ESCUELA POLITÉCNICA NACIONAL

FACULTAD DE CIENCIAS

ON THE ASSOCIATION BETWEEN WATER AND ECONOMIC GROWTH: EVIDENCE FROM LATIN AMERICA

TRABAJO DE TITULACIÓN PREVIO A LA OBTENCIÓN DEL TÍTULO DE INGE-NIERO EN CIENCIAS ECONÓMICAS Y FINANCIERAS

PROYECTO DE INVESTIGACIÓN

DIEGO ADRIÁN YEROVI OÑA

diego.yerovi@epn.edu.ec

DIRECTOR: JUAN PABLO DÍAZ, PhD.

juan.diaz@epn.edu.ec

Quito, abril del 2022

Yo, Diego Adrián Yerovi Oña, declaro bajo juramento que el trabajo aquí descrito es de mi autoría; que no ha sido previamente presentado para ningún grado o calificación profesional; y que he consultado las referencias bibliográficas que se incluyen en este documento.

La Escuela Politécnica Nacional puede hacer uso de los derechos correspondientes a este trabajo, según lo establecido por la Ley de Propiedad Intelectual, por su Reglamento y por la normatividad institucional vigente.

Diego Adrián Yerovi Oña

CERTIFICACIÓN

Certifico que el presente trabajo fue desarrollado por el Sr. Diego Adrián Yerovi Oña, bajo mi supervisión.

_____

Juan Pablo Díaz, PhD

DIRECTOR

AGRADECIMIENTO

Contents

List of Tables

List of Figures

Abstract

A wide and extensive literature has documented a negative association between rising global temperatures and economic growth. Much of these studies see temperature as a weather element that is both fixed and exogenous to economic outcomes in the short run. Under a similar point of view precipitation has also been analyzed, being an important element of weather too, but results from this research field are ambiguous. This study provides an investigation of precipitation by taking a different and novel approach, using a panel of 20 Latin American cuntries. Our empirical strategy relies on exogenous changes of rainfall seasonality, which is how much rainfall is temporally concentrated over the year within a particular location. This allows us to associate changes of precipitation with economic outcomes not only on the amount of rain for a given year, but on how evenly distributed it is across the year. To account its spatial distribution, we show that changes in relative rainfall seasonality are positively related with economic growth using disaggregated measures of production. On average, we find that when locations have changes in rainfall seasonality they experience as much as a 14% increase in economic growth as opposed to locations that do not. We also find that variability in the average amount of rainfall is consistently associated with a lower rate of economic growth.

1

1 Introduction

Recent literature has documented how fixed national characteristics may explain cross-country heterogeneity in economic outcomes. It hypothesizes that unchangeable elements in the environment of an economy (such as geography, climate, and colonial history) can explain why some countries grow faster than others.[1] With an empirical motivation that is different than that of traditional approaches, this new wave of studies found a way to be suggestive beyond simple correlations and cross-sectional estimates by exploiting the randomness in these fixed determinants of growth. As such, there is now more accurate evidence regarding the mechanisms behind development, and few experts would deny that geography or history are unrelated to the growth process.[2]

In a way, the branch of the literature that has focused on climate and weather as geographical elements has found temperature to be negatively associated with economic outcomes on several dimensions. With "climate" being the term reserved for the long run distribution of outcomes and "weather" used for the short run, on both of these dimensions temperature has been hypothesized to have effects through various mechanisms. For specific sectors, like agriculture, and through a more abstract relationship with individual productivity and labor, most of the evidence points to rising global temperatures as being detrimental to the economy.

---

[1]Some of the most notable contemporary analyses on development that hypothesize persistent effects of colonial history through institutions are the papers by Acemoglu et al. (2001), Glaeser & Shleifer (2002), Rodrik et al. (2004), Nunn (2008), Dell (2010), Michalopoulos & Papaioannou (2011) and Dell & Olken (2020). The relation between geography and growth has been analyzed by Gallup et al. (1998), Sachs (2001), Nordhaus (2006b), with climate being the center of examination in Deschênes & Greenstone (2007), Hsiang (2010), Jones & Olken (2010), Dell et al. (2012) and Burke et al. (2015). Much of these are discussed here. The incidence of disease due to geography is investigated in Gallup & Sachs (2001), Sachs & Malaney (2002) and Deschênes & Moretti (2009).

[2]Interestingly, a conclusion in Acemoglu et al. (2001) was that the relationship between geography and development is a spurious one, implying that regressions of output on climatic variables may be in fact capturing the effect of institutions. This discrepancy is illustrated in the debate between Rodrik et al. (2004) and Sachs (2003). Later on, evidence by Dell et al. (2012) and Burke et al. (2015) would strengthen the position of climate in the geographic side.

However, temperature conforms one part of the story regarding climate and weather. Precipitation, which is strongly correlated with temperature, has not been as scrutinized—often neglected to analyses that do not go beyond the obvious connections it has with agriculture. And while precipitation has served at least as a control variable in the climate literature, the relationship that it has with growth has not been successfully deciphered. Moreover, studies that include this topic have mixed results, and the ones that do find a relationship tend to differ in their implications.

Of the studies that find an association, some have suggested that precipitation plays an important role through a similar mechanism with productivity, similar to that of temperature, such as the individual leisure-labor decisions (Connolly, 2008; Damania et al., 2020). It has also been found that specific regions with a big agricultural composition such as the African continent are plausibly the ones most prone to suffering from rainfall variability, and not mid- to high-income regions like Latin America (Barrios et al., 2010). Others have also suggested that rainfall variation, if associated to an extreme event, is detrimental on an aggregate perspective (Deschênes & Moretti, 2009; Desbureaux & Damania, 2018). In this light, the present study investigates the combination of these seemingly unrelated elements: year-to-year fluctuations in precipitation and their effect on economic growth in Latin America.

In particular, our analysis diverges from previous ones in that we recognize that rainfall—and water in general—is economically characterized by its spatial and temporal distribution. In other words, the marginal effects of rainfall in an economy are heavily dependent on the location, the time the variations happen, and numerous unobserved characteristics that make difficult the disentanglement of causal effects. And while it is possible to model rainfall and water availability in different ways (e.g., General Equilibrium Models) to account for

spatial and temporal correlations, we opt for a more parsimonious approach, in line with the quasi-experimental estimation strategies in the studies mentioned above.

To account for the spatial factor statistically, we leave the macroeconomic realm and delve into the grayer area that lies between micro and macro, making use of *dissagregation*. This allows us to not think about how individual outcomes aggregate while having a fine level of detail regarding spatial activity. The concept, which can be traced to authors such as Nordhaus (2006b), captures the notion that while production and growth are often presented and most easily computed at the country level, there is heterogeneity across locations that is both useful and unobserved. Dissagregation involves computing a location-specific number that represents the level of production for a specific location in a given country. The earliest use of it, as originally proposed by Nordhaus, involved the use of population density data to obtain a geographical proxy of economic activity (e.g., we would expect that economic activity is concentrated the most in places where there are more people.) We use both population-dissagregated data (constructed "from scratch") and the dataset from Kummu et al. (2018), which contains data from more countries and is constructed using more sophisticated techniques.[3]

To account for the temporal distribution, we first recognize that year-to-year observations are not appropriate because rainfall effects are more likely felt in shorter intervals of time. And despite the existence of high-frequency and location-specific weather data, GDP and growth are constrained to a yearly periodicity. So, to harmonize the data coming from separate sources we borrow the concept of *relative rainfall seasonality* from the climatol-

---

[3]The dataset from Kummu et al. (2018) provides disaggregated measures of GDP and the HDI (Human Development Index). The authors use subnational GDP data where available, interpolating and extrapolating the GDP data with various methods in other cases ("filling the gaps"). They also provide estimates of the errors derived this process. See their paper for a better description.

ogy literature, which is simply an estimate of the dispersion of rainfall across the year—i.e., how much of the total yearly rainfall is accumulated across months. In practice this number, called the Seasonality Index (SI), varies between 0 and 1 and is an interpretative measure of temporal dispersion because it has clearly specified thresholds defining rainfall regimes. Like studies using average precipitation rates as the independent variable of interest, our identification strategy relies on changes in seasonality being exogenous.

Regarding identification, the estimation developed here also differs in that dissagregation allows us to define the counterfactual in terms of different locations (as opposed to the same locations at different points in time, which is a less convincing assumption that cannot be avoided when using country-level observations.) In the "potential outcomes" approach to causality in observational data (Rubin, 1978, 2005; Holland, 1986; Angrist & Imbens, 1991, 1994), the counterfactual is defined in terms of groups where observations are classified according to their treatment status. Hence, our estimates are derived from an approach that is designed to more closely resemble an experiment. At each particular year and country, locations are classified in seasonality categories and "treatment" is defined as a change in seasonality with respect to the previous period. At each baseline year, say t, we show that these locations are similar on average over some observable characteristics, so that the difference in outcomes across groups at $t + 1$ can be thought of as the effect of a seasonality change. We further discuss why this is also likely to hold for the unobserved characteristics, making the case of internal validity in a causal sense.

Our results are consistent with some of the previous findings supporting the existence of negative effects of average precipitation on growth. Our results also suggest that the biggest share of the burden posed by rainfall is due to temporal variability. The magnitude of the rel-

5

ative seasonality estimate is greater than the effect of average precipitation, and the estimates are significant when the regimes are extreme: changes in seasonality affect more growth when most rainfall is accumulated over just a few months.

To be more specific, our estimates suggest that any change in seasonality (up or down a category, either to a more equable or more extreme regime) is associated to a 14% decrease in growth, derived from a 0.3 percentage point decline. A 1/2 standard deviation increase in the average precipitation rate (an increase of about 36 mm^3 over the year) is associated to a 16% decrease in growth. However, we extend a word of caution for our average precipitation estimates since our main identification strategy revolves around seasonality. As such, we are less confident about the causality of our precipitation estimates when compared to the seasonality ones, specially since not all our identification tests hold for the former. Additionally, temperature seems to be significant across all specifications, consistent with the extensive findings documented on the matter.

From a policy perspective, our results provide evidence that precipitation is a two-dimensional element of weather with observable effects in the short run, independently of the share of agricultural activity in the economy as a whole. Unlike temperature, precipitation is a more nuanced factor in the equation when considering short intervals of time, and the implications of a greater seasonal variability may be lower rates of growth—at least in Latin America.

Although our analysis is limited in that we cannot point to the specific mechanisms that make this happen (e.g., the productivity-leisure decisions or the impossibility of weather adaptation) we believe the conclusions provided are interesting in their own right. As precipitation is not only an element of weather, but part of a much more complex cycle that is ultimately connected to water availability, the advent of climate change could amplify the de-

scribed effect of rainfall on growth. Thus, future research will have to analyze in detail both the mechanisms of transmission and the burden of weather on water availability in a more general way.

The remainder of the paper proceeds as follows. Section 2 provides a brief review of the literature on weather and economic outcomes, discussing previous findings and current limitations. Section 3 provides a theoretical framework that captures the notion of rainfall having a temporal and spatial dimension in a dissagregated setting. Section 4 describes the data and the summary statistics. Section 5 develops the main strategy and tests for identification. Section 6 provides the main results, and Section 7 concludes.

2 Literature review

In this section we briefly discuss some of the most important findings of the literature associating climatic and weather elements with economic outcomes. We start by analyzing the attempts at measuring the effects of climate in the economy, which started in a microeconomic setting and later progressed onto an active area of research in the macroeconomic field. The limitations posed by both the traditional and contemporary approaches are outlined, highlighting a few of the discrepancies found in the most relevant investigations of the climate-economy literature. With this, we also analyze how the use of dissagregated measures of a whole, complete, economy can be helpful in deciphering the association between precipitation and economic outcomes. This section finishes by providing an overview of some of the literature that consider rainfall as a form of water availabilty, remarking the importance of the literature relating climate and weather with the aggregated economy.

2.1 Measuring the relationship between climate and the economy

The traditional approach for measuring the effect of climate relied on the specification of a firm-level production function; that is, in a purely microeconomic setting. With cross-sectional techniques that considered a wide variety of regressors and functional forms, it was consistently seen that changes in precipitation, temperature, CO_2 levels and in the costs of irrigation are negatively associated with agricultural output [see, for example, Adams (1989) and Adams et al. (1990)]. However, estimates from this approach are potentially biased (upwards) as they do not allow for changes in the conditions of the agricultural environment, like technological change or productivity adaptations made by the farmer, that could mitigate climate effects. For instance, farmers could offset the losses of seasons with low precipitation rates by setting irrigation systems on their crops.

To circumvent the technological problem, seminal work by Mendelsohn et al. (1993) develops a methodological contribution regarding measurement, still in the microeconomic realm. Their main insight is that the *rent* of land is more likely to take into account the real effect on climate variability than output or crop yield, because land prices respond to both technological change and climate variation (under well-functioning markets). With rent having an operational role in their framework, their suitably-named "Ricardian approach" in principle corrects the bias. As Mendelsohn, Nordhaus and Shaw's main interest was measuring the effect of temperature, precipitation was interpreted as a control variable—its inclusion is nevertheless appropriate due to both temperature and rainfall being highly correlated variables.[4]

Using cross-sectional data from the United States, Mendelsohn et al. (1994) would use

---

[4]Meaning that if either temperature or precipitation is not included, the effect of one cannot be separately identified from the effect of the other in a regression framework.

the Ricardian approach to show thar higher temperatures seem to reduce agricultural revenue and farm values, while higher rates of precipitation increase them. The impact of climate is, as expected, lower in magnitude than the one documented in the literature using agricultural yield as the dependent variable. Interestingly, in an extension of this work, Mendelsohn & Dinar (2003) would further include irrigation and find that the value of irrigated land is not sensitive to precipitation and that land value increases with temperature; a complete change in direction for the results in Mendelsohn, Nordhaus and Shaw's 1993 and 1994 papers, which served as starting evidence that irriguation plays an important role offsetting climatic adversity in agriculture. As such, the most important contribution of the Ricardian approach to the field of climate economics is that systematic increases in temperature may play an observable role in agriculture and that investing in irrigation may function as insurance, suggesting a role for policy-making.

