Droughts and the growth of cities

Nicklas Nordfors*

December 9, 2024

Abstract

Some researchers and policymakers posit that climate change should increase city growth and urbanization as rising temperatures make rural livelihoods precarious, while others argue that climate change might trap rural households who cannot afford to migrate because of increasing poverty. Existing empirical evidence on the link between climate and urbanization is inconclusive. This paper exploits novel data mapping city growth for 7,000 cities in 108 low- to middle-income countries across 23 years to provide new evidence on the relationship between drought and urbanization. Cities experience large and persistent declines in growth rates after major drought events: after 11 years, cities are 0.7 percent smaller compared to a drought-free counterfactual. I show that fully accounting for dynamic effects is essential to correctly understand the relationship between drought and city growth and that a positive correlation between drought and contemporaneous city growth is misleading. Consistent with models that envision a drought-migration poverty trap, the negative effects on urbanization are more pronounced for the poorest, and most agricultural countries.

Key words: city; climate; drought; migration; urbanization

*Department of Economics, University of Gothenburg. E-mail: *nicklas.nordfors@economics.gu.se*.
I thank my advisors Anna Tompsett and Andreas Madestam for their guidance and support. This paper has benefited greatly from discussions with Niklas Bengtsson, Tessa Bold, Tillmann von Carnap, Sol Hsiang, Evelina Linnros, and feedback from seminar participants at the Stockholm University Development Tea, Stockholm University Brown Bag, UC Berkeley Development Lunch, IPSWD Columbia University, Nordic Conference in Development Economics, ICG-World Bank Young Economist Workshop, Aarhus University, and University of Gothenburg. The author alone is responsible for any remaining errors.

1 Introduction

Climate scientists predict that the frequency, intensity, and severity of droughts will increase with climate change (IPCC, 2022). How this affects the growth of cities is unknown, presenting a major problem for urban planners and policymakers who must make infrastructure and planning decisions with time horizons of several decades. Researchers and policymakers alike claim that more extreme weather events will lead to an increase in migration, projecting the number of internal climate migrants in developing countries to reach 200 million by 2050 (Clement et al., 2021), many of whom are likely to move to cities (Cattaneo & Peri, 2016). But adverse weather events might also decrease migration, especially for the poorest and agricultural households who are liquidity constrained, because it wipes out the funds needed to migrate (Bryan, Chowdhury, & Mobarak, 2014; Kleemans, 2023). Our understanding of this relationship has been constrained by the absence of disaggregated or high-frequency data on city growth, necessary to trace out the effects of specific weather events over time at lower spatial scales.

In this paper, I address these challenges by first assembling data on the annual growth of cities—as measured by city footprint, the horizontal built-up area—for 7,000 cities in 108 low- to middle-income countries, where the effects of droughts are likely to be most salient. I then combine these data with historical weather data to measure the effect of drought exposure on the growth of cities over the course of 23 years.

I study the reduced-form relationship between droughts in city hinterlands and city growth with an event study-distributed lag model framework, using it to trace out the dynamic response of cities following exposure to a drought. Fully accounting for these dynamic effects is critical for multiple reasons. Omitting the lagged effects of weather events could introduce bias in estimated coefficients, because of serial correlation in weather variables (Nath, Ramey, & Klenow, 2023; Newell, Prest, & Sexton, 2021). Moreover, the impact of weather shocks can be the result of temporal displacement, where effects observed in one time period are completely reversed in a later period, as highlighted by e.g. Deschenes and Moretti (2009) and Hsiang (2016). Lastly, estimating the dynamic effects allows me to both recover the cumulative impact of a drought and determine whether this impact persists in the long run.

I find that droughts have a large, negative, and persistent impact on the growth of cities. I trace the effects on city growth up to 11 years after drought exposure to cities and their hinterlands. I find that cities continue to diverge from their counterfactual growth trajectory over the course of seven years following a major drought, equivalent to a 1-in-50-year event. After this time cities remain 1.2 percent smaller compared to a drought-free counterfactual. While the effect fades over time, city growth does not recover completely and remains 0.7 percent smaller 11 years after a drought. The median city expands by 24 percent over the course of 11 years. Hence, a major drought effectively undoes around 3 percent of this total growth. While less intense droughts also have large effects after seven years, they are less persistent.

These results show that major droughts have long-term effects on the growth of cities, implying that a model that only accounts for contemporary effects is misleading. Indeed, the contemporaneous correlation between drought and city growth is weakly positive. However, these effects are overturned when accounting for the lagged effects. The results are robust to a battery of specification tests and alternative approaches to estimation and inference, including randomization-based inference based on reshuffling observed weather conditions across time, and showing that the impact is not the result of spatial spillovers from other cities.

Studies have highlighted that urbanization has progressed differently across continents and that climate change could have heterogenous effects on urbanization,[1] suggesting the impact of droughts could vary across geographical regions.[2] I run sub-sample analyses by continent. I find the largest and most persistent effects in Africa. After 11 years, the estimated effect of a major drought is 2 percent, which is substantially larger than the average effect. For Asia, comprising a large share of the sample, the results are similar to the average effects.[3] I find that the impact of droughts is statistically insignificant in North America, South America, and Oceania.

The heterogeneity across continents could suggest a larger impact on cities in poorer countries. Poorer countries have been found to be more vulnerable to disasters

[1] See, e.g., Henderson and Kriticos (2018) and Henderson and Turner (2020) for a review.

[2] Barrios, Bertinelli, and Strobl (2006) and Henderson, Storeygard, and Deichmann (2017).

[3] About two-thirds of the cities in the sample are located in Asia.

and weather shocks (Dell, Jones, & Olken, 2012; Kahn, 2005). I find that the impact of droughts is most pronounced in low and lower-middle-income countries, while the impact is statistically insignificant in middle-income countries, the richest countries in the sample. Low-income countries are more agriculturally dependent, and the agricultural sector is particularly vulnerable to weather shocks (Schlenker & Roberts, 2009). I therefore investigate the role of agriculture, as measured by the national share of employment in agriculture, dividing the sample into low, middle, and high shares. Cities in countries with the highest share of agricultural employment remain 3.5 percent smaller 11 years after a major drought. I find no impact for the cities in countries with the lowest share.

What do these results imply for the growth trajectory of cities? The median expansion for a city in poorer countries is 23 percent, a major drought undoes 10 percent of that growth, or the equivalent of a whole year. For cities in the most agricultural countries, it undoes the equivalent of almost two years of growth.

My results have two important implications. The first is that rather than inducing growth, droughts inhibit the growth of cities, which is contrary to the working hypothesis of many policymakers and researchers. This pattern holds across large parts of the sample, which together represent around 84 percent of the world's population. As the size of a city's footprint is highly correlated with its population,[4] my results also suggest that droughts decrease rather than induce city population growth. Moreover, city population growth is often driven by rural-to-urban migration (Brueckner & Lall, 2015), calling into question the often-made assumption about adverse weather shocks, such as droughts, driving rural-to-urban migration, or urbanization.

The larger and more persistent effect of drought exposure on cities in poorer and more agricultural countries is consistent with models which envision environment-migration poverty traps. These models predict that adverse weather shocks could effectively trap already resource-poor people in agriculture and further poverty since shocks worsen their liquidity constraints and reduce their ability to pay for migration costs (Cattaneo & Peri, 2016; Mayda, 2010). This chain of events could contribute to explaining the persistent gaps between rural and urban sectors in developing coun-

[4]The correlation between city population and footprint is high, ranging between 0.75 and 0.85. The results of these analyses are shown in Figure C1.1.

tries, where the higher overall amenities and income enjoyed by urban dwellers leads Henderson and Turner (2020) to question why urbanization is not happening at an even higher rate.

This paper relates to three strands of the literature. First, I build on a vast and still-growing literature that evaluates the socio-economic effects of climate change. For example, Dell, Jones, and Olken (2012), Hsiang and Jina (2014), and Nath, Ramey, and Klenow (2023) study the impact of temperature and tropical cyclones on GDP. Deschênes and Greenstone (2007) and Schlenker and Roberts (2009) examine the impact of temperature on agricultural yields, and Kahn (2005) study how natural disasters, for example floods, affect mortality.

My findings connect specifically to the literature on the links between climate change and urbanization. This literature has studied how long-run changes in the climate have affected the shares of urban population in Sub-Saharan African districts (Henderson, Storeygard, & Deichmann, 2017) or studied the effect of yearly weather variation at the country-level (Barrios, Bertinelli, & Strobl, 2006; Castells-Quintana, Krause, & McDermott, 2021). Contemporaneous work studies the effect on built-up area using global grid cells as the level of analysis (Chlouba, Mukim, & Zaveri, 2023). I add to this literature by extending the city-level evidence vastly, to five different continents.[5] Using city-level data, I study disaggregated weather shocks, which are important since there can be considerable variation in weather within larger administrative units, and even internal migration is largely local. Additionally, in contrast to the previous literature, I trace out the dynamic (and cumulative) effects of weather shocks and show that they are both persistent and different from contemporaneous correlations.

The climate and urbanization literature is closely related to the literature on climate change and migration, especially internal migration, an already vast literature. Recent reviews include Hoffmann, Šedová, and Vinke (2021) and Kaczan and Orgill-Meyer (2020). Research in this literature focusing on internal migration has largely studied single countries (Bohra-Mishra, Oppenheimer, & Hsiang, 2014; Gray

[5]Henderson, Storeygard, and Deichmann (2017) uses a city-level analysis to examine the impact of rainfall on city income, and Castells-Quintana, Krause, and McDermott (2021) limit their city sample to the largest city in each country.

& Mueller, 2012; Kleemans & Magruder, 2018), primarily using survey data. I build on this literature by showing that the effect holds across a large share of the world, and that the impact of droughts on city growth is strongest in poor and agricultural countries, results that are consistent with environment-migration poverty traps.

The paper proceeds as follows: Section 2 summarizes the data. Section 3 describes the empirical approach, section 4 presents the results, and section 5 presents the heterogeneity results. Section 6 concludes.

