results after re-introducing outliers are quantitatively very close to those obtained after excluding outliers. There is one exception, as the effect of mild temperatures is estimated to be only half as large and becomes insignificant. Finally, we address concerns about biases arising from measurement errors and the endogeneity of lagged variables. Our main first-differences specification is exposed to the risk that lags of the dependent variable are correlated with the unobserved fixed effects, making standard estimators inconsistent. Specifically, we implement the Arellano-Bond estimator that uses further lags of the dependent variable to instrument the lagged dependent variables. It requires that there is no autocorrelation in the idiosyncratic errors. Because the result of an Arellano-Bond test for autocorrelation is a borderline case, we also employ the Arellano-Bover/Blundell-Bond estimator that allows for idiosyncratic errors that follow a first-order moving average process and only use late lags as instruments. Further, since these GMM methods can suffer from weak/many instruments problem, particularly in cases where the time dimension is moderately long as in our case, we additionally use the half-panel jackknife FE-TE estimator showing little size distortions proposed by [Chudik et al. \(2016\)](#) and used in [Kahn et al. \(2021\)](#). Overall, our results remain essentially unchanged.

4.4. Comparisons with the empirical macro literature

Does the set of variables we select improve substantially our understanding of GDP variations? We examine this question by comparing our results with two central papers in the literature, [Burke et al. \(2015\)](#) and [Kahn et al. \(2021\)](#).

We estimate the two papers' respective baseline models using our sample and confirm their robustness.¹⁴ We follow the specification in [Burke et al. \(2015\)](#) and regress GDP growth on annual average temperature, precipitation, the two squares of these variables, and include two lags of GDP growth, country quadratic trends and year effects. Despite the fact that our sample is much smaller and other minor differences, we obtain very similar coefficient estimates.¹⁵ Results in column B in Table 2 feature an inverted U-shaped relationship between temperature and growth that is quantitatively close to that found by [Burke et al. \(2015\)](#), with optimal temperature estimated to be equal to 13.3 °C (instead of 13.1 °C).

We compare the performance of our approach with the model in [Burke et al. \(2015\)](#) by introducing climate variables sequentially. Relative to a specification without climate variables (column A in Table 2), the introduction of annual average temperature and precipitation and their squares improve the out-of-sample within R-squared by a factor 5.¹⁶ If we additionally include our four selected climate variables,

¹⁴Both papers examine real GDP per capita growth. We use our GDP variable to facilitate comparisons across specifications. For climate variables instead, we use the variables provided in the replication package of both papers.

¹⁵Our estimation sample starts in 1979 instead of 1960, and is 40 percent smaller with only 3,935 observations.

¹⁶The out-of-sample within R-squared increases from 0.001 to 0.005. It is small because of the large set of fixed effects and country-specific quadratic trends which absorb a lot of variations.

Table 2: Estimation of the effect of climate shocks on GDP growth: comparisons with the literature

	Burke et al. (2015)			Kahn et al. (2021)		
	(A) base	(B) unchanged	(C) augmented	(D) base	(E) unchanged	(F) augmented
Average Annual Temperature ^[W]		0.00909*** (0.00305)	0.00792** (0.00308)			
– squared		-0.000374*** (0.0000865)	-0.000276*** (0.0000921)			
Average Annual Precipitation ^[W]		0.0000186* (0.00000981)	0.00000275 (0.00000904)			
– squared		-5.97e-09** (2.52e-09)	-3.27e-09 (2.27e-09)			
Harsh Drought Prevalence ^[W]			-0.313*** (0.0777)			-0.294*** (0.0654)
Max T °C above 35 ^[W]			-0.200** (0.0979)			-0.228*** (0.0694)
Mean T °C in [9; 12]			0.198*** (0.0503)			0.197*** (0.0593)
Temperature Deviations from Trend ^[W]					-0.0275 (0.0264)	0.00543 (0.0272)
Precipitation Deviations from Trend ^[W]					-0.0757 (0.0663)	-0.0620 (0.0661)
Country fixed effects	yes	yes	yes	yes	yes	yes
Country quadratic trends	yes	yes	yes	no	no	no
Year effects	yes	yes	yes	no	no	no
Observations	4,214	4,214	4,214	4,500	4,500	4,500
AIC	-15,549	-15,567	-15,607	-15,141	-15,125	-15,170
BIC	-15,530	-15,523	-15,543	-15,122	-15,029	-15,055
Out-of-sample within R2	0.001	0.005	0.014	0.075	0.073	0.080

Notes: Each column corresponds to a fixed-effect regression of the first difference in log real GDP per capita on lags of the dependent variables and on climate variables for column B, C, E, and F. The within R-squared is calculated using out-of-sample observations as described in the text. The estimations in columns D-F include 4 lags of the deviation variables and are obtained with the half-panel jackknife FE estimator used in Kahn et al. (2021). [W] indicates population-weighted variables. Estimates of the coefficients of lag variables that are omitted in this table are reported in appendix Table A.5.

the within R-squared increases much more, by a factor 13 (column C in Table 2). Similarly, both the AIC and the BIC unambiguously improve with our selected variables, while neither support the introduction of annual averages and their squares on their own.

Turning our attention to Kahn et al. (2021), we adopt the same baseline specification focusing on the absolute deviations of annual average temperatures and precipitations relative to their respective 30-year moving average. Column E in Table 2 reports results that are again extremely similar to those originally reported despite the smaller size of our sample.¹⁷ Their specification also includes four lags for each of these variables that are mostly significant and reported in appendix Table A.1. In this case, the introduction of their climate variables relative to a basic case abstracting from climate fails to improve the out-of-sample within R-squared (columns D-E in Table 2). Furthermore, both the information criteria do not support the relevance of their climate variables. By contrast, additionally including our selected variables improves the out-of-sample within R-squared by 7 percent.¹⁸ In this sample, the relevance of

¹⁷Our estimation sample starts in 1979 instead of 1960, and is 25 percent smaller with only 4,917 observations.

¹⁸Increases in the within R-squares are much smaller compared to the increases in specifications with country trends a la Burke et al. (2015) because the within R-squared in the base case of no climate variable is much smaller

our selected variables is supported by a lower AIC but not by the BIC. Overall, a range of performance indicators suggests that our climate variables are much more relevant in explaining GDP growth variations than annual average temperature and precipitation.

Our findings imply that the magnitude of current weather impacts on GDP growth is larger than what is suggested in [Kahn et al. \(2021\)](#) and [Burke et al. \(2015\)](#). The effects captured by our variables are, at least partially, additional to the effects that are reported in these two papers. This is shown by the fact that the introduction of our weather variables does not fully remove the significance of the variables of these papers.

Furthermore, the effects captured by our variables alone have similar or greater impacts on GDP than previously reported. To compare the results in [Kahn et al. \(2021\)](#) and [Burke et al. \(2015\)](#) with ours, we focus on temperature variables because the effects of their precipitation variables are largely insignificant. We consider the following shocks. In the case of [Kahn et al. \(2021\)](#), a temperature deviation from country trends equal to one standard deviation ($0.029\text{ }^{\circ}\text{C}$) lowers GDP growth by 0.08 percentage points.¹⁹ This is at most half the size of the effects of any of our selected weather shocks. In the case of [Burke et al. \(2015\)](#), we calculate the GDP loss implied by the typical inter-annual variation in temperature. Specifically, we consider a hypothetical country with an annual temperature equal to the global average ($19.13\text{ }^{\circ}\text{C}$) and the variation in its GDP growth implied by an increase in annual temperature that is equal to one standard deviation of the inter-annual change in average temperature ($0.55\text{ }^{\circ}\text{C}$). Using the estimates from our replication exercise (column B), we find that GDP growth would be cut by 0.29 percent.²⁰ This effect is quantitatively comparable to the weather shocks that we identified.

4.5. The persistent effects of weather shocks

We examine GDP effects of weather variables in the years following the initial shock by using the local projection method. For each variable, we estimate the impulse response to a one standard deviation shock as described in Section 2.2. The flexibility of this approach allows us to investigate if a climate shock has a temporary effect, a persistent effect on GDP levels, or a persistent effect on GDP growth.

The distinction between “level” and “growth” effects often used in the literature can be confusing. In practice, a one percent reduction in the growth rate between any two years is equivalent to a one percent “level” loss of GDP.²¹ The precise question we address is whether weather shocks in one year can have persistent impacts, affecting GDP both instantaneously and in future years, by estimating the magnitude

with country trends. Country-specific quadratic trends absorb a large amount of variation, leaving little to be explained in the within R-squared.

¹⁹The loss is given by $-0.0275 * 0.029 = 0.0008$.

²⁰In other words, we calculate the change in GDP growth implied from having an average temperature increasing from $19.13\text{ }^{\circ}\text{C}$ to $19.68\text{ }^{\circ}\text{C}$. The loss is given by $(0.0091 * 19.68 - 0.00037 * 19.68^2) - (0.0091 * 19.13 - 0.00037 * 19.13^2) = 0.0029$. If we use the coefficient estimates reported in their paper, we obtain a loss of similar magnitude (0.37 percent).


²¹Similarly, over a longer time horizon, damage functions that give percentage losses of GDP from gradual climate change used in integrated assessment models (e.g., [Nordhaus 1993](#); [Tol 1997](#); [Bosetti et al. 2006](#); [Barrage and Nordhaus 2023](#)) can be given both a “level” and a “growth” interpretation. For a review of “growth” versus “levels” econometric models of the GDP-temperature relationship see [Newell et al. \(2021\)](#).

of these impacts at different horizons.

The results in Figure 4 show that our selected climate shocks have persistent effects on GDP levels that remain broadly significant and stable over the 7-year horizon we consider. Specifically, a one-year increase in mild temperatures has a positive effect that remains stable and significant throughout the period considered (Panel c). The respective adverse effects of an increase in harsh droughts and in hot temperatures are also stable, albeit subject to slightly more uncertainty (Panels a and b). The effects still remain significant over most of the horizons considered. The results in Figure 4 also show clearly that all effects do not affect GDP growth after the initial impact. If any of these variables had a persistent effect on GDP growth, the absolute value of the coefficient estimates would increase with the horizon considered.

The literature reaches similar empirical results when it tests the persistence of the effect of weather on GDP. Both [Dell et al. \(2012\)](#) and [Burke et al. \(2015\)](#) find that lagged effects of temperature are mostly non significant and with many sign reversals.²² [Kahn et al. \(2021\)](#) find that shocks have a persistent effect on GDP level peaking after a few years.

Figure 4: Persistence of selected weather shocks on GDP per capita



Notes: Each panel depicts the impulse response of per capita output in levels to a one standard deviation shock of the corresponding climate variable. Horizon 0 is the year of the shock. The shaded areas show the 90 percent confidence intervals around estimates. [W] indicates population-weighted variables.

4.6. Heterogeneity

We test if results obtained using the whole sample of countries are different from those obtained using sub-groups of homogeneous countries. This exercise provides both a robustness test and new insights on

²²[Burke et al. \(2015\)](#) examine dynamics effects of temperature at longer horizons with lags, find some evidence of persistence, but conclude that they cannot confirm if GDP growth rates are affected long after the initial shock (Section C.2 of Supplementary Information). [Dell et al. \(2012\)](#) find that the cumulative effect of temperature is significant for poor countries, albeit smaller than the contemporaneous effect. They interpret this as evidence of a “growth effect”, but our interpretation aligns more closely with [Newell et al. \(2021\)](#): small, non-significant lagged temperature coefficients suggest a permanent impact on the level of GDP rather than on growth.

the channels through which climate shocks affect GDP per capita.

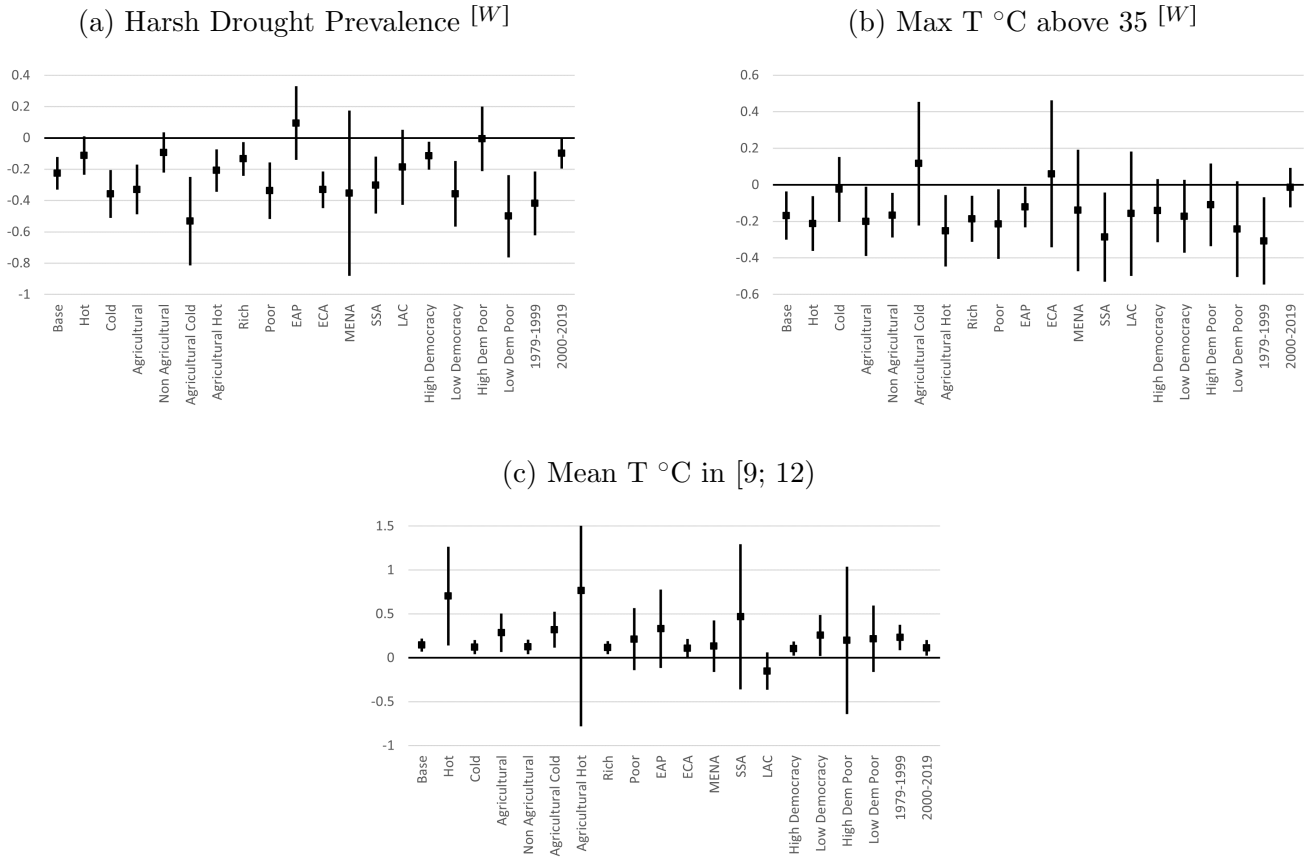
We differentiate between rich and poor, hot and cold, agricultural and non-agricultural, agricultural cold and agricultural hot, strong and weak democratic institutions. We further divide countries into six macro-regional groups to explore regional heterogeneity, and we compare earlier years of the panel (1979-1999) with more recent years (2000-2019). For each subgroup, we use the same variables selected by the LASSO for the baseline specification because we are primarily concerned with the robustness of estimates across groups. We leave selection of variables by group to future research. For a detailed description of each group see the notes to Figure 5.

We find that our main results are generally confirmed across groups (Figure 5). Coefficients tend to remain significant, there are no significant sign switches, and coefficients are almost never significantly different between groups. However, there is suggestive evidence that aggregate results might be driven by specific vulnerabilities in sub-groups of countries.