But, despite the fact that the Ricardian approach was more convincing than predecessor frameworks, it still suffers from a fundamental selection problem. The cross-sectional comparison of economic outcomes using climate is limited in its interpretation, in the sense that climate is a geographic element that does not vary much within provinces, states and even countries. As Dell et al. (2014) point out:

> A basic challenge in deciphering the relationship between climatic variables and economic activity is that the spatial variation in climate is largely fixed. Canada is colder on average than Cameroon, and it always has been. As such, while there can be large cross-sectional correlations between a country's climate and its economic outcomes, it is difficult to distinguish the effects of the current climate from the many other characteristics potentially correlated with it.

Only recently, the literature found a way to cope with this problem by using high-frequency *variations* in climate within locations using panel data. The basic idea is that by isolating variation in the climatic variables from elements that remain fixed across time (regions, countries, states, provinces, etc.), the estimate associated to the climatic variable in a regression framework will capture a more precise, clearer effect. At the risk of oversimplification, one could think of this approach the same way as a differences-in-differences estimation strategy, where variations in climate are correlated to variations in economic outcomes (instead of levels) controlling for individual characteristics.[5]

In practice, this surmounts to specifying a regression with fixed effects. The added benefit of this approach is that variations in weather can be considered to be exogenous to the economic system,[6] avoiding endogeneity, omitted variable bias and thus allowing the estimates to be causative.[7]

In this framework, studies using year-to-year fluctuations and intra-seasonal variations in weather [Deschênes & Greenstone (2007) and Fishman (2016) respectively] have found agriculture to be sensitive to both temperature and precipitation. Others have specifically examined the effects on rain-fed crops with similar results (Auffhammer et al., 2006; Sawano et al., 2008). Irrigation was also put into more scrutiny in an expanding literature that exploited a variety of estimation approaches, considering it as a mitigation against rainfall vulnerability in the form of water accumulation (Duflo & Pande, 2007; Blanc & Strobl, 2013). In these

---

[5]This is not quite right though. It is helpful to think of it this way to motivate intuition, but the intra-variation approach differs from a DiD in that *all* observations receive treatment. This is closely related to the causal methodology which will be discussed later on.

[6]Note that weather, and geography in general, by itself is not exogenous to economic outcomes. What is exogenous (or can be assumed to be so) are the *fluctuations* of weather in the form of high-frequency observations.

[7]This is not causation in the Granger-Sims sense. The implicit definition of causality underlying observational studies of this type is the one known as the "Rubin Causal Model" or the "potential outcomes approach" to causality.

investigations, precipitation has been found to play a consistent and important role in agricultural indicators and irrigation has been reasserted to offset the negative shocks of climate changes.

Leaving the microeconomic realm, macroeconomic analysis was motivated by other papers such as Sachs (2001) and Sachs & Malaney (2002), interested in the incidence of disease due to geography. A key hypothesis in much of the work of Sachs and his collaborators is that fixed conditions, such as a hot weather, are favorable to the development of diseases in much of the developing world—especially rural Africa and parts of Asia.

The macroeconomic analysis of climate and productivity (i.e., not associating it with other correlated mechanisms like health) surged with in the seminal work of Dell et al. (2012) (henceforth "DJO") and Burke et al. (2015) ("BHM"). These papers would estimate the effects of weather on an aggregate perspective using intra-country variation. In particular, DJO hypothesizes that temperature is a fundamental determinant of growth through its linkage to productivity, and that it may as well be a fixed element with persistent effects, just as institutions and colonial history. Their idea is to go beyond the widely-acknowledged correlation between higher temperatures and income—i.e., a cross-section—by instead exploiting within-country variations. On the opposite of a cross-section, the higher the temperature the least growth a particular country seems to experience. BHM, in a sense, builds upon DJO's paper and argues that more flexible, non-linear, functional forms of temperature may fit the picture better. Their results are somewhat consistent with DJO's. The (common and) most discussed result of these two investigations is that climate, particularly temperature, does seem to have a profound effect in the functioning economy *on average*.

2.1.1   The new limitations in the macroeconomic setting

Despite their identification strategies being largely the same and on similar data, the fact that DJO and BHM part from a very different set of methodological motivations make their results and implications differ for practical considerations. On one hand, DJO's main finding is that higher temperatures have largely negative effects but only on poor countries. They further report that hotter years affect both agricultural and industrial output. On the other hand, BHM find that non-linear functional forms of productivity on temperature seem to be consistent on a global scale; meaning that when non-linearity in temperature is allowed in the estimation, higher temperatures are associated with a significant decline in output for all countries. Neither find any effect of precipitation.

If we were to take DJO's story seriously, then the burden of climate—and more specifically hotter climates on the poor countries—can be escaped through growth. Since the poor are the ones affected by higher temperatures, be it by the agricultural composition of their economy or the lack of adaptation, the solution is to motivate a path of growth. On the other hand BHM's story has a more pessimistic implication, where no matter the degree of development or size of the economy, climate is associated to a generalized decline in productivity that will worsen with the progress of climate change.

As such, the crux of the problem is answering whether climatic adversity is only relevant for the poor and agriculturally-driven countries, or if this is in fact a phenomenon increasingly affecting economic outcomes for all countries. Is temperature a problem because of income heterogeneity? Or is it because of its non-linearity? This is not straightforward, however, as the empirical setting to answer such a question is, again, constrained by the lack of variation. But this time, the variation problem comes from the income variable—there

are only a handful of rich countries that also happen have hot climates (such as Shanghai). Hence, we cannot statistically discern between nonlinear effects and income heterogeneity even when using fixed effects, for if both are included in the same regression the parameter of the non-linearity will absorb some of the variation of the income-identifying variable.[8]

Moreover, even though temperature has been found to have clashing implications in the seminal investigations of DJO and BHM, in either case the estimate for temperature was statistically significant. Precipitation, on the other hand, has been found to have mixed results at best in these and other similar investigations found in the literature. Namely, while both DJO and BHM find no significant effect of precipitation, under similar strategies and different data sets the papers by Brown et al. (2013), Sadoff et al. (2015) and Damania et al. (2017) find the contrary.

Plausible explanations for these disparities are that the impacts of rainfall are heterogeneous for its spatial and temporal variability (Damania, 2020), and that the significance level of the estimates tend to depend on the level of data disaggregation (Damania et al., 2020). In economic parlance, the marginal utility of rainfall may differ significantly between geographical areas, for it is not the same having an additional inch of rainfall in the summer than in the winter, and that additional inch of rainfall would not have the same utility on a desert and

---

[8]It may be illustrative to pose this explicitly. DJO uses a global panel and a regression of the form

$$Y_{it} = \beta_0 + \beta_1 T_{it} + \beta_2 P_{it} + \beta_3 \mathbb{1}_{\{i \text{ poor}\}} + \beta_4 (T_{it} \times \mathbb{1}_{\{i \text{ poor}\}}) + \delta_{rt} + \varepsilon_{it},$$

where Y is the economic outcome of interest (the log of GDP per capita), T and P are the weather variables (temperature and precipitation respectively), $\mathbb{1}_{\{i \text{ poor}\}} = 1$ if the country is poor and 0 otherwise, δ_{rt} are time-region fixed effects, t indexes time and i indexes countries. BHM consider a more general form for temperature $g(T_{it})$, such as a quadratic. If we were to include both nonlinear effects in temperature and income heterogeneity, as in

$$Y_{it} = \alpha_0 + \alpha_1 T_{it} + \alpha_2 T_{it}^2 + \alpha_3 P_{it} + \alpha_4 \mathbb{1}_{\{i \text{ poor}\}} + \alpha_5 (T_{it} \times \mathbb{1}_{\{i \text{ poor}\}}) + \delta_{rt} + \mu_{it},$$

we would want to test H_0: $\hat{\alpha}_2 = 0$ and/or H_0: $\hat{\alpha}_5 = 0$. But, it is likely that $\hat{\alpha}_2$ is capturing the effect of $\hat{\alpha}_5$ given the small number of observations that satisfy $\mathbb{1}_{\{i \text{ poor}\}} = 0$ and have a high T_{it}. In other words, there is not enough statistical power. This also implies that we cannot estimate (robustly) how income is affected in hot and cold countries separately because there is not enough income variability in the group of the hotter ones. We cannot run a regression $Y_{it} = \beta_0 + \beta_1 T_{it} + \beta_2 P_{it} + \delta_{rt} + \varepsilon_{it}$ separately for cold and hot countries.

in a tropical region. Plus, the fact that the relevant economic data comes spatially aggregated could easily veil any relevant effects deriving from weather shocks. Moreover, precipitation cannot always be convincingly assumed to be a random phenomenon even when controling for previous periods to account for weather station effects. This argument is based on the idea that periodicity of changes in the weather could be anticipated with enough precision by individuals and firms, that the outcome variable and rainfall are only weakly correlated (Auffhammer et al., 2006).

Before discussing in more detail these econometric problems, it is worth reviewing evidence that does not use precipitation in the interest of weather and climate effects.

2.2 Precipitation as a measure of water variability

In a more general sense, rainfall has played an important role in the empirical literature through mechanisms other than climate. For instance, the notion that precipitation functions as an exogenous, transient phenomenon, has served to empirically corroborate the presence of consumption smoothing over temporary income shocks in agricultural families (Paxson, 1992). To the extent to which increments in income deriving from changes in rainfall are considered by farmers to be temporary, Paxson finds evidence that the marginal propensity to save is greater than on income increases that are regarded as permanent.

Similarly, the fact that precipitation can be thought of as an independent and identically distributed variable over time has motivated its use as an instrument in determining the effects of weather on conflict and political stability (Miguel et al., 2004; Hsiang et al., 2013). In Miguel et al., lower rainfall is found to lead onto more conflict, with negative shocks on economic growth increasing the likelihood of civil unrest.[9] Subsequent works using similar

---

[9]Something remarkable that could serve as a motivation for a more systematic analysis is that, for the

14

estimation strategies document negative economic consequences when precipitation declines (Hidalgo et al., 2010).[10]

Rainfall has also been included in studies about labor productivity with interesting results. A notable empirical paper on the matter was that by Connolly (2008). She examines the impact of rainfall on the individual trade-off between labor and leisure for the US finding that men, specifically, substitute on average thirty minutes a day away from leisure when it rains. This was an early attempt at using panel data to corroborate aggregate effects of a climatic variable with a causal framework. In fact, the notion that labor productivity is affected by weather outcomes greatly motivated the next major advancement with the use of panel data, allowing for the subsequent analysis on output level.

While these studies provide insights into the mechanisms over which precipitation affects individual decisions and economical environments, they do not provide clear evidence on how its variability translates on national economic measures. From a macroeconomic perspective they do not tell what are, if any, the *aggregate* effects of water variability. Specially because it may require a different set of analytical machinery, both from a theoretical and statistical point of view.

Probably the first to explicitly analyze water for its implications on per capita income in an aggregate perspective was Barbier (2004), through the concept of *water utilization*. Using the neoclassical endogenous growth model of Barro (1990) and Barro & Sala-I-Martin (1992), water is seen as a publicly provided commodity that is subject to congestion. Theoretically,

---

validity of their instrument, they rely on assuming precipitation does not affect conflict directly but through growth—meaning their results can only hold if rainfall has (*ex ante*) an effect on economic growth.

[10]An interesting aspect about the literature on civil conflicts, is that it is not free of ambiguity. It has been argued that these discrepancies can be explained by a number of identification omissions, with many of those studies not accounting for fixed effects, and relying on the inclusion of endogenous controls (Dell et al., 2014). Moreover, weather measures for rainfall can be correlated with the outcome of interest when serving as an instrument, with possible exceptions for these being regions with few weather stations (Auffhammer et al., 2011).

an inverted U relationship is formulated between growth and the rate of water use in an economy. Then, the cross-sectional estimates suggest that growth is not limited by utilization rates.[11] Despite of the geographical limitations in data and estimation, this first attempt by Barbier's 2004 paper addressed an empirical issue that would not have significant progress for quite some time.

The most common critique of Barbier's approach is that, by modeling water utilization rates, results are susceptible to endogeneity problems. On the opposite to precipitation, utilization rates are seldom exogenous in the economic system. Besides, as Damania (2020) notes, the interpretation of a utilization variable allows for odd conclusions: water shortages would lead to an increase in GDP under Barbier's original results. This means that the inclusion of water, in a complete sense of the word, is not statistically feasible to obtain robust and consistent estimators.

More recently, Russ (2020), while recognizing the heterogeneous effects attributed to space, argues that precipitation is a poor indicator of water availability because it does not account for all the factors that conform the total supply of water: for example, precipitation results in a significant amount of water that is absorbed in the soil and then goes thorough a natural cycle that ends up benefiting economic activity—specially agriculture. With this, it is proposed using water runoff as a measure of water availability, and satellite data on night-time lights is used as a proxy of economic activity to account for the spatial distribution.

An alternative method is modelling water through a computable general equlibrium model (CGE) and, more specifically, general algebraic modeling systems (GAMS). These allow for a very stylized specification of an economy that can be spatially disaggregated, incorporating

---

[11]However, the estimation results implied that this relationship is different for developing countries, something to be reviewed later in Barbier (2015) with the same approach. Instead, a U-shaped relationship was detected: for low- and middle-income economies, an increase in water use would first reduce economic growth.

specific economic mechanisms, like water sources, demand functions, supply functions, and water uses [see Hurd (2015)]. The biggest problem with CGEs is, again, fully specifying the many complex linkages of water in an economy, and rainfall is particularly complicated for its sequential nature—it is also the basis for an immense amount of externalities in the system—making linkages difficult to trickle down.

2.3 Disaggregation

The spatial distribution of rainfall makes it difficult to get unbiased estimators of the real effect on an economy. As production and growth are spatially correlated, we would expect that positive or negative effects in the short run are masked by aggregation. Moreover, considering individual variables (such as individual productivity, income or consumption) difficult the problem more because one has to think and clearly define how these variables aggregate—an impossible task considering the complex linkages between individual economic outcome. Some of the methodologies mentioned above, i.e., General Equilibrium Models (GEMs) or Integrated Assessment Models (IAMs), the number of parameters to estimate and equations to specify makes these inconvenient in a practical setting.

