THE GRADUATE INSTITUTE GENEVA

INSTITUT DE HAUTES ÉTUDES INTERNATIONALES ET DU DÉVELOPPEMENT

GRADUATE INSTITUTE OF INTERNATIONAL AND DEVELOPMENT STUDIES

From Climate Risk to Growth-at-Risk: An Empirical Model of Downside Risks to Growth from Temperature Shocks

Thesis submitted at the Graduate Institute of International and Development Studies in fulfillment of the requirements of the International Economics department

Joshua E. Ostry

July 2022

Geneva, Switzerland

Supervisor: Cédric Tille

Second Reader: Beatrice Weder di Mauro

Abstract

This paper investigates the degree to which the mean, variance and skewness of a country's temperature distribution worsen downside risks to GDP growth. Using a quantile regression framework for a sample of 124 countries over 37 years, I first document that while rising mean temperatures reduce mean GDP growth, mean temperature shocks have more pronounced effects on worst-case (left-tail) growth outcomes. For poor countries, rising mean temperatures lower left-tail GDP growth by reducing output in both the agricultural and industrial sectors. On the other hand, rising mean temperatures affect worst-case growth in rich countries only through the agricultural sector. Next, I highlight that higher moments of the temperature distribution are also salient predictors of left-tail growth. For rich countries, greater temperature volatility and more extreme cold temperatures — left-skewed temperature distributions — forecast weaker agricultural growth. On the other hand, more extreme hot temperatures — right-skewed distributions — predict significant deteriorations in agricultural growth in poor countries. Finally, using country case studies, I quantify the degree to which shocks to the moments of the temperature distribution affect the conditional distribution of (sectoral) growth, with an emphasis on the distribution's left tail.

1 Introduction

It is well-documented that we live in an age where human activity is the dominant force on the environment, a phenomenon known as anthropogenic climate change (e.g. Manabe and Wetherald 1967; Keeling et al. 1976; Held and Soden 2006). The central mechanism behind climate change is the emission of greenhouse gases (GHG) into the atmosphere from the burning of fossil fuels. A higher concentration of greenhouse gases in the atmosphere traps greater amounts of solar radiation in the air, which, in turn, increases the earth's surface temperature – global warming.

Global warming caused by climate change is only part of the problem. Climate change increases the variability of surface temperatures and the frequency of extreme weather events, floods, wildfires, droughts and severe storms. Moreover, some disastrous environmental outcomes, such as the extinction of innumerable wildlife species and the massive release of greenhouse gases from the melting of permafrost, are expected to occur along the planet's current emissions path. The Intergovernmental Panel on Climate Change forecasts the effects of emissions on global temperatures, and warns that average surface temperatures must not rise by more than 2° Celsius (C) relative to a long-term average if catastrophic outcomes are to be avoided. Soberingly, it assesses that the world's current "business-as-usual" trajectory will lead to temperature increases of about 4° C (Lee and Marotzke 2022).

In my thesis, I empirically assess the effects of climate change on *worst-case* economic growth outcomes. The existing macroeconomic literature on the effects of climate change considers how rising mean temperatures affect *mean* growth rates, and finds an adverse effect on mean growth for poor countries (Dell, Jones, and Olken 2012). The theoretical basis for this finding could be that higher temperatures adversely affect productivity in particular sectors of the economy, as well as productivity in the economy as a whole, or that countries' ability to mitigate the effect of climate change depends on their level of development, which acts as a constraint on the financial resources available for adaptation. A frequently-cited channel is the impact on the agricultural sector (which is likely to be more salient in less diversified, poorer economies), where rising temperatures tend to result in more droughts and shorter growing seasons, especially in poor countries that tend to be hotter. However, other channels related to the impact of higher temperatures on labor productivity in the industrial sector (Sudarshan and Tewari 2014) and the incidence of disease as temperatures climb (Sachs and Warner 1997), have also been used to explain additional mediating channels for reduced growth.

It is only recently that economic research has begun to investigate the relationship between higher moments of the temperature distribution and economic growth (Kotz et al. 2021; Angelova and Käbel 2019). Greater temperature variance may make crop yields more unpredictable, and more skewed temperature distributions – more extreme hot or cold temperatures – typically signal increases in the frequency and intensity of weather events and natural disasters, which will certainly affect expected aggregate and sectoral growth rates.

In addition, a recent contribution by Kiley 2021 considers how shocks to mean temperature affects left-tail growth. He finds that rising temperatures have larger effects on downside risks to growth than on mean growth. However, neither this paper, nor the existing literature that focuses on the effects of climate change on mean growth, address the central issue of my thesis: how do shocks to the entire temperature distribution, not just the mean, affect downside risks to growth? There is a strong prima facie case that the higher moments of the temperature distribution do not just impact mean growth, but also impact left-tail growth disasters. If policy makers are concerned about growth disasters and how to avert them, or at least make them less severe, this research question would seem to be very much of the essence.

Specifically, in this paper, I assess the empirical relationship between various moments of the temperature distribution, in particular its mean, variance and skewness, and the distribution of aggregate and sectoral growth, particularly the left-tail. While other empirical papers on the macroeconomic effects of climate change have considered certain links between moments of the temperature distribution and specific features of the GDP growth distribution, my approach is novel as it incorporates these links into a unified empirical framework. Moreover, I contribute to the literature by applying growth-at-risk techniques to the sectoral growth distribution as well, rather than to (aggregate) GDP alone.

To motivate my empirical framework, I first build a simple growth model that allows higher moments of the temperature distribution, not just the mean as assumed in Bond, Leblebicioglu, and Schiantarelli 2010, to impact a country's production function, and hence the country's growth rate. I then empirically test this growth-temperature relationship not only for mean growth, but for the entire growth distribution. This approach relies on quantile regression techniques that have been used extensively in the macro-finance literature to assess how shocks to a country's financial conditions affect its growth-at-risk, that is, its left-tail growth outcome (Adrian, Boyarchenko, and Giannone 2019).

Quantile models of growth-at-risk have become a mainstay of economic forecasts because they allow for shocks to have amplified effects on, for example, left-tail growth as compared to mean growth based on structural features of the economy (such as leverage and features of private sector balance sheets). The leap to the climate literature seems a priori plausible in that there may be amplification mechanisms that lead climate shocks to differentially affect different parts of the growth distribution, but it is an empirical matter to determine the salience of such mechanisms. It seems plausible, moreover, given the limited capacity of poorer countries to improve the resilience of their economies to climate shocks, that amplification mechanisms that worsen growth-at-risk would be larger in poorer countries.

My thesis presents three main findings. First, I find that in poor countries, an increase in mean temperature has a strong negative effect on the bottom 10th percentile of the GDP growth distribution (a standard measure of growth-at-risk), and that the effect of rising mean temperatures on growth-at-risk is mediated through both the agricultural and industrial sectors. Moreover, the effect of mean temperature is stronger for left-tail growth than for mean growth, which underscores the value of an empirical strategy that allows the effect of temperature on growth to vary along the growth distribution. I interpret the result that left-tail growth is more severely impacted than mean growth as indicating that mean temperature can help to predict growth disasters: this is an important result from a policy perspective.

Second, I find that rich countries are quite resilient to rising mean temperatures in terms of growth-at-risk impacts, except through the impact on the agricultural sector, which is not the dominant sector in the typical advanced-economy case (in contrast to agriculture's role in a number of poorer, developing countries).

Third, I find that, while the higher moments of the temperature distribution (standard deviation and skewness) may not significantly (in a statistical sense) affect left-tail GDP growth, these higher moments do have an appreciable effect on growth-at-risk for the agricultural sector. For richer countries, agricultural left-tail growth outcomes are more severe with higher temperature volatility and more extreme cold temperatures (left skewness). For poorer countries, on the other hand, more extreme hot temperatures (greater right skewness in temperature) have large adverse effects on left-tail agricultural growth.

The remainder of my paper is structured as follows. Section 2 discusses the existing literature related to temperature, macroeconomic growth and sectoral growth, and the literature on predicting downside risks to growth. Section 3 outlines the data sources used in my analysis and provides a brief overview of the methods employed to aggregate temperature data from the weather station-level to the country-level. Additionally, it presents summary statistics regarding temperatures and income levels for countries in the dataset. In section 4, I present the methodology used in my empirical analysis, where I build a growth regression framework with a set of standard determinants, but then adapt this framework to assess growth-at-risk using quantile regression. Section 5 presents the main results from my panel quantile growth regressions and quantifies the impact of temperature shocks on left-tail growth outcomes. In section 6, I discuss the possible channels through which temperature shocks can affect growth and present regression results for temperature variables on agricultural and industrial growth rates. Section 7 presents a number of country case studies to quantify the impact of shocks to the temperature distribution on conditional future left-tail growth. Finally, section 8 concludes and provides a direction for future research.

2 Literature Review

My thesis sits at the intersections of three distinct literatures in economics. The first is the literature on the drivers of economic development; the second is the empirical literature examining how climate change affects (sectoral and aggregate) economic growth; and the third is the literature employing novel empirical techinques to quantify downside risks to growth – the so-called growth-at-risk literature.

2.1 Literature on Drivers of Economic Development

In order to motivate the research surrounding the effects of temperature on growth, an important question to ask is whether a set of reasonably robust correlations between temperature and growth implies a causal relationship. While full consensus on the causality issue has not been reached among professional economists, empirical papers that identify these correlations without establishing causality have been quite influential. While it is outside the scope of this paper to establish causality between growth and temperature, other papers have offered some insight on this issue. One argument in favor of causality is that temperature has important implications for total factor productivity, agricultural production, tourism and other sectors that utilize natural and environmental resources. Additionally, warmer countries tend to be more prone to natural disasters and disease. Certainly, through these channels, changes in temperature affect the level of macroeconomic growth and development. On the other hand, it may be the case that the observed relationship between temperature and growth reflects underlying factors that have not been properly modeled. This question is typically discussed in the literature on temperature and economic development.

Sachs and Warner 1997 estimate a panel growth regression for African countries from 1965-1990 with several independent variables, including trade openness, institutional quality, life expectancy, and whether or not the country has a tropical climate, among others. They find that countries that have a tropical climate have significantly lower growth, which they attribute to two main factors. First, tropical countries tend to face various parasitic diseases that are not as common in temperate climates; and second, weather is less favorable for agricultural production in tropical climates. Sachs and Warner 1997 consider their results to be robust and do not shy away from interpreting them causally.

However, Acemoglu, Johnson, and Robinson 2001 argue that the observed relationship between temperature and development actually reflects a more fundamental cause: namely, that in the colonial period, better institutions were established in countries with colder climates (where the colonists would choose to settle). Thus, the causal relationship is between institutional quality and economic development, rather than between climate and economic development. When Europeans colonized African, North and South American countries, they created settlements in countries with lower mortality rates, which were often in cooler countries due to the lower incidence of disease. The institutions set up in these cooler countries tended to be higher quality as they were meant to support colonial settlements, while institutions set up in hotter countries were lower quality as colonialists would not settle there. Acemoglu, Johnson and Robinson argue that differences in the quality of institutions are responsible for differences in the levels of development, and as a result, the observed relationship between temperature and development conceals the underlying cause.