Harsh droughts are significantly harmful in almost all groups. A key vulnerability to harsh droughts comes from a large agricultural sector. Harsh droughts are not significant in hot countries, but they are significant in hot agricultural countries. Droughts are more harmful in agricultural cold countries than in cold countries. Poor countries are more affected than rich countries probably because they have a relatively larger agricultural sector. Across regions, droughts in Sub-Saharan Africa (SSA) and in the Europe and Central Asia region (ECA) have large and significant negative impacts. These are the regions with the largest exposure to droughts (see Figure 3). The effect of droughts is instead not significant in Latin American and the Caribbean (LAC), in Eastern Asia and Pacific (EAP), and in Middle East and North Africa (MENA). The small sample size of the MENA region may explain why the large negative impact is not significant. Droughts are less harmful in countries that have a high electoral democracy index relative to countries with a low index, but the confounding effect of income cannot be excluded. To further examine the importance of institutions in shaping the effect of droughts on the economy, we separate poor countries among those with high and low electoral democracy indexes. Results indicate that droughts are more damaging in poor countries that have low democratic institutions relative to poor countries. The effect of droughts in poor countries that rank high in democracy is not significant from zero. These results provide suggestive evidence that institutions play a role, similarly to what is found by (Kahn, 2005) when studying the impact of natural disasters on deaths. However, the results among groups are not statistically significant and a more in-depth analysis of the role of institutions is left for future research. Droughts are significantly harmful in both the first and the second part of the panel, but their impact has more than halved over time. This could be the outcome of economic and institutional development and the reduction of the share of agriculture in GDP in many countries.

The harmful effect of high temperatures is generally not significant in cold countries or in groups with a high share of cold countries, because maximum daily temperatures above 35 °C are rarely observed in these areas. Hot temperatures are equally harmful in poor and rich countries, agricultural and non-

Figure 5: Weather Coefficients Across Groups



Notes: Each panel depicts the estimated coefficient for each weather variable using our baseline specification (column A in Table 1) for different sub-groups. Climate variables are standardized and their definitions are detailed in appendix Table A.1. The vertical lines show the 95 percent confidence intervals using standard errors clustered by country. [W] indicates population-weighted variables. Hot ($N=3,315$): 1979-2019 average temperature > 22.8 °C. Cold ($N=3,338$): 1979-2019 average temperature ≤ 22.8 °C. Agricultural ($N=3,130$): share of “Agriculture, forestry, and fishing, value added (% of GDP)” in 2002 is above median across countries. Non Agricultural ($N=3,062$): countries that are not Agricultural. Agricultural Cold ($N=1,334$): agricultural and cold. Agricultural Hot ($N=1,785$): agricultural and hot. Rich ($N=3,823$): “High Income” and “Upper Middle Income” in WDI. Poor: “Low Income” and “Lower Middle Income” in WDI ($N=2,727$). SSA ($N=1,629$): Sub-Saharan Africa. MENA ($N=538$): Middle-East and North Africa. LAC ($N=1,372$): Latin America and the Caribbean. ECA ($N=1,638$): Europe and Central Asia. EAP ($N=1,013$): Eastern Asia and Pacific. None of the coefficients is significant for North America and South Asia, due to the low number of countries in these regions. High Democracy ($N=3,895$) and Low Democracy ($N=2,753$): countries with average value of the V-Dem Project Electoral Democracy Index from the the Varieties of Democracies (V-Dem) project [Coppedge et al. \(2023\)](#) above (below) the median in our sample. High Democracy and Poor ($N=950$), Low Democracy and Poor ($N=1,766$). 1981-1999 ($N=2,762$) and 2000-2019 ($N=3,891$). The V-Dem index aggregates indices measuring freedom of association, clean elections, freedom of expression, elected officials, and suffrage [Coppedge et al. \(2023\)](#). Table B.5 in the online appendix features summary statistics by sub-group for GDP per capita and climate variables.

agricultural countries. This suggests that the harmful impact of hot temperatures is not confined to the agricultural sector but may affect labor productivity at large, as found by the empirical sectoral literature. SSA, with mostly agricultural, poor and hot countries, and EAP are the only regions with a significant harmful effect. The effect of hot temperatures is never significant when we compare countries with weak and strong democratic institutions. The negative impact of high temperatures is much larger when estimated using the first part of the panel. As for droughts, the effect of high temperature in more recent years is smaller and, in this case, becomes non-significant possibly for similar reasons.

The effect of mild temperatures is not significant in countries that are agricultural and hot, and poor. This is due to the relatively low frequency of these temperatures in these groups. The positive impact is particularly large in agricultural cold countries and in hot countries. This suggests that the variable is picking the beneficial effect of cooler than average temperatures in hot countries and of warmer than average temperatures in cold agricultural countries, reducing heat stress on crops in the first group and a longer growing season either in spring or fall in the second group (Massetti et al., 2016). Among regional groups, the effect is positive and significant only in the ECA region, which comprises many cold agricultural countries. There are no significant differences among countries when we separate them using the Electoral Democracy Index. Also in this case there is an indication that the impact of weather on GDP becomes smaller as time goes by.

5. Macro-Fiscal Outcomes

In this last section, we investigate whether our selected weather shocks are typically associated with a fiscal policy response that amplifies or mitigate their effect on GDP. We also apply our approach to see what other weather shocks have a substantial impact on fiscal variables.

We consider three main fiscal indicators: government revenue, expenditure, and debt. We start by studying the effect of weather shocks on these variables measured by their GDP ratio as is standard practice in fiscal policy analysis. We supplement this analysis by considering the effect of weather shocks on the percent change of the fiscal variables expressed in constant 2011 USD. Examining variations in levels allows us to separately identify the effects of climate shocks on the numerator (fiscal variable) and denominator (GDP) of the ratios, as these can be of different magnitude and opposite sign.

We present empirical OLS estimates of the relationships between weather shocks and macro-fiscal outcomes in Table 3. Our goal is to study how climate shocks affect GDP and all fiscal outcomes systematically while keeping the selection of variables compact. Therefore, we restrict the selection of climate variables: we keep those three that were selected in our baseline study of GDP and only introduce three additional climate variables: the first variable that LASSO selects independently for each of the three fiscal ratios. The full list of the variables selected by LASSO using the BIC criteria and the associated results for every fiscal variables are reported in the online appendix.²³

²³The list of variables selected to optimize the BIC is reported in Table B.9. These variables systematically

Our sample becomes smaller when we introduce fiscal variables because of their narrower coverage. To allow for a meaningful comparison across macro-fiscal outcomes, we implement our analysis on the sample with non-missing values for all macro-fiscal variables. The smaller sample size and the addition of other weather variables explain the difference in results obtained for GDP growth compared to the previous section (column A in Table 1 versus column A in Table 3). Each weather effect separately tends to become smaller and the effect of hot temperatures becomes insignificant.

The three climate shocks previously identified have a significant impact on many fiscal policy ratios. Columns (B)-(D) in Table 3 show that more harsh drought prevalence leads to a significant increase of 0.11 percentage points in the expenditure-to-GDP ratio. Further, an increase in the occurrence of hot temperatures and a decrease in the occurrence of mild temperatures are associated with a significantly lower revenue-to-GDP ratio (respectively by 0.14 and by 0.08 percentage points). In these three cases, the effect is not entirely driven by a change in the denominator as we verify that there is a consistent change in the numerator (see Table A.6 in appendix). On the other hand, these three climate variables do not impact the debt ratio significantly.

The fiscal policy response to the three climate shocks that we identified as relevant for GDP mitigates direct GDP effects. In all three cases, the estimated effects of the weather shocks on the fiscal balance, which is the difference between the effect on revenue and spending, has the same sign as the effect on GDP. Government spending increases in response to droughts. This could come from relief measures (either built in automatic stabilizers or ad-hoc) or because droughts raise the cost of government purchases. In the case of revenue, its drop following a temperature shock could result from high-tax sectors, goods, or agents being disproportionately affected by the shock. Overall and assuming that the fiscal multiplier is non-zero, the results suggest that the direct effect on GDP of the weather shocks would be larger if fiscal policy were neutral (i.e., if the fiscal ratios remained constant).

Applying the LASSO to select new climate variables that are relevant for explaining fiscal ratios results in the selection of new variables with mixed effects. These three new variables have relatively small and insignificant effects on GDP per capita.

For government revenue, the LASSO selects the lag of the variable “Longest Day Cold Wave”, which counts the number of days in the longest period during which day temperatures are substantially below seasonal norms for at least three consecutive days.²⁴ We find that longer day cold waves in the previous

include the first variable that we report in the main text, meaning that the first variable to be selected does not get dropped at a later stage when additional variables are introduced to maximize the BIC. The algorithm performance as a function of the penalty parameter λ is reported in the online appendix in Figures B.6-B.5.

²⁴Seasonal norms are defined for every calendar day over the 15-day window centered on that day. We determine that temperatures are *substantially* below seasonal norms when temperatures are below the 10th percentile of the distribution of temperatures in the aforementioned windows in all years 1979-2019. See Table A.1 in appendix for mathematical definitions.

Table 3: Estimates of macro-fiscal effects of selected climate variables

	(A) $\Delta \ln \frac{\text{GDP}}{\text{POP}}$ (p.c.)	(B) $\Delta \frac{\text{Revenue}}{\text{GDP}}$ (p.p.)	(C) $\Delta \frac{\text{Expenditure}}{\text{GDP}}$ (p.p.)	(D) $\Delta \frac{\text{Debt}}{\text{GDP}}$ (p.p.)
Lag-1 dependent variable	0.211*** (0.0487)	-0.131*** (0.0205)	-0.118 (0.0861)	0.106*** (0.0306)
Lag-2 dependent variable	0.0966** (0.0411)	-0.137*** (0.0145)	-0.0911 (0.0760)	0.0248 (0.0267)
First difference in				
Harsh Drought Prevalence ^[W]	-0.180*** (0.555)	0.0625 (0.0431)	0.108** (0.0493)	-0.0765 (0.0957)
Max T °C above 35 ^[W]	-0.0452 (0.0629)	-0.135** (0.0525)	-0.0454 (0.0542)	-0.0603 (0.127)
Mean T °C in [9; 12]	0.129*** (0.0448)	0.0750* (0.0417)	-0.0636 (0.0654)	0.00017 (0.0975)
Lag-1 Longest Day Cold Wave ^[W]	0.0454 (0.0434)	-0.232* (0.137)	-0.0766 (0.0889)	0.200 (0.136)
Mean Wet Day PPT	-0.0501 (0.0608)	-0.0145 (0.0755)	-0.197*** (0.0695)	0.0674 (0.115)
Lag-1 PPT Minimum	0.132 (0.0858)	0.00557 (0.0464)	0.0302 (0.0715)	-0.313*** (0.102)
Constant	1.381*** (0.0912)	0.0782*** (0.00153)	0.0599*** (0.00879)	0.193*** (0.00488)
Observations	3,890	3,890	3,890	3,890
R-squared	0.304	0.118	0.0818	0.160

Note: The dependent variables are indicated in the column titles and are expressed in percentage points. We use the same three climate variables used for GDP growth and the first climate variables selected by the LASSO respectively for government revenue, expenditure, and debt. All climate variables are standardized with standard deviations equal to 100 to ease interpretation. ^[W] indicates population-weighted variables. Controls include the first two lags of the dependent variable (reported in the first two rows), and year and country fixed effects. Standard errors are clustered by country and reported in brackets.

year are associated with a significant 0.23 percentage point decline in the revenue ratio and a comparable (although insignificant) increase in the debt ratio. The expenditure ratio is left mostly unchanged.

For government spending and debt, the LASSO respectively selects the mean precipitation in wet days and the average across space of precipitation in the driest month.²⁵ Less precipitation in wet days is estimated to raise the expenditure ratio by 0.20 percentage points while less precipitation in the driest month of the previous year is estimated to raise the debt ratio by 0.19 percentage points. These variables seem to leave our other variables mostly unchanged.

In summary, we find that weather shocks can have rich and sizeable effects on fiscal aggregates, although their impacts are not always significant and can be hard to interpret. When the weather shocks have a clear effect on GDP, we find that the response of fiscal policy tends to act counter-cyclically. In other cases and generally, a more granular investigation would be needed to better understand the sectoral channels through which climate affects macro-fiscal outcomes.

²⁵See Table A.1 in appendix for mathematical definitions.

6. Conclusion

In this paper, we show how to leverage large global weather datasets with high-frequency and high-resolution data to estimate the impact of weather on economic outcomes. With few minor exceptions, the macro literature has traditionally focused on average annual temperature and total precipitation at the country level, leaving aside many potentially relevant weather phenomena.

The main empirical challenge of using dozens of billions of weather measurements is to reduce the high dimensionality of the data to a manageable set of variables that can be related to country-year macro-fiscal aggregates. We propose a method that relies on a mix of expert judgement and machine learning techniques, including the LASSO and Elastic-Net. This approach can identify a small set of variables that improves the ability to explain the outcome of interest without sacrificing the possibility to meaningfully interpret their effect. Applying this method leads to a substantial improvement in our ability to explain GDP growth.

We use our method to select a few variables capturing harsh droughts and hot and mild temperatures and find that a shock of one standard deviation to any of these variables has an affect approximately equal to 0.2 percent on GDP growth in the year of the shock. GDP level does not recover after each shock, but growth goes back to normal. However, “good” and “bad” shocks are symmetric by construction in all countries. Their average value is equal to zero over our estimation period and our specification implies that their average impact on average growth (and GDP level) is also equal to zero. Our results do not imply that GDP is unaffected by climate change. On the contrary, trends observed in some of our climate variables may have already negatively affected average GDP growth and would be expected to remain detrimental in the future.

Our most striking result is that country annual average temperature, the variable most frequently used in the literature, is never part of the core set of variables that is selected with our method. The weather variables we select are indeed far better in explaining GDP growth and macro-fiscal variations than average annual temperature alone, also when using data and methods of important papers in the literature. We find that weather extremes that had been shown by the literature to be harmful at sectoral level are also harmful at macro level.

This paper additionally contributes to the literature by identifying and estimating the effects of weather shocks on important fiscal aggregates. We find that fiscal policy tends to respond in a way that would counteract the direct effects of weather shocks on GDP. Government revenue declines following an increase in high temperatures or a decrease in mild temperatures while government spending increases in response to harsh droughts. A promising work avenue would be to look more precisely into the mechanisms underlying these policy responses.

Because we find that the impact of weather shocks on the macro economy is larger than previously mea-

sured, the main policy implication is that decision-makers should pay better attention to the shocks that we identified as relevant and ensure that they have adequate capacity to address them, including by creating fiscal space. We see our work as a starting point to learn more about the macro-fiscal vulnerability of countries that are exposed to weather shocks. It should also be useful to predict the macro-fiscal effects of future weather shocks.

Our method opens avenues for more advanced weather data analysis and applications to many potentially policy-relevant analysis. Future extensions could examine additional climate variables such as humidity and wind, and climate phenomena that we have not studied, such as tropical cyclones. It could also be applied to study other macroeconomic and sectoral outcomes, such as the impact of weather on health, agriculture, labor productivity, trade, and inflation. Future work could also examine other policy responses, like monetary policy. And by focusing on specific countries or regions, it could generate new useful insights and identify group-specific sets of relevant explanatory variables in specific contexts.

Our work highlights vulnerabilities to climate shocks but our empirical strategy cannot be easily extended to infer the long-term impacts of climate change. We rely on short-term weather fluctuations that are random, of relatively small magnitude compared to the long-term changes that we expect from climate change, and are not observed in all countries. Our method is thus inherently unable to estimate the impact of gradual but large changes in *average* conditions, and the effect of the appearance of new weather shocks in countries where they have never been observed. However, it should be possible to project our selected variables to estimate impacts from climate change in the near future. This would require constructing climate change scenarios with a spatial and temporal granularity similar to our historical data.

References

- Abatzoglou, J. T., S. Z. Dobrowski, S. A. Parks, and K. C. Hegewisch (2018). Terraclimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015. Scientific data 5(1), 1–12.
- Acevedo, S., M. Mrkaic, N. Novta, E. Pugacheva, and P. Topalova (2020). The Effects of Weather Shocks on Economic Activity: What are the Channels of Impact? Journal of Macroeconomics 65.
- Adebayo, E. and E. Fokoue (2019). An empirical demonstration of the no free lunch theorem. Mathematics for Applications 8, 173–188.
- Auffhammer, M., S. M. Hsiang, W. Schlenker, and A. Sobel (2013). Using weather data and climate model output in economic analyses of climate change. Review of Environmental Economics and Policy, ret016.
- Barrage, L. and W. D. Nordhaus (2023). Policies, projections, and the social cost of carbon: Results from the dice-2023 model. Technical report, National Bureau of Economic Research.