2 Data

I analyze the effects of drought on city growth using a city-level panel dataset, which I assemble from multiple sources. Below I summarize the main sources of data, variable construction, and how I assemble the panel data.

Countries I restrict my sample to include low- to middle-income countries, using the definition employed by the World Bank. I also make restrictions on a geographical basis and exclude all European countries from my analysis. In addition, I exclude small island states because data on drought conditions are missing, and jurisdictions that belong to high-income nations, e.g. the Dutch Caribbeans and Overseas France.[6]

City footprint I use the World Settlement Footprint (WSF) Evolution data (Marconcini et al., 2021) to calculate yearly city footprint growth measure. The data is provided in a raster format where the resolution of each pixel is approximately 30 meters around the equator. Pixels are classified as built-up or non-built, where the value of each pixel represents the year in which the pixel was classified as built-up, between 1985 and 2015. The main input of the classification is imagery from the Landsat satellites, where the spectral signature of the images is used to distinguish built-up area from non-built area. Because the availability of Landsat images varies across time, and between parts of the world, there are missing observations in the

[6]This includes departments and territories outside of mainland France, including Réunion and Guadeloupe.

data. Therefore, I exclude cities that have missing data in any year starting from 1992, the year in which the data becomes more consistent.[7]

To calculate the yearly city footprint growth, I use the WSF data to record the year in which each pixel was classified as built-up within city perimeters and subsequently calculate the sum of these pixels. To define cities, I use metropolitan area boundaries as defined by Moreno-Monroy, Schiavina, and Veneri (2021), called Functional Urban Areas (FUA). They represent estimated commuting zones of urban areas with a population of at least 50,000 in 2015. Using this definition, I avoid limiting the analysis to the political boundaries of a city, as these can both change over time and be endogenously determined. The associated data is provided by Schiavina et al. (2019). Figure **??** showcases an example of the FUA boundary and city growth in Jaipur, India.

An advantage of using these data is that the measures should be unaffected by differences in e.g. institutional quality, or definitions of what a city is, across countries.[8] In addition, the quality of the data should not be affected by the economic and political conditions in a country, which could introduce endogeneity (Hsiang & Jina, 2014).

Drought conditions I use the standardized precipitation evapotranspiration index (SPEI) (Vicente-Serrano, Begueria, & Lopez-Moreno, 2010) to construct a binary variable measuring whether a drought event has taken place in each year for cities in my sample. The weekly SPEI data is from Vicente-Serrano et al. (2022) and is provided in a raster format, with a resolution of 0.5 degrees (around 55 kilometers around the equator). The values of the index are standardized at the pixel-level with a mean of zero, and calibrated using data between 1979 and 2020. Hence, the values of the SPEI are interpreted as deviations from a long-term trend and do not capture a location's climate. For example, locations with drier climates will not have higher average SPEI values. As such, the values are comparable across different geographical

[7] This is also noted in Rentschler et al. (2023), who show that input data quality is worse before 1992.

[8] Administrative data may in some cases be manipulated, for example, because of political incentives, which has been highlighted for certain census data in Nigeria (Okafor, Adeleke, & Oparac, 2007).

regions.

The SPEI is an estimate of the climatic water balance in locations worldwide, and a combination of precipitation and temperature data is used. An advantage of the SPEI compared to other drought indices (e.g., the standardized precipitation index) is the inclusion of temperature, which has a documented effect on agricultural yields (Schlenker & Roberts, 2009).

The SPEI data from Vicente-Serrano et al. (2022) is calculated using ERA5 Climate Reanalysis data and is recommended by the authors over previous versions (SPEI-Base), which is calculated using weather station data. Using weather station data can be problematic in many settings (Auffhammer et al., 2013; Harari & Ferrara, 2018). The coverage of weather station data is sparse in many regions of the world, including Africa, which necessitates a large amount of spatial interpolation to create data that is disaggregated to a fine enough level. This procedure may introduce artificial spatial correlation (Harari & Ferrara, 2018). If weather station coverage is not consistent across time, estimates may be biased (Auffhammer et al., 2013). Schultz and Mankin (2019) show that weather station coverage is affected by civil conflict, which may introduce measurement error. The dataset they study is from the University of East Anglia's Climatic Research Unit (CRU), which is used to compute the original SPEI-Base dataset (Vicente-Serrano, Begueria, & Lopez-Moreno, 2010).

Growing season data I use data from Ortiz-Bobea et al. (2021) to define the agricultural growing seasons in the data. The authors use the normalized difference vegetation index (NDVI) data to calculate the month during which there is the most available biomass as a proxy for the growing season. I follow their approach and use the two months preceding and after this 'greenest' month as the main growing season.

Income and agricultural employment data I use data from the World Development Indicators of the World Bank (World Bank, 2023) to calculate the average agricultural employment and GDP per capita, at the country-level. I also use data from the United Nations Statistics Division (UNSD) to identify sub-regions, which can be sub-continental (for example Africa is divided into North Africa and Sub-Saharan Africa) or supra-continental (e.g. Latin America and the Caribbean). A map of the sub-regions can be found in Figure C1.2a.

The sample of countries used in this paper covers most low- to middle-income countries in the world. However, there is a large variation between these countries, for example in terms of income and levels of urbanization. Of the cities included in the sample, the median city had a population of around 139,000 in 2015 (the last year of the panel). The median growth rate of a city's horizontal footprint was around 2 percent per year. Summary statistics of the data can be found in Table D1.1. These are shown both for the whole sample and separately by continent.

The distribution of yearly city footprint growth rates is found in Figure C1.3. These are shown using the entire sample (panel A), by income categories (panel B), and by agricultural employment (panel C). As evidenced by the distributions, growth rates do not vary substantially across these categories.

3 Empirical strategy

This section details the empirical strategy of this paper. I begin by explaining the drought exposure variable. I then describe the specification in more detail. I also examine the potential pitfalls in omitting the lagged terms of drought, or weather shocks in general.

3.1 Drought exposure

I define exposure to a drought event using the SPEI index. To my knowledge, there are no established best practices in how a drought episode should be defined. While a large number of studies in economics make use of the SPEI to study various outcomes, the definition of a drought event varies.[9]

I construct the binary drought exposure variable as follows. First, I use the classifications from Wang et al. (2014) to define drought events of different intensities.

[9]Imbert and Ulyssea (2022) use crop value-weighted indices to study internal migration, Harari and Ferrara (2018) averages the values of drought indices over a dominant crop's growing season to study the impact on conflict, and Albert, Bustos, and Ponticelli (2021) use values above the mean to study internal migration and capital reallocation.

I aggregate the weekly SPEI data to the monthly level by taking the average. I then calculate a city-specific SPEI value by taking the area-weighted average of the values within a 100-kilometer buffer around each city. Lastly, I dichotomize the drought exposure measure, where the indicator D_{it} takes on the value 1 if the value in a year t in a city i is 1, 1.5, or 2 SD above the long-run city SPEI average. I vary the threshold to make sure the results are not driven by a specific cut-off and to understand whether the intensity of droughts matters for city growth. The following equation describes how drought exposure is defined:

$$D_{i,t} = \begin{cases} 1 & \text{if } S_{it} \geq \bar{S}_i + \lambda, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

where $\bar{S}_i$ is the 40-year average of the SPEI in a city and its hinterlands. λ represents the cutoffs multiplied with the standard deviation of the SPEI within a city and its hinterlands. Hence, for each year t, if the SPEI value for city i is higher than the mean and the standard deviation, I define it as a drought year. The cutoffs are used to characterize droughts as moderate (1 SD), severe (1.5 SD), or extreme (2 SD), following Wang et al. (2014). Accordingly, more intense droughts are less frequent. A 1 SD drought is, in expectation, a 1-in-6-year event, a 1.5 SD drought is a 1-in-16-year event, and a 2 SD drought is a 1-in-50-year event. Figure C1.4 shows the share of cities that are exposed to 1.5 and 2 SD droughts in any given year between 1980 and 2020.

I define the catchment area of a city using a 100-kilometer buffer (while keeping the buffer within country borders) since migration to a large extent is internal in developing countries (Jónsson, 2010), and decreases with distance (Bryan & Morten, 2019). This captures both the effect of droughts on the city as well as its rural hinterlands, where droughts are likely to be more salient because rural areas are more dependent on agriculture. Since i) the broader climate impacts literature has shown that weather shocks have larger effects on the agricultural sector (Burke, Hsiang, & Miguel, 2015; Zappalà, 2023), and ii) agriculture is the main economic sector in rural areas, I assume that the relevant temporal dimension is the agricultural growing season. Following Ortiz-Bobea et al. (2021), I approximate the main growing season for each pixel in the SPEI data using the month for which the NDVI is the highest. I then aggregate the SPEI values over the two preceding months, the highest value month, and the succeeding two months. Figures C1.5-C1.8 show the spatial variation in the growing season SPEI average, by year.

3.2 Specification

To estimate the causal effects of local droughts on city growth, I adopt an event study-distributed lag model approach,[10] where I model city growth rates (the first difference of the logarithm of city footprint) as a function of drought exposure $D_{i,t}$, out to a maximum lag length k. I estimate the following model using ordinary least squares:

$$\Delta ln(Y_{i,t}) = \sum_{l=-5}^{11} \gamma_l D_{i,t-l} + \alpha_i + \delta_{st} + \varepsilon_{d,t} \tag{2}$$

where cities are indexed by i and year by t. The parameters of interest are the coefficients γ.

This approach follows the general framework for identifying the effects of weather shocks (Deschênes & Greenstone, 2007; Hsiang, 2016). I include city fixed effects, α_i, to account for time-invariant properties of cities, which includes, for example, institutions or geography, and could lead to differences in average growth rates between cities. Since the outcome is defined in growth terms, it is the first derivative of city footprint. Hence, including city fixed effects is equivalent to controlling for a linear trend in city size.

I include subregion-by-year fixed effects, δ_{st}, to account for common nonlinear trends and shocks at the UNSD subregion-level.[11] This non-parametrically adjusts for all factors that are common across cities within a subregion by year, such as crop price levels. However, as highlighted in Deschênes and Greenstone (2007), if there exist local segmented markets, prices will not be held constant using this approach.