While it may be the case that, over a long time-horizon, development occurred in colder climates because higher quality institutions were established, this does not mean that empirical relationships between temperature and economic growth are not robust and of clear relevance from a policy standpoint. The legacy of settler mortality and institutional quality may well be in the background, but the question of how evolving climate shocks affect future economic growth still needs to be asked, and data brought to bear to produce an answer. The fundamental long-term influences identified by Acemoglu, Johnson, and Robinson would, in any event, be constant over the forecasting horizon that is relevant for policy analysis in most cases.

2.2 Empirical Literature on Climate Change and Growth

An important recent paper on the relationship between temperature shocks and economic growth was written by Dell, Jones, and Olken 2012. In this influential study, the authors perform a fixed effects panel regression analysis where they estimate the effect of mean temperature on mean growth rates for 125 countries between 1970 and 2006. They find that for rich countries, mean temperature does not affect mean growth rates. By contrast, they find a strong negative effect of mean temperature on mean growth in poor countries. The authors discuss the potential channels through which temperature may affect growth. They find that increasing mean temperatures result in decreasing agricultural and industrial growth rates in poor countries, but not in rich ones. They did not find that temperature had significant effects on investment or on political variables (other possible mediating channels for the observed effects on economic growth).

Kotz et al. 2021 estimate the effect of temperature variability, defined in terms of daily deviations from a monthly mean, on average GDP growth. Using a fixed effects regression model for 1537 regions over 40 years, they find that a one degree increase in temperature variability results in a 5 percentage point reduction in GDP growth. Additionally, they find that the effect of an extra unit of variability on growth rates is more negative in poorer countries.

Burke, Hsiang, and Miguel 2015, consider model specifications in which squared temperature and precipitation may affect mean growth rates. They believe such a methodology is needed in order to reconcile inconsistencies between the micro- and macroeconomic evidence of temperature's effects on growth. Microeconomic studies have found non-linear effects of temperature on production, while macroeconomic studies do not detect such nonlinearities. Their specification establishes that productivity peaks at approximately 13 $^{\circ}$ C and declines strongly at higher temperatures. They also find support for such a result in both rich and poor countries.

In a paper written by Michael Kiley 2021, the author builds upon the work of Dell, Jones, and Olken 2012 and Burke, Hsiang, and Miguel 2015 by employing a quantile regression framework to evaluate the effect of temperature on downside risks to growth and incorporating non-linear temperature effects into the model. Kiley finds that there is a stronger negative effect of rising temperatures on the left-tail of growth than on mean growth, and this result serves as a jumping-point for the specification in my thesis, which adds a number of new dimensions, including higher moments of the temperature distribution, an investigation of mediating channels, and quantified country case studies.

One issue to consider when evaluating the relationship between temperature and growth is what causes these two variables to be related—mediating channels. Determining the effects of temperature on certain sectors of the economy can shed light on the extent to which particular countries may be vulnerable to climate change. Angelova and Käbel 2019 estimate the effect of the mean and volatility of temperature and precipitation on European agricultural production. They find that a 1° C increase in mean temperature would reduce agricultural growth by 1.6 percentage points, while there is a negligible effect of mean precipitation on agriculture. Additionally, they find that increased rainfall and temperature variability increases agricultural output slightly.

Beyond the agricultural sector, there are several other mechanisms through which temperature may affect growth. Colacito, Hoffmann, and Phan 2018 estimate the effect of rising mean temperatures on country-level and state-level growth in the United States. They find that rising mean temperatures slightly reduce growth rates in the US. In their analysis, they consider how this relationship changes based on the season, and find that the strongest growth effects occur during the summer and autumn months, and that these effects have opposite signs, i.e., warmer temperatures are less favorable in the summer and more favorable in the autumn. They find, in addition, that in summer months, productivity falls as temperatures rise, while the opposite is true in the autumn months. However, employment rates do not change much with temperature. Finally, in terms of sectors, they find that rising temperatures have modest negative effects on the services, finance, insurance and real estate sectors. These sectoral effects are dwarfed by the effect of rising temperatures on agricultural production. However, because agriculture, forestry and fishing only comprise 1.1% of GDP in the United States, the transmission to GDP is weak.

Sudarshan and Tewari 2014 estimate the effect of rising temperatures on output from manufacturing plants in India using a multi-year panel regression. They find that a 1° C increase in mean temperature reduces manufacturing output by between 1-3 percentage points, and that this result is driven by declines in labor productivity.

2.3 Literature on Growth-at-Risk

While the previous literature has discussed the sources of mean growth, except for the paper by Kiley 2021, predicting downside risks to GDP growth is an important consideration for policymakers and climate economists. Adrian, Boyarchenko, and Giannone 2019 model the conditional distribution of future GDP growth in the U.S. as a function of financial conditions. They do so by running a series of quantile regressions that estimate different quantiles of h-year ahead growth rates, and then interpolate between these conditional quantiles by fitting them to a skew-t distribution. The authors find that, while the upper tail of the growth distribution is relatively stable in response to tighter financial conditions, the lower tail moves in tandem with tighter financial conditions. Specifically, tighter financial conditions make growth disasters more severe, worsening growth-at-risk.

Adrian et al.'s analysis relies on the use of quantile regression, which was developed by Koenker and Bassett 1978. The authors develop the quantile regression methodology by adapting the loss function used in ordinary least-squares regressions. This new loss function, which they call the sum of absolute deviations, can take any number of forms depending on the weights placed on observations above or below a best-fit line. Therefore, changing the weights (or quantiles), and then minimizing the loss function, will yield a regression equation that estimates the effect of a set of independent variables on a given quantile of the dependent variable. I use the quantile regression methodology in the empirical section of my thesis.

3 Data Sources and Summary Statistics

3.1 Data Sources

The empirical analysis performed in this paper uses both weather and GDP data. My historical average temperature and precipitation data comes from Dell, Jones, and Olken 2012. Their dataset is based on the Terrestrial Air Temperature and Precipitation: 1900–2006 Gridded Monthly Time Series, Version 1.01 (Matsuura and Willmott 2007). Matsuura and Willmott construct a worldwide 0.5 x 0.5 degree (approximately 56km at the equator) resolution grid of temperature and precipitation based on interpolating data from an average of 20 weather stations per grid node.

Average country-year temperatures are then calculated using geospatial software (such as ArcGIS). This is done by creating multi-layered maps, with layers corresponding to country borders, Matsuura and Willmott's temperature grid, and population density maps based on 1990 population data at a 30 arc-second resolution (approximately 1km at the equator) from the Global Urban-Rural Mapping Project (Balk et al. 2004). From these layers, average country-year temperatures are determined in two ways: first, by calculating a simple average temperature in all grid nodes within a country's borders, where incomplete nodes are weighted by a fraction of their 0.5 x 0.5 degree land-mass; and second, by calculating a weighted average of temperatures in all grid nodes based on the population living in each node belonging to a given country, where incomplete nodes are weighted in the same way as the first method. Dell, Jones, and Olken 2012 find that both the spatial- and population-weighted average temperatures yielded very similar results in their analysis, so I have chosen to employ population-weighted country-year temperatures in my analysis.

In addition to the country-year temperature observations, I also use country-month data from the Climate Change Knowledge Portal (World Bank Group 2021a). Their observations are calculated by weighting temperature nodes by land-mass in the same way as in Dell et al. This data is used to calculate within-year temperature variance and skewness. Country-year and country-month precipitation data from Dell, Jones, and Olken 2012 and the World Bank Climate Change Knowledge Portal are calculated in the same way as their respective temperature datasets.

The growth rates of aggregate, agricultural and industrial output come from the World Development Indicators (World Bank Group 2021b), and the response variables used in my analysis are the growth rates of GDP, agricultural, and industrial output (constant local currency units). While the country-level output dataset currently has observations for 217 countries, I am limited by the years for which I have weather data. Additionally, I have required that each country should have at least 20 years of output observations in order to be included in the sample as fewer observations may harm the significance of my regression results upon running country-level regressions and scenario analyses, as well as when I include country fixed effects. In the end, my dataset includes 124 countries with observations between 1970 and 2006.

3.2 Summary Statistics

Figures 1a and 1b plot overlapping histograms of average annual temperatures in the early and late stages of my sample in "rich" and "poor" countries, respectively. In these figures, as in the remainder of this paper, I define "rich" as a country having above median GDP PPP per capita in the first year of my sample, and "poor" as a country having below median GDP PPP per capita in the first year, which is the same classification used in Dell, Jones, and Olken 2012.

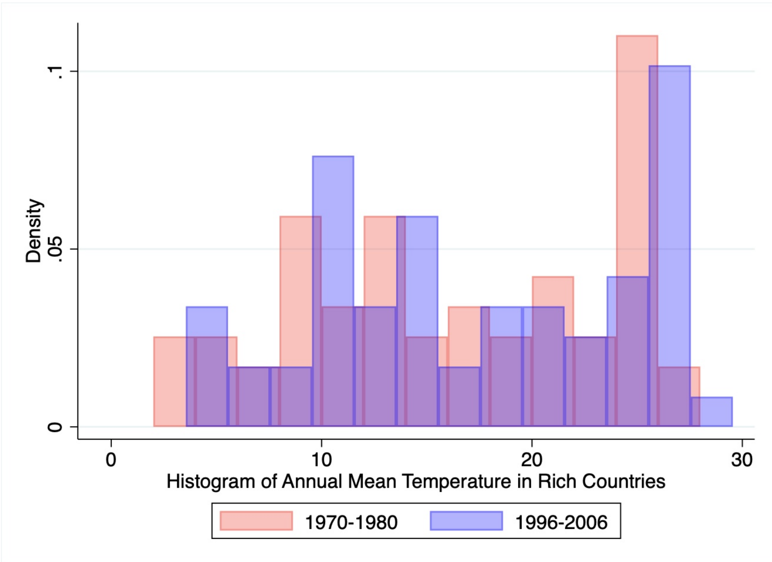

Figure 1a: Rich Countries

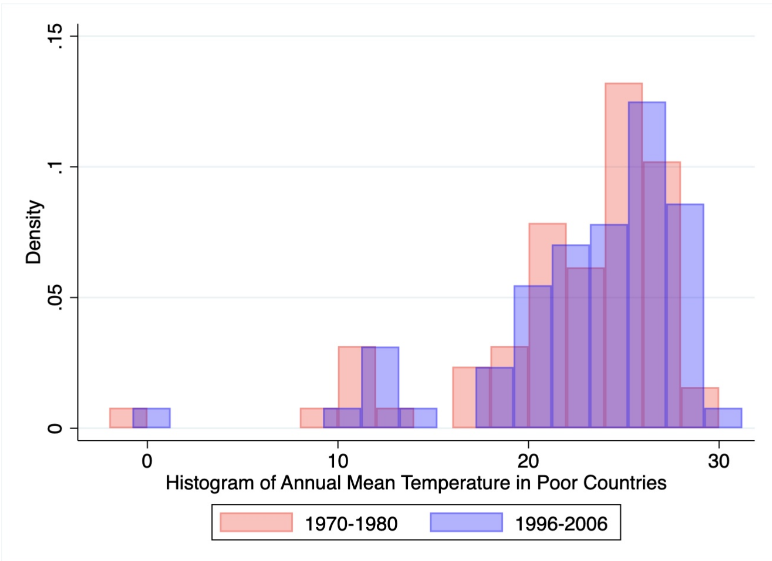

Figure 1b: Poor Countries

These figures show that in both rich and poor countries, average annual temperatures are rising during my sample. The histograms of average annual temperature in the late period (between 1996 and 2006), which are given by the blue areas, appear to have shifted to the right relative to the histogram in the early period (between 1970 and 1980) in both figures. This confirms the view that climate change is causing significant global warming in all countries.