- Belloni, A. and V. Chernozhukov (2013). Least squares after model selection in high-dimensional sparse models. Bernoulli 19(2), 521–547.
- Belloni, A., V. Chernozhukov, and C. Hansen (2014). Inference on treatment effects after selection among high-dimensional controls. The Review of Economic Studies 81(2), 608–650.
- Blanc, E. and W. Schlenker (2017). The use of panel models in assessments of climate impacts on agriculture. Review of Environmental Economics and Policy.
- Bosetti, V., C. Carraro, M. Galeotti, E. Massetti, and M. Tavoni (2006). A world induced technical change hybrid model. The Energy Journal (Special Issue# 2).
- Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. Nature 527(7577), 235–239.
- Cantelmo, A., G. Melina, and C. Papageorgiou (2019). Macroeconomic outcomes in disaster-prone countries. (217).
- Christensen, O. B., C. M. Goodess, and J.-C. Ciscar (2012). Methodological framework of the peseta project on the impacts of climate change in europe. Climatic Change 112(1), 7–28.
- Chudik, A., K. Mohaddes, M. H. Pesaran, and M. Raissi (2016). Long-run effects in large heterogeneous panel data models with cross-sectionally correlated errors. Emerald Group Publishing Limited.
- Coppedge, M., J. Gerring, C. H. Knutsen, S. I. Lindberg, J. Teorell, D. Altman, M. Bernhard, A. Cornell, M. S. Fish, L. Gastaldi, H. Gjerløw, A. Glynn, A. Good God, S. Grahn, A. Hicken, K. Kinzelbach, J. Krusell, K. L. Marquardt, K. McMann, V. Mechkova, J. Medzihorsky, P. Paxton, D. Pemstein, J. Pernes, O. Rydén, J. von Römer, B. Seim, R. Sigman, S.-E. Skaaning, J. Staton, A. Sundström, E. Tzelgov, Y.-t. Wang, T. Wig, S. Wilson, and D. Ziblatt (2023). V-dem country-year dataset v13.
- Coppedge, M., J. Gerring, C. H. Knutsen, S. I. Lindberg, J. Teorell, D. Altman, M. Bernhard, A. Cornell, M. S. Fish, L. Gastaldi, H. Gjerløw, A. Glynn, S. Grahn, A. Hicken, K. Kinzelbach, K. L. Marquardt, K. McMann, V. Mechkova, A. Neundorf, P. Paxton, D. Pemstein, O. Rydén, J. von Römer, B. Seim, R. Sigman, S.-E. Skaaning, J. Staton, A. Sundström, E. Tzelgov, L. Uberti, Y.-t. Wang, T. Wig, and D. Ziblatt (2023). V-dem codebook v13.
- Cubasch, U., D. Wuebbles, D. Chen, M. Facchini, D. Frame, N. Mahowald, and J.-G. Winther (2013). Introduction. In T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley (Eds.), Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, pp. 119–158. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.
- Dai, A., K. E. Trenberth, and T. Qian (2004). A Global Dataset of Palmer Drought Severity Index for 1870-2002: Relationship with Soil Moisture and Effects of Surface Warming. Journal of Hydrometeorology 5, 1117–1130.

- Dell, M., B. F. Jones, and B. A. Olken (2012). Temperature Shocks and Economic Growth: Evidence from the Last Half Century. American Economic Journal: Macroeconomics 4, 66–95.
- Deryugina, T. and S. Hsiang (2014). Does the environment still matter? daily temperature and income in the united states. NBER Working Paper (w20750).
- Deschênes, O. and M. Greenstone (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. American Economic Journal: Applied Economics 3(4), 152–185.
- Gleditsch, N. P., P. Wallensteen, M. Eriksson, M. Sollenberg, and H. Strand (2002). Armed conflict 1946-2001: A new dataset. Journal of peace research 39(5), 615–637.
- Gorelick, N., M. Hancher, M. Dixon, S. Ilyushchenko, and R. Moore (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. Remote Sensing of Environment 202, 18–27.
- Graff Zivin, J. and M. Neidell (2014). Temperature and the allocation of time: Implications for climate change. Journal of Labor Economics 32(1), 1–26.
- Gruss, B. and S. Kebhaj (2019). Commodity terms of trade: A new database. IMF Working Papers 2019(021).
- Hersbach, H., B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, et al. (2018). Era5 hourly data on single levels from 1979 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS) 10.
- Hope, C., J. Anderson, and P. Wenman (1993). Policy analysis of the greenhouse effect: an application of the PAGE model. 21(3), 327–338.
- Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the caribbean and central america. Proceedings of the National Academy of sciences 107(35), 15367–15372.
- IMF (2020). Adapting to climate change in sub-saharan africa. In Regional Economic Outlook, Sub-Saharan Africa, Regional Economic Outlook. International Monetary Fund, African Development.
- IPCC (2021a). Atlas. In Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.
- IPCC (2021b). Summary for policymakers. In Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.
- IPCC (2021c). Technical summary. In Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.

- Jorda, O. (2005). Estimation and Inference of Impulse Responses by Local Projection. American Economic Review 95, 161–182.
- Kahn, M. E. (2005). The death toll from natural disasters: the role of income, geography, and institutions. Review of economics and statistics 87(2), 271–284.
- Kahn, M. E., K. Mohaddes, R. N. Ng, M. H. Pesaran, M. Raissi, and J.-C. Yang (2021). Long-term macroeconomic effects of climate change: A cross-country analysis. Energy Economics, 105624.
- Kalkuhl, M. and L. Wenz (2020). The impact of climate conditions on economic production. evidence from a global panel of regions. Journal of Environmental Economics and Management 103, 102360.
- Kim, Y.-H., S.-K. Min, X. Zhang, J. Sillmann, and M. Sandstad (2020). Evaluation of the CMIP6 multi-model ensemble for climate extreme indices. Weather and Climate Extremes 29.
- Lai, Peiyu and Zhang, M., Z. Ge, B. Hao, Z. Song, J. Huang, M. Ma, H. Yang, and X. Han (2020). Responses of Seasonal Indicators to Extreme Droughts in Southwest China. Remote Sensing 12.
- Letta, M. and R. S. Tol (2019). Weather, climate and total factor productivity. Environmental and Resource Economics 73(1), 283–305.
- Massetti, E., R. Mendelsohn, and S. Chonabayashi (2016). How well do degree days over the growing season capture the effect of climate on farmland values? Energy Economics 60, 144–150.
- Mbaye, S., M. Moreno Badia, and K. Chae (2018). Global debt database: Methodology and sources. IMF Working Paper No. 18/111.
- Mérel, P. and M. Gammans (2021). Climate econometrics: Can the panel approach account for long-run adaptation? American Journal of Agricultural Economics 103(4), 1207–1238.
- Newell, R. G., B. C. Prest, and S. E. Sexton (2021). The gdp-temperature relationship: implications for climate change damages. Journal of Environmental Economics and Management 108, 102445.
- Nordhaus, W. D. (1993). Optimal greenhouse-gas reductions and tax policy in the” dice” model. The American Economic Review 83(2), 313–317.
- Nordhaus, W. D. and Z. Yang (1996). A regional dynamic general-equilibrium model of alternative climate-change strategies. pp. 741–765.
- Palmer, W. C. (1965). Meteorological Drought. Research Paper, U.S. Weather Bureau.
- Perkins, S. E. (2015). A review on the scientific understanding of heatwaves—their measurement, driving mechanisms, and changes at the global scale. Atmospheric Research 164, 242–267.
- Perkins, S. E. and L. V. Alexander (2013). On the Measurement of Heatwaves. Journal of Climate 26.

- Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to US crop yields under climate change. 106(37), 15594–15598.
- Somanathan, E., R. Somanathan, A. Sudarshan, and M. Tewari (2021). The impact of temperature on productivity and labor supply: Evidence from indian manufacturing. Journal of Political Economy 129(6), 1797–1827.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological) 58(1), 267–288.
- Tol, R. S. (1997). On the optimal control of carbon dioxide emissions: an application of FUND. 2(3), 151–163.
- Tol, R. S. (2019). A social cost of carbon for (almost) every country. Energy Economics 83, 555–566.
- Tol, R. S. (2021). The economic impact of weather and climate. CESifo Working Paper (8946).
- Zou, H., T. Hastie, and R. Tibshirani (2007). On the” degrees of freedom” of the lasso. The Annals of Statistics, 2173–2192.

A. Appendix

Table A.1: Definitions of climate variables

Variable	Model	Descriptor	Spatial	Time	Definition	Ref
<i>Temperature</i>						
$T_{j,d}$		Mean Temperature	Grid	Day		
T_d		Mean Temperature	Country	Day	$\sum_j T_{j,d}/J$	
T	L	Mean Temperature	Country	Year	$\sum_d T_d/365$	
$TVar$	L	Temperature Variance	Country	Year	$\sum_d (T_d - T)^2 / (365 - 1)$	
$TN_{j,d}$		Daily Temperature Minimum	Grid	Day	Note: temperature minimums almost always occur at night	
$TX_{j,d}$		Daily Temperature Maximum	Grid	Day	Note: temperature maximums almost always occur in daytime	
TN_d		Daily Minimum T °C	Country	Day	$\sum_j TN_{j,d}/J$	
TX_d		Daily Maximum T °C	Country	Day	$\sum_j TX_{j,d}/J$	
DTR_d		Diurnal T °C Range	Country	Day	$TX_d - TN_d$	
DTR	L	Diurnal T °C Range	Country	Year	$\sum_d DTR_d/365$	
$TNp(k)_d$		Percentile of Daily Minimum Temperature	Country	Day	p^{th} percentile of the 1979-2019 distribution of TN_d in a k -day window centered on d	
$TXp(k)_d$		Percentile of Daily Maximum Temperature	Country	Day	p^{th} percentile of the 1979-2019 distribution of TX_d in a k -day window centered on d	
$CN10$	L	# of Cold Nights	Country	Year	$\sum_d [TN_d < TN10(5)_d]$	a
$CD10$	L	# of Cold Days	Country	Year	$\sum_d [TX_d < TX10(5)_d]$	a
$WN90$	L	# of Warm Nights	Country	Year	$\sum_d [TN_d > TN90(5)_d]$	a
$WD90$	L	# of Warm Days	Country	Year	$\sum_d [TX_d > TX90(5)_d]$	a
TNn	L	Night T °C Minimum	Country	Year	Minimum of minimum daily temperature, $\sum_j \min_d \{TN_{j,d}\} / J$	d
TXx	L, R	Day T °C Maximum	Country	Year	Maximum of maximum daily temperature, $\sum_j \max_d \{TX_{j,d}\} / J$	d

Continued on next page

Notes: The letter L in the second column indicates whether a variable is used in the LASSO procedure. The letter R indicates the variables used in any of the OLS regressions. In the last column, a refers to [Kim et al. \(2020\)](#), b refers to [Perkins and Alexander \(2013\)](#), c refers to [Palmer \(1965\)](#), and d refers to [IPCC \(2021c\)](#).

Table A.1 (Continued): Definitions of climate variables

Variable	Model	Descriptor	Spatial	Time	Definition	Ref
<i>Heat Waves</i>						
<i>DDHW</i>	L	Heat Wave Days	Country	Year	Number of days in which $TX_d > TX90(15)_d$ for at least 3 consecutive days	b
<i>DNHW</i>	L	Heat Wave Nights	Country	Year	Number of nights in which $TN_d > TN90(15)_d$ for at least 3 consecutive days	b
<i>LDHW</i>	L	Longest Day Heat Wave	Country	Year	Number of days in the longest period during which $TX_d > TX90(15)_d$ for at least 3 consecutive days	b
<i>LNHW</i>	L	Longest Night Heat Wave	Country	Year	Number of days in the longest period during which $TN_d > TN90(15)_d$ for at least 3 consecutive days	b
<i>NDHW</i>	L	# of Day Heat Waves	Country	Year	Number of intervals of at least 3 consecutive days in which $TX_d > TX90(15)_d$	b
<i>NNHW</i>	L	# of Night Heat Waves	Country	Year	Number of intervals of at least 3 consecutive days in which $TN_d > TN90(15)_d$	b
<i>TDHW</i>	L	Day Heat Wave T °C	Country	Year	Average TX_d during day heat waves (intervals of at least 3 consecutive days in which $TX_d > TX90(15)_d$)	b
<i>TNHW</i>	L	Night Heat Wave T °C	Country	Year	Average TN_d during night heat waves (intervals of at least 3 consecutive days in which $TN_d > TN90(15)_d$)	b
<i>Cold Waves</i>						
<i>DDCW</i>	L	Cold Wave Days	Country	Year	Number of days in which $TX_d < TX10(15)_d$ for at least 3 consecutive days	b
<i>DNCW</i>	L	Cold Wave Nights	Country	Year	Number of days in which $TN_d < TN10(15)_d$ for at least 3 consecutive days	b
<i>LDCW</i>	L	Longest Day Cold Wave	Country	Year	Number of days in the longest period during which $TX_d < TX10(15)_d$ for at least 3 consecutive days	b
<i>LNCW</i>	L	Longest Night Cold Wave	Country	Year	Number of days in the longest period during which $TN_d < TN10(15)_d$ for at least 3 consecutive days	b
<i>NDCW</i>	L	# of Day Cold Waves	Country	Year	Number of intervals of at least 3 consecutive days in which $TX_d < TX10(15)_d$	b
<i>NNCW</i>	L	# of Night Cold Waves	Country	Year	Number of intervals of at least 3 consecutive days in which $TN_d < TN10(15)_d$	b
<i>TDCW</i>	L, R	Day Cold Wave T °C	Country	Year	Average TX_d during day heat waves (intervals of at least 3 consecutive days in which $TX_d < TX10(15)_d$)	b
<i>TNCW</i>	L, R	Night Cold Wave T °C	Country	Year	Average TN_d during night heat waves (intervals of at least 3 consecutive days in which $TN_d < TN10(15)_d$)	b

Continued on next page

Notes: The letter *L* in the second column indicates whether a variable is used in the LASSO procedure. The letter *R* indicates the variables used in any of the OLS regressions. In the last column, *a* refers to [Kim et al. \(2020\)](#), *b* refers to [Perkins and Alexander \(2013\)](#), *c* refers to [Palmer \(1965\)](#), and *d* refers to [IPCC \(2021c\)](#).

Table A.1 (Continued): Definitions of climate variables

Variable	Model	Descriptor	Spatial	Time	Definition	Ref
<i>Cold and Warm Spells</i>						
<i>CSD</i>	L, R	Cold Spell Duration	Country	Year	Number of days in which $TN_d < TN10(5)_d$ is observed in intervals of at least 6 consecutive days	a
<i>WSD</i>	L	Warm Spell Duration	Country	Year	Number of days in which $TX_d > TX90(5)_d$ for at least 6 consecutive days	a
<i>Temperature Absolute Thresholds</i>						
<i>TN0</i>	L	Frost prevalence	Country	Year	Share of grid-days with frost, $\sum_d \sum_j [TN_{j,d} < 0] / (J \times 365)$	d
<i>TX35</i>	L, R	Max T °C above 35	Country	Year	Share of grid-days with maximum daily temperature above 35 °C, $\sum_d \sum_j [TX_{j,d} > 35] / (J \times 365)$	d
<i>TX40</i>	L	Max T °C above 40	Country	Year	Share of grid-days with maximum daily temperature above 40 °C, $\sum_d \sum_j [TX_{j,d} > 40] / (J \times 365)$	d
<i>TS_{<-9}</i>		Mean T °C below 9	Country	Year	Share of grid-days with mean temperature below -9 °C, $\sum_d \sum_j [T_{j,d} < -9] / (J \times 365)$	
<i>TS_{[x₁,x₂)}</i>	L,R	Mean T °C in $[x_1, x_2)$	Country	Year	Share of grid-days with mean temperature in the interval $[x_1, x_2)$, $\sum_d \sum_j [x_1 \leq T_{j,d} < x_2] / (J \times 365)$. We use increments of 3 °C from -9 °C to 30 °C for x_1, x_2 .	
<i>TS_{≥30}</i>	L	Mean T °C above 30	Country	Year	Share of grid-days with mean temperature above 30 °C, $\sum_d \sum_j [T_{j,d} \geq 30] / (J \times 365)$	
<i>Precipitation</i>						
<i>P_{j,d}</i>		Precipitation (PPT)	Grid	Day		
<i>P_d</i>		Precipitation (PPT)	Country	Day	$\sum_j P_{j,d} / J$	
<i>P</i>	L	Mean Precipitation	Country	Year	$\sum_d P_d / 365$	
<i>PW_d</i>		Wet Day Precipitation	Country	Day	$P_d [P_d \geq 1]$	
<i>PWT</i>		Wet Day Precipitation	Country	Year	$\sum_d P_d [P_d \geq 1]$	
<i>W</i>	L	# of Wet Days	Country	Year	$\sum_d [P_d \geq 1]$	
<i>PWA</i>	L	Mean Wet Day PPT	Country	Year	Average daily precipitation in wet days, PTW/W	
<i>PVar</i>	L	Precipitation Variance	Country	Year	$\sum_d (P_d - P)^2 / (365 - 1)$	
<i>PWVar</i>	L	Wet Day PPT Variance	Country	Year	$\sum_d (P_d - PWA)^2 [P_d \geq 1] / (W - 1)$	
<i>PW_{pj}</i>		Percentile of Daily Precipitation in Wet Days	Grid	1979-2019	p^{th} percentile of the 1979-2019 daily distribution of PW_d (using only wet days) in grid cell j	
<i>PW_p</i>		Percentile of Daily Precipitation in Wet Days	Country	1979-2019	p^{th} percentile of the 1979-2019 daily distribution of PW_d (using only wet days)	
<i>P95WT</i>	L	Very Wet Day PPT	Country	Year	$\sum_d P_d [P_d \geq 1 \text{ and } PW_d > PW95]$	a

Continued on next page

Notes: The letter *L* in the second column indicates whether a variable is used in the LASSO procedure. The letter *R* indicates the variables used in any of the OLS regressions. In the last column, *a* refers to [Kim et al. \(2020\)](#), *b* refers to [Perkins and Alexander \(2013\)](#), *c* refers to [Palmer \(1965\)](#), and *d* refers to [IPCC \(2021c\)](#).