[10] The parameters recovered from an event-study specification and distributed lag models are identical under certain assumptions, as discussed in Schmidheiny and Siegloch (2023)

[11] Subregions are large; for example, Africa is divided into North Africa and Sub-Saharan Africa. A map of the subregions is found in Figure C1.2a.

While this could potentially be accounted for by introducing more spatially fine fixed effects, for example, country-by-year fixed effects, a downside is that finer fixed effects absorb a great deal of the variation in weather and climate (Fisher et al., 2012).

I assume that the disturbance term, ε_{it}, may exhibit both spatial correlation and autocorrelation within a city over time. To account for this possibility, I estimate standard errors clustered by ADM1-by-year and cities. ADM1 denotes the first administrative division level in each country, for example, states in India or provinces in Indonesia. A map delineating these first administrative levels is found in Figure C1.9b.

Conditional on each city's average climate and trend in drought conditions, which is absorbed by city fixed effects and the time fixed effects, the timing and intensity of drought exposure should be unpredictable and stochastic across years. Hence, the specified model allows me to assume that drought exposure, D_{it}, is plausibly exogenous and uncorrelated with other unobserved factors that could influence city growth. As an additional test to verify these assumptions, I include five years of leads in the specification. This is added as a placebo test (Hsiang, 2016), and to rule out potential anticipatory effects (Schmidheiny & Siegloch, 2023). For these assumptions to be plausible, the coefficients of the leads should be close to zero.

Interpretation of the coefficients The effect of interest in this paper is the long-run effects of droughts on the growth of cities, i.e. the dynamic (or cumulative) treatment effects.[12] The coefficients γ capture the incremental, or marginal, changes in the dynamic treatment effects. The dynamic effects can be calculated by summarizing the marginal effects over multiple years. Let β_l be the dynamic treatment effect up to k years after drought exposure:

$$\beta_k = \sum_{l=0}^{k} \gamma_l \tag{3}$$

In the interest of brevity and clarity, I only report the dynamic treatment effects,

[12]Dynamic treatment effects, cumulative treatment effects, cumulative effects, and intertemporal treatment effects are used interchangeably in the economics literature (see, e.g., de Chaisemartin and d'Haultfoeuille (2022)).

β_k throughout the paper, and do not report the estimates of γ_l.[13]

There is an explicit choice to be made in how many lags and leads should be included in the model. This choice reflects an assumption on when effects have been fully materialized. To the best of my knowledge, there is no empirical precedent or theoretical model to directly inform this choice with regard to the relationship between city growth and droughts, or other climate indicators more generally. Therefore, I choose the number of lags based on the available data, where eleven is roughly half the length of the outcome data (city growth). For the number of leads, I choose the maximum number available in the drought conditions data which does not change the effective sample for estimation. Finally, I make two additional assumptions. First, the effects are assumed to be additively separable, i.e., the effect of drought exposure in each year is independent of the effect in a previous year. Second, the treatment effects are homogenous across cohorts and years. I explore the second assumption in a robustness test and show that the results do not appear substantially different when accounting for heterogenous treatment effects.

Short vs. long-run effects The empirical modeling of the relationship between city growth and droughts is not obvious, since the functional form has little precedent in the literature. However, some insights can be gained from the climate impacts literature. The dynamic effects of climate indicators or weather shocks are important in understanding the true effects. As highlighted in Hsiang (2016), weather shocks can result in temporal displacement effects, where a shock in period t brings an event that would otherwise occur in a future time period, e.g. $t + j$, forward in time. Deschênes and Greenstone (2007) illustrate this phenomenon by showing that a substantial number of deaths seemingly related to heat would have taken place in the near future, even if a heat wave had not occurred. If that were the case in this paper, the dynamic treatment effect as expressed in Equation 3 would be equal to zero.

In this setting, it is plausible that the effects could take time to materialize. City footprint expands through the construction of new built-up areas. While this

[13] The marginal effect, γ_l can be recovered from the dynamic treatment effects, since $\beta_{k=l} = \beta_{l-1} + \gamma_l$.

construction could occur at a faster pace in developing countries, where there is often a lack of urban planning measures or the institutional capacity to enforce them, it will not be instantaneous. For this reason, lagged effects are of interest to capture the true effect of droughts on city growth.

Previous research has emphasized that climate indicators often exhibit considerable serial correlation (Nath, Ramey, & Klenow, 2023; Newell, Prest, & Sexton, 2021). Because of this, including a sufficient number of lags of the climate indicators or weather events is important in order to recover unbiased causal estimates of their effects.

4 Results

In this section, I first examine the effects of drought exposure on city growth for the entire sample and establish that the effect is indeed substantial, statistically significant, and persistent. I then explore whether there is a significant impact in geographical sub-samples, more specifically, by continent.

4.1 Main results

The results from estimating Equation 2 using the full sample of cities in the panel are presented in Figure 1 (see Table **??** for the point estimates of each lag). The figure presents the long-run effects of drought exposure on city growth relative to a drought-free city counterfactual. Each graph in the figure shows the dynamic treatment effects, β_l, from five years before (year -5) up to 11 years after drought exposure. The estimates are the result of using the three different intensities of drought exposure in separate analyses, where 1 SD represents the least intense drought measure, and 2 SD the most intense drought.

Following exposure to a drought, cities remain smaller for up to 11 years after exposure, compared to a scenario with no drought exposure. However, the result varies substantially across drought intensities. The results from estimating Equation 2 using the least intense measure of drought (1 SD) do not show either statistically or economically significant results over the course of 11 years. More extreme drought

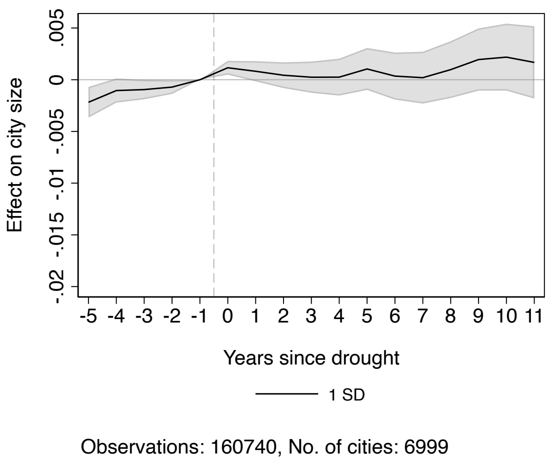

(a) 1 SD.

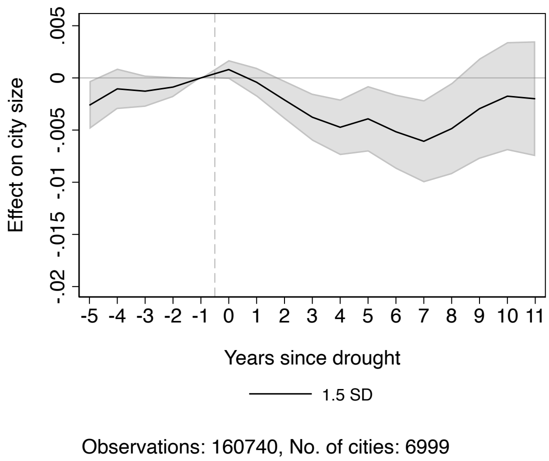

(b) 1.5 SD.

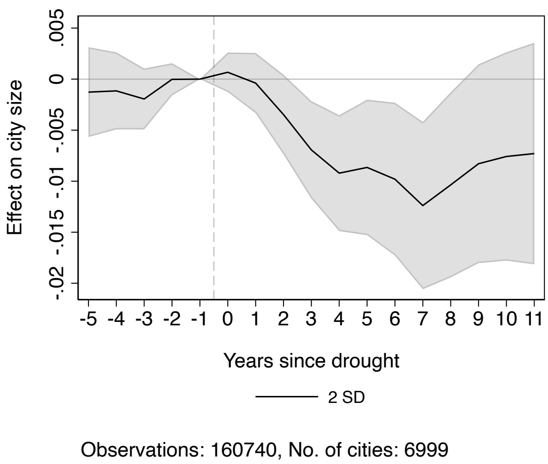

(c) 2 SD.

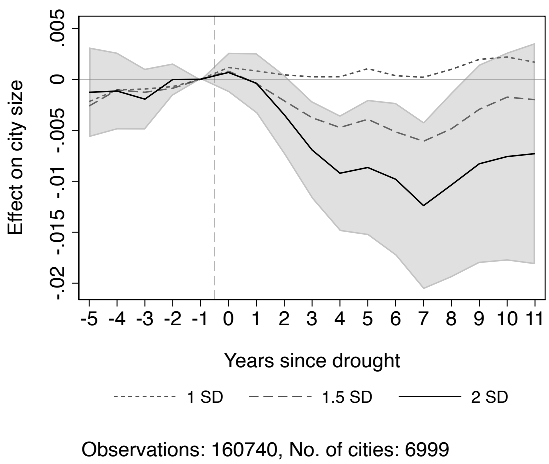

(d) 1, 1.5, 2 SD.

Figure 1: Estimates show the dynamic treatment effects of drought on city size. Panels a)-c) show the effects by drought intensity. Panel d) shows the analyses of a)-c) in the same plot. City and subregion-by-year fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Lines show the point estimates for each lag (or lead) and the shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

events, i.e., using a 1.5 or 2 SD cut-off, have statistically significant effects on city growth and are also larger in magnitude.

Focusing on the 1.5 and 2 SD measures, it is also evident that they show different patterns of persistence. While the 1.5 SD events are estimated to decrease city size by 0.5 percent 7 years after the exposure, this effect dissipates substantially over the course of the effect window and seems to revert back to trend after 11 years. The pattern is similar for 2 SD events, but the estimated effects are almost twice as large in terms of magnitude. After 7 years, the effect is 1.2 percent - this effect decreases somewhat 11 years after but does not entirely dissipate, although the confidence intervals become wider.

The fact that results differ considerably across drought intensity also suggests that there may be some degree of non-linearity in the effects of droughts on city growth. This might be explained by the in-sample probability of droughts, which differ depending on severity. Droughts which occur as often as every 5-6 years could be easier to cope with, while shocks which happen once every 15-50 years likely represent events that are both more damaging to agricultural production and unexpected.