Next, figure 2 plots the average annual temperatures and the natural logs of GDP per capita in 1970 of countries in my sample.

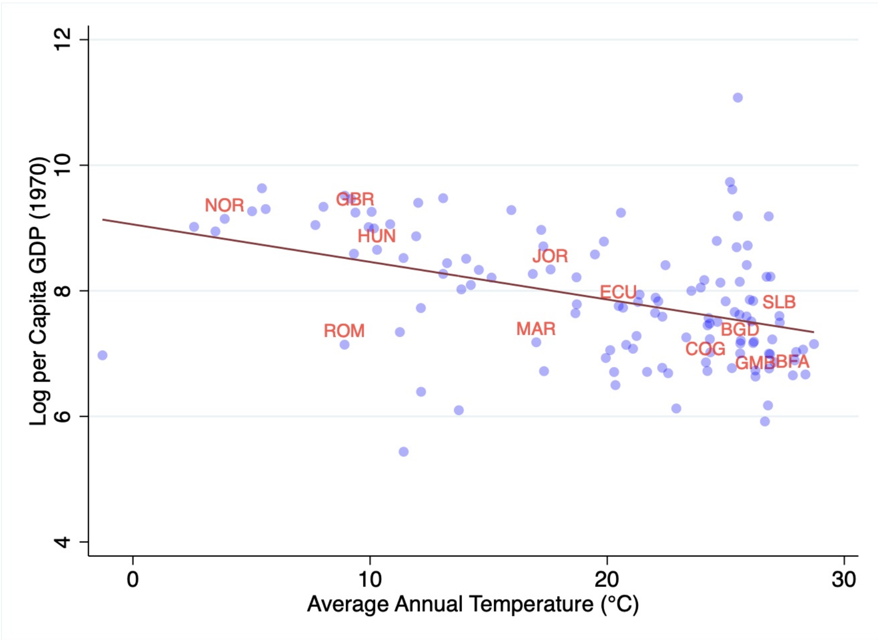

Figure 2: Differences in Average Temperatures and Income per Capita

From this figure, there appears to be a negative correlation between per capita GDP and mean temperatures at the cross-section of my panel. Richer and colder countries, such as Norway (NOR) and Great Britain (GBR), are found in the top-left of this figure, while poorer and hotter countries, like Gambia (GMB) and Bangladesh (BGD), are found in the bottom right of the figure. However, this simple negative correlation does not imply that rising temperatures make countries poorer, nor that countries will become poorer as a result of climate change. In the empirical section of this paper, however, I test this relationship more rigorously with both country and time controls. Additionally, climate change is more than just rising mean temperatures and will affect the wider temperature distribution. I also consider these effects in the subsequent sections.

A table of descriptive statistics concerning the mean, standard deviation and skewness of temperature and GDP growth rates for rich and poor countries can be found in Table 5 in the appendix. It highlights that poor countries have temperature distributions that are hotter and more left-skewed than rich countries, which have temperature distributions that are more variable. Additionally, they illustrate that rich countries tend to have GDP growth distributions with a greater mean and more left-skewness. Poor countries have more volatile growth and their growth rates are more right-skewed.

4 Methodology

4.1 Theoretical Motivation

In this section, I begin by adapting a simple model economy as done in Bond, Leblebicioglu, and Schiantarelli 2010, and later, Dell, Jones, and Olken 2012. While they include mean temperature as a determinant of total factor productivity (TFP) in their production function, it is certainly possible that both within-year temperature variance and skewness matter for TFP.[1]

In the remainder of this paper, I define $\bar{T}_{i,t} = \frac{1}{12}\sum_{m=1}^{12} T_{i,t,m}$ as the average temperature in year t, where $T_{i,t,m}$ is the temperature of country i in year t and month m. Additionally, I refer to $\tilde{\mu}^j_{i,t} = E[(T_{i,t,m} - \bar{T}_{i,t})^j]$ as the j'th sample central moment of the temperature distribution in country i at year t, where $\tilde{\mu}^2_{i,t}$ is the within-year variance (i.e. $\sqrt{\tilde{\mu}^2_{i,t}}$ is the standard deviation), and $\tilde{\mu}^3_{i,t}$ is the within-year skewness. Therefore, a standard growth model that seeks to estimate the impact of various moments of the temperature distribution on growth would take the form:

$$g_{i,t} = g_i + \bar{T}_{i,t} + \sqrt{\tilde{\mu}^2_{i,t}} + \tilde{\mu}^3_{i,t} \tag{1}$$

where $g_{i,t}$ is the growth rate of country i in time t. In a panel context with controls, this equation becomes:

$$\begin{aligned} g_{i,t} = \alpha_i + \alpha_t + \beta_0 \mathbf{1}_{Poor} \beta_1 \textbf{Mean}_{i,t} + \beta_2 \textbf{Standard Deviation}_{i,t} + \beta_3 \textbf{Skewness}_{i,t} + \\ \beta_4 \textbf{Mean}_{i,t} * \mathbf{1}_{Poor} + \beta_5 \textbf{Standard Deviation}_{i,t} * \mathbf{1}_{Poor} + \beta_6 \textbf{Skewness}_{i,t} * \mathbf{1}_{Poor} + \epsilon_{i,t} \end{aligned} \tag{2}$$

where α_i are country fixed effects and α_t are time fixed effects. I include these country fixed effects in order to account for differences in mean growth rates between countries, and include time fixed effects to account for global drivers of growth common to all countries. For example, if a country with high temperatures on average has lower growth rates, then I cannot conclude that rising temperatures reduce growth as there may be some other characteristics shared by hot countries that cause growth rates to fall. However, by including country fixed effects, I demean growth rates and temperatures which allows me to determine whether marginal increases in temperature reduce growth. Additionally, I choose to include time fixed effects

[1] The derivation of the growth equation can be found in section 1 of the appendix.

in order to account for time-varying effects that systematically affect growth for all countries by demeaning my variables in each period. Finally, I include interactions between my temperature moments and a poor country dummy ($\mathbf{1}_{Poor}$), as well as the poor country dummy itself, which takes a value of 1 if the country has below-median purchasing power parity-adjusted GDP per capita in their first year of the sample, to account for potential growth differences between rich and poor countries.

It should be emphasized that the dependent variable in equation 1 would, in the standard empirical growth literature, be mean growth; but as a matter of logic, the dependent variable could be a specific quantile or alternative measure of central tendency of the growth distribution. The underlying relationships linking growth and various moments of the temperature distribution might be the same or different depending on the exact definition of the dependent variable: whether it be mean growth, median growth, or a specific quantile of the growth distribution. The purpose of quantile regression is to assess the degree to which the impact of temperature is homogeneous or heterogeneous across the growth distribution. Equations 1 and 2 do not, of themselves, speak to the issue of whether quantile regression is or is not called for in the case of my specific research question. The motivation for equations 1 and 2 is to link up with the standard growth literature, and also because I begin my empirical results section with an estimation of mean growth.

4.2 Empirical Approach: Quantile Regression

To further determine the effect of temperature on growth, I estimate not only the effects of the first three moments of the temperature distribution on mean growth, but also estimate the effect of these temperature variables on the distribution of growth outcomes. For example, I am interested in the impact of temperature on the median, right-tail, and especially, the left-tail of a country's growth distribution. In the final case, evidence that temperature has stronger effects on the left-tail of growth when compared to the mean supports the view that climate change predicts economic disasters, and merely using point estimates of mean growth rates may conceal important downside risks. My empirical approach to estimating the impact of temperature changes on the distribution of growth outcomes is based on Adrian, Boyarchenko, and Giannone 2019.

Consider a panel dataset of country-year GDP growth rate and temperature observations. Performing an ordinary least-squares regression with two-way fixed effects on this dataset would yield a coefficient estimating the effect of temperature on mean growth by minimizing the sum of squared distances between a best-fit line and my observations. Instead, I am interested in estimating the effects of temperature on different parts of the conditional growth distribution, which I do by estimating equation 2 by quantile regression:

$$g_{i,t}^{\tau} = \alpha_i^{\tau} + \alpha_t^{\tau} + \beta_0^{\tau}\mathbf{1}_{Poor} + \beta_1^{\tau}\mathbf{Mean}_{i,t} + \beta_2^{\tau}\mathbf{Standard\ Deviation}_{i,t} + \beta_3^{\tau}\mathbf{Skewness}_{i,t} + \\ \beta_4^{\tau}\mathbf{Mean}_{i,t} * \mathbf{1}_{Poor} + \beta_5^{\tau}\mathbf{Standard\ Deviation}_{i,t} * \mathbf{1}_{Poor} + \beta_6^{\tau}\mathbf{Skewness}_{i,t} * \mathbf{1}_{Poor} + \epsilon_{i,t}^{\tau} \tag{3}$$

where $\tau \in (0,1)$ represents the quantile of the growth distribution conditional on the temperature variables. In order to compute the coefficients β_k^{τ}, I must introduce a loss function for each value of τ that I am interested in, which is given by the sum of absolute distances between each observation and the best fit line, weighting observations above the best fit line by τ and weighting observations below the best fit line by $(1-\tau)$. Formally, the estimated coefficient for a given quantile τ is given by:

$$\hat{\beta^{\tau}} = argmin_{\beta^{\tau} \in \mathbb{R}^m} \sum_{t=1}^{T} \sum_{i=1}^{I} \tau * \mathbf{1}_{g_{i,t} \geq \beta x_{i,t}} |g_{i,t} - \beta^{\tau} x_{i,t}| + (1-\tau) * \mathbf{1}_{g_{i,t} \leq \beta x_{i,t}} |g_{i,t} - \beta^{\tau} x_{i,t}| \tag{4}$$

By changing the value of τ, the quantile, I adjust the weighting scheme of the loss function to estimate the marginal effect of temperature variables $x_{i,t}$ on different quantiles of GDP (or sectoral) growth, $g_{i,t}^{\tau}$. For example, if I set $\tau = 0.5$, $\hat{\beta}^{0.5}$ is the estimated effect of a unit increase in my temperature variables on the median growth outcome..