The idea of disaggregation, which is defined as a number capturing production in a sub-region of a country, may eliminate these concerns and provide a useful way to obtain unbiased estimators. The reasoning is that we can aggregate production within a sub-region, and use this as a measure of economic activity in a finer way than country-aggregated GDP. It was originally conceived, to our knowledge, by Nordhaus and his collaborators (Nordhaus, 2006a,b), who also proposed the simplest way to do it in a practical setting: using population data and computing disaggregated GDP in terms of the proportion of people that leave in a

particular sub-region. In other words, this number is simply the country's per capita GDP multiplied by the number of people living in the sub-region.

Other methods, some more sophisticated and more computationally demanding, have been proposed since trying to account for spatial correlation across sub-regions or to compute other measures that are only available on a country level [e.g., Kummu et al. (2018); Thomas et al. (2019)].

3 Theoretical Framework

Here we present a framework that illustrates the connection between disaggregated effects of weather, specifically precipitation, and the aggregated measures which are observed as data. Empirical investigations have identified a negative association between: (*i*) rain and the allocation of leisure chosen by workers (Connolly, 2008), (*ii*) temperature and the allocation of time and leisure (Graff Zivin & Neidell, 2014), and (*iii*) hourly/daily temperature and agricultural yields (Schlenker & Roberts, 2006, 2009); however none consider explicitly both precipitation and the transition of disaggregated output measures (wages, profits, crop yields, etc.) onto macroeconomic data (GDP).

Borrowing from Burke et al. (2015),[12] consider a country with a fixed number of industries indexed by $i \in I$ with homogeneous firms, and assume that all firms respond equally to changes in weather conditions in short intervals of time, it being indexed by $t \in \tau$ (t can be thought of as a "moment" in time happening on a continuum of points, and τ as being the larger observed unit of time such as a quarter or a year). With respect to the spatial dimension of production, let s index the unit of spatial allocation so that all $s \in S$ jointly conform

---

[12]Readers are referred to the supplementary paper of Burke et al. (2015), which briefly illustrates the first half of this model only considering temperature.

the space within the boundaries of the country. We say that some particular industry i is established in some, plausibly not all, locations s. When this happens we write $i \in s$.

For illustration purposes, we consider a standard Cobb-Douglas aggregate production function with constant returns to scale and parameter α, and abstract productivity as being a function of weather. Namely,

$$A_i^K(P_{st}, T_{st}) \cdot K_{ist} \quad \text{and} \quad A_i^L(P_{st}, T_{st}) \cdot L_{ist}$$

determine the productive units resulting from each input, capital and labor respectively, at location s and moment t. The processes behind productivity are said to differ only across industries, and productivity differs within industries only due to local and transitory weather experienced (hence the subscripts for P and T). The parameter L_{ist} is to be interpreted as the amount of labor employed by industry i at location s and moment t. Also,

$$\partial A_i^j / \partial P_{st} < 0 \quad \text{and} \quad \partial A_i^j / \partial T_{st} < 0 \quad \text{for } j \in \{K, L\} \text{ with } s \text{ and } t \text{ fixed}$$

consistent with (i) and (ii).

In reality, there is the possibility that the quantities of capital and labor themselves are a function of P_{st} and T_{st}. However, by considering short intervals of time $t \in \tau$ it is possible to assume, not unrealistically, that input reallocations cannot be immediately done by firms in response to changes in weather. Furthermore, we shall regard K_{ist} and L_{ist} as exogenous (scalar) parameters in the model determined in the competitive equilibrium, which satisfies $K_{ist}/L_{ist} = \alpha/(1 - \alpha)$.

Letting p_i denote the unitary price of output, the value of production of industry i in the

location s at time t is

$$Y_{ist} = p_i \left[A_i^K(P_{st}, T_{st}) K_{ist} \right]^\alpha \left[A_i^L(P_{st}, T_{st}) L_{ist} \right]^{1-\alpha}.$$

It is convenient to further substitute

$$w_i(P_{st}, T_{st}) = \left[A_i^K(P_{st}, T_{st}) \right]^\alpha \left[A_i^L(P_{st}, T_{st}) \right]^{1-\alpha} \quad \text{and} \quad U_{ist} = p_i K_{ist}^\alpha L_{ist}^{1-\alpha}$$

so as to write $Y_{ist} = w_i(P_{st}, T_{st}) U_{ist}$. In essence, $w_i(.)$ describes the immediate effect of local precipitation and temperature (weather) on productivity (heterogeneous across industries), and U_{ist} is the scalar measure of resources in industry i applied at location s. Evidently, if i is not located at, say $\hat{s}$, then $U_{i\hat{s}t} = 0$ so that $Y_{i\hat{s}t} = 0$.

Assuming that the economy is additively separable in the spatial and time dimensions, the value of aggregate production Y (GDP) for the larger time interval τ is obtained by summing the production across all locations and industries and integrating for all moments,

$$Y = \sum_{i \in I} \sum_{s \in S} \int_{t \in \tau} w_i(P_{st}, T_{st}) U_{ist} \, dt.$$

The key insight is that the number of productive units across al moments U_{is} can be expressed in terms of the joint distribution of the time spent at all possible temperatures and precipitation rates. Let $g_i(P_s, T_s \mid i \in s)$ express the amount of time each productive unit of i in s $\left(\text{i.e., } p_i K_{is}^\alpha L_{is}^{1-\alpha} \right)$ experiences some determined temperature and precipitation rate.

Explicitly, define g_i such that[13]

$$U_{is} = \int_{t \in \tau} U_{ist}\, dt = \int_0^\infty \int_{-\infty}^\infty g_i(P_s, T_s \mid i \in s)\, dT_s\, dP_s$$

$$= \int_0^\infty \int_{-\infty}^\infty g_i(P, T) \cdot \mathbb{1}\{i \in s\}\, dT dP.$$

This function g_i can be thought of as a density distribution. (But not a probability density because it integrates to U_{is} instead of 1 conditional on $i \in s$.) It follows that using g_i we can express production spatially in terms of temperature and precipitation:

$$Y_s = \sum_{i \in I, i \in s} \int_0^\infty \int_{-\infty}^\infty w_i(P, T)\, g_i(P, T)\, dT dP. \tag{1}$$

Intuitively, the production at s is the sum on productivity times output (in market prices) for all possible weather combinations (P, T). Hence, under the conditions so far stated one can write spatially-dissagregated production without knowing in detail the temporal distribution of production (how much is produced at each particular moment in time). This is different—and less stringent—from what is embedded in the frameworks and findings of Dell et al. (2012), Burke et al. (2015) and other papers that do not account for spatial dissagregation. In particular, if we further assume that there exists a function density $h_i(.)$ that describes the number of productive units over time for each possible values of (P, T), as these papers do, then *we do not need to know the spatial distribution of production neither* because under the same procedure as above it follows that total aggregate production can be expressed as

$$Y = \sum_{i \in I} \int_0^\infty \int_{-\infty}^\infty w_i(P, T)\, h_i(P, T)\, dT dP. \tag{2}$$

---

[13]Note how precipitation varies in the bounded-below interval $[0, \infty)$ as opposed to temperature.

It is plausible—and likely, as the evidence of those papers show—that (2) is not unrealistic if we are interested in temperature, which is more prone to have systematic changes across large regions that scale up to the country level. But, as seen in section 1, if GDP is modeled this way, then one is likely to find that precipitation does not have a discernible effect because it is required *ex ante* that it have large enough shocks or systematic changes across time periods τ for any effect to be traceable up to the country level.[14]

4   Data

Monthly temperature and precipitation data comes from Willmott & Matsuura (2001), a global database in $0.5° \times 0.5°$ cells of which the territorial administrative limits of the Latin American countries considered lie under 6795 cells. This database is updated yearly, and observations are available for the entire period corresponding to 1900–2017. As per the availability of the rest of the data, the time span used throughout the analysis is the period 2000–2015.

We consider two different datasets of dissagregated GDP. The first, and the one used to present the results, is the gridded GDP data of Kummu et al. (2018) at the 5 arcmin resolution. To merge it with the weather data, the observations from this dataset are centered and aggregated to the $0.5° \times 0.5°$ level. The second dataset, which is used to replicate the analysis in the Appendix, is instead constructed "from scratch" using the economic data from the World Bank's Global Economic Monitor Database (GEM) and dissagregating it with data on yearly population density—which corresponds to the original methodology proposed by Nordhaus

---

[14]On the other side, one could argue the opposite; that the level of dissagregation in (1) would not allow for the identification of temperature effects. But this is only plausible if capital and labor can be quickly reallocated over space due to sudden changes in temperature, which violates a core assumption of the model presented and generally may be deemed as unrealistic.

Table 1: Summary statistics

Country	Grid cells	Precipitation	Temperature	Disaggregated log(GDP) Kummu et al.	Pop. disaggregated	Pop. Density	HDI
Argentina	1207	52.945 (36.259)	13.964 (6.373)	17.831 (2.127)	17.080 (2.035)	44068.17	26.719 (6.940)
Belize	14	172.025 (60.524)	25.059 (1.511)	19.218 (0.985)		21838.59	24.255 (6.128)
Bolivia	434	92.789 (57.232)	21.233 (7.807)	16.881 (1.837)	15.883 (1.507)	26199.38	22.039 (5.669)
Brazil	2954	149.448 (57.402)	24.845 (2.652)	17.350 (3.464)	17.773 (2.476)	72658.14	23.108 (6.057)
Chile	374	66.412 (70.325)	7.566 (5.275)	16.404 (4.688)	16.946 (2.287)	55183.65	25.924 (7.224)
Colombia	429	209.125 (88.180)	25.043 (4.124)	18.337 (2.555)	17.883 (2.280)	120680.10	23.104 (5.963)
Costa Rica	20	251.056 (96.831)	23.992 (3.400)	20.572 (1.451)	19.743 (2.059)	220530.15	23.389 (6.398)
Dominican Republic	19	108.158 (39.430)	24.913 (2.865)	21.747 (0.993)		448246.40	20.428 (6.254)
Ecuador	101	161.915 (92.756)	22.714 (5.371)	19.236 (2.147)	18.356 (2.432)	155563.72	23.215 (6.378)
El Salvador	10	124.672 (30.402)	24.677 (2.456)	21.672 (0.935)	20.899 (0.981)	670674.58	19.437 (5.800)
Guatemala	51	180.425 (72.750)	24.092 (2.911)	20.205 (1.740)	19.105 (1.960)	318484.02	20.528 (5.374)
Honduras	50	125.802 (47.941)	25.510 (2.365)	19.674 (1.615)	18.467 (1.721)	173348.67	19.575 (5.294)
Mexico	740	60.821 (49.522)	20.776 (4.196)	19.695 (2.026)	18.930 (2.398)	176577.44	24.322 (6.676)
Nicaragua	55	164.356 (73.790)	26.788 (2.100)	18.937 (1.517)	17.910 (1.587)	115941.08	20.321 (5.222)
Panama	28	205.344 (52.998)	25.579 (2.561)	19.250 (1.667)		100709.38	23.070 (6.572)
Paraguay	177	95.130 (37.153)	24.013 (1.709)	17.006 (3.034)	16.506 (2.028)	43855.45	22.856 (5.867)
Peru	482	142.079 (93.426)	20.409 (8.730)	17.761 (2.038)	17.772 (2.008)	51950.85	23.664 (6.106)
Puerto Rico	3	187.068 (42.329)	23.718 (1.032)	24.258 (0.450)		1295622.70	24.481 (7.244)
Uruguay	83	109.208 (26.554)	17.979 (1.016)	18.955 (1.358)	18.079 (1.738)	56442.51	25.714 (6.603)
Venezuela	337	155.304 (68.710)	25.475 (3.349)	17.807 (3.337)		79518.82	23.765 (6.204)
All	6965	120.224 (77.099)	21.298 (6.938)	17.804 (3.068)	17.625 (2.389)	81950.49	23.894 (6.497)

Note: Each column presents the grid cell average and standard deviation (in parentheses, if applicable) for each country. Each grid cell is a ~55 km^2 "square" determined by a longitude-latitude pair. The unit of measurement for precipitation are mm^3 and temperature is measured in centigrade degrees. Precipitation and temperature data is from the most recent version of Willmott and Mutsuura (2001). The grid average of log(GDP) is presented for the two datasets mentioned: one is from Kummu et al. (2018); the other one is constructed using data from the World Bank's GEM and disaggregated with population density data from the World Pop project. Grid population density (rightmost column) is the average number of people per grid cell (the average number of people per approx. 55 km^2).

(2006b). The population data comes from World Pop's database (WorldPop, 2020) at a 0.1°
precision (around 1 km at the equator), which is subsequently aggregated to the 0.5° level to
be harmonized with the weather data. It is seen that results do not differ significantly across
datasets, implying that population-dissagregated GDP works well enough as an indicator of
spatial economic activity, consistent with the statements of Nordhaus and his collaborators.

Figure 1 is divided into 6 subfigures, each presenting the empirical spatial distribution
and dispersion of our 3 main variables of interest: precipitation, temperature and GDP. Each
colored gridcell on the map represents the mean value of the variable across all periods of
observation (2000–2015). Interestingly, as shown in subfigures (e) and (f), the spatial distri-
bution Kummu et al. GDP and the population dissagregated GDP are virtually identical (ex-
cept for the availability of some data; in particular, there is no WB GDP data for Venezuela,
Cuba, Belize, and certain subregions, which are painted white in subfigure (d)). Figure 1
also displays in an intuitive way how precipitation tends to be more variable along forestall
Amazonian regions, and along certain parts of the Andean Mountains. Temperature, on the
other hand, is less variable around the equator.

Table 1 presents summary statistics of the main variables in the dataset. Each column
presents the grid cell average and standard deviation (in parentheses, if applicable) for each
country. Each grid cell is a $\sim$55 km^2 "square" determined by a longitude-latitude pair. The
unit of measurement for precipitation are mm^3 and temperature is measured in centigrade
degrees. The grid average of log(GDP) is presented for the two datasets mentioned above.
GDP data for some of the countries is not available from the World Bank's GEM, so not all
are presented under the population-disaggregated column. Grid population density (rightmost
column) is the average number of people per grid cell. Figure 1 shows instead, graphically,

the measures of the main variables in a spatially-disaggregated way. It allows us to see spatial concentration in precipitation, temperature, GDP and their variability.