4.2 Robustness of the main results

In order to ensure the robustness of the results described above, I perform a number of robustness tests. These include a randomization test, changing the main specification, using alternative statistical inference adjustments, and using alternative estimators.

Randomization inference test I use randomization inference tests to understand how likely it is that my estimated effects are observed by chance, in the spirit of Fisher (1935). I randomize drought conditions by shuffling years while keeping the cross-section constant, maintaining the spatial correlation across the sample intact (Heß, 2017). I then compare the point estimate β_7, which I recover from the true data to the distribution of point estimates recovered from the randomized placebo assignments. The results of this exercise are shown in Figure C1.10. While the results are weaker for the 1.5 SD droughts, with a p-value slightly above 0.1, the 2 SD is highly unlikely to be observed by chance, with a p-value far below 0.05. Hence, the results are unlikely to be spurious.

Alternative non-parametric time controls I replace the subregion-by-year fixed effects with intermediate regions-by-year, a more disaggregated definition of a region, to verify that the results are not driven by the choice of fixed effects. I show that the results are insensitive to changing the specification to include these finer-scaled fixed effects, as they remain almost identical to the results using subregion-by-year fixed effects. The results can be found in Figure C1.11a.

Influential observations I investigate whether extreme values affect the results. The results remain very similar if I winsorize city growth at the 1st and 99th percentiles. The results are found in Figure C1.11b.

Alternative standard errors In the main specification, Equation 2, I control for spatial correlation by clustering at ADM1-by-year. I show that the results remain unchanged when accounting for spatial correlation in the residuals using a different approach, by estimating the standard errors using Conley (1999). I vary the cut-off between 250 and 500 kilometers. In these specifications, I also allow for serial correlation up to 20 years using Newey and West (1986) standard errors. The resulting confidence intervals are similar to my main specification, as shown in Figure C1.12.

Alternative SPEI measure The SPEI can be calculated using different time lengths. In my main specification, I use SPEI-3, i.e. the SPEI is calculated using the water balance from the three previous months. I show that the results are not substantially different when instead choosing another time period, using SPEI-6, which uses the six previous months, and construct the drought exposure measures using the same approach as previously. This is the preferred time scale of Harari and Ferrara (2018), who argue that it captures the effect on agriculture. Since the estimates are not substantially different from using SPEI-3, this suggests that the precise time length chosen for the SPEI values does not substantially affect the results.

As described in Section 3, I calculate the drought variable over the agricultural growing season (which is approximated following Ortiz-Bobea et al. (2021)). I verify that the results hold when calculating the drought measure using the entire calendar year. The results do not change drastically - the one important difference is that the effects of a 2 SD are somewhat smaller after 11 years, at 0.5 percent (compared to 0.7 percent in the main analysis). The results can be found in Figure C1.13.

Spatial spillovers I assess whether spatial spillovers across cities drive the results by estimating a spatial lag model. The spatial lag model, incorporates the effect of drought events within pre-defined annuli from the centroid of each city, following the procedure outlined in, e.g., Hsiang and Jina (2014). City i's growth is modeled as a function of drought exposure within 100 kilometers as in Equation 2 and also includes all temporal lags of neighbors j whose centroids fall within concentric annuli (around i) with 100-kilometer increments. The results show very little evidence of spatial spillovers, as the growth in city i is not affected by droughts in neighboring cities j: the point estimates are close to zero and are statistically insignificant. The spatial lag regression model is described further in Appendix B, and the results can be found in Figure C1.14.

Lag structure I examine whether the lag length k changes the results by estimating the model in Equation 2, and varying the lag length, so that $k = \{3, ..., 10\}$. Each regression includes five leads, as in the original specification. The pattern of the result is remarkably stable across the number of included lags. The point estimate at, e.g., lag 7 is almost identical regardless of whether 7 or 11 lags are included in the regression. Hence, I conclude that the choice of the number of included lags does not meaningfully change the estimates and results. The results of this exercise are found in Figure C1.15. The results are presented omitting the confidence intervals, including only the point estimates.

Auto-regressive lag model The main specification does not explicitly control for potential serial correlation in the outcome variable. If there is a strong serial correlation in the dependent variable, the estimates may be biased, as argued by Nath, Ramey, and Klenow (2023). I amend the main specification in Equation 2 to include one or two lagged terms of the dependent variable. The main difference compared to the regressions without any autoregressive controls is that the point estimates are somewhat larger, and the results, especially for the 1.5 SD measure, dissipate less towards lag 11. As such, the main conclusions remain unchanged. The specification is detailed in Appendix B, and the results are shown in Figure C1.16.

Local projections Local projections are frequently used to estimate impulse response functions in time series (Jordà, 2005). Local projections has also been used to identify causal dynamic treatment effects (Miller, 2023), for example in Tran and Wilson (2020) to study the local effects of disasters. Colmer, Evans, and Shimshack (2023) argue that "under plausible conditions, the local projections estimator is the most consistent for identifying the dynamic effects of repeat transitory shocks". I amend Equation 2 to a local projections framework. I limit the exercise to estimating the effect of the most intense droughts, 2 SD. I find that the estimated effects are somewhat larger than those estimated using OLS (the main specification), and confidence intervals are smaller, especially toward the last lags. The specification is detailed in Appendix B, and the results are found in Figure C1.17a.

Heterogenous treatment effects The recent literature on heterogeneous treatment effects has shown that traditional estimators used in the broader event study and difference-in-differences empirical framework can be problematic if treatment effects are heterogenous. To ensure my results are robust to these concerns, I use the newly developed estimator from de Chaisemartin and d'Haultfoeuille (2022) and show that the results are in line with those recovered using my main specification. The main difference is that estimates are somewhat smaller until lag 8 – after 11 years, the point estimate is 0.7 percent, which is similar to the main results. This suggests that heterogenous treatment effects are not a main threat to identification. The results can be found in Figure C1.17b.

4.3 Results by continent

Having established that droughts have a substantial and persistent impact on cities over the course of several years, I examine whether this pattern is driven by certain regions in the sample. Urbanization patterns vary across countries in the sample. There are stark differences between, for example, South America, where the population is largely urban at the beginning of the sample period, and Africa. It is also the case that while the countries in the sample are low-income compared to the world average, there are some countries that are much richer than others.

I turn to examining whether the impact of drought is the same across the different continents in my sample, or whether specific continents drive the average effects. I run sub-sample analyses for each continent at a time, still using the specification in Equation 2. Since Africa and Asia have subregions within the continent, subregion-by-year fixed effects will be included, whereas for North America, and South America, the subregion is larger than these continents. Hence, the analysis for these countries includes a year fixed effect. I omit results from Oceania since the sample is very small.[14]

I find that the results vary substantially across continents. The results of the separate analyses for Asia and Africa are found in Figure 2. The pattern of the results for Asia is strikingly similar to those using the entire sample, which is plausibly explained by the fact that cities in Asia represent around two-thirds of the cities in the entire sample. The largest difference is that the estimates are less persistent - there is a marked reversion to zero 11 years after a drought. Additionally, the confidence intervals are also wider. This suggests that the effects of droughts are less persistent in Asia as a whole.

The differences are more pronounced when considering the sub-sample results for Africa. Here, the results are both very persistent and larger in magnitude, compared to the other continents as well as the entire sample. The impact of a major (2 SD) drought increases with every lag and only shows signs of dissipating somewhat after 8 years. In contrast, the 1.5 SD drought estimates do not dissipate at all after 11 years, and remain at an effect of 2 percent 11 years after a drought. Interestingly, the impact of a 1.5 or 2 SD drought is approximately the same in Africa.

The results for the analysis in South and North America are found in Figure C1.18. The results for South America are statistically insignificant with small point estimates throughout all lags. This suggests that drought does not have an impact on city growth in these countries. The estimates for the analysis of North American cities are somewhat more complicated. While all the point estimates are statistically insignificant, they are negative until lag 4, where they change sign after lag 6, and become positive. However, the irregular pattern and the large confidence intervals

[14]There are a total of 6 cities in Oceania. This is largely because the SPEI data is not available for many of the island states, or territories, in Oceania.

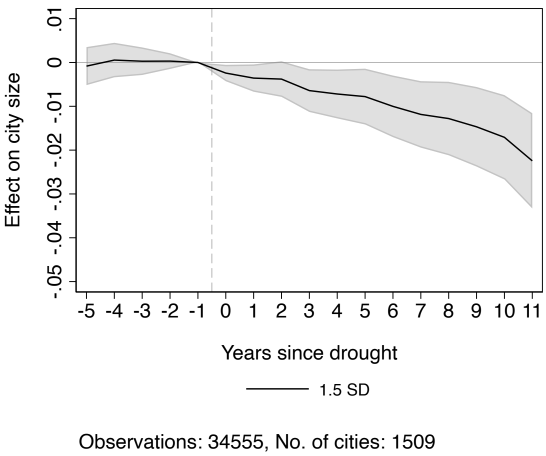

(a) Africa.

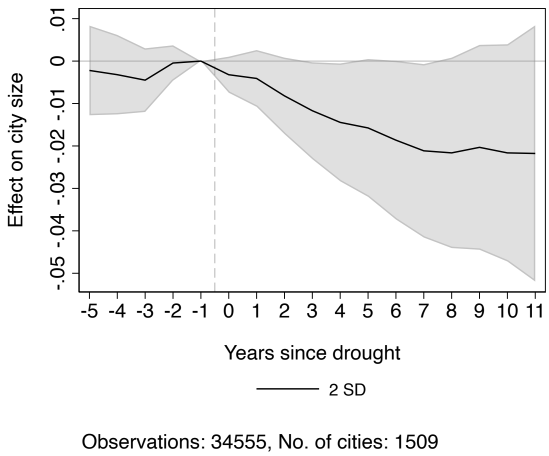

(b) Africa.