5 Empirical Results

In this section, I begin by motivating my application of quantile regression to the Dell, Jones, and Olken 2012 ordinary least-squares regression of mean temperature on mean growth by replicating their results, and then comparing them to quantile regressions performed at the median, as well as the bottom and top 10th percentile of the growth distribution. Then, I augment the panel quantile regressions of mean temperature on the 10th percentile growth rate by including the within-year variance and skewness of the temperature distribution as well, arriving at the full regression equation 3.

Table 1 presents the effect of a 1°C increase in mean temperature on various points along the GDP growth distribution. Column 1 replicates the findings of Dell, Jones, and Olken 2012, which estimates this effect for mean growth. As shown in the first row, an increase in temperature has a small, positive and statistically insignificant effect on mean growth in rich countries. Row 2 shows the effect of a 1°C increase in temperature interacted with a "poor" country dummy on mean growth. This value represents the relative effect of mean temperature on mean growth in rich and poor countries. This coefficient, -1.074, is large, negative, and statistically significant, indicating a clear difference between the effect of mean temperature on growth in rich and poor countries. The point estimate, standard error and p-value of the effect of mean temperature on mean growth in poor countries themselves is computed by conducting a t-test on the sum of the coefficients from rows 1 and 2. Here, I find that the temperature effect on poor countries' mean growth is -0.815 percentage points and is statistically significant at the 1% level.

Table 1: Effect of Change in Mean Temperature on Mean, Median, Bottom and Top 10% GDP Growth Rates

VARIABLES	(1) Mean	(2) Median	(3) Bottom 10%	(4) Top 10%
Temperature Mean	0.259	0.094	0.442	0.461
	(0.270)	(0.184)	(0.491)	(0.348)
Temperature Mean x Poor	-1.074***	-0.265	-1.907***	-1.119*
	(0.391)	(0.281)	(0.642)	(0.635)
Temp mean effect on poor countries	-0.815***	-0.171	-1.464***	-0.659
	(0.310)	(0.234)	(0.520)	(0.488)

Cluster robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.

Next, I test whether changes in mean temperature have an effect on median growth, another measure of central tendency, in order to determine the robustness of the previous result. This is done using quantile regression, where I choose a quantile of $\tau = 0.5$, and the results are shown in column 2. Surprisingly, I find that the coefficients in rows 1 and 2 of column 2 are smaller and less significant than in column 1, indicating that the relationship does not hold at the median of growth. In the case of an asymmetric growth distribution, the median of the distribution would not be affected by large outliers, while the mean of the distribution is shifted toward the outliers, or tail, of the distribution. Therefore, it is worth investigating whether the result found using OLS are due to outlier effects. To do so, I again perform quantile regression, but change my values of τ to 0.1 and 0.9 to see whether changes in temperature can predict left- and right-tail growth outcomes, respectively.

Column 3 presents the coefficients calculated using quantile regression where $\tau = 0.1$. In rich countries, there is a similar story to the one found in Dell, Jones, and Olken 2012, where the effect of increases in temperature are somewhat small, positive, and statistically insignificant. However, I find that temperature interacted with the poor country dummy is nearly twice as large as the value from OLS, while also remaining statistically significant at the 1% level. The net effect of an increase in temperature on the bottom 10th percentile of growth in poor countries is the sum of these two coefficients, or -1.464 percentage points. This indicates that in poor countries, rising temperatures can predict disastrous growth outcomes, and suggests that the negative effect of mean temperature on mean growth may actually be due to its effects on left-tail growth.

Finally, I estimate the effect of temperature on the right-tail of the growth distribution. To do so, I change my value of τ to 0.9, and present the results in Column 4. In rich countries, an increase in temperature has a small, positive and statistically insignificant effect on the 90th percentile of growth. In the second row, the sign is reversed, but unlike the OLS and quantile regression (τ=0.1) results, the coefficient is only significant at the 10% level. Moreover, the effect of temperature mean on the top 10th percentile of growth in poor countries is the sum of rows 1 and 2, which equals -0.659. This value is lower than the effect on poor countries in both columns 1 and 3.

Altogether, these results lead to two important conclusions: first, the effect of changes in mean temperature on mean growth rates can be partially explained by changes in the left tail of the growth distribution; and second, rising temperatures can predict disastrous growth outcomes in poor countries.

While increases in average temperature have a large, negative effect on the bottom 10th percentile of the growth distribution, it might be the case that this coefficient is picking up the effect of higher moments of the temperature distribution. As climate change worsens, mean temperatures increase, but temperature distributions may become more volatile and asymmetric. It is not unreasonable to hypothesize that greater within-year temperature volatility, or some months of very extreme temperatures, might affect growth rates in that year. I test this hypothesis in Table 2, where I regress the 10th percentile growth rate against the mean, standard deviation and skewness of the temperature distribution, as well as their interactions with the poor dummy. Column 1 presents the same result as in Column 3 of Table 1, and is meant to serve as a baseline for comparing quantile regression models with models containing higher moments of the temperature distribution.

Column 2 in Table 2 presents the regression equation of column 1, but with the inclusion of temperature standard deviation as an independent variable. The effect of changes in mean temperature in both rich and poor countries remains quite similar to the baseline model in terms of magnitude, direction and significance. In rich countries, I find that a 1 unit increase in the within-year standard deviation of temperature decreases the bottom 10th percentile growth rate by 0.79 percentage points. This result follows the intuition that countries do not experience a growth benefit from large variations in monthly temperatures. For example, greater variation in temperature might cause crop failures, discourage investment due to the unpredictability of weather, damage ecological communities, or reduce tourism in the greater number of months with undesirable temperatures. However, I find that the effect of a unit increase in temperature standard deviation interacted with the poor country dummy has a positive effect on the bottom 10th percentile of growth, such that the net effect on poor countries is positive (0.600), though the net effect in poor countries is statistically insignificant. This result runs counter to my previous explanations as poor countries seem to do better when temperature is more volatile. One potential explanation for this result is gambling for resurrection, a phenomenon observed in financially-distressed banks or politically-distressed countries. This theory posits that

Table 2: Effect of Change in Moments of Temperature Distribution on Bottom 10% GDP Growth Rate

Dependent variable is $g_{i,t}^{\tau=0.1}$	(1)	(2)	(3)	(4)
Temperature Mean	0.442	0.267	0.329	0.231
	(0.491)	(0.532)	(0.618)	(0.493)
Temperature Mean x Poor	-1.907***	-1.789***	-1.762***	-1.723***
	(0.642)	(0.626)	(0.679)	(0.606)
Temperature Standard Deviation		-0.786**		-0.584
		(0.400)		(0.433)
Temperature Standard Deviation x Poor		1.386**		1.039
		(0.622)		(0.645)
Temperature Skewness			1.825**	1.517
			(0.858)	(1.015)
Temperature Skewness x Poor			-2.102	-1.535
			(1.348)	(1.544)
Temp mean effect on poor countries	-1.464***	-1.522***	-1.432***	-1.492***
	(0.520)	(0.464)	(0.472)	(0.519)
Temp S.D. effect on poor countries		0.600		0.455
		(0.460)		(0.459)
Temp skewness effect on poor countries			-0.277	-0.018
			(0.973)	(1.155)

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.

these distressed countries or financial institutions will take part in excessive risk-taking. If the risky behavior has a bad outcome, then the country or bank will not lose much compared to if they had not taken on the risk. However, if the risk ends up having a positive outcome, it may be enough to get the bank or country out of its bad situation, hence it is gambling for its resurrection. In the context of temperature volatility, the analogy is that poor countries, which are the countries where higher temperatures worsen left-tail growth (Column 1), may gain from an extreme cold-temperature draw, without being hurt much by an extreme hot-temperature draw. The rationale is similar to the financial literature case. Poor countries may already be experiencing sufficiently hot temperatures that the cost of a warmer draw affects them little, just as a bad draw for an insolvent business will not materially affect its financial condition. But such poor countries may benefit quite a bit from a cold temperature draw, just as a financially-distressed corporation may be able to exit from insolvency with a sufficiently favorable shock.

So far I have explored the impact of the first two moments of the temperature distribution. But it is possible that the temperature distribution is not symmetric, and that I also need to explore the effect of within-year temperature skewness. This is shown in column 3, where I regress the bottom 10th percentile growth rate against temperature mean and skewness. Here, the results are quite interesting and intuitive. In rich countries, greater skewness, which means a greater propensity to experience hot temperatures, improves left tail growth outcomes, which may relate to the fact that rich countries tend to be colder, so some hot months may create growth opportunities for those countries. Similarly, lower skewness, or the existence of several very cold months, would be harmful in rich countries who already live in a colder climate. Conversely, in poor countries, greater skewness very slightly reduces left-tail growth outcomes, though this result is not significant at the 10% level. As poor countries tend to be warmer, the existence of some very hot months may make production and living conditions unbearable, while some cold months would improve growth outcomes. However, it is important to note that the point estimate of the effect of skewness in poor countries is not statistically significant.

Finally, I regress bottom 10th percentile growth rate against the first three moments of the temperature distribution, and their interaction with the poor dummy (equation 3), and I present these results in Column 4. In this case, I find that an increase in mean temperature reduces growth at risk in poor countries, similar to my previous three models. However, the results for temperature standard deviation and skewness become less significant, indicating that for aggregate GDP growth, these variables may not matter.

While the higher moments of the temperature distribution do not appear to significantly affect aggregate growth at risk, they may be masking important underlying effects on sectoral growth. For example, it may be the case that some temperature-vulnerable sectors, such as agriculture and industry, are affected by changing temperatures in one way, and these effects are offset by growth in other sectors. [2]

6 Channels

In this section, I estimate the relationship between the within-year mean, standard deviation and skewness of the temperature distribution and the bottom 10th percentile growth rates in certain sectors, using the specification of equation 3. My motivation for this approach is to determine the channels through which temperature affects growth in order to tell which sectors and countries are the most vulnerable to climate change.

6.1 Agriculture

Sachs and Warner 1997, Burke, Hsiang, and Miguel 2015, Dell, Jones, and Olken 2012, and Angelova and Käbel 2019 argue that weather is an integral component for agricultural production. For example, rising temperatures may make growing certain crops unfeasible, and in the short term, adapting to these changes by planting different crops is difficult. Greater within-year temperature variance and skewness signal ex-

[2] For reference, I include panel regression result that estimate the effect of the temperature variables employed in this section on mean GDP growth in Table 6 of the appendix.

treme weather events and shorter growing seasons for certain crops. In order to estimate the relationship between the left-tail of agricultural growth outcomes and temperature, I present a quantile regression of the first three moments of the temperature distribution on the 10th percentile agricultural growth rate in Table 3.