Table A.1 (Continued): Definitions of climate variables

Variable	Model	Descriptor	Spatial	Time	Definition	Ref
<i>P99WT</i>	L	Extremely Wet Day PPT	Country	Year	$\sum_d P_d [P_d \geq 1 \text{ and } PW_d > PW99]$	a
<i>CDD</i>	L	Cont'd Dry Days	Country	Year	Largest number of consecutive days with $P_d < 1mm$	a
<i>CWD</i>	L	Cont'd Wet Days	Country	Year	Largest number of consecutive days with $P_d \geq 1mm$	a
<i>PCWD</i>	L	Cont'd Wet Day PPT	Country	Year	Total precipitation during the longest period of consecutive wet days with $P_d \geq 1$	
<i>C95WD</i>	L	Cont'd Very Wet Days	Country	Year	Largest number of consecutive wet days with $PW_d > PW95$	
<i>C99WD</i>	L	Cont'd Extra Wet Days	Country	Year	Largest number of consecutive wet days with $PW_d > PW99$	
<i>PC95WD</i>	L	Cont'd Heavy PPT	Country	Year	Total precipitation during the longest period of consecutive very wet days with $PW_d \geq PW95$	
<i>PC99WD</i>	L	Cont'd Extreme PPT	Country	Year	Total precipitation during the longest period of consecutive extremely wet days with $PW_d \geq PW99$	
<i>PX(5)</i>	L	5-Day PPT Maximum	Country	Year	Maximum 5-day precipitation, $\max_d \left\{ \sum_{i=0}^4 P_{d-i} \right\}$	a
<i>PX(1)</i>	L	1-Day PPT Maximum	Country	Year	Maximum 1-day precipitation, $\max_d \{P_d\}$	a
$P_{j,m}$		Monthly Precipitation	Grid	Month	$\sum_d P_{j,d}$	
<i>PXM</i>	L	PPT Maximum	Country	Year	Max 1-month precipitation, $\sum_j \max_m \{P_{j,m}\}/J$	
<i>PNM</i>	L	PPT Minimum	Country	Year	Min 1-month precipitation, $\sum_j \min_m \{P_{j,m}\}/J$	
Precipitation Absolute Thresholds						
<i>LLDS_{.5}</i>	L	Longest Dry Spell (.5)	Country	Year	Maximum number of days in a Dry Spell, defined as an uninterrupted series of days in which $\sum_j [P_{j,d} < 1]/J > 0.50$	
<i>LLDS_{.65}</i>	L	Longest Dry Spell (.65)	Country	Year	Maximum number of days in a Dry Spell, defined as an uninterrupted series of days in which $\sum_j [P_{j,d} < 1]/J > 0.65$	
<i>LLDS_{.80}</i>	L	Longest Dry Spell (.80)	Country	Year	Maximum number of days in a Dry Spell, defined as an uninterrupted series of days in which $\sum_j [P_{j,d} < 1]/J > 0.80$	
<i>LLDS_{.95}</i>	L	Longest Dry Spell (.95)	Country	Year	Maximum number of days in a Dry Spell, defined as an uninterrupted series of days in which $\sum_j [P_{j,d} < 1]/J > 0.95$	
<i>PS_{≤1}</i>	L	Less than 1 mm PPT	Country	Year	Share of grid-days with precipitation less than 1 mm, $\sum_d \sum_j [P_{j,d} \leq 1]/(J \times 365)$	
<i>PS₁₋₁₀</i>		1 to 10 mm PPT	Country	Year	Share of grid-days with precipitation between 1 and 10 mm, $\sum_d \sum_j [1 < P_{j,d} \leq 10]/(J \times 365)$	
<i>PS₁₀₋₂₀</i>	L	10 to 20 mm PPT	Country	Year	Share of grid-days with precipitation between 10 and 20 mm, $\sum_d \sum_j [10 < P_{j,d} \leq 20]/(J \times 365)$	

Continued on next page

Notes: The letter *L* in the second column indicates whether a variable is used in the LASSO procedure. The letter *R* indicates the variables used in any of the OLS regressions. In the last column, *a* refers to [Kim et al. \(2020\)](#), *b* refers to [Perkins and Alexander \(2013\)](#), *c* refers to [Palmer \(1965\)](#), and *d* refers to [IPCC \(2021c\)](#).

Table A.1 (Continued): Definitions of climate variables

Variable	Model	Descriptor	Spatial	Time	Definition	Ref
$PS_{>20}$	L	Above 20 mm PPT	Country	Year	Share of grid-days with precipitation above 20 mm, $\sum_d \sum_j [P_{j,d} > 20] / (J \times 365)$	
$MaxP_{>10}$	L	Heavy PPT Maximum	Country	Year	$\max_d \left\{ \sum_j [P_{j,d} > 10] / J \right\}$	
$MaxP_{>20}$	L	Extreme PPT Maximum	Country	Year	$\max_d \left\{ \sum_j [P_{j,d} > 20] / J \right\}$	
$BP_{1-10}(0.5)$	L	Balanced PPT Indicator	Country	Year	$-\left \sum_d \sum_j [1 < P_{j,d} \leq 10] / (J \times 365) - 0.5 \right $	
Droughts						
$PDSI_{j,m}$		Palmer Drought Severity Index	Grid	Month		c
$PDSI_{<-3}$	L	Drought Prevalence	Country	Year	$\sum_m \sum_j [PDSI_{j,m} < -3] / (J \times 12)$	c
$PDSI_{<-4}$	L, R	Harsh Drought Prevalence	Country	Year	$\sum_m \sum_j [PDSI_{j,m} < -4] / (J \times 12)$	c
$PDSI_{>3}$	L	Wet Conditions Prevalence	Country	Year	$\sum_m \sum_j [PDSI_{j,m} > 3] / (J \times 12)$	c
$PDSI_{>4}$	L	Very Wet Conditions Prevalence	Country	Year	$\sum_m \sum_j [PDSI_{j,m} > 4] / (J \times 12)$	c
$MPDSI_{<-3}$	L, R	Drought Intensity	Country	Year	$\max_m \left\{ \sum_j [PDSI_{j,m} < -3] / J \right\}$	c
$MPDSI_{<-4}$	L	Harsh Drought Intensity	Country	Year	$\max_m \left\{ \sum_j [PDSI_{j,m} < -4] / J \right\}$	c
$MPDSI_{>3}$	L	Wetness Intensity	Country	Year	$\max_m \left\{ \sum_j [PDSI_{j,m} > 3] / J \right\}$	c
$MPDSI_{>4}$	L	High Wetness Intensity	Country	Year	$\max_m \left\{ \sum_j [PDSI_{j,m} > 4] / J \right\}$	c

Notes: The letter *L* in the second column indicates whether a variable is in the set of climate variables used in the LASSO procedure. The letter *R* indicates the variables used in any of the OLS regressions. In the last column, *a* refers to [Kim et al. \(2020\)](#), *b* refers to [Perkins and Alexander \(2013\)](#), *c* refers to [Palmer \(1965\)](#), and *d* refers to [IPCC \(2021c\)](#).

Table A.2: The effect of changes in selected climate variables on GDP per capita growth

	(A)	(B)	(C)	(D)	(E)	(F)	(G)
Lag-1 GDP p.c. Growth	0.215*** (0.0332)	0.213*** (0.0333)	0.213*** (0.0333)	0.213*** (0.0334)	0.223*** (0.0327)	0.126*** (0.0366)	0.236*** (0.0292)
Lag-2 GDP p.c. Growth	0.0896*** (0.0201)	0.0899*** (0.0200)	0.0900*** (0.0204)	0.0891*** (0.0203)	0.0941*** (0.0197)		0.0929*** (0.0202)
World GDP p.c. Growth					0.753*** (0.0585)		
First difference in							
Harsh Drought Prevalence ^[W]	-0.229*** (0.0541)	-0.262*** (0.0535)			-0.247*** (0.0531)	-0.207*** (0.0538)	-0.310*** (0.0622)
Max T °C above 35 ^[W]	-0.169** (0.0676)		-0.214*** (0.0668)		-0.156** (0.0635)	-0.174** (0.0679)	-0.210** (0.0923)
Mean T °C in [9; 12]	0.145*** (0.0380)			0.162*** (0.0386)			
PPT Minimum ^[W]							-0.193** (0.0777)
Constant	1.220*** (0.0627)	1.221*** (0.0628)	1.221*** (0.0627)	1.223*** (0.0630)	-1.091*** (0.205)	1.514*** (0.0618)	1.118*** (0.0504)
Observations	6653	6653	6653	6653	6653	6653	4644
Year fixed effects	Yes	Yes	Yes	Yes	No	Yes	Yes
World GDP growth	No	No	No	No	Yes	No	No
Country quadratic trends	No	No	No	No	No	Yes	No
Balanced	No	No	No	No	No	No	Yes
R-squared	0.264	0.261	0.261	0.260	0.253	0.349	0.224
Within R-squared	0.0997	0.0967	0.0965	0.0950	0.150	0.0360	0.0790

Notes: This is the full table corresponding to the main text summary Table 1. All regressions include country fixed effects. The dependent variable is the first difference of log real GDP per capita expressed in constant 2015 USD. Climate variables are standardized and their definitions are detailed in appendix Table A.1. (W) indicates population-weighted variables. Standard errors are clustered by country.

Table A.3: Weather shocks on GDP per capita growth: robustness to additional controls

	(A)	(B)	(C)	(D)	(E)	(F)	(G)
L.GDP p.c. Growth	0.215*** (0.0332)	0.207*** (0.0366)	0.207*** (0.0367)	0.213*** (0.0335)	0.205*** (0.0327)	0.179*** (0.0337)	0.193*** (0.0331)
L2.GDP p.c. Growth	0.0896*** (0.0201)	0.0866*** (0.0212)	0.0863*** (0.0215)	0.0877*** (0.0197)	0.0842*** (0.0303)	0.0742*** (0.0247)	0.0797** (0.0338)
L3.GDP p.c. Growth		0.0431** (0.0207)	0.0431** (0.0206)				0.0368* (0.0198)
L4.GDP p.c. Growth		-0.00867 (0.0138)	-0.00864 (0.0137)				-0.00244 (0.0161)
Harsh Drought Prevalence (W)	-0.225*** (0.0532)	-0.217*** (0.0538)	-0.196*** (0.0581)	-0.226*** (0.0531)	-0.222*** (0.0529)	-0.256*** (0.0608)	-0.236*** (0.0643)
L.Harsh Drought Prevalence (W)			0.0731 (0.0547)				0.0878 (0.0649)
L2.Harsh Drought Prevalence (W)			-0.0261 (0.0536)				-0.0787 (0.0588)
Max T °C above 35 (W)	-0.168** (0.0670)	-0.160** (0.0710)	-0.169** (0.0670)	-0.172** (0.0665)	-0.156** (0.0629)	-0.189*** (0.0688)	-0.108* (0.0625)
L.Max T °C above 35 (W)			-0.0167 (0.0818)				0.0664 (0.0748)
L2.Max T °C above 35 (W)			-0.00436 (0.0605)				0.0386 (0.0521)
Mean T °C in [9; 12)	0.144*** (0.0377)	0.124*** (0.0375)	0.216*** (0.0527)	0.144*** (0.0379)	0.134*** (0.0394)	0.146*** (0.0445)	0.154*** (0.0588)
L.Mean T °C in [9; 12)			0.178** (0.0772)				0.0933 (0.0817)
L2.Mean T °C in [9; 12)			0.0539 (0.0473)				0.0112 (0.0552)
D.Violent conflict indicator				-0.686* (0.408)			-0.626* (0.354)
LD.Violent conflict indicator				-0.884** (0.434)			-0.797* (0.455)
L2D.Violent conflict indicator				-1.283*** (0.402)			-0.816* (0.423)
D.Inflation (%)					-0.0130 (0.0116)		-0.00152 (0.00995)
LD.Inflation (%)					-0.00781 (0.0128)		0.00390 (0.0111)
L2D.Inflation (%)					-0.000365 (0.0103)		0.00466 (0.0100)
D.Log exchange rate						-1.043*** (0.254)	-0.602** (0.270)
LD.Log exchange rate						-0.121 (0.177)	-0.212 (0.151)
L2D.Log exchange rate						0.103 (0.132)	0.0857 (0.127)
D.Log terms of trade						5.805** (2.248)	3.371 (2.646)
LD.Log terms of trade						9.176*** (1.975)	9.715*** (2.142)
L2D.Log terms of trade						2.694 (2.251)	4.469** (2.156)
Observations	6653	6253	6253	6653	5494	5653	4819

Notes: The dependent variable is GDP per capita growth. The violent conflict indicator is from the Uppsala Conflict Dataset v23.1 (Gleditsch et al., 2002). Inflation and the exchange rate are from the World Bank WDI dataset. The terms are trade are from (Gruss and Kebhaj, 2019). See additional notes in Table 1 in the main text.

Table A.4: Weather shocks on GDP per capita growth: robustness to estimation choices

	(A)	(B)	(C)	(D)	(E)
Lag-1 GDP p.c. Growth	0.215*** (0.0332)	0.160* (0.0825)	0.138 (0.0857)	0.0918 (0.120)	0.233*** (0.0403)
Lag-2 GDP p.c. Growth	0.0896*** (0.0201)	0.0773*** (0.0220)	0.117*** (0.0253)	0.0223 (0.0362)	0.115*** (0.0299)
Lag-3 GDP p.c. Growth			0.0755*** (0.0169)	0.102*** (0.0144)	
Harsh Drought Prevalence (W)	-0.225*** (0.0532)	-0.182*** (0.0572)	-0.248*** (0.0598)	-0.239*** (0.0615)	-0.229*** (0.0514)
Max T °C above 35 (W)	-0.168** (0.0670)	-0.172** (0.0733)	-0.143* (0.0807)	-0.129 (0.0794)	-0.175*** (0.0605)
Mean T °C in [9; 12)	0.144*** (0.0377)	0.0709 (0.0660)	0.117** (0.0464)	0.117** (0.0467)	0.129*** (0.0474)
Constant	0.0122*** (0.000627)	0.0127*** (0.00139)	-0.0175** (0.00740)		
Observations	6,653	6,680	6,253	6,253	6,604
Arellano-Bond test for zero autocorrelation in first-differenced errors					
p-values (H0: No autocorrelation)					
Order 1			.0013	.0078	
Order 2			.0519	.9915	

Notes: The dependent variable is GDP per capita growth and all specifications include country and year fixed effects. Column A refers to the same estimation as in the first column in Table 1 in the main text (see additional notes there for more details). Column B shows the result for the same specification but without the exclusion of outliers. Column C reports results from the Arellano-Bond GMM estimator. Results in column D are from the Arellano-Bover/Blundell-Bond estimator with lagged dependent variables of order 4 or higher as instruments. The bottom panel reports the results of Arellano-Bond tests. Rejecting the null hypothesis of no serial correlation in the first-differenced errors at order zero does not imply model misspecification because the first-differenced errors are serially correlated by assumption. Rejecting the null hypothesis of no serial correlation in the first-differenced errors at an order greater than one implies model misspecification, but the hypothesis is accepted at the 5% level in columns C and D. Column E reports the results obtained with the half-panel jackknife FE-TE estimator that is robust to small size distortions as proposed by Chudik et al. (2016) and used in Kahn et al. (2021). Standard errors in column A are clustered by country and column B-C-D are robust to unspecified heteroskedasticity.