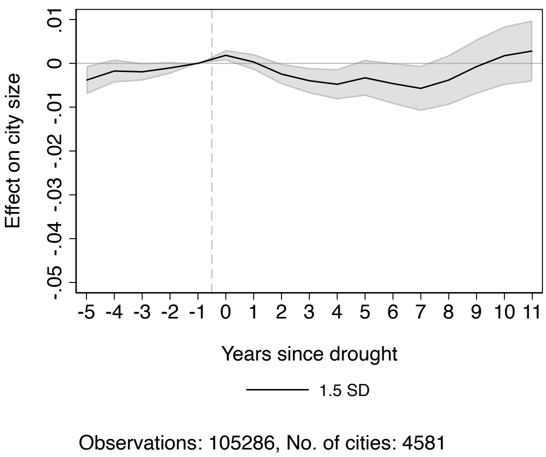

(c) Asia.

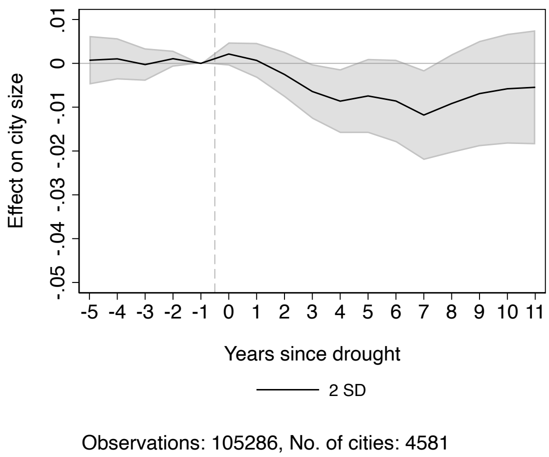

(d) Asia.

Figure 2: Estimates show the dynamic treatment effects of drought on city size. Panels a)-b) show the effects for cities in Africa, by 1.5 and 2 SD droughts. Panels c)-d) show the effects for cities in Asia, by 1.5 and 2 SD droughts. City and subregion-by-year fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Lines show the point estimates for each lag (or lead) and the shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

suggest that droughts have little impact on city growth in North American cities.

4.4 Relation to previous results

The results differ from previous research studying the effect of climate change or weather shocks on urbanization, or built-up area. Barrios, Bertinelli, and Strobl (2006) and Castells-Quintana, Krause, and McDermott (2021) find that less rainfall or increased temperatures lead to increased urbanization, Henderson, Storeygard, and Deichmann (2017) find no impact of drier conditions on urbanization in Sub-Saharan Africa, and Chlouba, Mukim, and Zaveri (2023) find that drier conditions lead to an increase in built-up area. Two main differences in this paper compared to the mentioned studies may explain the differences in the results.

First, the two first-mentioned studies use country-level data. As Henderson, Storeygard, and Deichmann (2017) points out, such data rely heavily on interpolation between years. The analysis using such data also ignores within-country variation, which is often considerable. However, the same critique can be applied to the district-level analysis in Henderson, Storeygard, and Deichmann (2017), who use census data. These are infrequently collected and are only available for a subset of countries in Sub-Saharan Africa. The definition of urbanization, and what constitutes a city, also varies across countries.

Second, these studies focus mainly on the contemporary effects, apart from Henderson, Storeygard, and Deichmann (2017). These effects may be misleading if effects take time to materialize, or if there are temporal displacement effects, as outlined in Section 3. Hence, my results are not directly comparable to the contemporaneous effects found in Barrios, Bertinelli, and Strobl (2006), Castells-Quintana, Krause, and McDermott (2021), and Chlouba, Mukim, and Zaveri (2023).

5 Heterogeneity and channels

This section discusses the results based on heterogeneity analysis by income and agricultural dependency, as well as a discussion regarding potential channels of the results.

5.1 Agriculture and income

The continent-wise analyses suggest large spatial heterogeneities in the effect of droughts on city growth. The large and persistent impact of drought in Africa, the smaller but also negative impact in Asia, and the lack of an effect in the Americas, warrants further investigation into the relationship between drought and city growth. Macroeconomic models of climatic effects have often emphasized the importance of agriculture and income (Dell, Jones, & Olken, 2012). This has been corroborated by more recent empirical studies on the relationship between, e.g., temperature and GDP, for example, Burke, Hsiang, and Miguel (2015), Nath, Ramey, and Klenow (2023), and Zappalà (2023), where the impact is larger on either agricultural GDP or the agricultural sector.

Therefore, I investigate whether agriculture and income are potential channels that could explain the differences in the impact between countries. I partition the sample by income categories, as defined by the World Bank. I do the same using the share of employment in agriculture and estimate Equation 2 separately by categories.

By country income I estimate Equation 2 separately for each income category of countries, as defined by the World Bank; low-income, lower-middle-income, and middle-income countries. The results of these analyses can be found in Figure 3.

The results for the low-income countries are presented in Figure 3a. The magnitude of the effect is larger than the effects in the entire sample, but only for the most intense drought events (2 SD). Additionally, the effect is highly persistent, remaining around 1.5 percent 11 years after drought exposure. For the medium-intensity droughts (1.5 SD), I find close to no effect.

The results for the lower-middle-income countries are presented in Figure 3b. These effects are highly significant, and large in magnitude. The effect of drought exposure results in a steady decline in city size across both the 1.5 SD and 2 SD droughts, although they show some sign of leveling off towards the end of 11 years. The magnitudes are large: at around 2 percent for the 1.5 SD measure, and around 2.5 percent for the 2 SD measure.

Finally, Figure 3c shows the results for the middle-income countries. The estimated effects are close to zero for the 1.5 SD droughts across all lags. The 2 SD

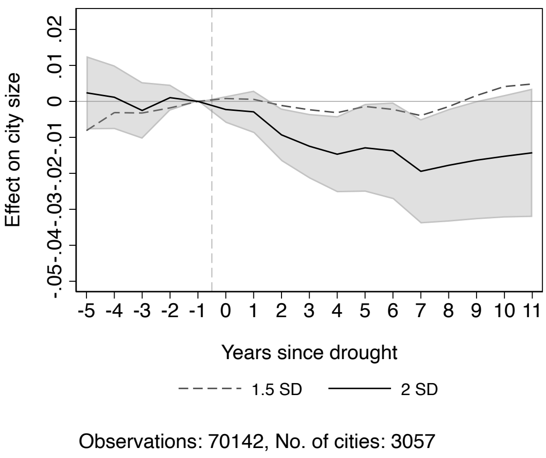

(a) Low-income countries.

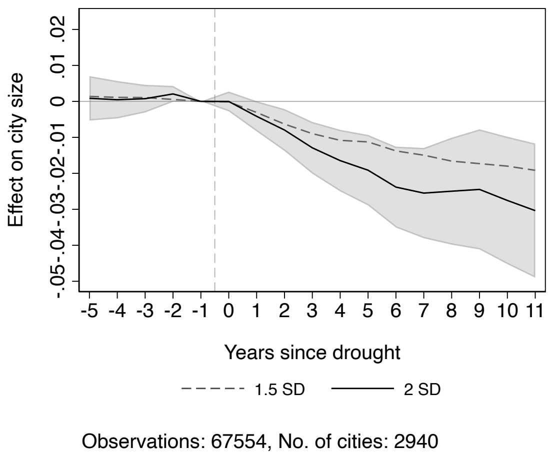

(b) Lower-middle-income countries.

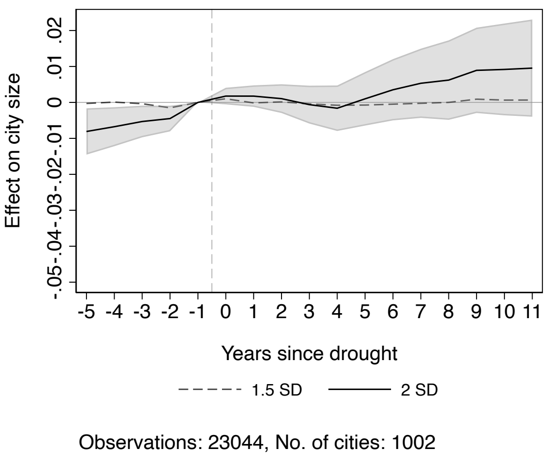

(c) Middle-income countries.

Figure 3: Estimates show the dynamic treatment effects of drought on city size. Panel a) shows the effects on cities in low-income countries, panel b) in lower-middle-income countries, and panel c) in middle-income countries. Standard errors are clustered by ADM1-by-year and city. Dashed lines show the point estimates for each lag (or lead) of a 1.5 SD drought, and solid lines the estimates for a 2 SD drought. The shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

droughts have a positive effect after 6 years. However, the results are statistically insignificant across all lags for both 1.5 and 2 SD droughts. In addition, the leads are statistically significant for the 2 SD droughts, indicating that these results may not be interpreted as a causal effect.

Altogether, my findings suggest that the cities which are most impacted by local droughts are cities in lower and lower-middle-income countries. Somewhat surprisingly, the effect is not the largest in the poorest category of countries, i.e., the low-income countries. However, this could be explained by the fact that certain natural resource-rich countries, e.g. Angola have higher levels of GDP while the general population is poor.

By share of agricultural employment In order to better understand potential driving forces behind the results, I examine whether cities in more agriculturally dependent countries are more impacted compared to cities in less agriculturally dependent countries. I proxy agricultural dependence using the agricultural share of employment in 1992.

I divide the cities into three groups based on the distribution of agricultural employment share at the country-level, creating three terciles of the distribution. I present the results from estimating Equation 2 for each of these sub-samples in Figure 4.

The largest effect is found in the third tercile, which are the cities in countries with the highest share of employment in agriculture. In fact, the effect of the most intense drought events, 2 SD, is almost 4 percent after 7 years and does not dissipate substantially after 11 years, suggesting that the effect is both large and highly persistent. Likewise, the 1.5 SD shows a steady decline in city size, resulting in a 2.2 percent decrease after 11 years.

The cities in the second most agriculturally dependent countries are also heavily impacted by drought. While the results are larger than those of the global sample, the effect for the 1.5 SD droughts is less persistent. While the effect is almost 1 percent after 6-7 years, cities do recover after 11 years. However, the 2 SD droughts do not show any sign of recovery and instead remain at 2 percent after 7 to 11 years.