Table 3: Effect of Change in Moments of Temperature Distribution on Bottom 10% Agricultural Growth Rate

Dependent variable is $agr_{i,t}^{\tau=0.1}$	(1)	(2)	(3)	(4)
Temperature Mean	-1.361**	-2.146*	-1.525**	-2.650***
	(0.571)	(1.242)	(0.684)	(0.908)
Temperature Mean x Poor	-1.285	-0.319	-1.024	0.032
	(0.875)	(1.240)	(0.910)	(1.022)
Temperature Standard Deviation		-3.295		-3.306**
		(2.193)		(1.636)
Temperature Standard Deviation x Poor		3.612		3.154*
		(2.304)		(1.825)
Temperature Skewness			2.749**	2.992***
			(1.369)	(1.040)
Temperature Skewness x Poor			-5.098**	-5.998***
			(2.287)	(1.965)
Temp mean effect on poor countries	-2.646***	-2.464***	-2.549***	-2.618***
	(0.822)	(0.829)	(0.856)	(0.738)
Temp S.D. effect on poor countries		0.317		-0.152
		(1.314)		(1.064)
Temp skewness effect on poor countries			-2.349**	-3.006***
			(1.408)	(1.113)

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.

Column 1 estimates the effect of mean temperature on the bottom 10th percentile agricultural growth rate. In rich countries, a 1°C increase in mean temperature reduces growth rates by 1.361 percentage points. In poor countries, however, this effect is nearly twice as strong, where a 1°C increase in temperature reduces growth rates by 2.646 percentage points. There are several factors that may account for the difference in the effect of mean temperature on agricultural growth between rich and poor countries. First, rich countries have greater access to sophisticated irrigation systems and to stable water supplies than poor countries, making crops more resilient to droughts and heat waves. Second, large-scale farms use pesticides and genetically-modified seeds which prevent insects from eating their crops. These farms are typically found in rich countries, while poor countries typically employ small-scale and subsistence farming practices. Third, reductions in crop output reduce the supply of animal feed, and because the agricultural sector includes both crops and livestock, any decrease in crop output results in an amplified decrease in agricultural output due to a decrease in livestock.

Column 2 adds within-year temperature standard deviation into the model. Surprisingly, in this model, greater variance is not especially significant for agricultural growth rates, though the point estimate shows that in rich countries, greater within-year temperature variation reduces agricultural growth-at-risk, while in poor countries, this effect is negligible. Additionally, with the inclusion of standard deviation into the model, rising mean temperatures adversely affects agricultural growth at risk in rich countries more than in the model from Column 1.

In column 3, I replace temperature standard deviation with temperature skewness. In this model, I find a very intuitive result: a one unit increase in temperature skewness improves left-tail agricultural growth in rich countries by 2.749 percentage points. This follows the same reasoning as the earlier GDP growth regressions, where rich countries tended to be cold, so the emergence of a right-skewed temperature distribution would be favorable for growing crops. Conversely, a one unit decrease in temperature skewness – leading to a more left-skewed temperature distribution – would reduce agricultural growth-at-risk as very cold temperatures would not provide adequate conditions for farming in an already-cold country. In poor countries, the opposite is the case, as a 1 unit increase in temperature skewness deteriorates poor countries' agricultural growth at risk by 2.349 percentage points, with this effect significant at the 5% level. This is likely because poor countries tend to be hotter on average, so some extremely hot monthly temperatures (in line with an increase in skewness) would tend to deteriorate agricultural growth-at-risk.

Finally, column 4 considers a model with temperature mean, standard deviation and skewness. In this model, I find that a 1°C increase in mean temperature reduces the bottom 10th percentile agricultural growth rate by 2.650 percentage points in rich countries, and by 2.618 percentage points in poor countries. Next, in rich countries, a one unit increase in temperature standard deviation reduces the bottom 10th percentile agricultural growth by -3.306 percentage points, while it has a negligible effect on poor countries. Finally, a one unit increase in temperature skewness increases the bottom 10th percentile agricultural growth rate in rich countries by 2.992 percentage points but reduces the bottom 10th percentile agricultural growth rate in poor countries by 3.006 percentage points. Altogether, this highlight the salience of the higher moments of the temperature distribution for explaining left-tail agricultural growth outcomes in both rich and poor countries.

6.2 Industry

Another channel through which temperature may affect growth at risk is the industrial sector. Higher temperatures may decrease labor productivity, leading to declines in manufacturing output as described in Sudarshan and Tewari 2014. Furthermore, it is worth investigating whether the higher moments of the temperature distribution can predict left-tail industrial growth rates. One might expect that with greater temperature volatility comes more frequent natural disasters that can lead to the destruction of machinery and manufacturing plants, as well as the displacement of workers in the industrial sector. Therefore, I run quantile regressions of the first three moments of the temperature distribution on the bottom 10th percentile industrial growth rate, and present these results in Table 4, below.

Table 4: Effect of Change in Moments of Temperature Distribution on Bottom 10% Industrial Growth Rate

Dependent variable is $ind_{i,t}^{\tau=0.1}$	(1)	(2)	(3)	(4)
Temperature Mean	0.856	0.643	0.795	0.660
	(0.565)	(0.601)	(0.605)	(0.621)
Temperature Mean x Poor	-3.003***	-2.904***	-3.047***	-3.036***
	(0.872)	(0.893)	(0.857)	(0.914)
Temperature Standard Deviation		-0.295		-0.113
		(0.588)		(0.712)
Temperature Standard Deviation x Poor		2.057		1.717
		(1.492)		(1.509)
Temperature Skewness			1.288	1.051
			(2.607)	(3.683)
Temperature Skewness x Poor			-1.197	-0.952
			(2.845)	(3.839)
Temp mean effect on poor countries	-2.147***	-2.260***	-2.251***	-2.376***
	(0.793)	(0.730)	(0.742)	(0.737)
Temp S.D. effect on poor countries		1.762		1.604
		(1.343)		(1.372)
Temp skewness effect on poor countries			0.092	0.099
			(1.470)	(1.537)

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.

Columns 1-4 present the effects of temperature shocks on industrial growth-at-risk. In each of these models, I find a very constant relationship between mean temperature and the bottom 10th percentile industrial growth rate. Rich countries' industrial output appears to be fairly resilient to rising mean temperatures. This is likely due to the fact that rich countries tend to be colder, so rising temperatures may not be much of an issue for labor productivity. Additionally, manufacturing plants in rich countries will often contain amenities such as air conditioning and water fountains to keep workers cool and hydrated. However, in poor countries, a 1 °C increase in mean temperature results in between a 2.1 and 2.4 percentage point decrease in industrial growth rates. This effect follows from the intuition that, unlike rich countries, poor countries often cannot mitigate reductions in labor productivity resulting from higher temperatures.

In terms of the effects of the higher moments of the temperature distribution on industrial growth at risk, there appears to be a very minimal and statistically insignificant relationship. None of the coefficients of temperature standard deviation and skewness in both rich and poor countries are significant at the 10% level, so I conclude that the central mechanism through which temperature affects industrial growth-at-risk is rising mean temperature.

To summarize, I find that in poor countries, rising mean temperatures worsen downside risks to agricultural and industrial growth. Conversely, for rich countries, only mean temperatures affect left-tail agricultural growth. Finally, agricultural growth-at-risk outcomes worsen in rich countries when faced with more extreme cold temperatures or more temperature volatility, while growth-at-risk outcomes worsen in poor countries when faced with more extreme hot temperatures.[3]

7 Country Case Studies

So far, I have estimated the effects of the within-year mean, standard deviation and skewness of the temperature distribution on the bottom 10th percentile aggregrate, agricultural and industrial growth rates. However, it is helpful to visualize how changing these temperature variables impacts the entire growth distribution, not just the bottom 10th percentile. In this section, I use a counterfactual exercise to assess the effects of temperature shocks on a country's growth distribution using country case studies. The code used in this section comes from the Growth-at-Risk tool developed in Prasad et al. 2019, and more technical details are available in Lafarguette 2019 and the associated GitHub page.[4]

7.1 Methodology

In order to perform this analysis, I begin by estimating the empirical predicted quantiles of a country's growth distribution conditional on the within-year mean, standard deviation and skewness of the country's temperature distribution. This is done in two stages. First, I estimate the quantile marginal effects, β^τ, from equation 3 by solving equation 4 for a particular country, but without country fixed effects or poor dummy interactions as this is a pure time-series regression.[5] Second, I use these β^τ coefficients to compute the empirical predicted quantiles according to:

$$\hat{y}_t^\tau | x_t = \hat{\beta}^\tau * x_t \tag{5}$$

for $\tau \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$, where x_t are the temperature moments for a given country at time t.

[3]For reference, I include panel regression result that estimate the effect of the temperature variables employed in this section on mean agricultural and industrial growth in Tables 7 and 8 of the appendix.

[4]Link to the Growth-at-Risk tool and mathematical appendix: `https://github.com/IMFGAR/GaR`

[5]Similarly, the quantile regression minimization in equation 4 no longer includes the summation with respect to countries i, as I am in a pure time-series space.

Next, I interpolate between the five estimated empirical quantile functions by fitting them to theoretical quantiles of a skew-t distribution, as developed by Azzalini and Capitanio 2003. The probability distribution function takes the form:

$$f(y;\mu,\sigma,\alpha,v) = \frac{2}{\sigma}\ \mathbf{t}\left(\frac{y-\mu}{\sigma};v\right)\ \mathbf{T}\left(\alpha\frac{y-\mu}{\sigma}\sqrt{\frac{v+1}{v+\left(\frac{y-\mu}{\sigma}\right)^2}};v+1\right) \tag{6}$$

where $\mathbf{t}(.)$ and $\mathbf{T}(.)$ are the probability distribution function and cumuluative distribution function of the Student t-distribution. In order to fit the quantiles to this distribution, the code employ matching quantile estimation (Sgouropoulos, Yao, and Yastremiz 2015), which minimzes the following objective function:

$$\{\hat{\mu}_t,\hat{\sigma}_t,\hat{\alpha}_t,\hat{v}_t\} = argmin_{\{\mu,\sigma,\alpha,v\}}\sum^{\tau}[\hat{y}_t^{\tau}|x_t - Q^{\tau}(\mu,\sigma,\alpha,v)]^2 \tag{7}$$

where $Q^{\tau}(\mu,\sigma,\alpha,v)$ denotes the theoretical quantile function of the skew-t distribution in equation 6. Essentially, minimizing equation 7 chooses a set of values for the location (μ), scale (σ), shape (α) and degrees of freedom (v) parameters that minimize the sum of squared distances between the estimated empirical quantile functions and the theoretical quantiles of the skew-t distribution. The benefit of fitting the data to the skew-t distribution is that its distribution can be fully characterized by only four parameters, allowing me to reduce potential overfitting of the data and making it less vulnerable to noise and outliers, as opposed to a non-parametric fit. Additionally, the skew-t distribution has been used extensively in financial econometrics in order to model tail events, and this approach is equally useful for my temperature analysis. In this analysis, I have chosen to fit the skew-t distribution in the final year of my sample, and this year varies by country.