Figure 2 is composed of scatterplots, where each point represents a pair of variable observations for each grid cell in the dataset. It displays graphically that there is mostly a negative relationship between mean precipitation and GDP, HDI and population density. In particular, GDP seems to be increasing in precipitation at low levels (the blue line is a non-parametric fit) and increasing at high levels. This non-linearity, however, must be seen with caution because of external validity concerns: the better the fit in the curve for Latin American countries, the more bias it is introduced if we want to extrapolate this to a global scale. Moreover, it can be seen that the cloud of points for GDP (data from Kummu et al.) and population density with precipitation is virtually identical, just as their linear and non-linear fits, supporting our notion that it does not make a difference how we disaggregate economic activity. Population-disaggregated GDP works as well for our purposes.

Figure 1: Spatial distribution of variables

(a) Mean precipitation (mm^3)

(b) Mean temperature (°C)

(c) Standard deviation of precipitation (mm^3)

(d) Standard deviation of temperature (°C)

(e) Mean grid cell GDP (data from Kummu et al.)

(f) Mean grid cell GDP (disaggregated with pop. data)
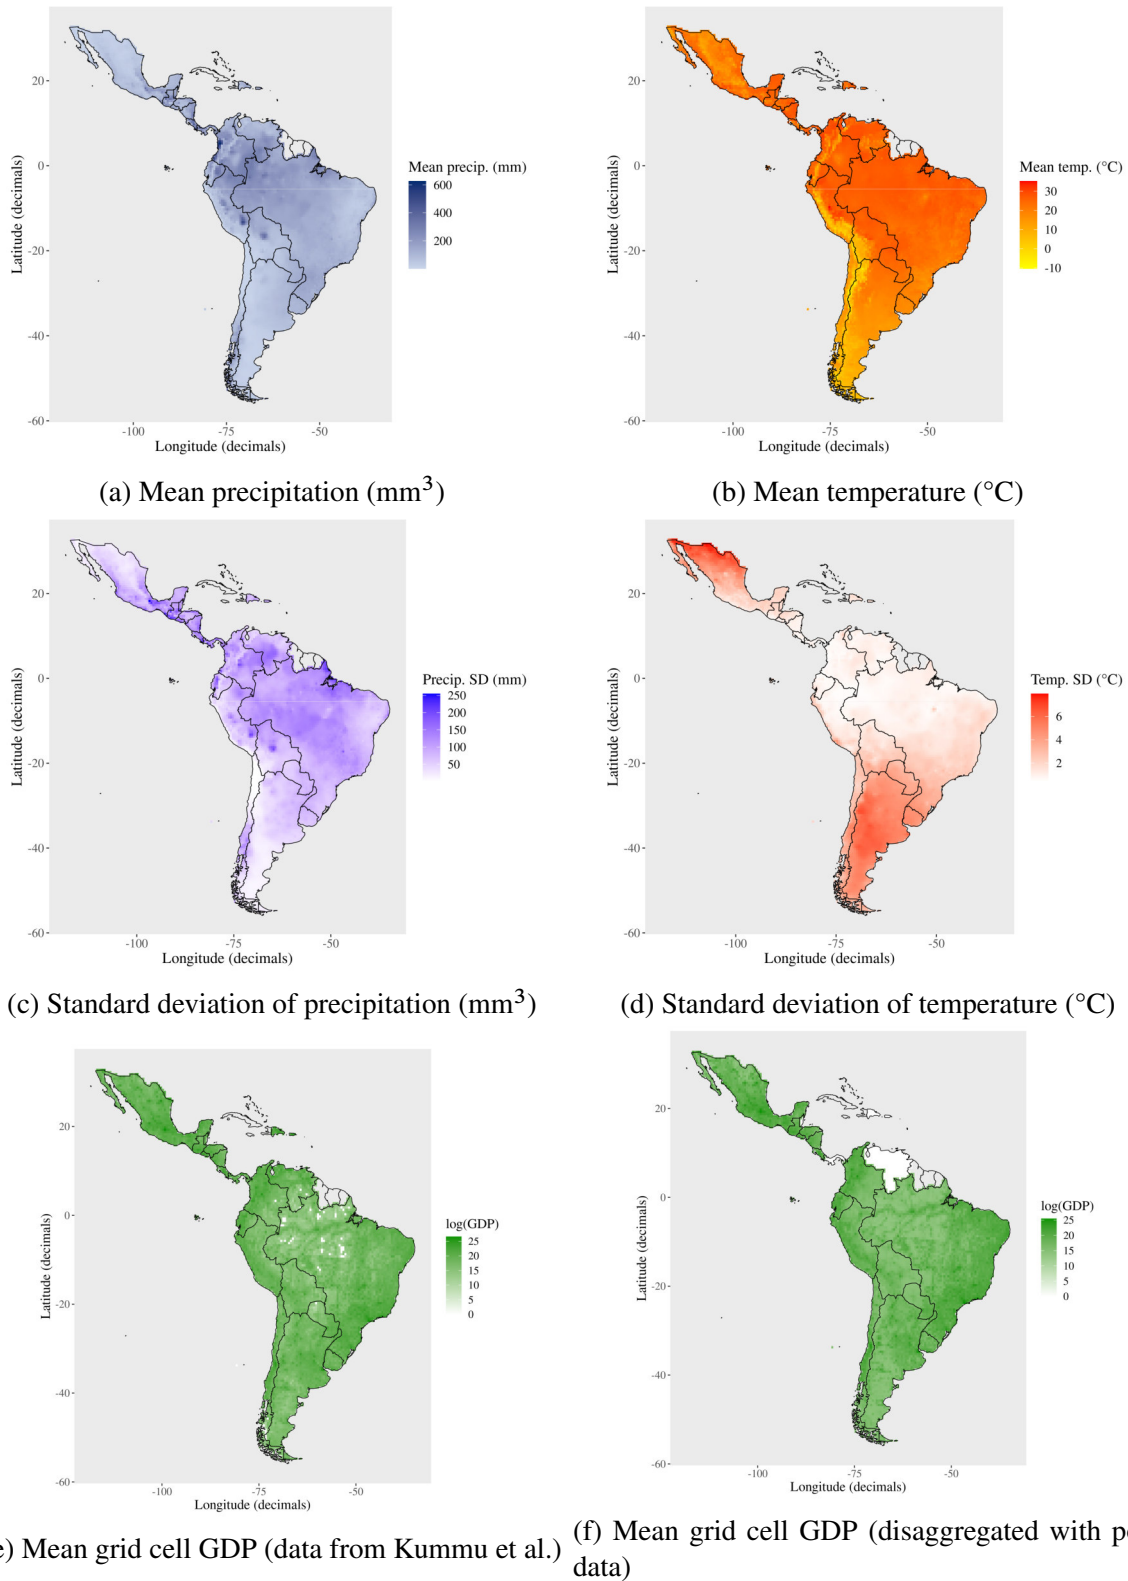
Note: This figure displays the mean and standard deviation per spatial unit of observation on the climatic variables and log(GDP) across the entire period (2000–2015).

26

Figure 2: Scatterplots at the grid cell level
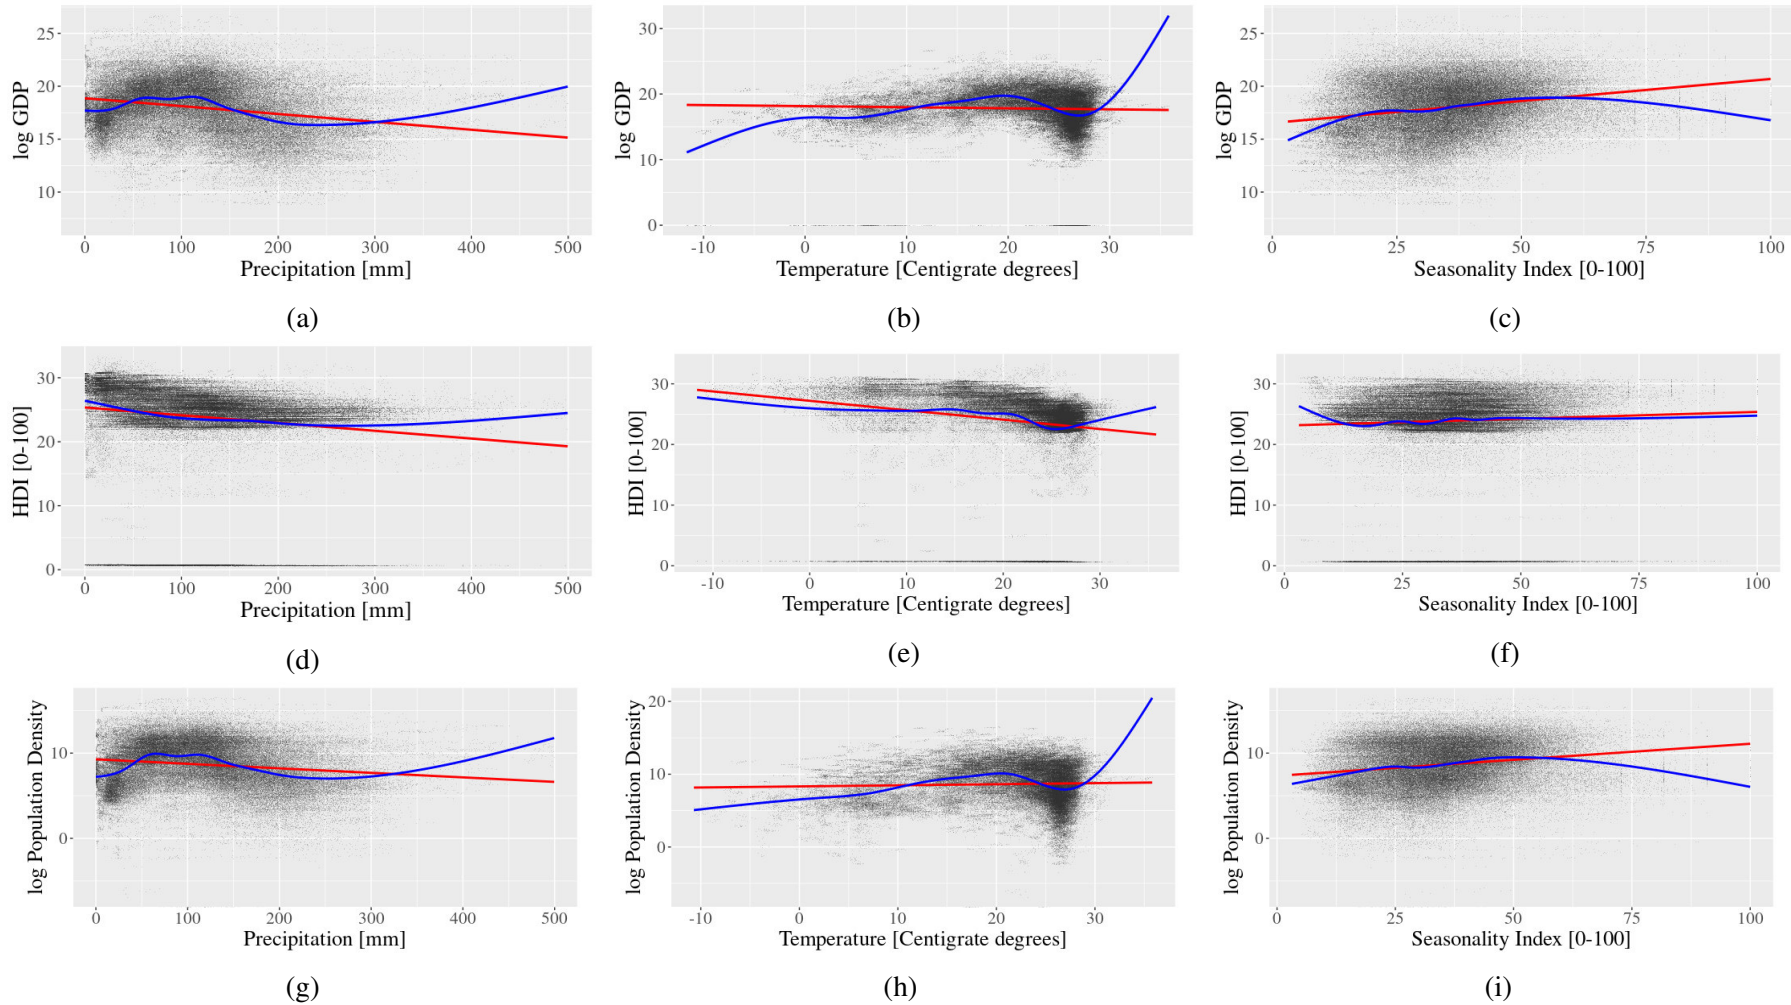
Note: Each point is represents an observation pair of variables associated to a grid cell in the dataset. The yearly mean precipitation, temperature and Seasonality Index are used when indicated. The red line in each plot represents the best fit on a linear regression. The blue line is the best non-parametric fit with Epaechnikov curve at each bin.

5 Estimation

The first part of the approach taken for the estimation of weather effects is akin to the recent literature exploiting intra-country variaton in the weather variables. Additionally, we propose a quasi-expermiental methodology to identify the effects of precipitation on a dissagregated economy. This latter approach is based on the definition of causality for observational data provided by what is known as the "Rubin Causal Model" or the "potential outcomes" framework (Rubin, 1978, 2005, 2019; Holland, 1986).[15] This section specifies the empirical equations to be estimated and test some implications associated with the variables in the model. The main results are presented and discussed in section 6.

Using the same notation as above, we start by estimating the following regression

$$\log(Y_{s\tau}) = \beta_0 + \beta_1 P_{s\tau} + \beta_2 T_{s\tau} + \mathbf{X}'_{s\tau}\boldsymbol{\beta} + \delta_{\mathscr{C}\tau} + \varepsilon_{s\tau} \tag{3}$$

where s indexes location (the center of a longitude-latitude grid cell), $\mathscr{C}$ indexes countries, τ indexes time (a year), $Y_{s\tau}$ is the dissagregated observed measure of GDP, $P_{s\tau}$ and $T_{s\tau}$ are the *average* precipitation and temperature experienced in s at time τ, $\delta_{\mathscr{C}\tau}$ are country-year fixed effects,[16] and $\mathbf{X}_{s\tau}$ is a column vector of control variables. One could also estimate these equations adding the variable $T_{s\tau}^2$ to account for the possible non-linear effect of temperature mentioned before, to reduce the variance of $\hat{\beta}_1$. The reasoning behind the inclusion of fixed effects at the year-country level is similar in that we seek to reduce the variance of the estimator. Specifically, it accounts for the fact that GDP is a variable whose value depends on the country of belonging—a correlation within groups—for grid cells at specific points in time.