Table A.5: Estimation of the effect of climate on GDP growth: comparisons with the literature

	Burke et al. (2015)			Kahn et al. (2021)		
	(A) base	(B) unchanged	(C) augmented	(D) base	(E) unchanged	(F) augmented
Lag-1 GDP p.c. Growth	0.0742* (0.0423)	0.0747* (0.0418)	0.0786* (0.0415)	0.234*** (0.0465)	0.234*** (0.0467)	0.236*** (0.0465)
Lag-2 GDP p.c. Growth	-0.0159 (0.0203)	-0.0162 (0.0203)	-0.0150 (0.0203)	0.145*** (0.0362)	0.145*** (0.0364)	0.145*** (0.0367)
Lag-3 GDP p.c. Growth				0.000686*** (0.000211)	0.000666*** (0.000213)	0.000643*** (0.000215)
Lag-4 GDP p.c. Growth				-0.0000984 (0.000189)	-0.0000797 (0.000189)	-0.0000588 (0.000190)
Average Annual Temperature		0.00909*** (0.00305)	0.00792** (0.00308)			
– squared		-0.000374*** (0.0000865)	-0.000276*** (0.0000921)			
Average Annual Precipitation		0.0000186* (0.00000981)	0.00000275 (0.00000904)			
– squared		-5.97e-09** (2.52e-09)	-3.27e-09 (2.27e-09)			
Harsh Drought Prevalence (W)			-0.313*** (0.0777)			-0.294*** (0.0654)
Max T °C above 35 (W)			-0.200** (0.0979)			-0.228*** (0.0694)
Mean T °C in [9; 12]			0.198*** (0.0503)			0.197*** (0.0593)
Temperature Deviations from Trend				-0.0275 (0.0264)		0.00543 (0.0272)
– first lag				-0.0444 (0.0331)		-0.0484 (0.0329)
– second lag				-0.104*** (0.0396)		-0.109*** (0.0392)
– third lag				-0.132*** (0.0362)		-0.128*** (0.0357)
– fourth lag				-0.0526* (0.0299)		-0.0511* (0.0297)
Precipitation Deviations from Trend				-0.0757 (0.0663)		-0.0620 (0.0661)
– first lag				-0.0707 (0.0730)		-0.0493 (0.0727)
– second lag				-0.0401 (0.0774)		-0.0231 (0.0772)
– third lag				-0.0448 (0.0756)		-0.0377 (0.0751)
– fourth lag				-0.0485 (0.0626)		-0.0518 (0.0623)
$\hat{\theta}_{\Delta \bar{T}_{it}(m)}$					-0.581*** (0.217)	-0.536** (0.215)
$\hat{\theta}_{\Delta \bar{P}_{it}(m)}$					-0.451 (0.419)	-0.363 (0.419)
Observations	4,214	4,214	4,214	4,500	4,500	4,500

Notes: $\hat{\theta}_{\Delta|\bar{T}_{it}(m)}$ and $\hat{\theta}_{\Delta|\bar{P}_{it}(m)}$ are the estimated long-term effects of temperature and precipitation deviations from trend as calculated in Kahn et al. (2021). See additional notes in Table 2 in the main text.



Table A.6: Additional estimates of macro-fiscal effects of selected climate variables

	(A) $\Delta \ln$ Revenue (p.c.)	(B) $\Delta \ln$ Expenditure (p.c.)	(C) $\Delta \ln$ Debt (p.c.)
Lag-1 dependent variable	-0.0829** (0.0348)	-0.0704** (0.0274)	0.146*** (0.0293)
Lag-2 dependent variable	-0.0515* (0.0281)	-0.0358 (0.0335)	0.0390** (0.0151)
First difference in			
Harsh Drought Prevalence (W)	0.0230 (0.195)	0.0822 (0.178)	-0.224 (0.205)
Max T °C above 35 (W)	-0.730** (0.287)	-0.152 (0.225)	-0.141 (0.216)
Mean T °C in [9; 12)	0.264** (0.118)	-0.0100 (0.123)	-0.0520 (0.205)
Lag-1 Longest Day Cold Wave (W)	-0.368* (0.197)	0.0439 (0.210)	0.172 (0.241)
Mean Wet Day PPT	-0.351 (0.240)	-0.560*** (0.190)	0.188 (0.264)
Lag-1 PPT Minimum	0.146 (0.179)	-0.0680 (0.177)	-0.677*** (0.209)
Constant	4.406*** (0.205)	4.287*** (0.197)	3.450*** (0.130)
Observations	3,890	3,890	3,890
R-squared	0.120	0.0733	0.201

Note: The dependent variables are indicated in the column titles and are expressed in percentage points. We use the same three climate variables used for GDP growth and the first climate variables selected by the LASSO respectively for government revenue, expenditure, and debt. All climate variables are standardized with standard deviations equal to 100 to ease interpretation. ^[W] indicates population-weighted variables. Controls include the first two lags of the dependent variable (reported in the first two rows), and year and country fixed effects. Standard errors are clustered by country and reported in brackets.



Figure A.1: Elastic Net climate variable selection impacting (GDP baseline FE specification)

(a) Variations of the BIC with ϕ (fixed $\lambda = 0.04$) (b) Variations of the BIC with λ (fixed $\phi = 0.80$)





(c) Variations of the AIC with ϕ (fixed $\lambda = 0.03$)

(d) Variations of the AIC with λ (fixed $\phi = 0.70$)



(e) Variations of the out-of-sample within R2 with ϕ (fixed $\lambda = 0.027$)

(f) Variations of the out-of-sample within R2 with λ (fixed $\phi = 0.648$)



--- BIC (LHS) --- AIC (LHS) Out of Sample Within R-Squared (LHS) — # of Selected Variables (RHS)

Note: The figures show the results of implementing the EN for different penalty parameters α and λ . In each graph, we fix one of these two parameters, setting it to its optimal value, and vary the other one. The red lines show how the number of selected variables vary with each parameter. The various grey dashed lines show the variation of different criteria with each parameter. The within R-squared is calculated on a sub-sample of countries (evaluation set) based on coefficients estimated on the rest of countries (training set) as explained in the main text. The estimated model has GDP per capita growth as the dependent variable and includes country and year effects. The dots indicate the different selection outcomes given by the local optimum for each criteria respectively.

B. ONLINE Appendix

B.1. Source Data

We use weather from the ERA5 dataset from 1979 to 2021.²⁶ The original ERA5 dataset has hourly data but we use data aggregated at daily level by Google Earth Engine (GEE).²⁷ This includes daily mean temperature in each day d and grid cell j , calculated using ERA5’s 24 measures per day ($T_{j,d}$), the minimum of those 24 measures within a day ($TN_{j,d}$), and the maximum of those 24 measures within a day ($TX_{j,d}$). Total daily precipitation ($P_{j,d}$) is calculated by summing all the hourly precipitation measures within a day. From these daily grid-cell data points we construct all our variables.

Number of observations in the original databases. The resolution of ERA5 data is 0.25 degrees. A global map has 180 degrees along the North-South dimension and 360 degrees along the West-East dimension: the total number of cells is therefore equal to $(180/0.25) \times (360/0.25) = 1,036,800$. The percentage of Earth’s surface covered land, after excluding Antarctica and Greenland, is approximately equal to 27%. This means that we use approximately $1,036,800 \times 0.27 = 279,936$ cells on land. For each grid and each day of the 41 years from 1979 to 2019 we have four weather data points (T , TN , TX , and P). This means that we start with approximately $279,936 \times 365 \times 41 \times 4 = 16,756,968,960$ (≈ 17 billion) temperature and precipitation data points.

The Palmer Drought Severity Index (PDSI) is from [Abatzoglou et al. \(2018\)](#) and is accessed using GEE.²⁸ PDSI data comes at monthly intervals with spatial resolution equal to 0.0416 degrees. This corresponds to $(180 / .0416) \times (360 / 0.0416) \times 0.27 = 10,110,022$ cells on land excluding Antarctica and Greenland. Summing over all months from January 1979 to December 2019 we have a total of $10,110,022 \times 12 \times 41 = 4,974,130,917$ (≈ 5 billion) observations on PDSI from the Terra Climate data.

To sum up, we start with 21,715,195,392 (≈ 22 billion) data points on temperature, precipitation, and the PDSI.

Merging datasets and zonal statistics. We merge the ERA5 and PDSI datasets into one single geospatial dataset that uses the higher resolution of PDSI data of approximately $5 \text{ km} \times 5 \text{ km}$ at the equator. This dataset is projected on a global map of countries to calculate zonal statistics at country level.²⁹ The whole process is managed using Google Earth Engine and delivers a total of 9,621,976 (≈ 10 million) country-matched grid cells for each one of our five core climate variables ($T_{j,d}$, $TN_{j,d}$, $TX_{j,d}$,

²⁶<https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels?tab=overview>

²⁷https://developers.google.com/earth-engine/datasets/catalog/ECMWF_ERA5_DAILY#description

²⁸https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE#bands

²⁹Zonal statistics are operations that calculate statistics of cell values of a dataset (raster) within boundaries defined by another dataset.

$P_{j,d}$, and $PDSI_{j,d}$). Each grid cell has daily data for 41 years. This means we develop our full set of climate variables using 580,705,495,552 (≈ 600 billion) data points.

Weighted variables. The resolution for the population data is $\approx 1 \text{ km} \times 1 \text{ km}$ at the equator,^{30,31} and hence for the weighted data we use $1 \text{ km} \times 1 \text{ km}$ grid cells during zonal statistics. By mixing population and weather data we obtain 25 additional points for each grid cell of the raw weather data. This adds $25 \times 580,705,495,552 = 14,517,637,388,800$ (≈ 15 trillion) data points to our dataset for zonal statistics.

B.2. Definition of weather variables

This Section describes all the weather variables we construct from raw precipitation and temperature data. We start by an overview of weather variables, then give a brief presentation of mathematical notations and concepts, and finally provide the full list of the variables we construct and their mathematical definitions in table A.1.

Temperature variables. For each day in a year and country, we calculate country-wide averages of daily average, minimum, and maximum temperature (respectively T_d , TN_d , and TX_d) from daily grid level data. We aggregate average daily temperatures to get annual average temperature (T), the variance of daily temperature ($TVar$). We calculate the average diurnal temperature range (DTR) from minimum and maximum daily temperatures. Using the 10th and 90th percentiles of the 1979-2019 distribution of TN_d and TX_d in a 5-day window centered on each day of the year, we calculate the number of cold nights ($CN10$), cold days ($CD10$), warm nights ($WN90$) and warm days ($WD90$), to characterize cold and heat extremes using relative thresholds.

To account for impacts from extended exposure to temperature extremes, we build variables to capture heatwaves and coldwaves based on the climate literature. We follow Kim et al. (2020) and we define cold (warm) spell duration (CSD , WSD) as the number of days in which TN_d (TX_d) is below (above) the 10th (90th) percentile of the 1979-2019 distribution in a 5-day window centered on each day, for at least six consecutive days. We follow Perkins and Alexander (2013) to define eight additional indicators of day (night) heat waves based on exceeding the 90th percentile of the 1979-2019 distribution of TX_d (TN_d) in a 15-day window centered on each day, for at least three consecutive days. We count the number of days with day (night) heat wave, the length of the longest day (night) heatwave, the number of day (night) heatwaves during a year, and the average maximum (minimum) temperature during day (night) heatwaves. Similarly, we use the 10th percentile of the distribution of TX_d and TN_d to measure the characteristics of day and night cold waves.

We construct country averages of grid-level annual minimum of minimum daily temperature (TNn) and

³⁰<https://sedac.ciesin.columbia.edu/data/set/gpw-v4-population-count-adjusted-to-2015-unwpp-country-totals-rev11>

³¹https://developers.google.com/earth-engine/datasets/catalog/CIESIN_GPWv411_GPW_UNWPP-Adjusted_Population_Count

of grid-level maximum of maximum daily temperature (TXx), both used in the climate literature.

We also define another set of extreme temperature variables using absolute temperature thresholds based on the climate literature (e.g., IPCC, 2021a). With absolute temperature thresholds, using the highest possible level of spatial resolution is essential to avoid missing the potentially harmful events that can get averaged out over large areas. For example, if two grid cells have maximum daily temperature equal to, respectively, 33 °C and 36 °C, their average is equal to 34.5 °C, lower than the frequently used 35 °C threshold. By first averaging and then checking if the threshold is crossed, we would record zero extreme events, while temperature in 50% of the grid cells exceeds the threshold. The same does not apply to extremes measured using relative thresholds.


Therefore, when we use absolute thresholds, we sum the number of times a threshold is crossed in each grid cell and in each day, across all days and grid cells in a country, and then divide that number by the total number of grid-day observations ($J \times 365$). We do so to find the share of grid-days with frost (minimum daily temperature below 0°C – $TN0$), with maximum temperature above 35 °C ($TX35$) and above 40°C ($TX40$).

Finally, to capture potential non-linear effects of temperature on macroeconomic variables, we divide the distribution of temperature into 3 °C-wide intervals and we measure the share of grid-day observations in each interval (e.g., Schlenker and Roberts, 2009). For example, Figure B.1 illustrates the calculation of the share of grid-days that experiences temperature levels between x_1 and x_2 degrees Celsius. By using 3 °C wide intervals we aim to balance flexibility in modeling the temperature response function and avoiding multicollinearity problems that would arise from using narrower temperature intervals (Mérel and Gammans, 2021). One of the intervals is omitted in our estimation process to avoid perfect collinearity among all interval indicators. As very low and very high average daily temperatures are rare, all the days with average temperature below -9 °C and at or above 30 °C are grouped in two terminal intervals.

Precipitation (rain or snow) variables. We start by calculating the average of total daily precipitation in each country across all grid cells (P_d). We use this variable to construct annual average precipitation (P) and the annual variance of daily precipitation ($PVar$) for every country. Following the climate literature, we focus on days that have more than 1 mm of precipitation, which are called “wet days”. We calculate the number of wet days (W), average daily precipitation in wet days (PWA), and wet days precipitation variance ($PWVar$). We calculate total precipitation in very wet ($PW95T$) and extremely wet days ($PW99T$) using the 95th and 99th percentiles of the distribution of wet days over all days and years from 1979 to 2019.

We build several variables to capture extended wet and dry periods. We count the largest number of consecutive dry days (days with precipitation less than 1 mm — CDD), the largest number of consecutive wet days (CWD) and total precipitation during the longest wet days period ($PCWD$). To focus on extreme conditions, we count the number of consecutive very ($PC95WD$) and extremely ($C99WD$) wet

Figure B.1: Computing the share of grid-days with weather conditions in a specific interval



Notes: This figure illustrates the calculation methodology for “Share of Grid-Days with Mean Temperature in the interval $[x_1, x_2]$ ” (Mean $T \text{ } ^\circ\text{C}$ in $[x_1, x_2]$ — $TS_{[x_1, x_2]}$) for a given day d in any country j . We also zoom on Algeria. The grid cells colored in red represent the locations where $x_1 \leq T_{j,d} < x_2$ and grid cells colored in gray represent the locations where $T_{j,d}$ (average daily temperature in country j on day d) is outside of this range. For our study, we later obtain country-year measures by averaging daily percentages over the 365 days of a year. Note that the grid cells are pictured as much bigger than they are in the original dataset for visualization purposes. For example, there are 50 grid cells belonging to Algeria in this figure. However, there are more than 105 thousand grid cells in Algeria in the dataset.

days in the longest periods with daily precipitation above the 95th and 99th percentiles of the distribution, respectively. Similarly, we calculate total precipitation in consecutive very ($PC95W$) and extremely wet days ($PC99WD$).

To capture intense precipitation that may cause floods, which are among the most destructive climate disasters, we use the maximum amount in a year of rainfall in 1-day ($PX1$) or 5-day ($PX5$) intervals. To capture extreme precipitation at the local level, we use total monthly precipitation in each grid cell and we calculate the country average of maximum ($PX(1Month)$) and minimum ($PN(1Month)$) monthly precipitation.

As for temperature, precipitation extremes can also be characterized using absolute thresholds but this requires calculations at the grid level. We calculate the length of the longest dry spell ($LLDS$) in a country as the uninterrupted series of days in which a minimum percent of the country area has daily total precipitation less than 1 mm (dry day). We use four thresholds to identify dry spells and consider spells affecting 50%, 65%, 80% and 95% of a country area. Similarly to what we do with temperature intervals, we calculate the share of total grid-days with total precipitation in four intervals: less than 1 mm, from 1 mm to 10 mm, from 10 mm to 20 mm, and above 20 mm. The maximum extent of heavy

($MaxP_{>10}$) and very heavy ($MaxP_{>20}$) precipitation is equal to the maximum daily share of the country with precipitation respectively greater than 10 mm and 20 mm. To capture deviations from conditions with balanced level of precipitation across time and space, we develop an indicator that measures the absolute deviation from having 50% of the grid-days observations of precipitation between 1 and 10 mm ($BP_{1,10}(0.5)$).