I then use this framework to study how temperature shocks affect the entire conditional distribution of (sectoral) growth. I perform this comparative statics exercise by shocks a moments of the temperature distribution, for example by setting the shocked mean temperature equal to $\widetilde{Mean_t} = Mean_t + shock$, where the shock is equal to one standard deviation of a given country's mean temperature over my sample, and then re-estimating the empirical quantile functions according to equation 5. I then fit these new empirical quantile functions to the skew-t using equation 7 and output a new "shocked" (sectoral) growth distribution, which I then compare with the pre-shock dsitribution to assess the effects of mean temperature shocks.

In this section, I construct conditional growth distributions for one rich (the Netherlands) and one poor country (Haiti), where rich and poor are defined as having above- and below-median GDP PPP per capita at the beginning of my sample, respectively, in order to maintain consistency with my regression results. Next, I study the effects of shocks to the moments of the temperature distribution on the aggregate, agricultural and industrial growth distributions. Finally, I describe how the shocks have affected the shape of the conditional growth distribution functions.

7.2 The Netherlands

In my results section, I presented a finding that rich countries' left tail aggregate and industrial growth are quite resilient to temperature shocks, but that their agricultural sectors were vulnerable to temperature shocks. In order to illustrate this result at the country-level, I begin by modelling the agricultural sector's conditional growth distribution in the Netherlands, which is shown in Figure 3a.

This figure shows the agriculture sector's conditional growth distribution in the Netherlands in 2005. The mean growth rate, as shown by the blue dotted line, is 7.5%, while the bottom 5th percentile growth rate, which I now refer to as the growth-at-risk, is -3.3%. The growth-at-risk is the threshold growth rate for which 5% of the growth distribution's probability density (the red shaded area) is to the left of this value. When I apply temperature shocks to this temperature distribution, one can see how growth-at-risk changes. The first shock that I impose on this model is a one standard deviation (0.688° C) increase in mean temperature, which is shown in Figure 3b.

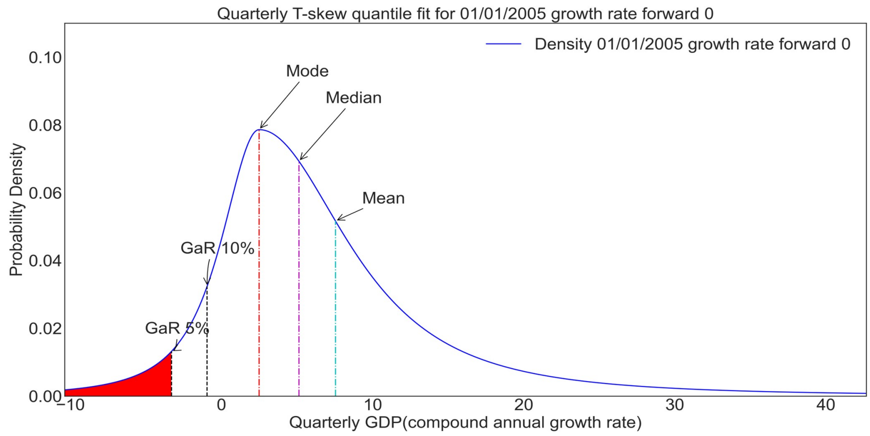

Figure 3a: Agricultural Sector's Conditional Growth Distribution - Netherlands

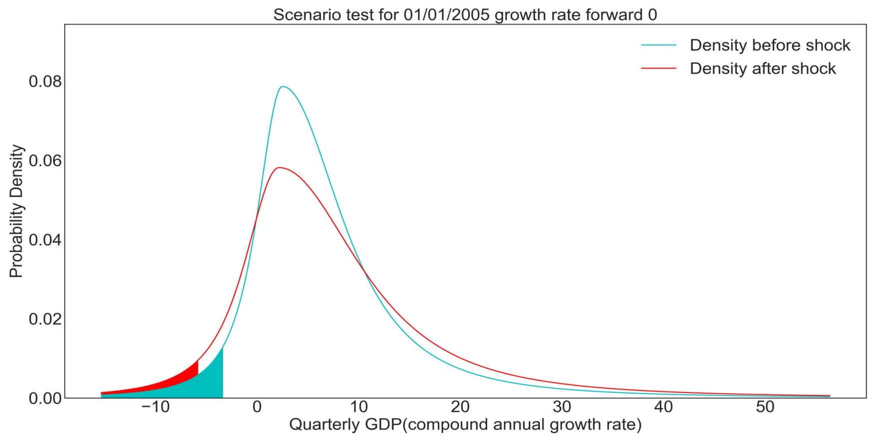

Figure 3b: Mean Temperature Shock in Agricultural Sector - Netherlands

As presented in this image, rising mean temperatures worsen left-tail growth. This is shown by the leftward shift in the growth-at-risk of the shocks distribution relative to the unshocked distribution. Moreover, it appears that changing mean temperatures had little effect on the mode of the growth distribution, which, like the mean, is a measure of central tendency. This agrees with my panel regression result which found that rising mean temperatures worsen left-tail agricultural growth in rich countries.

Next, I illustrate the effect of greater temperature volatility on the growth distribution. Figure 3c shows the effect of a one standard deviation increase (0.602° C) increase in within-year temperature standard deviation.

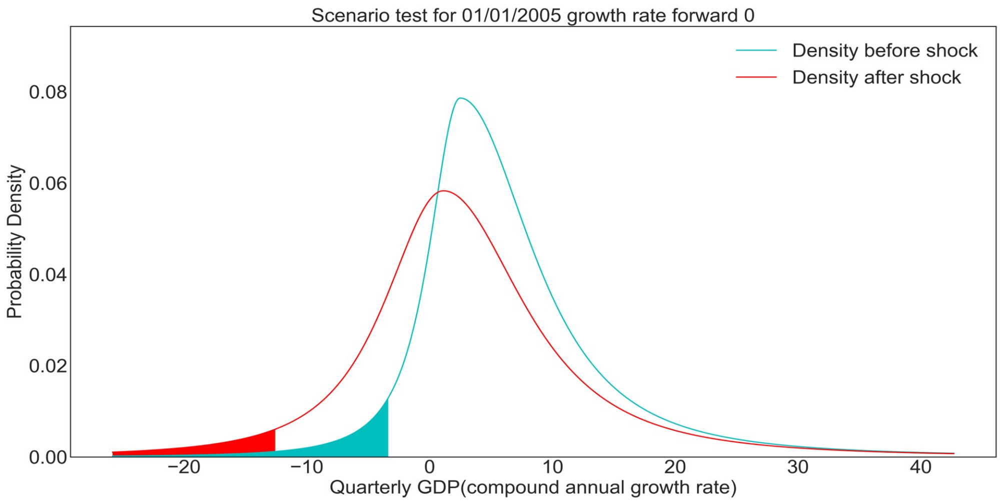

Figure 3c: Temperature Volatility Shock in Agricultural Sector - Netherlands

This image illustrates that increasing within-year temperature volatility worsens left-tail agricultural growth in the Netherlands. While growth-at-risk was -3.3% before the shock, this value decreased significantly to -12.6% after a 0.602° C increase in temperature standard deviation. This result agrees with my previous finding that rising temperature volatility worsens left-tail agricultural growth rates in rich countries.

Finally, I model the effect of a more left-skewed temperature distribution on the agricultural growth distribution. Figure 3d illustrates a one standard deviation (0.217 units) decrease in within-year temperature skewness on the Netherland's agricultural growth distribution.

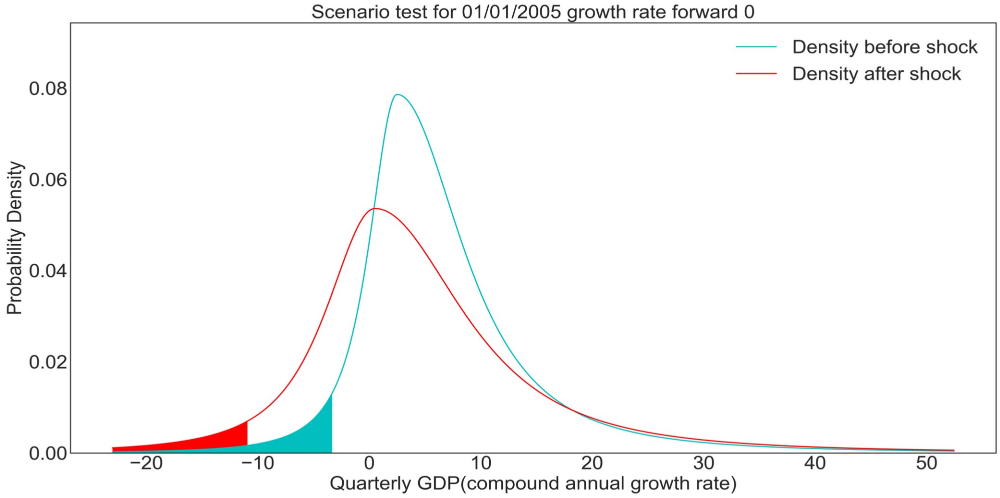

Figure 3d: Temperature Skewness Shock in Agricultural Sector - Netherlands

This figure shows that a more negatively skewed within-year temperature distribution, i.e. a greater number of very cold months, worsens left-tail growth in the Netherlands' agricultural sector. Moreover, this shock has a very similar effect to the temperature standard deviation shock in Figure 3c. This further supports my finding in Table 3, where the coefficients for temperature standard deviation and skewness in rich countries were very similar in magnitude (2.992 and -3.306).

To summarize, I modeled the effects of temperature shocks on the growth distribution in the Netherlands' agriculture sector, and I found similar results as the ones presented in my agricultural panel regressions. Specifically, I found that rising mean temperatures, greater temperature volatility, and more extreme cold temperatures make worst-case agricultural growth outcomes more severe in the Netherlands.

7.3 Haiti

While temperature shocks typically affect the agricultural sector in rich countries, poor countries were vulnerable to these shocks at the aggregate level, as well as in their agricultural and industrial sectors. In this section, I model the effects of various temperature shocks on the aggregate, agricultural and industrial growth distributions in Haiti.

To begin, I first construct the conditional GDP growth distribution for Haiti in 1992, which is their last year in my sample. This is shown in Figure 4a. Given this conditional growth distribution, I can illustrate the effects of temperature shocks on the distribution of GDP growth. Figure 4b shows the effect of a one standard deviation (0.635 °C) increase in mean temperature on the Haiti's conditional growth distribution.