---

[15]Readers are also referred to Chapter 1 of Imbens & Rubin (2015).

[16]I.e., $\delta_{\mathscr{C}\tau} = \mathbb{1}\{s \in \mathscr{C}\} \times \mathbb{1}\{\tau = j\}$. Since the time span of the data set is 2000-2017, as mentioned in the introduction, we have $j = 2000, \ldots, 2017$ or $j = 2000Q1, 2000Q2, \ldots, 2017Q4$.

Moreover, the dummy variables would capture the effect of a "good" or a "bad" year in GDP which is attributable to individual characteristics, isolating the effect of the weather variables. It is possible to further account for group heterogeneity by computing robust standard errors, or more conservatively, clustering the standard errors at the country level.

The thing about equation (3) is that $\mathbf{X}_{s\tau}$ may not perfectly control for individual characteristics at the location or grid cell level that also determine $Y_{s\tau}$. In such case this is an evident omitted variable bias problem that can be corrected by using a specification that resembles that of the Differences-in-Differences approach in an experimental setting. If we substract the first difference from the weather and outcome variables, the estimator associated to these variables will capture the intended effect accounting for all possible individual characteristics, because we are then regressing the change in GDP—*growth*—on the changes in weather. That is,

$$\log(Y_{s\tau}) - \log(Y_{s,\tau-1}) = \alpha_0 + \alpha_1 \gamma_{s\tau} + \alpha_2(P_{s\tau} - P_{s,\tau-1})$$

$$+ \alpha_3(T_{s\tau} - T_{s,\tau-1}) + \delta_{\mathscr{C}\tau} + \varepsilon_{s\tau}$$

$$\Delta \log(Y_{s\tau}) = \alpha_0 + \alpha_1 \Delta P_{s\tau} + \alpha_2 \Delta T_{s\tau} + \delta_{\mathscr{C}\tau} + \varepsilon_{s\tau} \tag{4}$$

where Δ represents the first difference. Note that (4) is also equivalent to estimating (3) with location and country-year FEs additively,

$$\log(Y_{s\tau}) = \omega_0 + \omega_1 P_{s\tau} + \omega_2 T_{s\tau} + \delta_s + \delta_{\mathscr{C}\tau} + \varepsilon_{s\tau}$$

where $\omega_1 = \alpha_1$. All in all, the estimates of equations (3) and (4) are intended to show the main hypothesis that precipitation has a significant effect when using spatially-dissagregated

measures of GDP.

Table 2 displays the estimates for equations (4) and (5). After controlling for altitude, latitude and ppopulation density, it can be observed that the estimates under robust standard errors are significant (in practice, they are not so different from classical standard errors; not shown in the table). To be conservative, we also present these results using clustered standard errors at the country-year level. While this may does not necessarily correspond to the assignment mechanism of variations in precipitation, it can be seen that the significance does not vanish. The estimates seem relatively small, consistent with previous literature, and the difference-in-difference estimator (column [3]) is half the magnitude of the ones in cols. [1] and [2]. However, there are several reasons why the estimates of these two regressions may not be interpreted as causal, even when clustering the standard errors.

First, there is the possibility that shocks are spatially correlated across administrative boundaries, which would lead to an underestimate of the true β_1. One would expect that a district located near the country border interacts with its neighbor at the other side of the fence, and likely experiences similar economic outcomes resulting from weather changes.

Second, we are not discerning between spatial economic activities (which is an inevitable constraint due to the lack of data availability). This implies that effects on agriculture, probably the most affected section of the economy regarding weather, may be offset by unrelated gains in industrial sector or others that conform a big share of GDP. In such a case, there would too be an underestimate.

Third, each grid cell (or country, for that matter) observed at different points in time is a different observation altogether, so we cannot be completely confident of what would have happened to the dissagregated GDP *had a weather shock not occurred*. In particular, the

Table 2: Panel FE Results

	log(GDP)		Growth
	[1]	[2]	[3]
Precipitation	−0.012	−0.010	
	(0.000)***	(0.000)***	
	[0.004]***	[0.004]***	
Temperature	−0.008	0.402	
	(0.003)***	(0.008)***	
	[0.097]	[0.059]***	
Temperature^2		−0.012	
		(0.000)***	
		[0.003]***	
ΔPrecipitation			−0.006
			(0.001)***
			[0.002]***
ΔTemperature			0.026
			(0.030)
			[0.294]
Constant	19.724	16.755	2.350
	(0.081)***	(0.098)***	(0.026)***
	[2.999]***	[1.640]***	[0.018]***
Observations	121088	121088	113520
R^2	0.232	0.271	0.160
adj. R^2	0.230	0.269	0.158
Prob. $> F$ (Robust SEs)	0.000	0.000	0.000
Prob. $> F$ (Clustered SEs)	0.000	0.000	0.024

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Note: The results of this table exclusively use the GDP data from Kummu et al. (2018). All specifications are calculated with country–year fixed effects. The specifications for [1] and [2] control for altitude, latitude and population density. Robust standard errors are reported in parenthesis and clustered standard errors to the country-year level are in brackets. The clustered SEs are vastly more conservative and may underrepresent the true assignment mechanism because seasonality changes would happen at the grid cell or sub-country level. Nevertheless, it is seen that the significance of the precipitation estimates is unchanged by this.

structural equation associated to (4) requires we assume that a country itself observed several periods where ΔP and ΔT are low is a good approximation of a "control" group, which is to be compared with periods where ΔP and ΔT are high—the "treatment". There are important reasons why this may not be a good resemblence, which are discussed in the next subsection.

And fourth, rainfall may not only be characterized by the mean but by its dispersion across time. It can be expected that economic outcomes are different when all rainfall is accumulated in one single month of the year, as opposed to it being evenly distributed across the year.

The first two limitations pose an interpretation problem rather than a methodological one, in the sense that the magnitude of the estimate may not causative but its sign would. The third one is methodological and is shared in common with all other papers using intra-country variation. It is related to the fact that we are comparing the same observations in different periods, assuming that a "normal" year in weather reflects what would have happened had a "bad" year not occurred. Hence, the intra-location approach to measuring the effects of weather consists of *assuming* that the normal instance of previous periods precisely resembles the *counterfactual* of the abnormal current period. The fourth limitation poses an identification problem, which implies that β_1 is not the effect of precipitation at large, but the only effect of the amount of rainfall averaged across the year.

5.1   The Seasonality Index

The climatology of precipitation can be exploited to analyze the same setting in a more experimental way, so as to circumvent the last two limitations mentioned above. In particular, we can characterize precipitation by the shape of its distribution over time in terms of its first

and second moments;[17] which are measures of the center and dispersion respectively. The latter is referred in the climatology literature as the relative seasonality,[18] and it defines the proportion of time the average amount of rainfall in the year happens at a certain number of months or sub-periods. Relative seasonality can be quantified using the Seasonality Index (SI) as defined by Ayoade (1970),

$$SI_\tau = \frac{100}{22P_\tau} \sum_{m=1}^{12} |12P_m - P_\tau| \times \%$$

where P_m and P_τ are the monthly and yearly precipitation averages respectively. The index is constructed so that $0 \leq SI_\tau \leq 1$. On one extreme, $SI_\tau = 1$ happens when all the rainfall within that particular year falls in a single month, implying an extreme regime. The opposite, $SI_\tau = 0$, happens instead when rainfall is evenly distributed across months so that the rainfall regime is perfectly equable. Using the thresholds of Walsh & Lawler (1981), intermediate cases $0 < SI_\tau < 1$ are defined (see Table 3).[19]

The SI is a preferable measure of dispersion in our setting as opposed to the empirical standard deviation because, first, it varies on a closed interval which allows for an intuitive interpretation, and second, the range of its possible values are clearly defined from the climatology literature. More importantly, the SI is a more natural way of measuring rainfall

---

[17]Theoretically, these may be the moments around the mean $[E(P - \mu_P)$ and $E((P - \mu_P)^2)]$, although recall that the temporal distribution of precipitation with respect to time is discontinuous at a rate of zero. One could also define this in terms of the moments around zero $[E(P)$ and $E(P^2)]$.

[18]The climatology literature defines two types of seasonality, absolute and relative. Absolute seasonality, in contrast, simply states weather a particular year has experienced more or less rainfall on average. Technically, absolute and relative seasonality jointly characterize precipitation in terms of its temporal distribution. However, it is evident that absolute seasonality is captured in the sign of the difference ΔP_τ.

[19]Walsh and Lawler's version of the seasonality index is virtually identical to that of Ayoade, which is

$$SI = \frac{1}{P_\tau} \sum_{m=1}^{12} \left| P_m - \frac{P_\tau}{12} \right| \times \% \in [0, 1.833].$$

Their reasoning is that this more computationally-efficient to compute than Ayoade's original index. Nevertheless, we opt to use Ayoade's version in conjunction with Walsh and Lawler's classes.

dispersion across time in a dissagregated setting because it was originally conceived to be used with climatic station data. Other studies at the aggregate level have opted to define rainfall dispersion in terms of the empirical standard deviation; some define it around the long run mean (the mean computed over the entire century or various decades), and others using high frequency (monthly) data define it around each year's mean (Deschênes & Greenstone, 2007; Damania et al., 2017, 2020). When using standard deviations, the customary thing to do is define abnormal instances as those where the precipitation rate is two standard deviations above or below the chosen mean.

Table 3: Walsh and Lawler's (1981) thresholds for the Seasonality Index

Classes	SI
0. Very equable	≤ 0.16
1. Equable with a wet season	0.17 – 0.33
2. Seasonal with a short dry season	0.34 – 0.49
3. Seasonal	0.50 – 0.66
4. Seasonal with a long dry season	0.67 – 0.83
5. Most rain happens in 3 months or less	0.84 – 0.89
6. Extreme regime, all rain in 1–2 months	$1 \geq 0.89$

Note: The original thresholds by Walsh and Lawler were proposed with respect of their own version of the Seasonality Index, which varied between 0 and 1.83; the class thresholds given in their work were converted to fit with Ayoade's version of the index which (more intuitively and better for our interpretation purposes) varies between 0 and 1.

5.2 Main strategy

While the estimates from Table 1 suggest significant associations of rainfall with the GDP levels and growth through the average amount of rainfall, we have cannot make inferences about their causal validity. To provide such validation, an alternative and more experimental approach is proposed.

The idea is to spatially identify locations within each country that have had changes in their seasonality category and others that have not in a manner of "treatment" and "control". Specifically, because there are various categories of seasonality and various periods, with the computation of the Seasonality Index we can classify each location, within a country and for each year, into one of the seven possible categories from Table A. Knowing which locations belong to each SI category at year τ, we can tell which locations experienced changes in their relative seasonality for year $\tau + 1$. Moreover, for a group of locations that belong to some category at τ, we can identify which them experience a change their seasonality at $\tau + 1$ and compare the difference in economic outcomes to obtain the effect of the seasonality change—provided that this change was exogenous. In essence, those locations which have not changed in seasonality for year $\tau + 1$ can be thought as the control, and those locations which have changed for $\tau + 1$ can be thought of as receiving treatment in the form of a different temporal distribution of rainfall starting at period τ. This comparison will be valid *only between those locations which belong to the same SI category at baseline or period τ*, and ideally, between locations that are as similar as possible before any seasonality change happens.

This last statement is important, because if the observed characteristics of treatment and control at period τ are equal on average, then the control group can be regarded as a precise resemble of the treatment group's counterfactual in the sense of Rubin (1978, 2005, 2019). In such case, the difference in economic outcomes from treatment and control (change in the SI vs. no change) reflects the causal effect of rainfall seasonality. The methodological nuances for a precise estimation under this strategy are straightforward, because we can compute the (real) Differences-In-Differences estimator, which is by nature more precise and efficient than

the simple difference in outcomes. Furthermore, we can expand this same idea to capture the marginal effect of rainfall given a change in the seasonal category. This specification captures precisely what we want to measure, which how the distribution of precipitation on both its temporal and quantitative dimension affect the dissaggregated measures of GDP.

Figure 3 illustrates this notion. We want to compare only locations that have both similar characteristics and the same seasonality category at baseline. This way, we get to statistically comparable groups of observations, and a difference-in-difference estimator would allow us to find an unbiased estimate of a seasonality change. We think of "treatment" as being subject to a change in seasonality in a quasi-experimental sense.

Figure 3: Illustration of the main identification strategy
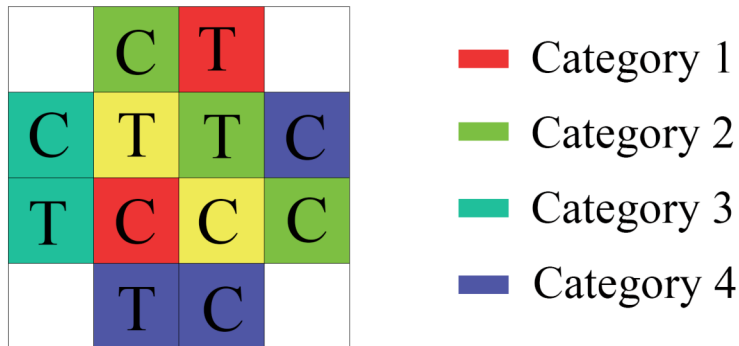
Note: This figure illustrates the comprasion between grid cells of the same seasonality index at baseline. Each color denotes a different seasonality category. "T" and "C" denoting treatment and control respectively, where treatment is a seasonality change (either up or down a category) and control is no seasonality change. The main identification assumption is that those grid cells represented by the same color at time t are similar for quasi-experimental comparison at time $t+1$, where the similarities hold in both the observed and unobserved characteristics. Then, comparing those grid cells marked "T" with those marked "C" would allow us to find the treatment effect of a seasonality change. In particular, we propose a differences-in-differences approach as opposed to a simple difference in the comparison of outcomes.