Wetness and drought variables. Finally, we use the Palmer Drought Severity Index (PDSI) (Palmer, 1965) to introduce a measure of dry and wet periods that combines temperature and precipitation data to estimate cumulative deviations in soil moisture from normal conditions (Dai et al., 2004; Abatzoglou et al., 2018; Lai et al., 2020).³² The PDSI ranges from -10 to +10, but values below -4 and above +4 are rare. We build variables measuring the share of total grid-months subject to extreme droughts ($PDSI < -4$), extreme and severe droughts ($PDSI < -3$), periods with extreme moisture ($PDSI > 4$), and periods with very high and extreme moisture ($PDSI > 3$). For each of these four categories and in every country, we also build variables reflecting the maximum extent of these events, that is the share of affected grid-cells in the month where the share is at its maximum.

Mathematical notations and concepts. We use d to denote calendar days, months with m , and $j = 1, \dots, J$ to denote grid cells in every country. For ease of notation, we do not index variables by country and year. In each year there are 12 months and for ease of notation we assume each year has the same number of days.

We use Iverson brackets in the definition of many variables. Iverson brackets map any statement inside brackets into a function that takes the value of the variables for which the statement is true, and take the value zero otherwise.³³ It is denoted by putting the statement inside square brackets:

$$[X] = \begin{cases} 1 & \text{if } X \text{ is true;} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, to count days in which a certain condition X is met we write: $\sum_d [X]$.

Some variables capture different percentiles of the long-term distribution of daily mean temperature and daily precipitation. We use the whole time horizon of our dataset for these distributions, from 1979 to 2019. This represents a 41-year time window that is well-suited to capture extreme realizations of temperature and precipitation.

For daily precipitation, we use all days of the calendar year as there are no obvious seasonal patterns

³²Data downloaded from Google Earth Engine. See <http://www.climatologylab.org/terraclimate.html> and https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE for a detailed description of the datasets.

³³Donald Knuth, "Two Notes on Notation" American Mathematical Monthly, Volume 99, Number 5, May 1992, pp. 403–422.

that apply to all countries. For temperature, there is a more marked seasonal cycle in most countries and deviations from normal conditions are more clearly dependent on the time of the year temperature is observed. For this reason, the distribution of temperature is restricted to moving windows centered on the day of interest. We use 5-day and 15-day windows following the literature [Kim et al. \(2020\)](#); [Perkins and Alexander \(2013\)](#). For example, consider August 16, 2000. To check whether precipitation is extreme, we compare daily precipitation with the distribution of precipitation over all days from 1979 to 2019. To check if temperature is extreme, we restrict the distribution of daily mean temperature to August 14, 15, 16, 17, and 18 (with a 5-day window) from 1979 to 2019.

B.3. Summary Statistics

Between and within variance. Our empirical and identification approach relies on inter-annual variation within country. Therefore, we use a standard approach to decompose the variance of variables into *between* and *within* components.

For any variable x , the variance across N countries and over T years can be decomposed by introducing country averages \bar{x}_i . The variance is equal to $\sum_{i,t} \frac{(x_{i,t} - \bar{x})^2}{NT} = \sum_{i,t} \frac{(x_{i,t} - \bar{x}_i)^2}{NT} + 2 \sum_{i,t} \frac{(x_{i,t} - \bar{x}_i)(\bar{x}_i - \bar{x})}{T} + \sum_{i,t} \frac{(\bar{x}_i - \bar{x})^2}{NT}$. It simplifies to $\sum_{i,t} \frac{(x_{i,t} - \bar{x})^2}{NT} = \sum_{i,t} \frac{(x_{i,t} - \bar{x}_i)^2}{NT} + \sum_i \frac{(\bar{x}_i - \bar{x})^2}{N}$ where the terms are respectively the within and between variance. We take the square roots of each component to obtain *between*- and *within*-country standard deviations.

The *between* standard deviation measures variation of average country weather around the global mean. The *within* standard deviation measures the average deviation from country averages.

Trends in weather variables. Table B.1 reports tests of trends in the levels of the weather variables. For each variable and each country we estimate a linear regression of the form $w_t = \alpha + \beta t + u_t$, where w_t is the value taken by the weather variable in year t , u_t is a random component and β is the country-specific trend coefficient.

Column A reports the average β across all countries. Our results are not truly indicative of global trends, because we use country-level observations instead of area-weighted averages. For an accurate assessment of climate trends, it is important to rely on conclusions from climate science ([IPCC, 2021b](#)). However, the positive trend for average annual temperature is equal to 0.03 °C per year, a value remarkably in line with the average decadal increase of temperature equal to 0.3 °C found by the IPCC WG I.

Column B shows the percentage of countries for which the trend is significantly different from zero at the 5 percent confidence level. We use this percentage value to rank variables in decreasing order. Most of the variables built using temperature show a significant trend consistent with global warming in the majority of countries, and in some cases in virtually all countries. Variables built using precipitation do not generally show a trend that is significant for the majority of countries and in most cases trends are

not significant for more than 2/3 of the countries.

Our model specification (see equation (2)) effectively removes trends in climate variables only if the trend is time invariant. To assess weather trends change over time, we conduct a test for a structural trend break with unknown break date in the time series of each climate variable, separately in each country. In column C we report the percentage of countries with both a significant trend and a significant break in the trend.³⁴ There is evidence of a trend with a structural break for more than 50 percent of the countries only for few variables. This suggests that our method, albeit imperfectly, helps to remove trends in weather variables.

Table B.1: Trends in weather variables

	(A)	(B)	(C)
	Average trend	Significant trend (% of countries)	Significant trend and break (% of countries)
Mean Temperature	0.0292	99%	50%
# of Warm Nights	1.4904	97%	71%
# of Warm Days	1.3516	96%	62%
# of Cold Days	-1.0564	95%	51%
# of Cold Nights	-1.2086	94%	64%
# of Day Cold Waves	-0.0858	90%	46%
# of Night Heat Waves	0.1175	89%	65%
Cold Wave Days	-0.5272	89%	53%
Heat Wave Nights	0.7510	89%	62%
# of Day Heat Waves	0.1094	88%	57%
Heat Wave Days	0.6824	88%	60%
Cold Wave Nights	-0.5965	87%	57%
# of Night Cold Waves	-0.0979	87%	56%
Longest Day Heat Wave	0.1574	81%	47%
Mean T °C in [27; 30)	0.0018	80%	52%
Longest Night Heat Wave	0.1737	78%	49%
Max T °C above 35	0.0008	76%	42%
Day T °C Maximum	0.0322	75%	42%
Warm Spell Duration	0.4343	74%	53%
Longest Night Cold Wave	-0.1439	73%	54%
Longest Day Cold Wave	-0.1313	72%	42%
Cold Spell Duration	-0.3679	72%	56%
Mean T °C in [24; 27)	-0.0008	70%	54%
Frost prevalence	-0.0009	69%	33%
Night Heat Wave T °C	0.0785	67%	54%
Mean T °C in [21; 24)	-0.0006	66%	43%
Mean T °C above 30	0.0006	66%	35%
Day Heat Wave T °C	0.0748	58%	54%
Max T °C above 40	0.0003	58%	36%
Mean T °C in [-6; -3)	-0.0002	58%	39%
Diurnal T °C Range	0.0051	53%	68%
Mean T °C in [15; 18)	-0.0001	53%	36%
Mean T °C in [18; 21)	-0.0001	52%	36%

Continued on next page

³⁴More precisely, we test if the null hypothesis of no structural break can be rejected at the 95 percent confidence level using a supremum Wald test which is the less restrictive of those commonly used.

Table B.1 (Continued): trends in weather variables

Mean T °C in [-3; 0)	-0.0002	48%	34%
Night T °C Minimum	0.0266	48%	39%
Mean T °C in [-9; -6)	-0.0002	45%	35%
Mean T °C in [0; 3)	-0.0002	44%	30%
Mean T °C in [12; 15)	-0.0001	41%	41%
Balanced PPT Indicator	0.0003	38%	63%
Mean T °C in [9; 12)	-0.0001	37%	37%
Night Cold Wave T °C	0.0039	36%	42%
Drought Intensity	0.0039	35%	75%
Less than 1 mm PPT	0.0006	35%	65%
Mean T °C in [3; 6)	-0.0001	35%	32%
# of Wet Days	-0.2077	34%	59%
Mean T °C in [6; 9)	0.0000	34%	33%
Day Cold Wave T °C	0.0155	32%	46%
Harsh Drought Intensity	0.0037	31%	63%
Drought Prevalence	0.0026	29%	61%
Above 20 mm PPT	0.0001	28%	50%
Very Wet Day PPT	1.6213	28%	46%
Wetness Intensity	0.0008	27%	73%
Precipitation Variance	0.1271	27%	42%
10 to 20 mm PPT	-0.0001	26%	57%
Harsh Drought Prevalence	0.0019	26%	48%
High Wetness Intensity	0.0016	26%	59%
Wet Day PPT Variance	0.1818	26%	38%
Wet Conditions	0.0011	25%	59%
Mean Wet Day PPT	0.0049	24%	48%
PPT Maximum	0.0003	24%	44%
Mean Precipitation	0.0008	23%	58%
Cont'd Wet Days	-0.1937	21%	41%
Very Wet Conditions	0.0012	21%	43%
Longest Dry Spell (.80)	0.0541	20%	40%
Cont'd Dry Days	0.0774	19%	35%
Longest Dry Spell (.65)	0.1109	19%	41%
Extremely Wet Day PPT	0.7905	19%	37%
1-Day PPT Maximum	0.1054	18%	37%
Cont'd Wet Day PPT	-0.6595	18%	37%
5-Day PPT Maximum	0.1418	17%	34%
PPT Minimum	0.0000	17%	31%
Longest Dry Spell (.95)	0.0241	15%	39%
Extreme PPT Maximum	0.0008	15%	36%
Longest Dry Spell (.5)	0.1123	15%	40%
Cont'd Heavy PPT	0.1693	14%	32%
Cont'd Very Wet Day PPT	0.0034	12%	32%
Temperature Variance	0.0046	12%	38%
Heavy PPT Maximum	0.0000	8%	33%
Cont'd Extreme PPT	0.1885	7%	18%
Cont'd Extra Wet Day PPT	0.0046	7%	20%

Macroeconomic variables. Summary statistics for the macro-fiscal variables used in our analysis are shown in Table B.2. The Table displays separately the growth rate of GDP per capita in the larger sample used for the analysis of weather impacts on GDP growth, and the growth rate of GDP per capita in the smaller sample used for the analysis of fiscal impacts.

Table B.2: Summary statistics of macro-fiscal variables

Summary Statistics of First Differences	N	Mean	St. Dev.	St. Dev. Between	St. Dev. Within
$\Delta \ln(GDP/POP)$ in GDP Growth Sample (p.c.)	6,653	1.726%	4.63%	1.77%	4.33%
$\Delta \ln(GDP/POP)$ in Fiscal Sample (p.c.)	3,890	2.005%	3.85%	1.55%	3.56%
Δ Revenue-to-GDP (p.p.)	3,890	0.064%	2.90%	0.71%	2.86%
Δ Expenditure-to-GDP (p.p.)	3,890	0.045%	3.61%	0.92%	3.56%
Δ Balance-to-GDP (p.p.)	3,890	0.019%	4.12%	0.71%	4.09%
Δ Debt-to-GDP (p.p.)	3,890	0.177%	8.1%	1.96%	7.9%
Δ Revenue (p.c.)	3,890	3.873%	11.6%	2.50%	11.4%
Δ Expenditure (p.c.)	3,890	3.867%	10.8%	2.51%	10.6%
Δ Debt (p.c.)	3,890	4.221%	16.3%	4.91%	15.8%

Notes: GDP per capita is measured by the difference of log GDP capita. Government revenue, Government expenditure and Government Debt growth are measured by the difference of log variables. All fiscal variables are measured as percentage of GDP and first differences are measured in percentage points.

Table B.3: Summary statistics of climate variables

Summary Statistics of First Differences	Mean	St. Dev. Between	St. Dev. Within
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0050	0.0097	0.1741
Max T above 35 °C (W) ($TX35$)	0.0006	0.0011	0.0205
Mean T in [9; 12] °C ($TS_{9,12}$)	0.0000	0.0007	0.0159
Longest Day Cold Wave ($LDCW$)	-0.1130	0.2584	5.903
Mean Wet Day PPT (PWA)	0.0025	0.0531	0.9897
PPT Minimum (PNM)	0.0000	0.0007	0.0161

Notes: Summary statistics of first differences of all weather variables used for either GDP analysis, including robustness tests, or for analysis of macro-fiscal outcomes. (W) indicates population-weighted variables. The sample of the baseline specification is used for all climate variables.

Correlation analysis. The analysis of raw correlations between GDP growth and the explanatory variables selected by the LASSO for our main specification is displayed in Table B.4. Correlations between GDP growth and first differences of weather variables are generally small. Correlation is negative for Max T °C above 35 and Harsh Drought Prevalence, and positive for Mean T °C in [9; 12]. The same relationships are confirmed in our baseline regression analysis (see Table 1).

We also display the correlation of GDP growth with both average annual temperature and annual precipitation even if these two variables are not selected by the LASSO because they are the only two weather variables typically used in the literature. The correlation between GDP growth and both temperature and precipitation is very low and much lower than for our selected weather variables. This is preliminary evidence that the literature may miss a large fraction of climate induced variation in GDP growth. Interestingly, the largest correlations among climate variables are between Average Temperature and Harsh Drought Prevalence and between Mean Temperature and Max T °C above 35, but the LASSO always selects Harsh Drought Prevalence and Max T °C above 35 instead of Mean Temperature to explain GDP growth.

Table B.4: Correlation Matrix Between Baseline Variables

	GDP Growth	Lag(1) of GDP Growth	Lag(2) of GDP Growth	Harsh Drought Prevalence (W)	Max T °C above 35 (W)	Mean T °C in [9; 12]	Mean Temperature
Lag(1) of GDP Growth	0.366						
Lag(2) of GDP Growth	0.248	0.281					
Harsh Drought Prevalence (W) ($PDSI < -4$)	-0.058	0.025	0.012				
Max T above 35 °C (W) ($TX35$)	-0.040	0.016	0.021	0.180			
Mean T in [9; 12] °C ($TS_{9,12}$)	0.040	-0.016	0.014	-0.032	-0.033		
Average T (T)	-0.015	0.020	-0.001	0.156	0.362	-0.053	
Mean Precipitation (P)	0.014	-0.006	-0.002	-0.196	-0.145	0.081	-0.087

Notes: These correlations are computed using first differences using the baseline regression sample. (W) indicates population-weighted variables.