The results for the cities least dependent on agriculture, the first tercile, are statistically insignificant across all lags, although they increase from lag 5 and onwards.

It should be noted that the number of cities (and observations) is unevenly distributed across sub-samples. However, despite the fact that the third tercile only has

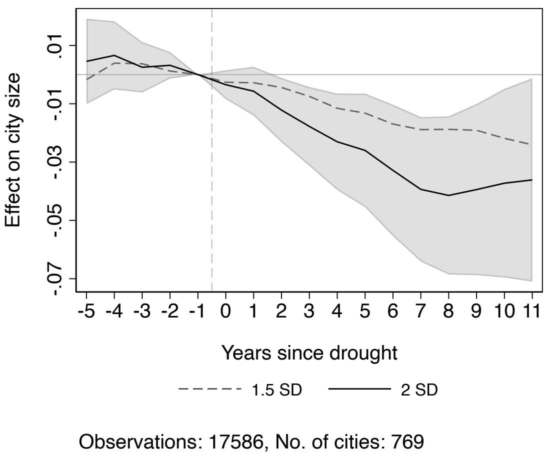

(a) Highest share, third tercile.

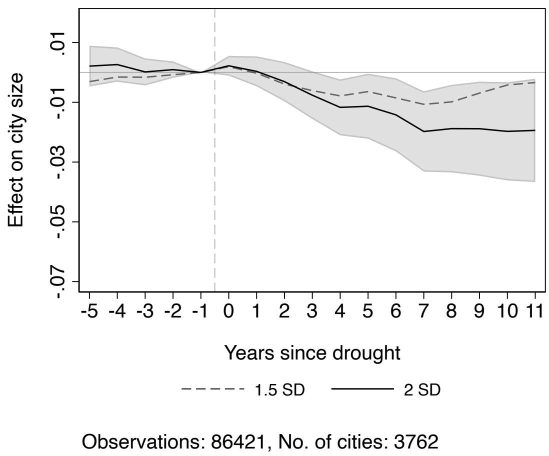

(b) Middle share, second tercile.

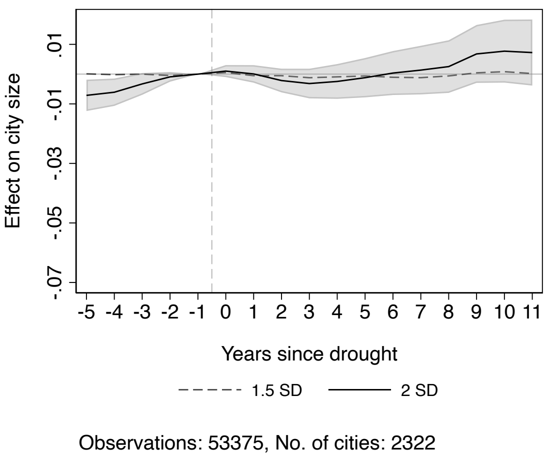

(c) Lowest share, first tercile.

Figure 4: Estimates show the dynamic treatment effects of drought on city size. Panel a) shows the effects on cities in countries with the highest share of agricultural employment (third tercile), panel b) on cities in the second tercile countries, panel c) shows the effects on cities in the first tercile countries, the lowest share. Standard errors are clustered by ADM1-by-year and city. Dashed lines show the point estimates for each lag (or lead) of a 1.5 SD drought, and solid lines the estimates for a 2 SD drought. The shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

784 cities, it represents a population from 27 countries totaling hundreds of millions.

The results suggest that both country-level income and the agricultural share of the economic sector are important channels in understanding the magnitude of the impact of drought on city growth. This is plausible, considering that agriculture is heavily impacted by adverse weather events.

5.2 Potential mechanisms

What mechanisms could explain the results of this paper? My reduced-form analysis calls for a certain degree of caution, as I am unable to pinpoint specific mechanisms with certainty. In this section, I will, however, discuss some potential mechanisms.

An important driver of urbanization, and city growth, in developing countries is rural-to-urban migration (Brueckner & Lall, 2015). I do not observe migration—there does not exist high-frequency population data at the city-level for most of the world—but my results are consistent with migration as a plausible explanation behind my findings. City footprint growth, as measured by built-up area, is highly correlated with measures of population. First, I show this using data of city populations in Sub-Saharan collected by Jedwab and Storeygard (2022). In panel A, Figure C1.1, I plot the log values of built-up area against the log values of city population. The estimated correlation between the two measures is high, at 0.76. Second, I use data from Schiavina, Freire, and MacManus (2019). These data are estimated, on a global scale, using a combination of built-up area and administrative data. The correlation is high, 0.85, which may partially be a result of the data being constructed using built-up area as an input. This plot is found in panel B, Figure C1.1.

The strong correlation between population and built-up area may indicate that a negative effect on city footprint could be associated with a negative effect on city population. As city population growth is driven partly by rural-to-urban migration, this could suggest that migration toward cities is decreasing. An alternative explanation could be that the data mainly captures built-up area in the form of non-residential buildings. I argue that this explanation is unlikely for two reasons. First, built-up area is often used in estimating population in the absence of other reliable data (Freire et al., 2016). Second, the correlation between built-up area and population is high, as shown previously in this paper.

If migration is negatively affected by droughts, it would be consistent with a model where households face credit and liquidity constraints which makes them less likely to migrate. People in developing countries often face substantial risks and liquidity constraints which may impact their decision to migrate, especially if moving requires a substantial up-front cost (Kleemans, 2023). Liquidity constraints have been shown to prevent migration in, for example, Bangladesh (Bryan, Chowdhury, & Mobarak, 2014), Indonesia (Bazzi, 2017), and China (Cai, 2020). In addition, my results echo the model in Cattaneo and Peri (2016), where liquidity constraints are more binding in poorer countries compared to richer countries.

The persistence of the results may be explained by household responses to shocks. For example, people who choose not to migrate in one period may invest more in agriculture, and forgo migration in subsequent periods. Nath, Ramey, and Klenow (2023) finds that increasing temperatures in hotter and low-income areas may induce additional specialization in agriculture, as a means to meet food demand under productivity losses. Evidence from Peru shows a similar pattern, where farmers respond to heat shocks by increasing their agricultural land and changing their product mix, rather than migrating (Aragón, Oteiza, & Rud, 2021).

I emphasize that there are other potential mechanisms that could explain the results. These include direct effects on mortality or fertility in cities, which could cause their populations to decline, and therefore expand less. While droughts are mostly damaging to the agricultural sector, cities may still be affected. If droughts decrease economic growth, it could also lead to less urban expansion. However, this does not necessarily preclude a migration channel, as rural-to-urban migration may be more appealing during periods of urban economic growth.

6 Conclusion

This paper studies the effects of drought on the growth of cities. I find that drought exposure leads to a large, negative, and persistent effect on the growth of cities. Following a major drought, city growth is lower than it would have been compared to a drought-free counterfactual. The effect lasts up to 11 years, suggesting that the effects are persistent. The impact is largest in Africa and Asia, which host several countries and regions undergoing rapid urbanization. I find that the most important predictor of where drought exposure will have a large effect is agricultural dependency, as measured by the national share of employment in agriculture.

Together, these results suggest that adverse weather shocks, such as droughts, can have an impact on the spatial allocation of people. The size of the effect varies across sub-samples, but the cities in countries with a large agricultural sector remain 4 to 5 percent smaller after exposure to a local drought, compared to a counterfactual of not having experienced a drought episode.

I utilize and build a novel data set to construct measures of the annual growth of cities at the city-level. These data allow me to overcome some of the data constraints and limitations facing previous research, including Barrios, Bertinelli, and Strobl (2006), Castells-Quintana, Krause, and McDermott (2021), and Henderson, Storeygard, and Deichmann (2017), who find positive or null effects. I show that the contemporaneous effects may be misleading, as they are either close to null effects or positive, suggesting that delayed effects are important when studying the impact of weather shocks on cities, and potentially on migration as well.

I am unable to speak directly on mechanisms in this paper. I show that city growth, as measured by the growth in horizontal built-up area, is correlated with population. Hence, a potential explanation of the results is that droughts in city hinterlands lead to less rural-to-urban migration. However, there are other potential explanations that I cannot rule out in this paper. These could include a direct economic effect on cities, which could lead to less construction of new urban areas. Isolating the mechanisms would require a more structurally oriented analysis, which could be an avenue for future work.

In the broader discussion of the effects of climate change, a prominent view argues that climate change will lead to inevitable urbanization, putting pressure on already strained cities. My results suggest that this is less of a concern for a large share of cities, which are hindered rather than induced to grow. However, I note that these results do not speak directly to other climate hazards that many people in developing countries face, for example, cyclones and floods. These events could have a different effect compared to that of droughts.

My findings have some policy implications. As climate change is expected to increase the number and intensity of droughts, more could be done to help the rural and agricultural sector to cope with droughts. This might include technologies such as irrigation, or cultivating other types of crops, which are the two main margins of adaptation. The results also suggest that it may not be adverse weather shocks to the agricultural sector which will increase internal climate migration.

References

Albert, C., Bustos, P., & Ponticelli, J. (2021). The effects of climate change on labor and capital reallocation. National Bureau of Economic Research.

Aragón, F. M., Oteiza, F., & Rud, J. P. (2021). Climate change and agriculture: Subsistence farmers' response to extreme heat. *American Economic Journal: Economic Policy*, *13*(1), 1–35.

Auffhammer, M., Hsiang, S. M., Schlenker, W., & Sobel, A. (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*.

Barrios, S., Bertinelli, L., & Strobl, E. (2006). Climatic change and rural–urban migration: The case of sub-Saharan Africa. *Journal of Urban Economics*, *60*(3), 357–371.

Bazzi, S. (2017). Wealth heterogeneity and the income elasticity of migration. *American Economic Journal: Applied Economics*, *9*(2), 219–255.

Bohra-Mishra, P., Oppenheimer, M., & Hsiang, S. M. (2014). Nonlinear permanent migration response to climatic variations but minimal response to disasters. *Proceedings of the National Academy of Sciences*, *111*(27), 9780–9785.