To write the empirical equations to be estimated, let $C_{s\tau}$ denote the Seasonal Index category of grid cell s at year τ, where it takes one of seven possible values corresponding to each class shown in Table A. As before, $\Delta C_{s\tau}$ would denote the change in category with

respect to the previous period, $C_{s\tau} - C_{s,\tau-1}$. This is a natural way of representing changes in the category that go up in the scale of table A, in which case the sign is positive, and changes that go down the scale, where the sign is negative. The idea is that $\Delta C_{s\tau}$ allows for a better approximation of treatment on a particular year, as opposed to changes in the average amount of rainfall. Like with all other papers analyzing precipitation, a categorical change of seasonality may be regarded as an exogenous phenomenon to economic activity.

The Differences-In-Differences strategy mentioned above would involve estimating the following

$$\Delta \log(Y_{s\tau}) = \alpha_0 + \alpha_1 \gamma_{s\tau} + \alpha_2 \Delta P_{s\tau} + \alpha_3 \Delta T_{s\tau} + \delta_{\mathscr{C}c\tau} + \zeta_{s\tau} \tag{5}$$

where $\gamma_{s\tau} = \mathbb{1}\{\Delta C_{s\tau} \neq 0\}$ is a dummy variable specifying whether there is change in the seasonal index category with respect to the previous period. We now take advantage of the use of fixed effects at the country-year-category level, $\delta_{\mathscr{C}c\tau}$. By including a dummy variable that equals 1 for every country-category-year combination, we are isolating treatment and control as defined above. The estimate α_1 will be the difference between treatment (a change in the seasonal category with respect to the previous period) and control (no change), substracted from individual characteristics. The Differences-in-Differences estimator in (5) more closely resembles an experimental setting as opposed to the one in (4) because we are explicitly separating the locations which have had an SI change from those that have not, at each year. Hence, we leave the assumption that each location is its own counterfactual at different points in time. At the same time, because we have defined treatment in terms of the temporal distribution of rainfall and not $P_{s\tau}$ itself, we can look at the marginal effect of precipitation *given relative seasonality*. That is, we can analyze the effect of precipitation on both of its dimensions: the amount of rain and the temporal distribution of rain each year. We

can estimate

$$\Delta \log(Y_{s\tau}) = \alpha_0 + \alpha_1 \gamma_{s\tau} + \alpha_2(\gamma_{s\tau} \times \Delta P_{s\tau}) + \alpha_3 \Delta P_{s\tau} + \alpha_4 \Delta T_{s\tau} + \delta_{\mathscr{C}c\tau} + \zeta_{s\tau} \qquad (6)$$

where the parameter associated to the interaction term, α_2, captures the marginal effect.

5.3 Identification & the SI-Precipitation rate relationship

Evidently, the identification assumption required for these estimates to be causal is that the variation in the Seasonality Index and precipitation rates are exogenous to the outcome variable of interest, growth, which ideally would imply that the observed and unobserved characteristics of those locations experiencing seasonality changes are equal (on average) than those locations that remain on the same category. Notably, most (if not all) of similar studies on the subject, have regarded rainfall as a variable that fluctuates without anything to do with the economic system in the short run. As it was mentioned in the Introduction, rainfall has even been used as a "benchmark" instrument in a wide variety of empirical studies, because it confidently provides variation that is seldom related to GDP or growth. Nevertheless, for our purposes and in our dissagregated setting, it is possible to provide tests for identification.

We start by estimating specifications that intend to look at the difference in outcomes across locations. Because we expect changes to happen exogenously, we can test whether there are systematical differences in GDP of locations that experience seasonality changes with locations that do not, in a particular category. If on average GDP is higher or lower moving up or down in the scale, then it would be likely that the assignment mechanism is not exogenous. This would, intuitively, mean that economic activity is more established on places that do not experience variations in relative seasonality, so any estimates would be

38

biased. Mathematically we have,

$$\log(Y_{s\tau}) = \beta_0 + \beta_1 \gamma_{s\tau} + \beta_2 P_{s\tau} + \beta_3 T_{s\tau} + \mathbf{X}'_{s\tau}\boldsymbol{\beta} + \delta_{\mathscr{C}c\tau} + \varepsilon_{s\tau} \qquad (7)$$

where β_1 is intended to capture the difference in outcomes of locations that have had *any* change seasonality with respect to the previous period *within countries and across time*. The test surmounts to not rejecting H_0: $\hat{\beta}_1 = 0$.

Because we can think of the treatment assignment repeating itself every year, we can see if locations are different before (*ex ante*) the change occurs within categories. This tests the assumption that baseline conditions in the outcome variable are similar on average before treatment, which would tell us if the dummy variable $\gamma_{s\tau}$ provides a valid way of comparison between regions that do not have experienced seasonality changes and those that have. In other words, this is what would allow to make the estimate α_1 in (5) valid in the sense that it would capture the effect of treatment. We can test this by running

$$\log(Y_{s\tau}) = \beta_0 + \beta_1 \gamma_{s,\tau+1} + \beta_2 P_{s\tau} + \beta_3 T_{s\tau} + \mathbf{X}'_{s\tau}\boldsymbol{\beta} + \delta_{\mathscr{C}c\tau} + \varepsilon_{s\tau} \qquad (8)$$

where $\gamma_{s,\tau+1}$ is the next periods' assignment. If $\gamma_{s,\tau+1} = 1$, then location s experiences a seasonality change at period $t + 1$, and for the comparison of the events[20]

$$[Y_{s',\tau+1} \mid \gamma_{s',\tau+1} = 1] \quad \text{and} \quad [Y_{s'',\tau+1} \mid \gamma_{s'',\tau+1} = 0] \quad \text{for} \quad s' \neq s''$$

to be valid, we want that $\mathrm{E}\,[Y_{s'\tau}] = \mathrm{E}\,[Y_{s''\tau}]$. In the regression, we need to not reject H_0: $\hat{\beta}_1 = 0$.

---

[20]Here we use the conditional notation from Imbens & Rubin (2015).

Table 4: Differences at baseline

Variable	Variation in SI	No variation in SI	Difference	Diff. p-value
log(GDP)	16.938	16.957	0.019	0.271
Growth	2.669	2.737	0.068	0.227
HDI	26.758	26.774	0.016	0.165
Population density	83,002.98	86,291.29	3,288.31	0.216
Latitude	−18.670	−18.621	0.058*	0.097

$*p < 0.10, **p < 0.05, ***p < 0.01$

Note: All specifications are computed using country-year-category fixed effects. Column 2 ("Variation in SI") shows the mean value at time t (within country $\mathscr{C}$ and category c) for locations that did not experience a change in seasonality at time $t + 1$, for each variable on the leftmost column. Likewise, Column 3 shows the mean value at time t for places that experienced seasonality changes in $t + 1$. A change in seasonality is defined as a dummy variable that takes the value of 1 whenever the seasonality is different according to the categories above. Column 4 shows the difference between columns 2 and 3, and column 5 shows the p-value associated to the difference. All regressions control for precipitation, temperature and temperature squared. HDI refers to the Human Development Index, and the data comes from Kummu et al. (2018)

Table 4 shows the baseline characteristics and the differences across regions that experience seasonality changes, within a country and a particular year, with regions that do not. It is seen that the differences in averege is mostly not significant. Furthermore, in Table 5 the rightmost columns, [3] and [4], show that H_0 from eqs. (7) and (8) can be confidently rejected, so that *within categories* locations assigned to seasonality changes before and after do not have different GDP levels, which is at the center of our main identification assumption. The two leftmost columns, [1] and [2], instead show that *within countries* we cannot confidently reject (at a 10% level at least) that assignment to seasonality variation does not differ across observations of GDP. The main findings of Table 5 as are that (*i*) baseline GDP is not different before the changes occur, as shown in col. [4], and (*ii*) locations experiencing seasonality changes are not necessarily different in their level of GDP, as shown in col. [3]. Interestingly, Table 5 also shows that precipitation is different across all these dimensions on

Table 5: Panel FE Results of the main identification assumptions

	log(GDP$_\tau$)			
	[1]	[2]	[3]	[4]
ΔC_τ dummy	0.029*		0.001	
	(0.016)		(0.017)	
$\Delta C_{\tau+1}$ dummy		0.023		0.017
		(0.016)		(0.017)
Precipitation$_\tau$	−0.010***	−0.010***	−0.007***	−0.007***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	16.740***	16.776***	17.346***	17.366***
	(0.099)	(0.102)	(0.090)	(0.093)
Country-Year FE	Yes	Yes		
Country-Category-Year FE			Yes	Yes
Observations	121088	113520	121088	113520
R^2	0.271	0.272	0.332	0.332
adj. R^2	0.269	0.270	0.323	0.324
Prob. $> F$	0.000	0.000	0.000	0.000

$^*p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Note: All specifications control for temperature, temperature^2, altitude, latitude and population density (not shown). The unit of measurement for precipitation are mm^3. The variable "ΔC_τ dummy" equals one if there is any change in the categorical Seasonality Index with respect to the previous period, and "$\Delta C_{\tau-1}$ dummy" equals one if there is a categorical change for the subsequent period with respect to the current one. Robust standard errors are reported in parentheses.

average. Precipitation rates are lower on those locations that do not have seasonality changes.

Regarding this seemingly negative correlation of precipitation with GDP levels when accounting for relative seasonality, one can further use the full variability in the category dimension to get a similar table showing heterogeneous correlations. Specifically, if we include interaction terms of the SI category with the rate of precipitation, the associated parameters will capture the additional effect of a mm^3 of rainfall at each category level. Letting $\gamma_{js\tau} = \mathbb{1}\{C_{s\tau} = j\}$ be a dummy variable that takes the value of one when the SI category of s at year τ is j, where $j \in \{0, 1, \ldots, 6\}$, we can write this explicitly as

$$\log(Y_{s\tau}) = \beta_0 + \sum_{j=1}^{6} \beta_j \left(\gamma_{js\tau} \times P_{s\tau}\right) + \beta_8 P_{s\tau} + \beta_9 T_{s\tau} + \mathbf{X}'_{s\tau}\boldsymbol{\beta} + \delta_{\mathscr{C}c\tau} + \varepsilon_{s\tau}. \quad (9)$$

The omitted category is the first one, where seasonality is very equable across months. Here $\hat{\beta}_{j+7} = 0$ for all j would imply that there is no attributable effect to more or less rainfall at any given seasonality level, contradictory with what is seen in Table 2. Similarly, $\hat{\beta}_8 = 0$ and $\hat{\beta}_j \neq 0$ for $j = 1, \ldots, 6$ would mean that the only attributable effect of precipitation on GDP is due to the SI—i.e., precipitation would only matter because of its relative, and not absolute, seasonality. A constistent result with our previous findings would be that $\hat{\beta}_j \neq 0$ for at least one j and $\hat{\beta}_8 \neq 0$. Note that because of the use of country-category-year fixed effects, $\delta_{\mathscr{C}c\tau}$, the estimate $\hat{\beta}_j$ reflects the marginal effect of precipitation given seasonality $j \neq 0$, in contrast with seasonality $j = 0$. In other words, $\hat{\beta}_j$ is to be interpreted as the additional effect perceived on the outcome variable with respect to a very equable rainfall

regime. Also note that this is a more precise way of estimating

$$\log(Y_{s\tau}) = \theta_0 + \sum_{j=1}^{6} \theta_j \gamma_{js\tau} + \sum_{j=1}^{6} \theta_{j+6} \left(\gamma_{js\tau} \times P_{s\tau} \right) + \theta_{14} P_{s\tau} + \cdots + \delta_{\mathscr{C}\tau} + \varepsilon_{s\tau}.$$

for the interaction terms $\hat{\theta}_7, \dots, \hat{\theta}_{12}$ using country-year fixed effects. For comparison, the estimates of both equations are shown in Table 3.

Table 6 shows the results of equation (9) in conjunction with the country-year FE specification with individual dummies for each category (columns [3] and [4]). Note that in the more isolated specification these terms would be collinear. Columns [1] and [2] show the same results from Table 3 for additional comparison and adding the terms for temperature. It is seen that seasonality categories are significant and positive across all categories in column [3], in contrast with a greater magnitude of the constant in col. [4]. This implies that (on average) more equable regimes tend to have a greater GDP level. A reason for this may be that most of the spatial economic activity is established in locations with relatively equable rainfall; it is also plausible that a higher share of GDP is obtained in the most seasonal areas. Because more equable regimes are the most uniformly distributed ones, this would support the notion that the predictability of rainfall is benefitial—contrasting with the idea that predictability offsets any changes in economic outcomes due rainfall variation (Auffhammer et al., 2006). Interestingly, the interactions are significant across all categories, and go from negative to positive going in the direction of less equable regimes. Since less equable regimes are the less uniformly distributed, they could also bear a high predictability rate for economic agents (e.g., agents would know which months of the year have rain, so they can plan ahead). So all this would support the notion that when the seasonality rainfall is predictable, *then* the amount of rainfall matters positively for the level of GDP. More explicitly, the interaction

Table 6: Panel FE Results, comparison of precipitation effects among SI categories

	log(GDP)			
	[1]	[2]	[3]	[4]
Constant	16.740***	17.346***	11.678***	17.249***
	(0.099)	(0.090)	(0.243)	(0.090)
ΔC dummy	0.029*	0.001	0.007	−0.010
	(0.016)	(0.017)	(0.016)	(0.017)
Precipitation	−0.010***	−0.007***	0.006***	0.003**
	(0.000)	(0.000)	(0.001)	(0.001)
Temperature	0.402***	0.348***	0.381***	0.346***
	(0.008)	(0.007)	(0.007)	(0.007)
Temperature^2	−0.012***	−0.011***	−0.012***	−0.011***
	(0.000)	(0.000)	(0.000)	(0.000)
SI category dummies:				
Wet season			5.030***	
			(0.233)	
Short dry season			5.480***	
			(0.222)	
Seasonal			5.035***	
			(0.221)	
Long dry season			4.478***	
			(0.221)	
all rain in 3 ≤ months			4.381***	
			(0.223)	
Extreme, all rain in 1-2 months			4.417***	
			(0.226)	
Precipitation interacted with:				
Wet season			−0.016***	−0.015***
			(0.001)	(0.001)
Short dry season			−0.017***	−0.012***
			(0.001)	(0.001)
Seasonal			−0.011***	−0.009***
			(0.001)	(0.001)
Long dry season			0.000	0.003**
			(0.001)	(0.001)
All rain in 3 ≤ months			0.009***	0.016***
			(0.001)	(0.002)
Extreme, all rain in 1- months			0.010***	0.026***
			(0.002)	(0.003)
Country-Year FE	Yes		Yes	
Country-Category-Year FE		Yes		Yes
Observations	121088	121088	121074	121074
R^2	0.271	0.332	0.298	0.340
adj. R^2	0.269	0.323	0.296	0.332
Prob. $> F$	0.000	0.000	0.000	0.000

$*p < 0.10, **p < 0.05, ***p < 0.01$

Note: All specifications control additionally for altitude, latitude and population density. The unit of measurement for precipitation are mm^3; and for temperature it is centigrade degrees. The variable "ΔC dummy" equals one if there is any change in the categorical class of the Seasonality Index with respect to the previous period. These results were computed using the dissagregated GDP data from Kummu et al. (2018). Robust standard errors are reported in parentheses.