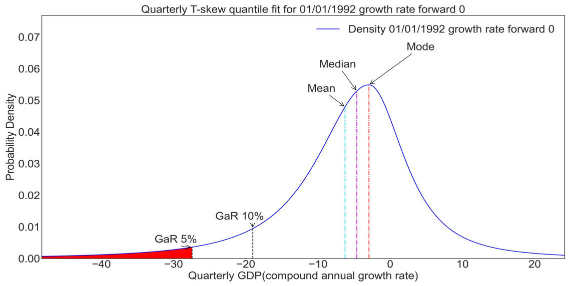

Figure 4a: Conditional GDP Growth Distribution - Haiti

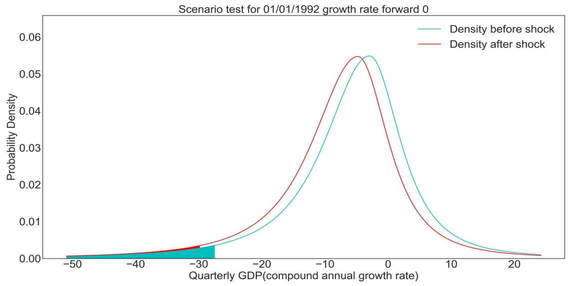

Figure 4b: Mean Temperature Shock and Aggregate Growth - Haiti

From these figures, it is evident that rising mean temperatures reduce left-tail GDP growth in Haiti. However, it is important to note that this temperature shock appears to have shifted the entire distribution left by a fairly uniform amount, indicating that, at least in Haiti, rising mean temperatures do not have a stronger growth effect in the left-tail than at the mean. This result disagrees with my result in Table 1, which found more negative growth effects in the left-tail than the mean due to rising temperatures in poor countries.

Next, I consider the effects of rising mean temperatures and more extreme hot months on the conditional agricultural growth distribution in Haiti. Figure 5a depicts the conditional growth distribution in Haiti's agricultural sector, while Figures 5b and 5c model the effects of a one standard deviation increase in temperature mean and skewness on the growth distribution, respectively.

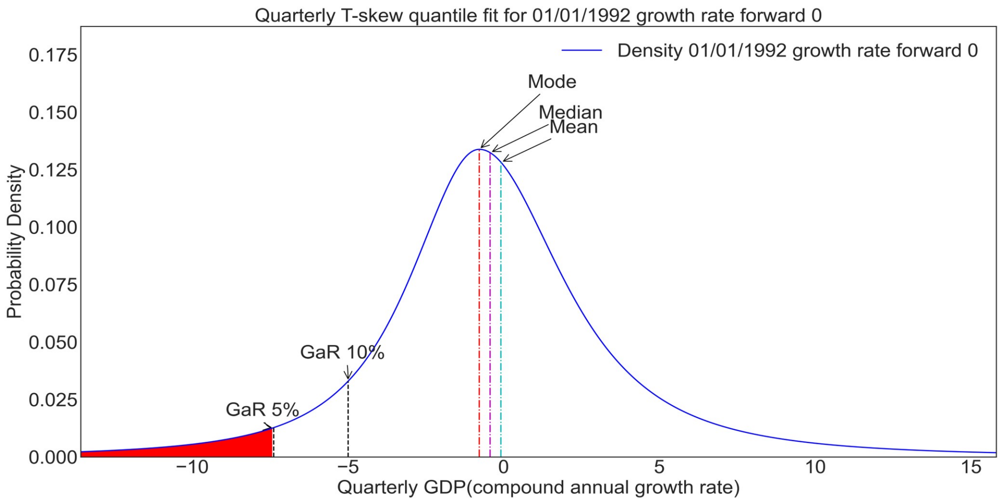

Figure 5a: Conditional Agricultural Growth Distribution - Haiti

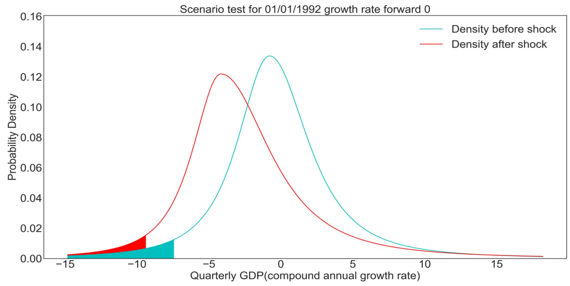

Figure 5b: Mean Temperature Shock in Agricultural Sector - Haiti

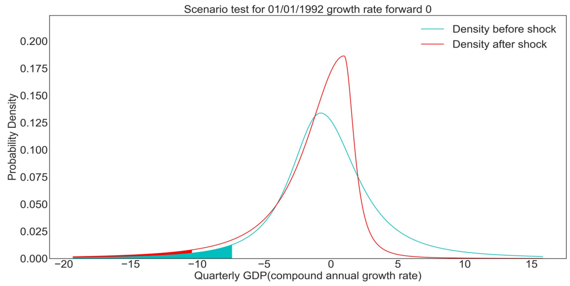

Figure 5c: Temperature Skewness Shock in Agricultural Sector - Haiti

Figures 5b and 5c indicate that rising mean temperatures, as well as more extreme hot temperatures, worsen downside risks to agricultural growth in Haiti. A one standard deviation (0.635 °C) increase in mean temperature reduces the bottom 5% growth rate from -7.4% to -9.4%. Additionally, Figure 5b illustrates that a mean temperature shock in Haiti shifts the mode and the left-tail of the distribution, while the right-tail remains somewhat the same. This would indicate that rising mean temperatures do not affect upside risks to agricultural growth.

Next, a one standard deviation (0.123 units) increase in temperature skewness reduces the bottom 5th percentile growth rate, the growth-at-risk, to -10.4%. Unlike the mean temperature shock, the mode of the growth distribution in Figure 5b does not move in the same direction as the left-tail. This indicates that more extreme hot months do not affect the expected growth outcomes much, but are salient for predicting the worst-case outcomes.

Finally, I assess the effect of rising mean temperatures on the industrial growth distribution in Haiti. Figure 6a depicts the industrial sector's conditional growth distribution in Haiti, while Figure 6b illustrates the effect of a positive temperature shock on this distribution.

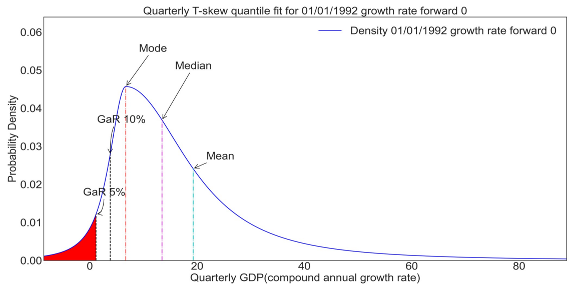

Figure 6a: Mean Temperature Shock in Industrial Sector - Haiti

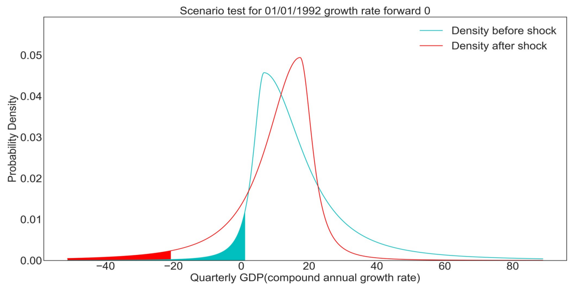

Figure 6b: Mean Temperature Shock in Industrial Sector - Haiti

The conditional growth distribution in Figure 6a is right-skewed, and the bottom 5th percentile growth rate is fairly high (1.1%). However, a one standard deviation (0.635 °C) increase in mean temperature has a strong effect on this distribution. First, the bottom 5th percentile growth rate falls to -20.8%. This result agrees with my panel regression results, which finds that rising mean temperatures reduce industrial growth-at-risk in poor countries. Second, the growth distribution becomes significantly more left-skewed, indicating that in Haiti, rising mean temperatures worsen downside risks more than the expected value and upside risks to industrial growth.

When performing these country case studies, it is important to consider that the shape of the conditional growth distributions and their behavior in response to temperature shocks will vary between countries. However, this exercise is nonetheless useful for two reasons: first, it allows me to show the extent to which my panel regression results hold at the country-level; and second, it allows me to create a conceptual map between the temperature distribution and the entire growth distribution, not just its left-tail or mean. To summarize my findings from these country case studies, I found that the Netherlands agricultural sector responds similarly to rich countries in my panel regressions when faced with temperature shocks. Likewise, Haiti's GDP, agricultural and industrial growth distributions behave in a similar way as poor countries in response to shocks to the first three moments of the temperature distribution.

8 Conclusion

In this paper, I first estimated the effect of rising mean temperature on the mean, median, bottom and top 10% of the GDP growth distribution using ordinary least-squares and quantile regression. I found that, in poor countries, mean temperature has a much larger negative effect on left-tail growth than on mean growth, which motivated the remainder of my panel quantile regressions on the bottom 10th percentile growth rate. Second, I incorporated within-year temperature standard deviation and skewness into my regressions, though I found that these variables are not significant in predicting left-tail GDP growth. Third, I estimated the effects of the first three moments of the temperature distribution on two mediating channels through which temperature may affect GDP growth: the agricultural and industrial sectors. I found that in poor countries, rising mean temperatures worsened left-tail industrial growth. However, all three moments of the temperature distribution mattered for left-tail agricultural growth in both rich and poor countries. Specifically, greater temperature mean and volatility, as well as more extreme cold temperatures (left-skewed temperature distributions), reduced left-tail agricultural growth in rich countries. In poor countries, rising mean temperatures and more extreme hot temperatures (right-skewed temperature distributions) reduced left-tail agricultural growth. Finally, through country case studies in Haiti and the Netherlands, I constructed conditional aggregate and sectoral growth distributions and compared their shape, and especially, their left-tail, to counterfactual distributions estimated by applying shocks to my temperature variables. I found that for these two countries, applying temperature shocks affected the left-tail of their growth distributions in similar ways to my panel regression findings.

The results and methodologies used in my paper suggest two directions for future research. First, while structural and computational models of the macroeconomic effects of climate change (such as integrated assessment models) include mean temperatures as their temperature variable of choice, adding the higher moments of the temperature distribution would be a worthwhile addition to these models. Second, in my analysis, I argue that changes in the higher moments of the temperature distribution signal natural disasters resulting from climate change. Thus, including natural disasters themselves as regressors in my model could offer a more complete picture of the empirical link between climate change and macroeconomic growth.

References

Acemoglu, Daron, Simon Johnson, and James A. Robinson (Dec. 2001). "The Colonial Origins of Comparative Development: An Empirical Investigation". In: *American Economic Review* 91.5, pp. 1369–1401.

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (Apr. 2019). "Vulnerable Growth". In: *American Economic Review* 109.4, pp. 1263–1289.

Angelova, Denitsa and Jan Käbel (2019). "Weather Volatility and Production Efficiency". In: *Multidisciplinary Digital Publishing Institute*.

Azzalini, Adelchi and Antonella Capitanio (2003). "Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t -distribution". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65, pp. 367–389.

Balk, Deborah et al. (2004). *Global Rural-Urban Mapping Project (GRUMP), v1*. Center for International Earth Science Information Network.

Bond, Stephen, Asli Leblebicioglu, and Fabio Schiantarelli (2010). "Capital accumulation and growth: a new look at the empirical evidence". In: *Journal of Applied Econometrics* 25.7, pp. 1073–1099.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (Oct. 2015). "Global non-linear effect of temperature on economic production". In: *Nature* 527, pp. 235–239.