Brueckner, J. K., & Lall, S. V. (2015). Cities in developing countries: Fueled by rural–urban migration, lacking in tenure security, and short of affordable housing. *Handbook of Regional and Urban Economics*, *5*, 1399–1455.

Bryan, G., Chowdhury, S., & Mobarak, A. M. (2014). Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh. *Econometrica*, *82*(5), 1671–1748.

Bryan, G., & Morten, M. (2019). The aggregate productivity effects of internal migration: Evidence from Indonesia. *Journal of Political Economy*, *127*(5), 2229–2268.

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, *527*(7577), 235–239.

Cai, S. (2020). Migration under liquidity constraints: Evidence from randomized credit access in china. *Journal of Development Economics*, *142*, 102247.

Castells-Quintana, D., Krause, M., & McDermott, T. K. (2021). The urbanising force of global warming: The role of climate change in the spatial distribution of population. *Journal of Economic Geography*, *21*(4), 531–556.

Cattaneo, C., & Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, *122*, 127–146.

Chlouba, V., Mukim, M., & Zaveri, E. (2023). After Big Droughts Come Big Cities: Does Drought Drive Urbanization? World Bank Policy Research Working Paper.

Clement, V., Rigaud, K. K., De Sherbinin, A., Jones, B., Adamo, S., Schewe, J., Sadiq, N., & Shabahat, E. (2021). Groundswell Part 2.

Colmer, J., Evans, M. F., & Shimshack, J. (2023). Environmental citizen complaints. CEP Discussion Papers.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, *92*(1), 1–45.

de Chaisemartin, C., & d'Haultfoeuille, X. (2022). Difference-in-differences estimators of intertemporal treatment effects. National Bureau of Economic Research.

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.

Deschenes, O., & Moretti, E. (2009). Extreme weather events, mortality, and migration. *The Review of Economics and Statistics*, *91*(4), 659–681.

Deschênes, O., & Greenstone, M. (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review*, *97*(1), 354–385.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., & Schlenker, W. (2012). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. *American Economic Review*, *102*(7), 3749–3760.

Fisher, R. A. (1935). *The Design of Experiments*. Edinburgh: Oliver; Boyd, Ltd.

Freire, S., MacManus, K., Pesaresi, M., Doxsey-Whitfield, E., & Mills, J. (2016). Development of new open and free multi-temporal global population grids at 250 m resolution. *Population*, *250*.

Gray, C., & Mueller, V. (2012). Drought and population mobility in rural Ethiopia. *World Development*, *40*(1), 134–145.

Harari, M., & Ferrara, E. L. (2018). Conflict, climate, and cells: A disaggregated analysis. *Review of Economics and Statistics*, *100*(4), 594–608.

Henderson, J. V., & Kriticos, S. (2018). The development of the African system of cities. *Annual Review of Economics*, *10*, 287–314.

Henderson, J. V., Storeygard, A., & Deichmann, U. (2017). Has climate change driven urbanization in Africa? *Journal of Development Economics*, *124*, 60–82.

Henderson, J. V., & Turner, M. A. (2020). Urbanization in the developing world: Too early or too slow? *Journal of Economic Perspectives*, *34*(3), 150–173.

Heß, S. (2017). Randomization inference with Stata: A guide and software. *The Stata Journal*, *17*(3), 630–651.

Hoffmann, R., Šedová, B., & Vinke, K. (2021). Improving the evidence base: A methodological review of the quantitative climate migration literature. *Global Environmental Change*, *71*, 102367.

Hsiang, S. (2016). Climate econometrics. *Annual Review of Resource Economics*, *8*, 43–75.

Hsiang, S. M., & Jina, A. S. (2014). The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones. National Bureau of Economic Research.

Imbert, C., & Ulyssea, G. (2022). Firm dynamics and the labor market effects of rural-urban migration.

IPCC. (2022). *Climate change 2022: Impacts, adaptation and vulnerability.* Cambridge, UK; New York, USA: Cambridge University Press,

Jedwab, R., & Storeygard, A. (2022). The average and heterogeneous effects of transportation investments: Evidence from Sub-Saharan Africa 1960–2010. *Journal of the European Economic Association*, *20*(1), 1–38.

Jónsson, G. (2010). The environmental factor in migration dynamics - a review of African case studies.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182.

Kaczan, D. J., & Orgill-Meyer, J. (2020). The impact of climate change on migration: A synthesis of recent empirical insights. *Climatic Change*, *158*(3-4), 281–300.

Kahn, M. E. (2005). The death toll from natural disasters: The role of income, geography, and institutions. *Review of Economics and Statistics*, *87*(2), 271–284.

Kleemans, M. (2023). Migration choice under risk and liquidity constraints.

Kleemans, M., & Magruder, J. (2018). Labour market responses to immigration: Evidence from internal migration driven by weather shocks. *The Economic Journal*, *128*(613), 2032–2065.

Marconcini, M., Metz-Marconcini, A., Esch, T., & Gorelick, N. (2021). Understanding current trends in global urbanisation - the World Settlement Footprint suite. *GI_Forum*, *9*(1), 33–38.

Mayda, A. M. (2010). International migration: A panel data analysis of the determinants of bilateral flows. *Journal of Population Economics*, *23*(4).

Miller, D. L. (2023). An introductory guide to event study models. *Journal of Economic Perspectives*, *37*(2), 203–230.

Moreno-Monroy, A. I., Schiavina, M., & Veneri, P. (2021). Metropolitan areas in the world. Delineation and population trends. *Journal of Urban Economics*, *125*, 103242.

Nath, I. B., Ramey, V. A., & Klenow, P. J. (2023). How much will global warming cool global growth? Working Paper.

Newell, R. G., Prest, B. C., & Sexton, S. E. (2021). The GDP-temperature relationship: implications for climate change damages. *Journal of Environmental Economics and Management*, *108*, 102445.

Newey, W. K., & West, K. D. (1986). A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix.

Okafor, R., Adeleke, I., & Oparac, A. (2007). An appraisal of the conduct and provisional results of the Nigerian Population and Housing Census of 2006. *Proceedings of American Statistical Association: Survey Research Methods Section*, 2199–2205.

Ortiz-Bobea, A., Ault, T. R., Carrillo, C. M., Chambers, R. G., & Lobell, D. B. (2021). Anthropogenic climate change has slowed global agricultural productivity growth. *Nature Climate Change*, *11*(4), 306–312.

Rentschler, J., Avner, P., Marconcini, M., Su, R., Strano, E., Vousdoukas, M., & Hallegatte, S. (2023). Global evidence of rapid urban growth in flood zones since 1985. *Nature*, *622*(7981), 87–92.

Schiavina, M., Freire, S., & MacManus, K. (2019). GHS population grid multitemporal (1975, 1990, 2000, 2015) R2019A. *European Commission, Joint Research Centre (JRC)*, *10*.

Schiavina, M., Moreno-Monroy, A., Maffenini, L., Veneri, P., et al. (2019). GHSL-OECD Functional Urban Areas. Publications Office of the European Union.

Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences*, *106*(37), 15594–15598.

Schmidheiny, K., & Siegloch, S. (2023). On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics*, *38*(5), 695–713.

Schultz, K. A., & Mankin, J. S. (2019). Is temperature exogenous? The impact of civil conflict on the instrumental climate record in Sub-Saharan Africa. *American Journal of Political Science*, *63*(4), 723–739.

Tran, B. R., & Wilson, D. J. (2020). The local economic impact of natural disasters. Federal Reserve Bank of San Francisco.

Vicente-Serrano, S. M., Begueria, S., & Lopez-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index. *Journal of Climate*, *23*(7), 1696–1718.

Vicente-Serrano, S. M., Domínguez-Castro, F., Reig, F., Tomas-Burguera, M., Peña-Angulo, D., Latorre, B., Beguería, S., Rabanaque, I., Noguera, I., Lorenzo-Lacruz, J., et al. (2022). A global drought monitoring system and dataset based on ERA5 reanalysis: A focus on crop-growing regions. *Geoscience Data Journal*.

Wang, Q., Wu, J., Lei, T., He, B., Wu, Z., Liu, M., Mo, X., Geng, G., Li, X., Zhou, H., et al. (2014). Temporal-spatial characteristics of severe drought events and their impact on agriculture on a global scale. *Quaternary International*, *349*, 10–21.

World Bank. (2023). *World Development Indicators 2023*.

Zappalà, G. (2023). Sectoral impact and propagation of weather shocks. Working Paper.

Appendices

A List of countries in sample

Afghanistan, Algeria, Angola, Argentina, Armenia, Azerbaijan, Bangladesh, Belize, Benin, Bolivia, Botswana, Brazil, Burkina Faso, Burundi, Cambodia, Cameroon, Central African Republic, Chad, Chile, China, Colombia, Costa Rica, Cote d'Ivoire, Cuba, Democratic Republic of the Congo, Djibouti, Dominican Republic, Ecuador, Egypt, El Salvador, Equatorial Guinea, Eritrea, Ethiopia, Gabon, Gambia, Georgia, Ghana, Guatemala, Guinea, Guinea Bissau, Guyana, Haiti, Honduras, India, Indonesia, Iran, Iraq, Jamaica, Jordan, Kazakhstan, Kenya, Kyrgyzstan, Laos, Lebanon, Lesotho, Liberia, Libya, Madagascar, Malawi, Malaysia, Mali, Mauritania, Mexico, Mongolia, Morocco, Mozambique, Myanmar, Namibia, Nepal, Nicaragua, Niger, Nigeria, North Korea, Pakistan, Palestine, Panama, Papua New Guinea, Paraguay, Peru, Philippines, Republic of Congo, Rwanda, Senegal, Sierra Leone, Somalia, South Africa, South Sudan, Sri Lanka, Sudan, Suriname, Swaziland, Syria, Tajikistan, Tanzania, Thailand, Timor Leste, Togo, Tunisia, Turkey, Turkmenistan, Uganda, Uruguay, Uzbekistan, Venezuela, Vietnam, Western Sahara, Yemen, Zambia, Zimbabwe.