44

estimates show that an additional mm^3 of rainfall is associated with a ($\sim$ 1 percentage point) decrease in the level of GDP at intermediately equable seasons; and an additional mm^3 of rainfall is associated with a ($\sim$ 1-2 percentage point) increase in the level of GDP in years with less uniformly distributed rain.

Finally, it is important to mention that because the assignment to seasonality categories themselves may not be random (as the significance of the estimates show), we are not attributing any causal validity to the results from tables 2 and 3. They simply intent to show the average differences in GDP across seasonal categories and their relationship with precipitation. And, as the assignment to the changes in seasonality does not seem to be significantly associated with GDP levels at the category level, we can say that the data does not support the idea that changes in seasonality happen *because* GDP is different across locations.

6   Main Results

Table 7 displays the main results from the estimation of equations (5) and (6) from section 5.2 in columns [1] and [2] respectively. Column [3] breaks down the dummy variable for any categorical change into positive (up in the scale) and negative (down in the scale). Column [4] breaks down instead the additional effect of variations in precipitation for each category in the Seasonality Index as opposed to the binary interaction for any treatment—similar to Table 6.

It can be observed that the associated estimate to $\gamma_{s\tau}$ ("ΔC dummy") is relatively unchanged across all three relevant specifications, suggesting that changes in seasonality drive growth up by 0.3 percentage points. To put this result in contrast, the average growth rate across all locations considered (20 countries in the entire Latin American region) is 2.2 per-

centage points—which is captured in the magnitude of the constant term. Meaning that a change in the seasonality scale, either up or down, on any given particular year will lead, on average, to a growth rate that is $(0.3/2.2)\% = 13.6\%$ bigger than it would have been otherwise. In other words, for a group of locations within a country that start at the same category of seasonality, those experiencing changes in their category for the subsequent period will have a bigger rate of growth.

Interestingly, across all specifications written above on both the level of GDP and the rate of growth, the amount of precipitation has been consistently found to have a negative associated effect of 0.007–0.010 percentage points, and Table 7 is no exception. Notably, these are not necessarily small effects neither when put into perspective. The variable for precipitation is measured in mm^3, and the standard deviation of precipitation across the most populated areas in the entire region is 72 mm^3 (see the summary statistics in the data section). Hence, a one-half standard deviation increase in such case would lead to approximately a 0.36 percentage point decrease in growth, implying that the end-of-year growth rate would be around 16% smaller had the increase not happened. Moreover, if a change in the seasonality of these areas is accompanied by a half-standard-deviation increase in the amount of rainfall, then the effect of seasonality on growth will be completely offset by it.

Even if such a variation on the amount of precipitation is unrealistic, the fact that we cannot reject the null hypothesis of a zero effect of precipitation on growth (and GDP) over a variety of specifications tells us that there is a statistical association that cannot be neglected. More importantly, although consistent with previous findings in the literature, the results from Table 4 go on to show that the temporal distribution of rain across the year is associated to growth in a way that may be similar in magnitude to that of average precipitation documented

Table 7: Main panel FE Results

	Growth			
	[1]	[2]	[3]	[4]
Constant	2.204***	2.206***	2.205***	2.204***
	(0.037)	(0.037)	(0.037)	(0.037)
ΔC dummy	0.314***	0.308***		0.310***
	(0.057)	(0.057)		(0.057)
$\Delta C > 0$ dummy			0.048	
			(0.090)	
$\Delta C < 0$ dummy			0.573***	
			(0.067)	
ΔPrecipitation $\times \Delta C$ dummy		−0.010***		
		(0.003)		
ΔPrecipitation	−0.007***	−0.003*	−0.007***	−0.016***
	(0.001)	(0.001)	(0.001)	(0.005)
ΔTemperature	0.183**	0.175*	0.206**	0.171*
	(0.090)	(0.089)	(0.088)	(0.089)
Δ*Precipitation interacted with*:				
Wet season				0.015***
				(0.005)
Short dry season				0.011**
				(0.006)
Seasonal				0.001
				(0.006)
Long dry season				0.011*
				(0.006)
All rain in $3 \leq$ months				0.018**
				(0.007)
Extreme, all rain in 1-2 months				0.049***
				(0.010)
Observations	113520	113520	113520	113506
R^2	0.203	0.203	0.203	0.203
adj. R^2	0.193	0.193	0.193	0.193
Prob. $> F$	0.000	0.000	0.000	0.000

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Note: All specifications were calculated with country-category-year fixed effects. The precipitation variables are measured in mm^2 and temperature is measured in centigrade degrees. The variable "ΔC dummy" equals one if there has been a change in the Seasonality Index category with respect to the precious period; "$\Delta C > 0$ dummy" equals one if the change occurs in the direction of a more extreme (less equable) rainfall regime; and "$\Delta C < 0$ dummy" equals one if the change is in the direction of a more equable regime. These results were computed using the disaggregated GDP data from Kummu et al. (2018). Robust standard errors are reported in parentheses.

in similar studies.

Column [3] shows an interesting relationship between the sign if the change in relative seasonality and the associated effect on growth. It suggests that the effect is not different from zero when the category goes on the direction of a more extreme seasonality. Instead, the non-negligible effect is attributable to only those categorical changes going into a more equable type of rainfall seasonality. In plain words, growth only seems to be affected, positively, when rainfall becomes more evenly distributed across months. This is again consistent with the hypothesis that more predictable rainfall is beneficial for growth, because uniformly-distributed rainfall tends to be the less variable. In the same light, this is also consistent with the fact that the variation in the average amount of precipitation is negatively associated with growth.

Nevertheless, some caution must be drawn into the estimates of col. [3]. On the contrary to the simple dummy variable for ΔC (this variable equals 1 for any change, no matter the direction), the assignment to a binary relationship such that $\Delta C > 0$ or $\Delta C < 0$ may not be close to randomly assigned. In fact, it is intuitive that some regions would tend to change their rainfall seasonality only in a particular direction; for instance, locations near the shores would tend to only have seasonality changes in the direction of more equable regimes, because the majority of the time such locations already have rainfall unevenly distributed across months. They cannot have a less equable seasonality if they are already in the least equable of categories.

The fact that directionality may not be exogenous is related to the fact that geography, as aluded in the Literature section, is indeed fixed across observations. Hence, the interpretation of the estimates from the dummy variables in column [3] is limited and we cannot confidently

assert the validity of their magnitudes. However, we can suggest there is a relevant relationship regarding directionality on the grounds of significance. We can reject that the estimate associated to the dummy variable $\Delta C > 0$ is zero, and this is consistent with the hypothesized relationship between rainfall predictability and growth mentioned in the literature. All that we can confidently assert, is that regions experiencing rainfall changes going into more equable regimes experience higher rates of growth as compared to locations within the same baseline categories that do not have any seasonality changes.

Column [4] displays a similar specification as the ones in [1] and [2] but adding interaction terms for each seasonal category with variations in precipitation; the omitted term is the one associated to category 0, or the most equable of regimes. This allows us to have an idea of the additional effect of rainfall for all seasonality categories independently of the treatment status. It is seen that additional amounts of rainfall are more significant at the extreme, opposite regimes, which suggests that there may not be any such marginal effect when the economy experiences seasonal years. The important aspect is that the estimates are positive, while the term for precipitation variation remains negative and has increased in magnitude. Thus, this implies that for a fixed amount of precipitation variation, there are benefits under the most extreme and equable regimes, while the same amount would be contractive for the seasonal ones. So overall, col. [4] shows that the positive effect associated to changes in precipitation is heterogeneous in the seasonal category dimension.

7 Concluding remarks

Precipitation matters for both its temporal and spatial distribution. Recognizing this, the present paper developed an estimation strategy around the Seasonality Index, a number that

captures the dispersion of rainfall across months in a year, and using GDP dissagregation to the grid cell level. As observations are on a finer level of geographical detail, we were able to isolate locations with the same baseline sasonality (within a country and within a particular year), and showed that locations that experience seasonality changes for the next period are similar ex ante, on average, over a number of covariates. With this, the developed empirical strategy is akin to a diff-in-diff approach, where we compared the variations in seasonality with variations in the GDP level, or growth. The identification assumption we needed was that locations that experienced seasonality variations would not have had different trends in GDP had the change not happened. Because baseline characteristics are mostly similar, we argued why this assumption is not unreasonable. We found that locations that experience seasonality changes, either to more equable or more extreme regimes, are better-off than the locations that do not. This suggests that while greater variations in the average rainfall rates are detrimental on growth, when the temporal distribution changes growth is bigger.

Although we were not able to identify explicitly the mechanism behind this positive effect, we can suggest a number of possible elements that are consistent with previous literature (see Section 2, Literature Review). Firstly, the accumulation of yearly rainfall on a few months could influence the capability of adaptation, so that individual labor-leizure decisions are taken in a more efficient setting and in turn increasing productivity. Secondly, an evenly distributed rainfall across the year (coming from an extreme regime) could have positive effects because it changes the production possibility frontier (PPF). In particular, it could improve agricultural yields and the production of other activities dependent on precipitation rates. It is uncluar how the trade-off between more equable or more extreme regimes work, but our estimates suggest that any negative effects are offset by positive ones in either case.

Finally, our results are also consistent with previous studies analyzing precipitation rates, in the sense that an increased average precipitation is negatively associated to growth. While our estimation strategy did not revolve around yearly (average) precipitation rates—thus lowering the confidence of the causality in this particular topic—the sign of our estimates is consistent, significant, and robust across all specifications.

References

Acemoglu, D., Johnson, S., & Robinson, J. A. (2001, December). The Colonial Origins of Comparative Development: An Empirical Investigation. *American Economic Review*, *91*(5), 1369–1401. doi: 10.1257/aer.91.5.1369

Adams, R. M. (1989, December). Global Climate Change and Agriculture: An Economic Perspective. *American Journal of Agricultural Economics*, *71*(5), 1272–1279. doi: 10.2307/1243120

Adams, R. M., Rosenzweig, C., Peart, R. M., Ritchie, J. T., McCarl, B. A., Glyer, J. D., . . . Allen, L. H. (1990). Global climate change and US agriculture. *Nature*, *345*(6272), 219–224.

Angrist, J. D., & Imbens, G. W. (1991, December). *Sources of Identifying Information in Evaluation Models* (Tech. Rep. No. t0117). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/t0117

Angrist, J. D., & Imbens, G. W. (1994, March). Identification and Estimation of Local Average Treatment Effects. *Econometrica*, *62*(2), 467. doi: 10.2307/2951620

Auffhammer, M., Hsiang, S. M., Schlenker, W., & Sobel, A. (2011). *Global Climate Models: A User Guide for Economists.*

Auffhammer, M., Ramanathan, V., & Vincent, J. R. (2006). Integrated model shows that atmospheric brown clouds and greenhouse gases have reduced rice harvests in India. *Proceedings of the National Academy of Sciences*, *103*(52), 19668 LP – 19672.

Ayoade, J. O. (1970, September). THE SEASONAL INCIDENCE OF RAINFALL. *Weather*, *25*(9), 414–418. doi: 10.1002/j.1477-8696.1970.tb04132.x

Barbier, E. B. (2004). Water and Economic Growth. *Economic Record*, *80*(248), 1–16. doi: 10.1111/j.1475-4932.2004.00121.x

Barbier, E. B. (2015). Water and growth in developing countries. In A. Dinar & Schwabe, Kurt (Eds.), *Handbook of water economics.* Cheltenham: Edward Elgar.

Barrios, S., Bertinelli, L., & Strobl, E. (2010, May). Trends in Rainfall and Economic Growth in Africa: A Neglected Cause of the African Growth Tragedy. *Review of Economics and Statistics*, *92*(2), 350–366. doi: 10.1162/rest.2010.11212

Barro, R. J. (1990). Government Spending in a Simple Model of Endogeneous Growth. *Journal of Political Economy*, *98*(S5), 103–125. doi: 10.1086/261726

Barro, R. J., & Sala-I-Martin, X. (1992). Public Finance in Models of Economic Growth. *The Review of Economic Studies*, *59*(4), 645–661.

Blanc, E., & Strobl, E. (2013). *Is Small Better? A Comparison of the Effect of Large and Small Dams on Cropland Productivity in South Africa.* Washington D.C.

Brown, C., Meeks, R., Ghile, Y., & Hunu, K. (2013, November). Is water security necessary? An empirical analysis of the effects of climate hazards on national-level economic growth. *Philosophical Transactions of the Royal Society*, *371*(2002), 20120416.

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, *527*, 235–239. doi: 10.1038/nature15725

Connolly, M. (2008). Here Comes the Rain Again: Weather and the Intertemporal Substitution of Leisure. *Journal of Labor Economics*, *26*(1), 73–100. doi: 10.1086/522067

Damania, R. (2020). The economics of water scarcity and variability. *Oxford Review of Economic Policy*, *36*(1), 24–44. doi: 10.1093/oxrep/grz027

Damania, R., Desbureaux, S., Hyland, M., Islam, A., Moore, S., Rodella, A.-S., … Zaveri, E. (2017). *Shocks in the Cities: The Economic Impact of Water Shocks in Latin American Metropolitan Areas* (Tech. Rep.). Washington D.C.: World Bank.

Damania, R., Desbureaux, S., & Zaveri, E. (2020). Does rainfall matter for economic growth? Evidence from global sub-national data (1990–2014). *Journal of Environmental Economics and Management*, *102*, 102335. doi: 10.1016/j.jeem.2020.102335

Dell, M. (2010). The Persistent Effects of Peru's Mining Mita. *Econometrica*, *78*(6), 1863–1903. doi: 10.3982/ECTA8121

Dell, M., Jones, B. F., & Olken, B. A. (2012). Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics*, *4*(3), 66–95. doi: 10.1257/mac.4.3.66

Dell, M., Jones, B. F., & Olken, B. A. (2014). What Do We Learn from the Weather? The New Climate–Economy Literature. *Journal of Economic Literature*, *52*(3), 740–798.