Table B.5: Summary Statistics for Sub-Groups

	Mean	St. Dev. Between	St. Dev. Within	Mean	St. Dev. Between	St. Dev. Within
		Hot (N=3,315)			Cold (N=3,338)	
Δ GDP p.c.	1.38	1.85	4.44	2.07	1.59	4.22

Continued on next page

Table B.5 (Continued): summary statistics for Sub-Groups

Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0032	0.0079	0.1668	0.0068	0.0108	0.1811
Max T above 35 °C (W) ($TX35$)	0.0008	0.0014	0.0257	0.0004	0.0007	0.0137
Mean T in [9; 12] °C ($TS_{9,12}$)	-0.00006	0.0003	0.0041	0.0002	0.0009	0.0220
	Agricultural (N=3,119)			Non Agricultural (N=3,107)		
Δ GDP p.c.	1.71	1.77	4.55	1.75	1.46	3.89
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0043	0.0085	0.1833	0.0055	0.0091	0.1632
Max T above 35 °C (W) ($TX35$)	0.0009	0.0011	0.0257	0.0004	0.0008	0.0140
Mean T in [9; 12] °C ($TS_{9,12}$)	0.00003	0.0005	0.0103	0.0001	0.0008	0.0193
	Agricultural Hot (N=1,785)			Agricultural Cold (N=1,334)		
Δ GDP p.c.	1.37	1.64	4.10	2.16	1.84	5.08
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0025	0.0062	0.1803	0.0067	0.0101	0.1872
Max T above 35 °C (W) ($TX35$)	0.0011	0.0012	0.0310	0.0006	0.0009	0.0160
Mean T in [9; 12] °C ($TS_{9,12}$)	0.0000	0.0001	0.0033	0.0001	0.0007	0.0153
	Rich (N=3,936)			Poor (N=2,717)		
Δ GDP p.c.	1.95	1.77	4.27	1.41	1.75	4.41
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0055	0.0102	0.1697	0.0042	0.0088	0.1803
Max T above 35 °C (W) ($TX35$)	0.0004	0.0008	0.0141	0.0010	0.0014	0.0273
Mean T in [9; 12] °C ($TS_{9,12}$)	0.0001	0.0009	0.0194	0.0000	0.0002	0.0083
	EAP (N=1,013)			ECA (N=1,632)		
Δ GDP p.c.	2.42	2.23	3.89	2.31	1.48	4.49
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0017	0.0060	0.1683	0.0100	0.0112	0.1892
Max T above 35 °C (W) ($TX35$)	0.0006	0.0014	0.0240	0.0002	0.0004	0.0089
Mean T in [9; 12] °C ($TS_{9,12}$)	-0.00004	0.0001	0.0054	0.0004	0.0011	0.0283
	MENA (N=620)			SSA (N=1,656)		
Δ GDP p.c.	0.97	1.86	5.16	1.07	1.52	4.66
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0003	0.0064	0.1961	0.0048	0.0096	0.1924
Max T above 35 °C (W) ($TX35$)	0.0011	0.0018	0.0246	0.0010	0.0012	0.0285
Mean T in [9; 12] °C ($TS_{9,12}$)	-0.0003	0.0009	0.0196	-0.0001	0.0003	0.0052
	LAC (N=1,372)			Base (N=6,550)		
Δ GDP p.c.	1.36	1.47	3.87	1.73	1.77	4.33
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0056	0.0086	0.1388	0.0050	0.0097	0.1741
Max T above 35 °C (W) ($TX35$)	0.0005	0.0009	0.0122	0.0006	0.0011	0.0205
Mean T in [9; 12] °C ($TS_{9,12}$)	-0.00002	0.0001	0.0056	0.0000	0.0007	0.0159
	High Democracy (N=3,895)			Low Democracy (N=2,753)		
Δ GDP p.c.	1.92	2.13	3.47	1.45	2.61	5.12
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0049	0.0180	0.1723	0.0049	0.0193	0.1758
Max T above 35 °C (W) ($TX35$)	0.0005	0.0031	0.0148	0.0008	0.0019	0.0265
Mean T in [9; 12] °C ($TS_{9,12}$)	0.00006	0.0011	0.0186	0.0000	0.0013	0.0109
	High Democracy and Poor (N=950)			Low Democracy and Poor (N=1,766)		
Δ GDP p.c.	1.87	1.57	3.46	1.15	2.41	4.73
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0026	0.0208	0.1785	0.0051	0.0148	0.1809
Max T above 35 °C (W) ($TX35$)	0.0013	0.0049	0.0232	0.0008	0.0021	0.0292
Mean T in [9; 12] °C ($TS_{9,12}$)	-0.00002	0.0009	0.0054	-0.00003	0.0004	0.0096
	1979-1999 (N=2,762)			2000-2019 (N=3,891)		

Continued on next page

Table B.5 (Continued): summary statistics for Sub-Groups

Δ GDP p.c.	2.13	1.94	3.68	2.13	1.94	3.68
Harsh Drought Prevalence (W) ($PDSI < -4$)	0.0061	0.0162	0.1827	0.0061	0.0162	0.1827
Max T above 35 °C (W) ($TX35$)	0.0009	0.0020	0.0190	0.0009	0.0020	0.0190
Mean T in [9; 12) °C ($TS_{9,12}$)	0.00011	0.0007	0.0171	0.0001	0.0007	0.0171

Note: Summary statistics of first difference of weather variables and GDP growth in percentage. Coefficients of weather variables are reported in Figure 5 and groups are described in the Notes to the Figure.

B.4. Additional Result Tables

Table B.6: Optimal LASSO selection of variables affecting GDP per capita growth under different fit criteria (baseline FE specification)

	BIC	AIC	OOS-Countries- R^2	OOS-Observations- R^2	
Selected variables	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12)	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Longest Night Heat Wave ^[W] , Lag-1 Mean T °C in [0,3) ^[W] , Lag-1 Cold Spell Duration, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days, Longest Night Cold Wave ^[W] , 1-Day PPT Maximum, Drought Intensity, Lag-1 PPT Minimum ^[W] , Lag-1 10 to 20 mm PPT, Lag-1 Day T °C Maximum, Lag-2 Balanced PPT Indicator, Lag-2 Mean T °C in [3,6), Cont'd Heavy PPT ^[W] , Heavy PPT Maximum ^[W] , Longest Dry Spell (.80) ^[W] , Lag-1 Mean T °C in [24,27) ^[W] , PPT Minimum ^[W] , PPT Maximum ^[W] , Cont'd Extreme PPT, Lag-1 Cont'd Wet Days ^[W] , Lag-1 Harsh Drought Prevalence ^[W] , Lag-1 Longest Dry Spell (.65) ^[W] , Lag-2 Day Heatwave T °C ^[W] , Lag-2 Longest Dry Spell (.80) ^[W] , Lag-2 Very Wet Conditions Prevalence ^[W] , Lag-2 Mean T °C in [0,3) ^[W] , Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Longest Dry Spell (.95), Lag-2 Longest Dry Spell (.5), Lag-2 Mean T °C in [21,24)	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Longest Night Heat Wave ^[W] , Lag-1 Mean T °C in [0,3) ^[W] , Lag-2 Cold Spell Duration, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days, Longest Night Cold Wave ^[W] , 1-Day PPT Maximum, Drought Intensity, Lag-1 PPT Minimum ^[W] , PPT Maximum ^[W] , Cont'd Extreme PPT, Lag-1 Cont'd Wet Days ^[W] , Lag-1 Harsh Drought Prevalence ^[W] , Lag-1 Longest Dry Spell (.65) ^[W] , Lag-2 Day Heatwave T °C ^[W] , Lag-2 Longest Dry Spell (.80) ^[W] , Lag-2 Very Wet Conditions Prevalence ^[W] , Lag-2 Mean T °C in [0,3) ^[W] , Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Longest Dry Spell (.95), Lag-2 Longest Dry Spell (.5), Lag-2 Mean T °C in [21,24)	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Longest Night Heat Wave ^[W] , Lag-1 Mean T °C in [0,3) ^[W] , Lag-2 Cold Spell Duration, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days, Longest Night Cold Wave ^[W] , 1-Day PPT Maximum, Drought Intensity, Lag-1 PPT Minimum ^[W] , PPT Maximum ^[W] , Cont'd Extreme PPT, Lag-1 Cont'd Wet Days ^[W] , Lag-1 Harsh Drought Prevalence ^[W] , Lag-1 Longest Dry Spell (.65) ^[W] , Lag-2 Day Heatwave T °C ^[W] , Lag-2 Longest Dry Spell (.80) ^[W] , Lag-2 Very Wet Conditions Prevalence ^[W] , Lag-2 Mean T °C in [0,3) ^[W] , Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Longest Dry Spell (.95), Lag-2 Longest Dry Spell (.5), Lag-2 Mean T °C in [21,24)	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Longest Night Heat Wave ^[W] , Lag-1 Mean T °C in [0,3) ^[W] , Lag-2 Cold Spell Duration, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days, Longest Night Cold Wave ^[W] , 1-Day PPT Maximum, Drought Intensity, Lag-1 PPT Minimum ^[W] , PPT Maximum ^[W] , Cont'd Extreme PPT, Lag-1 Cont'd Wet Days ^[W] , Lag-1 Harsh Drought Prevalence ^[W] , Lag-1 Longest Dry Spell (.65) ^[W] , Lag-2 Day Heatwave T °C ^[W] , Lag-2 Longest Dry Spell (.80) ^[W] , Lag-2 Very Wet Conditions Prevalence ^[W] , Lag-2 Mean T °C in [0,3) ^[W] , Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Longest Dry Spell (.95), Lag-2 Longest Dry Spell (.5), Lag-2 Mean T °C in [21,24)
Number of Selected Variables	5	36	18	18	
Optimal Penalty Weight (λ)	.0328	.0139	.019	.019	

Notes: This table shows some results of the implementation of the LASSO to select the climate variables that are best to explain GDP per capita variations after accounting for country and year fixed effects. Each column corresponds to a different fit criteria and refers to the outcomes of implementing the LASSO after setting λ to optimize that specific fit criteria. For each column, the second row shows the list of the climate variables selected by the LASSO. The optimal value of λ is presented in the last row.

Table B.7: Optimal EN selection of variables affecting GDP per capita growth under different fit criteria (baseline FE specification)

	BIC	AIC	OOS-Countries- R^2	OOS-Observations- R^2
Selected variables	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12)	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Longest Night Heat Wave ^[W] , Drought Intensity, Lag-1 Mean T °C in [0,3) ^[W] , Lag-1 Cold Spell Duration, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Longest Night Heat Wave ^[W] , Longest Night Cold Wave ^[W] , Drought Intensity, Cont'd Heavy PPT ^[W] , 1-day PPT ^[W] , 1-day PPT Maximum, Lag-1 Mean T °C in [0,3) ^[W] , Lag-1 Cold Spell Duration, Lag-1 PPT Minimum ^[W] , Lag-1 10 to 20 mm PPT, Lag-1 Day T °C Maximum, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days, Lag-2 Balanced PPT Indicator, Lag-2 Mean T °C in [3,6), Lag-2 Mean T °C in [21,24)	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , Mean T °C in [9,12), Drought Intensity, Cont'd Heavy PPT ^[W] , Longest Night Heat Wave ^[W] , Longest Night Cold Wave ^[W] , Heavy PPT Maximum ^[W] , Longest Dry Spell ^[W] , Mean T °C in [24,27) ^[W] , Mean T °C above 30 ^[W] , PPT Maximum ^[W] , PPT Minimum ^[W] , 1-day PPT Maximum, Very Wet Day PPT, Cont'd Extreme PPT, Mean T °C in [6,9), Mean T °C in [18,21), Lag-1 Mean T °C in [0,3) ^[W] , Lag-1 Cold Spell Duration, Lag-1 Wet Day PPT Variance, Lag-1 Cont'd Wet Days ^[W] , Lag-1 Harsh Drought Prevalence ^[W] , Lag-1 PPT Minimum ^[W] , Lag-1 10 to 20 mm PPT, Lag-1 Longest Dry Spell (.65), Lag-1 Mean T °C in [0,3), Lag-1 Mean T °C in [9,12), Lag-1 Day T °C Maximum, Lag-2 Mean T °C in [3,6) ^[W] , Lag-2 Cold Wave Days, Lag-2 Day Heat Wave T °C ^[W] , Lag-2 # of Night Cold Waves ^[W] , Lag-2 Longest Dry Spell (.80) ^[W] , Lag-2 High Wetness Intensity ^[W] , Lag-2 Very Wet Conditions Prevalence ^[W] , Lag-2 Mean T °C in [0,3) ^[W] , Lag-2 Mean T °C in [21,24) ^[W] , Lag-2 Balanced PPT Indicator, Lag-2 Longest Dry Spell (.95), Lag-2 Longest Dry Spell (.80), Lag-2 Longest Dry Spell (.50), Lag-2 Mean T °C in [3,6), Lag-2 Mean T °C in [21,24)
Number of Selected Variables	5	11	21	46
Optimal Penalty Weight (λ)	.04	.03	.027	.062
Optimal LASSO Ratio (ϕ)	.8	.7	.648	.215

Notes: This table shows some results of the implementation of the Elastic-Net (EN) to select the climate variables that are best to explain GDP per capita variations after accounting for country and year fixed effects. Each column corresponds to a different fit criteria and refers to the outcomes of implementing the EN after setting λ and ϕ in equation (5) to optimize that specific fit criteria. For each column, the second row shows the list of the climate variables selected by the EN. The optimal values of λ and ϕ are presented in the last two rows, consecutively.

Table B.8: Optimal LASSO selection of variables under different model specifications and BIC

	Without Year Effects	With Quadratic Trends	Balanced Sample with FE
BIC	World GDP Growth, Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W]	Lag-1 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W]	Lag-1 GDP p.c. growth, Lag-2 GDP p.c. growth, Harsh Drought Prevalence ^[W] , Max T °C above 35 ^[W] , PPT Minimum ^[W]
Number of Selected Variables	5	3	5
Optimal Penalty Weight (λ)	.039	.0285	.0372

Notes: This table shows some results of the implementation of the LASSO to select the climate variables that are best to explain GDP per capita variations after accounting for different fixed effects different specification. Each column corresponds to a different fixed-effect specification and refers to the outcomes of implementing the LASSO after setting λ to optimize the BIC. For each column, the second row shows the list of the climate variables selected by the LASSO. The optimal value of λ is presented in the last row.

Table B.9: Optimal LASSO selection affecting fiscal variables under BIC fit criteria

	Revenue	Expenditure	Debt
BIC	Lag-1 Revenue, Lag-2 Revenue, Harsh Drought Prevalence ^[W] , Mean T °C in [24,27] ^[W] , Cont'd Dry Days, Heavy PPT Maximum, PPT Minimum, Lag-1 Extremely Wet Day PPT ^[W] , Lag-1 Longest Day Cold Wave ^[W] , Lag-1 Harsh Drought Prevalence ^[W] , Lag-1 Extremely Wet Day PPT, Lag-1 Cold Wave Days, Lag 1 Longest Dry Spell (.80), Lag-1 Mean T °C in [24,27)	Lag-1 GDP p.c. growth, Lag-1 Expenditure, Lag-2 Expenditure, Mean T °C in [-3,0] ^[W] , Mean Wet Day PPT, Lag-1 Harsh Drought Prevalence ^[W]	Lag-1 GDP p.c. growth, Lag-1 Debt, Lag-1 PPT Minimum
Number of Selected Variables	14	6	3
Optimal Penalty Weight (λ)	.0318	.0419	.0339

Notes: This table shows some results of the implementation of the LASSO to select the climate variables that are best to explain different fiscal variables after accounting for country and year fixed effects. Each column corresponds to a different fiscal variable and refers to the outcomes of implementing the LASSO after setting λ to optimize the BIC. For each column, the second row shows the list of the climate variables selected by the LASSO. The optimal value of λ is presented in the last row.

Table B.10: Climate variables selected by the LASSO using the AIC and their GDP effect under the baseline specification

Variable	Est.	Variable	Est.	Variable	Est.
1 Lag-1 GDP per Capita Growth	0.220*** (0.0325)	13 Max T °C above 35 (W)	-0.133** (0.0673)	25 Longest Night Cold Wave (W)	-0.0870 (0.0530)
2 Lag-2 GDP per Capita Growth	0.0898*** (0.0197)	14 1-Day PPT Maximum	-0.127** (0.0560)	26 Lag-2 Longest Dry Spell (.5)	-0.0814 (0.0499)
3 PPT Maximum (W)	0.229*** (0.0828)	15 Lag-2 Longest Dry Spell (.95)	0.120 (0.0741)	27 Heavy PPT Maximum (W)	0.0783 (0.0524)
4 Lag-1 Mean T °C in [0; 3] (W)	-0.219*** (0.0601)	16 Lag-2 Mean T °C in [21; 24]	0.115** (0.0508)	28 Lag-2 Cold Wave Days	-0.0773 (0.0531)
5 Lag-2 Longest Dry Spell (.80) (W)	-0.219** (0.0869)	17 Longest Night Heat Wave (W)	-0.108*** (0.0350)	29 Lag-2 Very Wet Conditions (W)	-0.0765* (0.0436)
6 Lag-2 Longest Dry Spell (.80)	0.187** (0.0851)	18 Lag-1 Longest Dry Spell (.65)	0.108* (0.0639)	30 Lag-2 Mean T °C in [3; 6]	-0.0713 (0.0810)
7 Cont'd Heavy PPT (W)	-0.165** (0.0701)	19 Lag-1 Day T °C Maximum	0.103** (0.0439)	31 Drought Intensity	-0.0646 (0.0601)
8 Mean T °C in [9; 12]	0.165*** (0.0418)	20 PPT Minimum (W)	-0.101 (0.0769)	32 Lag-1 Harsh Drought Prevalence (W)	0.0549 (0.0530)
9 Harsh Drought Prevalence (W)	-0.164*** (0.0613)	21 Longest Dry Spell (.80) (W)	0.0977 (0.0673)	33 Lag-2 Balanced PPT Indicator	0.0546 (0.0541)
10 Lag-2 Mean T °C in [0; 3] (W)	-0.163*** (0.0608)	22 Lag-2 Day Heat Wave T °C (W)	0.0917 (0.0583)	34 Cont'd Extreme PPT	-0.0472 (0.0630)
11 Lag-1 Cold Spell Duration	0.160** (0.0658)	23 Mean T °C in [24; 27] (W)	0.0900* (0.0528)	35 Lag-1 PPT Minimum (W)	0.0455 (0.0687)
12 Lag-1 10 to 20 mm PPT	0.138** (0.0607)	24 Lag-1 Cont'd Wet Days (W)	-0.0872* (0.0510)	36 Lag-2 Mean T °C in [3; 6] (W)	-0.0455 (0.0734)
Observations	6,653	R-squared	0.281	Within R-squared	0.114

Notes: The table lists the variables selected by the LASSO after a random search for lambda to maximize the AIC. The Est. column indicates the coefficient estimates from a linear regression with country and year fixed effects. All climate variables are first-differenced and standardized. (W) indicates population-weighted variables. Standard errors are clustered by country and reported in brackets.