Colacito, Riccardo, Bridget Hoffmann, and Toan Phan (2018). "Temperature and Growth: A Panel Analysis of the United States". In: *Journal of Money, Credit and Banking* 51, pp. 313–368.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (July 2012). "Temperature Shocks and Economic Growth:Evidence from the Last Half Century". In: *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

Group, World Bank (2021a). *Climate Change Knowledge Portal*.

— (2021b). *World Development Indicators*.

Held, Isaac M. and Brian J. Soden (Nov. 2006). "Robust Responses of the Hydrological Cycle to Global Warming". In: *Journal of Climate* 19.21, pp. 5686–5699.

Keeling, Charles D. et al. (Dec. 1976). "Atmospheric carbon dioxide variations at Mauna Loa Observatory, Hawaii". In: *Tellus* 28.6, pp. 538–551.

Kiley, Michael T. (Aug. 2021). "Growth at Risk From Climate Change". In: *Finance and Economics Discussion Series* 54.

Koenker, Roger and Gilbert Bassett (1978). "Regression Quantiles". In: *Econometrica* 46.1, pp. 33–50.

Kotz, Maximilian et al. (Feb. 2021). "Day-to-day temperature variability reduces economic growth". In: *Nature Climate Change* 11.4, pp. 319–325.

Lafarguette, Romain (Feb. 2019). "Growth at Risk Tool: Technical Appendix". In: *International Monetary Fund* 36.

Lee, June-Yi and Jochem Marotzke (2022). “Future Global Climate: Scenario-based Projections and Near-term Information”. In: *IPCC Sixth Assessment Report*. Ed. by Masahide Kimoto Krishna Kumar Kanikicharla Vladimir Kattsov. Cambridge University Press, pp. 553–672.

Manabe, Syukuro and Richard T. Wetherald (May 1967). “Thermal Equilibrium of the Atmosphere with a Given Distribution of Relative Humidity”. In: *Journal of The Atmospheric Sciences* 24.3, pp. 241–259.

Matsuura and Willmott (2007). *Terrestrial Air Temperature and Precipitation: 1900–2006 Gridded Monthly Time Series, Version 1.01.* University of Delaware.

Prasad, Ananthakrishnan et al. (Feb. 2019). “Growth at Risk: Concept and Application in IMF Country Surveillance”. In: *IMF Working Papers* 36.

Sachs, Jeffrey D. and Andrew M. Warner (Oct. 1997). “Sources of Slow Growth in African Economies”. In: *Journal of African Economies* 6.3, pp. 335–376.

Sgouropoulos, Nikolaos, Qiwei Yao, and Claudia Yastremiz (Aug. 2015). “Matching a distribution by matching quantiles estimation”. In: *Journal of the American Statistical Association* 110.510, pp. 742–759.

Sudarshan, Anant and Meenu Tewari (July 2014). “The Economic Impacts of Temperature on Industrial Productivity: Evidence from Indian Manufacturing”. In: *Indian Council for Research on International Economic Relations* 278.

9 Appendix

9.1 Descriptive Statistics of Growth and Temperature Distributions

Table 5: Mean, Standard Deviation and Skewness of Growth and Temperature Distributions

	Rich	Poor
Temperature (°C)		
Mean	16.712	22.568
Standard Deviation	4.810	3.052
Skewness	-0.074	-0.208
GDP Growth (%)		
Mean	1.494	1.134
Standard Deviation	1.527	2.197
Skewness	-1.485	.345

Notes: In the Temperature category, the summary statistics are calculated by taking the mean, standard deviation, and skewness of within-year temperature (i.e. using monthly temperature observations), and then computing the average of these within-year moments. However, in the GDP Growth category, the summary statistics are calculated using annual real GDP growth observations.

9.2 Derivation of the Growth Equation

Consider a simple model economy with a production function and productivity growth rate as done in Bond, Leblebicioglu and Schiantarelli (2010).

$$Y_{i,t} = e^{\beta T_{i,t}} A_{i,t} L_{i,t} \tag{8}$$

$$\frac{\Delta A_{i,t}}{A_{i,t}} = g_i + \gamma T_{i,t} \tag{9}$$

where i, t denotes country i at time t, T is temperature, A is labor productivity and L is the total population.

Equation (1) models the level effects of temperature on output, perhaps due to more preferable or detrimental temperatures in a certain year affecting crop yields or tourism in that year. However, equation (2) models the effect of temperature on the growth of productivity. This, unlike the level effect, has long term consequences for the growth rate of the economy, and may be caused by changes in methods of production, policy and institutions.

In per capita terms, equation (1) is

$$y_{i,t} = (e^{\beta T_{i,t}} A_{i,t} L_{i,t})/L_{i,t} = e^{\beta T_{i,t}} A_{i,t}$$

Taking the natural logarithm of both sides and first differencing, this becomes:

$$ln(y_{i,t}) - ln(y_{i,t-1}) = (\beta T_{i,t} - \beta T_{i,t-1}) + [ln(A_{i,t}) - ln(A_{i,t-1})]$$

where the left-hand side is the growth rate of per capita GDP, and $ln(A_{i,t}) - ln(A_{i,t-1}) = \frac{\Delta A_{i,t}}{A_{i,t}}$. Therefore, by substituting in equation (2), my growth equation is:

$$g_{i,t} = g_i + (\beta + \gamma)T_{i,t} - \beta T_{i,t-1}$$

This equation models the growth rate of output as a function of current and lagged temperature. Dell, Jones and Olken (2012) use this growth equation to model their regression specification, given by:

$$g_{i,t} = \theta_i + \theta_{r,t} + \sum_{j=0}^{L} \rho_j T_{i,t-j} + \epsilon_{i,t} \tag{10}$$

where θ_i are country fixed effects, $\theta_{r,t}$ are time fixed effects interacted with region and poor country dummies, $T_{i,t}$ is the mean temperature in year t with up to L lags included, and $\epsilon_{i,t}$ is an error term clustered by country and year.

However, I consider the possibility that temperature variance and skewness matter for economic production. Therefore, I adapt Bond, et al.'s model economy to consider the effect of the higher central moments of the temperature distribution on growth rates. Let $\tilde{\mu}^j_{i,t} = E[(T_{i,t,m} - \bar{T}_{i,t})^j]$ be the jth sample central moment of the temperature distribution in country i at year t and month m, and $\bar{T}_{i,t} = \frac{1}{12}\sum_{m=1}^{12} T_{i,t,m}$ is the average temperature in country i at year t. Consider the following model setup:

$$Y_{i,t} = A_{i,t} L_{i,t} \tag{11}$$

$$\frac{\Delta A_{i,t}}{A_{i,t}} = g_i + \bar{T}_{i,t} + \tilde{\mu}^2_{i,t} + \tilde{\mu}^3_{i,t} \tag{12}$$

where $\tilde{\mu}^2_{i,t}$ and $\tilde{\mu}^3_{i,t}$ are the sample variance and skewness, respectively.

$$y_{i,t} = A_{i,t} L_{i,t}/L_{i,t} = A_{i,t}$$

$$ln(y_{i,t}) - ln(y_{i,t-1}) = ln(A_{i,t}) - ln(A_{i,t-1})$$

Substituting equation (5) into the growth rate of productivity, we have:

$$g_{i,t} = g_i + \bar{T}_{i,t} + \tilde{\mu}^2_{i,t} + \tilde{\mu}^3_{i,t} \tag{13}$$

9.3 Panel Regression Results using Mean Growth Rates as Dependent Variable

Table 6: Effect of Change in Moments of Temperature Distribution on Mean GDP Growth Rate

	(1)	(2)	(3)	(4)
Dependent variable is $g_{i,t}$				
Temperature Mean	0.259	0.321	0.260	0.320
	(0.270)	(0.289)	(0.272)	(0.290)
Temperature Mean x Poor	-1.074***	-1.088***	-1.075***	-1.087***
	(0.391)	(0.396)	(0.392)	(0.396)
Temperature Standard Deviation		0.276		0.279
		(0.313)		(0.319)
Temperature Standard Deviation x Poor		0.016		0.014
		(0.600)		(0.605)
Temperature Skewness			-0.020	0.043
			(0.648)	(0.660)
Temperature Skewness x Poor			0.022	-0.030
			(0.814)	(0.827)
Temp mean effect on poor countries	-0.815***	-0.766***	-0.814***	-0.767***
	(0.310)	(0.308)	(0.311)	(0.308)
Temp S.D. effect on poor countries		0.293		0.294
		(0.534)		(0.537)
Temp skewness effect on poor countries			0.002	0.013
			(0.501)	(0.504)

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.

Table 7: Effect of Change in Moments of Temperature Distribution on Mean Agricultural Growth Rate

Dependent variable is $agr_{i,t}$	(1)	(2)	(3)	(4)
Temperature Mean	-0.598	-1.128***	-0.714*	-1.199***
	(0.376)	(0.384)	(0.376)	(0.387)
Temperature Mean x Poor	-1.124**	-0.787	-1.040*	-0.732
	(0.540)	(0.532)	(0.542)	(0.536)
Temperature Standard Deviation		-2.225***		-2.139***
		(0.487)		(0.497)
Temperature Standard Deviation x Poor		2.576**		2.496**
		(1.206)		(1.224)
Temperature Skewness			1.880*	1.464
			(1.034)	(1.022)
Temperature Skewness x Poor			-2.202	-1.765
			(1.415)	(1.413)
Temp mean effect on poor countries	-1.722***	-1.915***	-1.755***	-1.931***
	(0.494)	(0.492)	(0.497)	(0.493)
Temp S.D. effect on poor countries		0.352		0.356
		(1.101)		(1.118)
Temp skewness effect on poor countries			-0.321	-0.301
			(0.977)	(0.996)

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.

Table 8: Effect of Change in Moments of Temperature Distribution on Mean Industrial Growth Rate

Dependent variable is $ind_{i,t}$	(1)	(2)	(3)	(4)
Temperature Mean	0.523	0.534	0.568	0.571
	(0.410)	(0.421)	(0.414)	(0.420)
Temperature Mean x Poor	-2.011***	-1.977***	-2.060***	-2.021***
	(0.626)	(0.642)	(0.624)	(0.639)
Temperature Standard Deviation		0.016		-0.020
		(0.384)		(0.404)
Temperature Standard Deviation x Poor		0.694		0.759
		(1.042)		(1.057)
Temperature Skewness			-0.611	-0.606
			(1.206)	(1.234)
Temperature Skewness x Poor			1.439	1.467
			(1.472)	(1.497)
Temp mean effect on poor countries	-1.488***	-1.443**	-1.492***	-1.449**
	(0.555)	(0.583)	(0.558)	(0.583)
Temp S.D. effect on poor countries		0.710		0.740
		(0.977)		(0.986)
Temp skewness effect on poor countries			0.827	0.860
			(0.865)	(0.872)

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Model uses country and year FE, and clusters errors by country. Poor is defined as a dummy with a value 1 if a country is below the median GDP PPP per capita in their first year in the sample.