Dell, M., & Olken, B. A. (2020, January). The Development Effects of the Extractive Colonial Economy: The Dutch Cultivation System in Java. *The Review of Economic Studies*, *87*(1), 164–203. doi: 10.1093/restud/rdz017

Desbureaux, S., & Damania, R. (2018). Rain, forests and farmers: Evidence of drought induced deforestation in Madagascar and its consequences for biodiversity conservation. *Biological Conservation*, *221*, 357–364.

Deschênes, O., & Greenstone, M. (2007, February). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *American Economic Review*, *97*(1), 354–385. doi: 10.1257/aer.97.1.354

Deschênes, O., & Moretti, E. (2009, November). Extreme Weather Events, Mortality, and Migration. *Review of Economics and Statistics*, *91*(4), 659–681. doi: 10.1162/rest.91.4 .659

Duflo, E., & Pande, R. (2007, May). Dams. *The Quarterly Journal of Economics*, *122*(2), 601–646.

Fishman, R. (2016). More uneven distributions overturn benefits of higher precipitation for crop yields. *Environmental Research Letters*, *11*(2), 24004.

Gallup, J. L., & Sachs, J. D. (2001). The Economic Burden of Malaria. *American Journal of Tropical Medicine and Hygiene*, *64*(1), 85–96.

Gallup, J. L., Sachs, J. D., & Mellinger, A. (1998, December). *Geography and Economic Development* (Tech. Rep. No. w6849). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/w6849

Glaeser, E. L., & Shleifer, A. (2002, November). Legal Origins. *The Quarterly Journal of Economics*, *117*(4), 1193–1229. doi: 10.1162/003355302320935016

Graff Zivin, J., & Neidell, M. (2014, January). Temperature and the Allocation of Time: Implications for Climate Change. *Journal of Labor Economics*, *32*(1), 1–26. doi: 10.1086/671766

Hidalgo, F. D., Naidu, S., Nichter, S., & Richardson, N. (2010). Economic Determinants of Land Invasions. *The Review of Economics and Statistics*, *92*(3), 505–523.

Holland, P. W. (1986, December). Statistics and Causal Inference. *Journal of the American Statistical Association*, *81*(396), 945–960. doi: 10.1080/01621459.1986.10478354

Hsiang, S. M. (2010, August). 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. doi: 10.1073/pnas.1009510107

Hsiang, S. M., Burke, M., & Miguel, E. (2013). Quantifying the Influence of Climate on Human Conflict. *Science*, *341*(6151). doi: 10.1126/science.1235367

Hurd, B. H. (2015). Concepts and methods for assessing economic impacts from climate change on water resources. In *Handbook of Water Economics* (pp. 56–68). Edward Elgar Publishing. doi: 10.4337/9781782549666.00009

Imbens, G. W., & Rubin, D. B. (2015). *Causal inference for statistics, social, and biomedical sciences: An introduction*. Cambridge: Cambridge University Press.

Jones, B. F., & Olken, B. A. (2010, May). Climate Shocks and Exports. *American Economic Review*, *100*(2), 454–459. doi: 10.1257/aer.100.2.454

Kummu, M., Taka, M., & Guillaume, J. H. A. (2018, December). Gridded global datasets for Gross Domestic Product and Human Development Index over 1990–2015. *Scientific Data*, *5*(1), 180004. doi: 10.1038/sdata.2018.4

Mendelsohn, R., & Dinar, A. (2003). Climate, Water, and Agriculture. *Land Economics*, *79*(3), 328–341.

Mendelsohn, R., Nordhaus, W. D., & Shaw, D. (1993). The impact of climate on agriculture: A Ricardian approach. In Y. Kaya, N. Nakicenovic, W. Nordhaus, & F. Toth (Eds.), *Costs, impacts and benefits of CO2 mitigation* (pp. 173–207). Luxemburg: International Institute for Applied Systems Analysis.

Mendelsohn, R., Nordhaus, W. D., & Shaw, D. (1994). The Impact of Global Warming on Agriculture: A Ricardian Analysis. *The American Economic Review*, *84*(4), 753–771.

Michalopoulos, S., & Papaioannou, E. (2011, June). *Divide and Rule or the Rule of the Divided? Evidence from Africa* (Tech. Rep. No. w17184). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/w17184

Miguel, E., Satyanath, S., & Sergenti, E. (2004). Economic Shocks and Civil Conflict: An Instrumental Variables Approach. *Journal of Political Economy*, *112*(4), 725–753.

Nordhaus, W. (2006a). The G-Econ Database on Gridded Output:. , 30.

Nordhaus, W. (2006b, March). Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences*, *103*(10), 3510–3517. doi: 10.1073/pnas.0509842103

Nunn, N. (2008, February). The Long-Term Effects of Africa's Slave Trades. *Quarterly Journal of Economics*, *123*(1), 139–176. doi: 10.1162/qjec.2008.123.1.139

Paxson, C. H. (1992, August). Using weather variability to estimate the response of savings to transitory income in thailand. *The American Economic Review*, *82*(1), 15–33.

Rodrik, D., Subramanian, A., & Trebbi, F. (2004, June). Institutions Rule: The Primacy of Institutions Over Geography and Integration in Economic Development. *Journal of Economic Growth*, *9*(2), 131–165. doi: 10.1023/B:JOEG.0000031425.72248.85

Rubin, D. B. (1978). Bayesian Inference for Causal Effects: The Role of Randomization. *The Annals of Statistics*, *6*(1), 34–58.

Rubin, D. B. (2005, March). Causal Inference Using Potential Outcomes: Design, Modeling, Decisions. *Journal of the American Statistical Association*, *100*(469), 322–331. doi: 10 .1198/016214504000001880

Rubin, D. B. (2019, January). Essential concepts of causal inference: A remarkable history and an intriguing future. *Biostatistics & Epidemiology*, *3*(1), 140–155. doi: 10.1080/ 24709360.2019.1670513

Russ, J. (2020). Water runoff and economic activity: The impact of water supply shocks on growth. *Journal of Environmental Economics and Management*, *101*, 302–322.

Sachs, J. D. (2001). Tropical Underdevelopment. doi: 10.7916/D8M90GJP

Sachs, J. D. (2003, February). *Institutions Don't Rule: Direct Effects of Geography on Per*

Capita Income (Tech. Rep. No. w9490). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/w9490

Sachs, J. D., & Malaney, P. (2002, February). The economic and social burden of malaria. *Nature*, *415*(6872), 680–685. doi: 10.1038/415680a

Sadoff, C., Hall, J., Grey, D., Aerts, J., Ait-Kadi, M., Brown, C., ... Wiberg, D. (2015). *Securing Water, Sustaining Growth: Report of the GWP* (Tech. Rep.). University of Oxford.

Sawano, S., Hasegawa, T., Goto, S., Konghakote, P., Polthanee, A., Ishigooka, Y., ... Toritani, H. (2008). Modeling the dependence of the crop calendar for rain-fed rice on precipitation in Northeast Thailand. *Paddy and Water Environment*, *6*(1), 83–90.

Schlenker, W., & Roberts, M. (2006, January). Nonlinear Effects of Weather on Crop Yields: Implications for Climate Change. *Rev Agric Econ*, *28*.

Schlenker, W., & Roberts, M. J. (2009, September). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, *106*(37), 15594–15598. doi: 10.1073/pnas.0906865106

Thomas, T. S., You, L., Wood-Sichra, U., Ru, Y., Blankespoor, B., & Kalvelagen, E. M. F. (2019). Generating gridded agricultural gross domestic product for brazil: A comparison of methodologies. *World Bank Policy Research Working Paper*(8985).

Walsh, R. P. D., & Lawler, D. M. (1981, July). RAINFALL SEASONALITY: DESCRIPTION, SPATIAL PATTERNS AND CHANGE THROUGH TIME. *Weather*, *36*(7), 201–208. doi: 10.1002/j.1477-8696.1981.tb05400.x

Willmott, C. J., & Matsuura, K. (2001). *Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1950 - 1999).*

WorldPop. (2020). *Individual Countries 1km Population Density (2000-2020).* University of Southampton. doi: 10.5258/SOTON/WP00674

Appendix

This section replicates the results from Tables 2, 5 and 6 using population-disaggregated GDP data. We find that the significance of the seasonality estimates and their sign do not change. Interestingly, we also see that the magnitude of these estimates is lower. We suggest two plausible explanations: (i) because spatial activity is dependent on population and thus migratory movements, the estimates resulting from the use od population-disaggregated data are biased below; and (ii) because this data completely relies on an aggregated measure (the WB's Development Indicators Database) there is some localized activity that is not captured and also biases the estimates below.

This analysis also demonstrates how average precipitation rates vary across datasets. It is found that under our estimation strategy and population-disaggregated GDP, identification is weaker. Specifically, in Table 10 the precipitation estimates are no longer significant for specifications [3] and [4]. While this does not intervene with our main findings, it suggests that average precipitation rates are more difficult to observe on a yearly basis; suggesting that further research to be made on the subject should rely on a finer level of detail in the temporal dimension.

Table 8: Replication Panel FE Results

	log(GDP)		Growth
	[1]	[2]	[3]
Precipitation	−0.008	−0.006	
	(0.0001)***	(0.0001)***	
	[0.005]	[0.005]	
Temperature	−0.017	0.357	
	(0.002)***	(0.005)***	
	[0.079]	[0.057]***	
Temperature^2		−0.010	
		(0.0001)***	
		[0.002]***	
ΔPrecipitation			-2.41×10^{-5}
			(8.56×10^{-6})***
			$[2.19 \times 10^{-5}]$
ΔTemperature			−0.004
			(0.001)***
			[0.003]
Constant	4.825	2.300	0.030
	(0.081)***	(0.065)***	(0.0002)***
	[2.999]***	[1.640]***	[0.0001]***
Observations	110917	110917	103750
R^2	0.263	0.313	0.148
adj. R^2	0.262	0.311	0.147
Prob. $> F$ (Robust SEs)	0.000	0.000	0.000
Prob. $> F$ (Clustered SEs)	0.000	0.000	0.341

$*p < 0.10, **p < 0.05, ***p < 0.01$

Note: The results of this table use population-disaggregated GDP data. All specifications are calculated with country–year fixed effects. The specifications for [1] and [2] control for altitude, latitude and population density. Robust standard errors are reported in parenthesis and clustered standard errors to the country-year level are in brackets. The clustered SEs are vastly more conservative and may underrepresent the true assignment mechanism because seasonality changes would happen at the grid cell or sub-country level.

Table 9: Replication Panel FE Results of identification assumptions

	[1]	[2]	[3]	[4]
		$\log(\text{GDP}_\tau)$		
ΔC_τ dummy	0.036*		-0.002	
	(0.013)		(0.013)	
$\Delta C_{\tau+1}$ dummy		0.035***		0.025*
		(0.013)		(0.013)
Precipitation$_\tau$	-0.006***	-0.006***	-0.003***	-0.003***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	2.281***	2.279***	2.193***	2.196***
	(0.066)	(0.068)	(0.070)	(0.070)
Country-Year FE	Yes	Yes		
Country-Category-Year FE			Yes	Yes
Observations	121088	113520	121088	113520
R^2	0.313	0.313	0.366	0.366
adj. R^2	0.311	0.312	0.359	0.360
Prob. $> F$	0.000	0.000	0.000	0.000

$*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Note: All specifications control for temperature, temperature^2, altitude, latitude and population density (not shown). The unit of measurement for precipitation are mm^3. The variable "ΔC_τ dummy" equals one if there is any change in the categorical Seasonality Index with respect to the previous period, and "$\Delta C_{\tau-1}$ dummy" equals one if there is a categorical change for the subsequent period with respect to the current one. Robust standard errors are reported in parentheses. These results were computed using population-disaggregated data.

Table 10: Replication Panel FE Results, effects within SI categories

	log(GDP)			
	[1]	[2]	[3]	[4]
Constant	2.281***	2.193***	0.240**	2.077***
	(0.066)	(0.070)	(0.115)	(0.069)
ΔC dummy	0.036***	−0.002	−0.017	−0.010
	(0.013)	(0.013)	(0.013)	(0.013)
Precipitation	−0.006***	−0.003***	0.001	0.003***
	(0.000)	(0.000)	(0.001)	(0.001)
Temperature	0.357***	0.343***	0.354***	0.342***
	(0.005)	(0.006)	(0.005)	(0.006)
Temperature^2	−0.010***	−0.010***	−0.011***	−0.010***
	(0.000)	(0.000)	(0.000)	(0.000)
SI category dummies:				
Wet season			2.226***	
			(0.106)	
Short dry season			2.009***	
			(0.102)	
Seasonal			1.762***	
			(0.102)	
Long dry season			1.440***	
			(0.103)	
all rain in $3 \leq$ months			0.943***	
			(0.108)	
Extreme, all rain in 1-2 months			0.515***	
			(0.109)	
Precipitation interacted with:				
Wet season			−0.009***	−0.013***
			(0.001)	(0.001)
Short dry season			−0.006***	−0.006***
			(0.001)	(0.001)
Seasonal			−0.001***	−0.004***
			(0.001)	(0.001)
Long dry season			0.008***	0.003***
			(0.001)	(0.001)
All rain in $3 \leq$ months			0.026***	0.024***
			(0.001)	(0.001)
Extreme, all rain in 1- months			0.038***	0.036***
			(0.002)	(0.003)
Country-Year FE	Yes		Yes	
Country-Category-Year FE		Yes		Yes
Observations	110917	110917	110903	110903
R^2	0.313	0.366	0.348	0.380
adj. R^2	0.311	0.359	0.347	0.373
Prob. $> F$	0.000	0.000	0.000	0.000

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Note: All specifications control additionally for altitude, latitude and population density. The unit of measurement for precipitation are mm^3; and for temperature it is centigrade degrees. The variable "ΔC dummy" equals one if there is any change in the categorical class of the Seasonality Index with respect to the previous period. These results were computed using the population-disaggregated data. Robust standard errors are reported in parentheses.