Table B.11: Climate variables selected by the LASSO using the out-of-sample within R-squared and their GDP effect under the baseline specification

Variable	Est.	Variable	Est.	Variable	Est.
1 Lag-1 GDP per Capita Growth	0.217*** (0.0333)	7 Lag-1 Mean T °C in [0; 3] (W)	-0.131*** (0.0381)	13 Lag-1 PPT Minimum (W)	0.0885 (0.0683)
2 Lag-2 GDP per Capita Growth	0.0897*** (0.0199)	8 Lag-1 Day T °C Maximum	0.106** (0.0412)	14 Lag-2 Cold Wave Days	-0.0795 (0.0542)
3 Harsh Drought Prevalence (W)	-0.169*** (0.0600)	9 Longest Night Heat Wave (W)	-0.105*** (0.0342)	15 Lag-2 Balanced PPT Indicator	0.0788 (0.0505)
4 Max T °C above 35 (W)	-0.162** (0.0709)	10 1-Day PPT Maximum	-0.104** (0.0477)	16 Drought Intensity	-0.0665 (0.0568)
5 Lag-1 Cold Spell Duration	0.158** (0.0676)	11 Lag-2 Mean T °C in [3; 6] (W)	-0.104 (0.0777)	17 Lag-1 10 to 20 mm PPT	0.0616 (0.0580)
6 Mean T °C in [9; 12]	0.151*** (0.0412)	12 Longest Night Cold Wave (W)	-0.0962* (0.0533)	18 Lag-2 Mean T °C in [3; 6]	-0.0450 (0.0818)
Observations	6,653	R-squared	0.281	Within R-squared	0.114

Notes: The table lists the variables selected by the LASSO after a random search for lambda to maximize the out-of-sample within R-squared. In this case, we obtain the same variable selection whether we randomly assign countries (and all their observations) to the test set or if we randomly assign observations to the test without accounting for the country panel structure. The Est. column indicates the coefficient estimates from a linear regression with country and year fixed effects. All climate variables are first-differenced and standardized. (W) indicates population-weighted variables. Standard errors are clustered by country and reported in brackets.


B.5. Additional Figures

Figure B.2: Selection of climate variables impacting GDP (specification without year effects)
 (a) using the AIC (b) using the BIC (c) using the OOS within R2



Note: GDP per capita growth is the dependent variable and the specification has country effects and world growth. See the notes of the following graph for more details.

Figure B.3: Selection of climate variables impacting GDP (specification with quadratic trends)
 (a) using the AIC (b) using the BIC (c) using the OOS within R2




Note: The figures show the results of implementing the LASSO for different penalty parameters λ . The red lines are similar in every panel and show how the number of selected variables vary with λ . The grey dashed lines in each panel show the variation of different criteria with λ . The out-of-sample (OOS) within R-squared is calculated on a sub-sample of countries (evaluation set) based on coefficients estimated on the rest of countries (training set) as explained in the main text. The dots indicate the different selection outcomes given by the local optimum for each criteria respectively. The estimated model has GDP per capita growth as the dependent variable and includes country quadratic trends and year effects.

Figure B.4: Selection of climate variables impacting GDP (balanced sample with FE)

(a) using the AIC

(b) using the BIC

(c) using the OOS within R2




Note: GDP per capita growth is the dependent variable and the specification with country and year effects was estimated on the balanced sample for 1984-2019. See the notes of the following graph for more details.

Figure B.5: Selection of climate variables impacting government revenue


(a) Selection using the AIC

(b) Selection using the BIC




Note: The figures show the results of implementing the LASSO for different penalty parameters λ . The red lines are similar in every panel and show how the number of selected variables vary with λ . The grey dashed lines in each panel show the variation of different criteria with λ . The within R-squared is calculated on a sub-sample of countries (evaluation set) based on coefficients estimated on the rest of countries (training set) as explained in the main text. The dots indicate the different selection outcomes given by the local optimum for each criteria respectively. The estimated model has the ratio of government revenue to GDP as the dependent variable and includes country and year effects.

Figure B.6: Selection of climate variables impacting government expenditure
 (a) Selection using the AIC
 (b) Selection using the BIC



Note: The estimated model has the ratio of government expenditure to GDP as the dependent variable and includes country and year effects. See the notes of the following graph for more details.

Figure B.7: Selection of climate variables impacting government debt
 (a) Selection using the AIC
 (b) Selection using the BIC



Note: The figures show the results of implementing the LASSO for different penalty parameters λ . The red lines are similar in every panel and show how the number of selected variables vary with λ . The grey dashed lines in each panel show the variation of different criteria with λ . The within R -squared is calculated on a sub-sample of countries (evaluation set) based on coefficients estimated on the rest of countries (training set) as explained in the main text. The dots indicate the different selection outcomes given by the local optimum for each criteria respectively. Government debt to GDP is the dependent variable and the specification has country and year effects.

B.6. LASSO and Elastic-Net Implementation

In this appendix, we detail our implementation of the various algorithms covered in the paper, with an emphasis on the technical steps and the specific software and functions that we use.

We define X as the matrix containing (first differenced) right hand side variables, excluding fixed effects, but including relevant lags of the dependent variable Δy . For details on the lags used in each regression model, please refer to the main article. We designate F as the matrix of fixed effects, which varies according to the model specification. Specifically, F can encompass:

- a) Country and year fixed effects, or
- b) Only country fixed effects, or
- c) Country fixed effects, year fixed effects, and country quadratic dummies.

The organization of these matrices is such that rows represent individual observations and columns correspond to variables. To ensure compatibility with our Python-based feature selection algorithm (Python version 3.9 or higher), we remove any missing observations.

Using this notation, the regression model incorporating all variables can be summarized as shown in equation (B.1), which is the same as equation (2) in the main text but using different notations to single-out and combine fixed effects under one matrix F . Note that we have omitted subscripts from the fixed effects matrix F to indicate its flexibility; depending on the specification, F may contain only country-related information (i) and/or year-related information (t).

$$\Delta y_{it} = X_{it}\beta + F\theta + \varepsilon_{it} \tag{B.1}$$

Before constructing the fixed effects matrix F , we first eliminate outlier observations for the dependent variable. Specifically, any observation Δy_{it} that deviates by more than 5 standard deviations from the mean $\overline{\Delta y_{it}}$ is removed as described in section 3.3 in the main text. Because the LASSO penalizes the value of the coefficients, the scales of the parameters can affect the selection. Therefore, we standardize each column of X_{it} to have 0 mean and a standard deviation of 1. After this preprocessing step, we proceed to generate the F matrix.

The variable selection algorithm focuses on the variables within X_{it} . However, theoretical considerations mandate the inclusion of fixed effects in the regression model. To reconcile these aspects, we force the presence of fixed effects in the regression. We do so by first subtracting $F(F'F)^{-1}F'\Delta y$ from both sides of the equation.

$$\begin{aligned}
\Delta y_{it} - F(F'F)^{-1}F'\Delta y &= X_{it}\beta + F\theta - \underbrace{F(F'F)^{-1}F'(X_{it}\beta + F\theta + \varepsilon_{it})}_{\Delta y} + \varepsilon_{it} \\
\underbrace{(I - F(F'F)^{-1}F')\Delta y_{it}}_{\Delta \tilde{y}_{it}} &= \underbrace{(I - F(F'F)^{-1}F')X_{it}}_{\tilde{X}_{it}}\beta + \underbrace{F\theta - F(F'F)^{-1}F'F\theta}_0 + \underbrace{(I - F(F'F)^{-1}F')\varepsilon_{it}}_{u_{it}}
\end{aligned}$$

$$\implies \Delta \tilde{y}_{it} = \tilde{X}_{it}\beta + u_{it} \tag{B.2}$$

The Frisch-Waugh-Lowell theorem implies that the estimations based on equations (B.1) and (B.2) result in the same estimate for β . Consequently, performing the selection algorithm after the above transformation effectively incorporates the fixed effects into the regression model.

As elaborated in the main text, our objective is to select a subset of columns from the matrix \tilde{X}_{it} . To achieve this, we employ LASSO and Elastic-Net methods, which are detailed in the subsequent sections. The analyses are conducted using version 1.2.2 of the Scikit-Learn package in Python. To ensure replicability due to the random sampling described later, we set the random seed using the numpy package, version 1.25.0. All computations are performed on a Windows 11 machine with a 13th Gen Intel(R) Core(TM) i7-13700 processor, operating at 2.10 GHz.

B.6.1 LASSO

As explained in Section 2.3 in the main text, the LASSO aims to solve equation (4), that is to minimize the following equation:

$$\min_{\beta} \Delta \tilde{y}_{it} - \tilde{X}_{it}\beta + \lambda \sum_{j=1}^K |\beta_j| \tag{B.3}$$

where the hyperparameter λ weighs the penalty term, which is the sum of the absolute values of the coefficients β_j . K denotes the number of columns in the matrix \tilde{X}_{it} .

The penalty term encourages some coefficients to shrink towards zero. As λ increases, the penalty term gains more weight, leading to more coefficients becoming zero. Conversely, a smaller λ results in fewer coefficients shrinking to zero. Coefficients that remain non-zero are those for which the reduction in standard error outweighs the penalty incurred by their inclusion in the regression.

To determine the optimal value of the hyperparameter λ , we explore four approaches:

1. Minimizing the Bayesian Information Criterion (BIC),
2. Minimizing the Akaike Information Criterion (AIC),

3. Maximizing the average out-of-sample R^2 using 5-fold cross-validation (and since we remove fixed-effects, this R^2 -metric corresponds to what is commonly defined as the *within* R^2). In this method, observations are randomly divided into five bins without considering the panel structure of the data,
4. Maximizing the average out-of-sample R^2 with a modified 5-fold cross-validation approach that respects the panel structure (and again, since we remove fixed-effects, this R^2 -metric corresponds to what is commonly defined as the *within* R^2). Specifically, countries are divided into five bins, and observations for these countries are used in each fold separately.

For the first two approaches, we employ the built-in `LassoLarsIC()` function available in Scikit-Learn. This function utilizes the Least Angle Regression (LARS) algorithm for LASSO variable selection, as opposed to Scikit-Learn’s main LASSO implementation, which relies on a gradient-descent algorithm. Both methods aim to solve the same optimization problem but take different computational routes.

For the third approach, we employ 5-fold cross-validation. In k -fold cross-validation, the dataset is randomly divided into k subsets of equal (or nearly equal) size called folds. One fold is reserved as the test set, and the model is trained on the remaining $k - 1$ folds. This process is repeated k times, each time with a different fold serving as the test set. The performance metric, in our case the out-of-sample R^2 , is then averaged across all k iterations.

We employ Scikit-Learn’s `RandomizedSearchCV()` function to conduct the 5-fold cross-validation. We perform the cross-validation for 200 distinct penalty weights, leading to a total of 1,000 model fits. These penalty weights are drawn from a half-normal distribution with a location parameter (`loc`) of 0.001 and a scale parameter of 0.05.

For the fourth approach, we use numpy’s `random.choice()` function to divide the countries into 5 folds: four folds contain 41 countries each, while the fifth contains 39 countries. We then proceed in a manner similar to the k -fold cross-validation described above. Specifically, for each of 200 distinct penalty weights, we fit the model using observations from four folds, reserving one fold as the test set to calculate the out-of-sample R^2 . Each penalty weight is evaluated five times—once for each fold serving as the test set—and the average R^2 is computed. The penalty weight yielding the highest average R^2 is then selected.

B.6.2 Elastic-Net

Using the notations of this section, the Elastic-Net optimization problem covered in equation (5) in the main text can be expressed as follows (note that ϕ would correspond to α in the Scikit-Learn package’s notations):

$$\min_{\beta} \frac{1}{2N} (\Delta \tilde{y}_{it} - \tilde{X}_{it} \beta)^2 + \lambda \phi \sum_{j=1}^K |\beta_j| + \lambda \frac{1 - \phi}{2} \sum_{j=1}^K \beta_j^2 \quad (\text{B.4})$$

In equation (B.4), we have two hyperparameters, λ and ϕ . The range of ϕ is between 0 and 1, and it determines the balance between the penalty terms associated with the LASSO and the Ridge. Increasing ϕ promotes sparsity in the solution. Likewise, increasing λ enhances sparsity, given that $\phi \neq 0$. However, the selection of variables may differ depending on the approach taken. To determine the optimal ϕ and λ combination, we employ the same approach as we did for the LASSO and separately consider 4 different fit criteria.

Unlike in the LASSO case, there is no built-in function available to minimize the BIC and AIC in the case of the Elastic-Net. As a result, we modify the source codes of the `LassoLarsIC()` to make it compatible with the Elastic-Net.³⁵ Since the faster LARS algorithm is not available for the Elastic-Net, we resort to the gradient descent algorithm. Consequently, we do not explore every possible combination of ϕ and λ . We consider 9 distinct values for ϕ (ranging from 0.1 to 0.9 with increments of 0.1) and 16 values for λ (0.0025, 0.005, 0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09, 0.1, 0.15, 0.2, 0.25, 0.3). This results in a total of $9 \times 16 = 144$ unique combinations. For each of these 144 combinations, we calculate the *AIC* and *BIC* values using the formulas:

$$AIC = N \times \log(2\pi\hat{\sigma}_u^2) + \frac{RSS}{\hat{\sigma}_u^2} + 2 \times \text{DoF} \quad (\text{B.5})$$

$$BIC = N \times \log(2\pi\hat{\sigma}_u^2) + \frac{RSS}{\hat{\sigma}_u^2} + \log(N) \times \text{DoF} \quad (\text{B.6})$$

We use the residual sum of squares (*RSS*) obtained after making predictions with the Elastic-Net, and the degrees of freedom (DoF) are equal to the number of non-zero coefficients after the Elastic-Net. N is the number of observations, and $\hat{\sigma}_u^2$ is the estimated variance of the error term in equation (B.2). The error term is estimated before the selection using all variables in the X_{it} matrix as in the source codes of `LassoLarsIC()` function.

For maximizing the out-of-sample within R^2 using k-fold cross-validation in the case when the observations are randomly allocated without considering the panel structure, we again utilize Scikit-Learn's `RandomizedSearchCV()` function. We sample the ϕ parameter from a uniform distribution ranging from 0.1 to 0.9, and the λ parameter from a half-normal distribution with a location parameter (`loc`) of 0.001 and a scale parameter of 0.5. We consider 200 distinct combinations, resulting in a total of 1,000 model fits across 5 folds.

Lastly, maximizing the out-of-sample within R^2 using k-fold cross-validation in the case when observations are randomly allocated factoring in the country panel structure, we employ numpy's `random.choice()` function to partition the countries into 5 folds. Four of these folds contain 41 countries each, and the fifth contains 39 countries. To prevent corner solutions, we use a smaller scale parameter for the half-normal distribution this time. Specifically, ϕ is sampled from a uniform distribution ranging between 0.1 and

³⁵The source codes can be found in https://github.com/scikit-learn/scikit-learn/blob/main/sklearn/linear_model/_least_angle.py#L2280, after lines 2089 as of November 2023.

0.9, while λ is drawn from a half-normal distribution with a location parameter (loc) of 0.001 and a scale parameter of 0.1 (as opposed to 0.5 used in previous exercises).

For the two implementation based on the out-of-sample within R^2 , we implement the EN five times for each combination of penalty weights, once with each fold serving as the test set. For each combination, we compute the average R^2 . The combination of penalty weights that maximizes this average R^2 is then determined to be the optimal combination.