B Robustness specifications

Autoregressive distributed lag model I amend the main specification, Equation 2, to include a lag of the outcome variable, which results in the following equation:

$$\Delta ln(Y_{i,t}) = \Delta ln(Y_{i,t-1}) + \sum_{l=-5}^{11} \gamma_l D_{i,t-l} + \alpha_i + \delta_{st} + \varepsilon_{d,t} \tag{4}$$

which is the AR(1) model. I augment the model further by including two lags:

$$\Delta ln(Y_{i,t}) = \Delta ln(Y_{i,t-1}) + \Delta ln(Y_{i,t-2}) + \sum_{l=-5}^{11} \gamma_l D_{i,t-l} + \alpha_i + \delta_{st} + \varepsilon_{d,t} \tag{5}$$

which results in an AR(2) model.

Local projections I amend Equation 2, following e.g. Colmer, Evans, and Shimshack (2023) and Tran and Wilson (2020), and estimate the effect of droughts on city growth for each time horizon h:

$$\Delta ln(Y_{i,t+h}) = \beta^h D_{i,t} + \alpha_i + \delta_{st} + \varepsilon_{d,t} \tag{6}$$

where $\Delta ln(Y_{i,t+h})$ is the change in city growth in the year $t+h$, relative to the year prior. The contemporaneous effect is captured by $\beta^h = 0$ and the effect after 11 years is captured by $\beta^h = 11$. These parameters all represent the dynamic treatment effects.

Spatial lag model I amend Equation 2 to include both spatial and temporal lags, following Hsiang (2016) and Hsiang and Jina (2014):

$$\Delta ln(Y_{i,t}) = \sum_{l=0}^{11} \sum_{\pi=0}^{11} \{D_{[j|dist(i,j)=\pi],t-1}\gamma_{l\pi}\} + \alpha_i + \delta_{st} + \varepsilon_{d,t} \tag{7}$$

where $D_{[j|dist(i,j)=\pi],t-1}\gamma_{l\pi}$ is the average drought exposure of all neighboring cities j, within a distance π from city i at time $t-l$, where i is the city-hinterland where the outcome is observed. $dist(i,j)$ is the distance from i to j.

C Additional figures

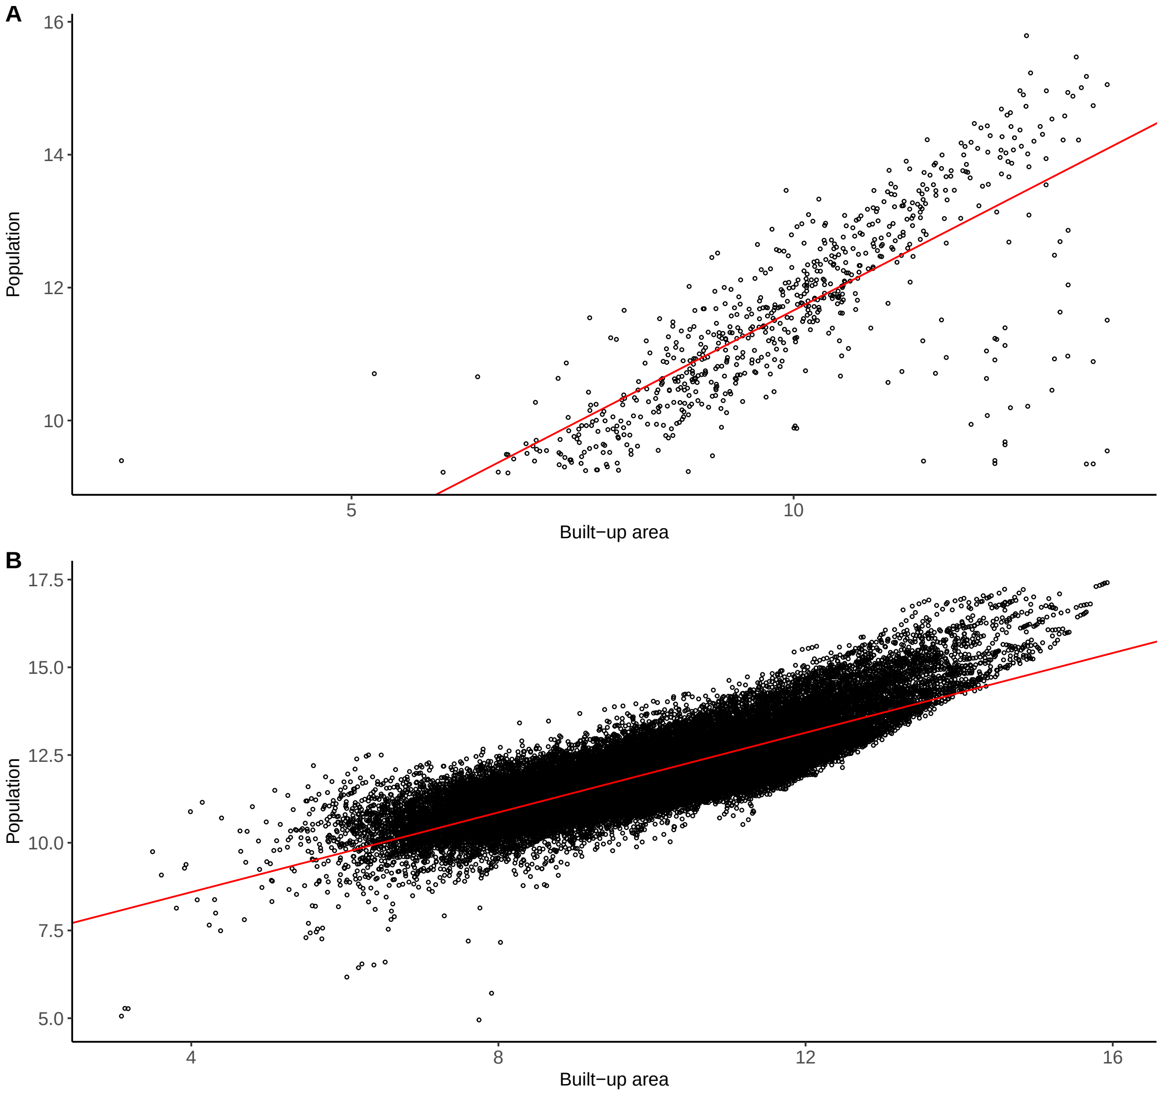

Figure C1.1: This figure shows the correlation between city population measures and the data on city footprint used to construct annual city growth, WSF (Marconcini et al., 2021). I plot the logarithm of the measures of city footprint (x-axis) against the logarithm two different sources of city population (y-axis). The solid red line show the estimates of a regression of the aforementioned measures. Panel A shows the results from comparing data from Jedwab and Storeygard (2022), for the years 1990, 2000, and 2010 with WSF data from the same years. Panel B shows the results from GHS Population (Schiavina, Freire, & MacManus, 2019), for the years 1990, 2000, and 2015.

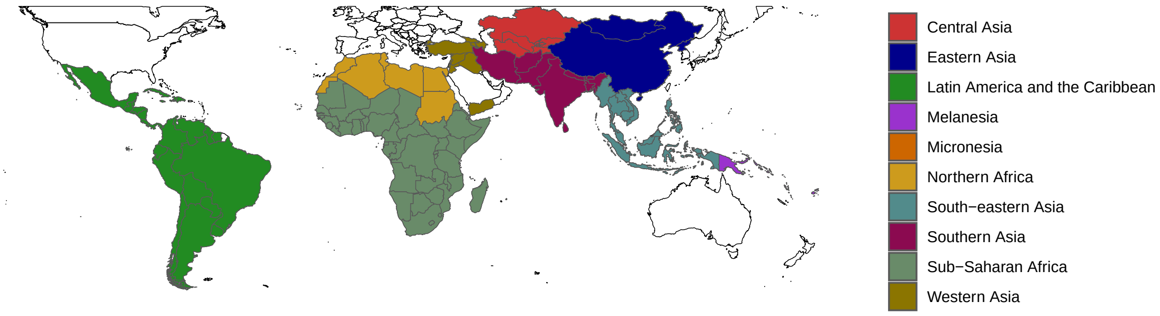

(a)

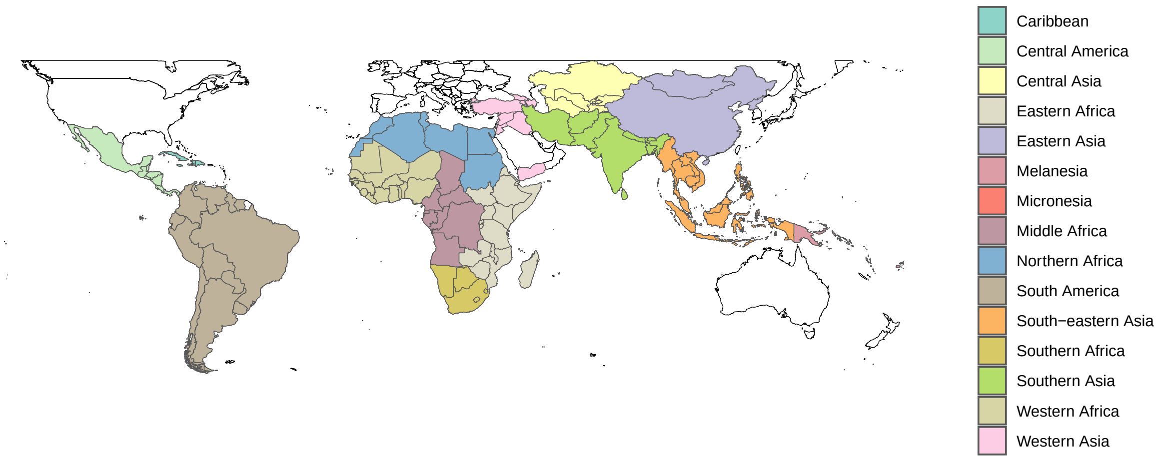

(b)

Figure C1.2: Panel a) shows the delineation of the UNSD subregions for the countries in the sample. Panel b) shows the delineation of the UNSD intermediate regions for the countries in the sample.

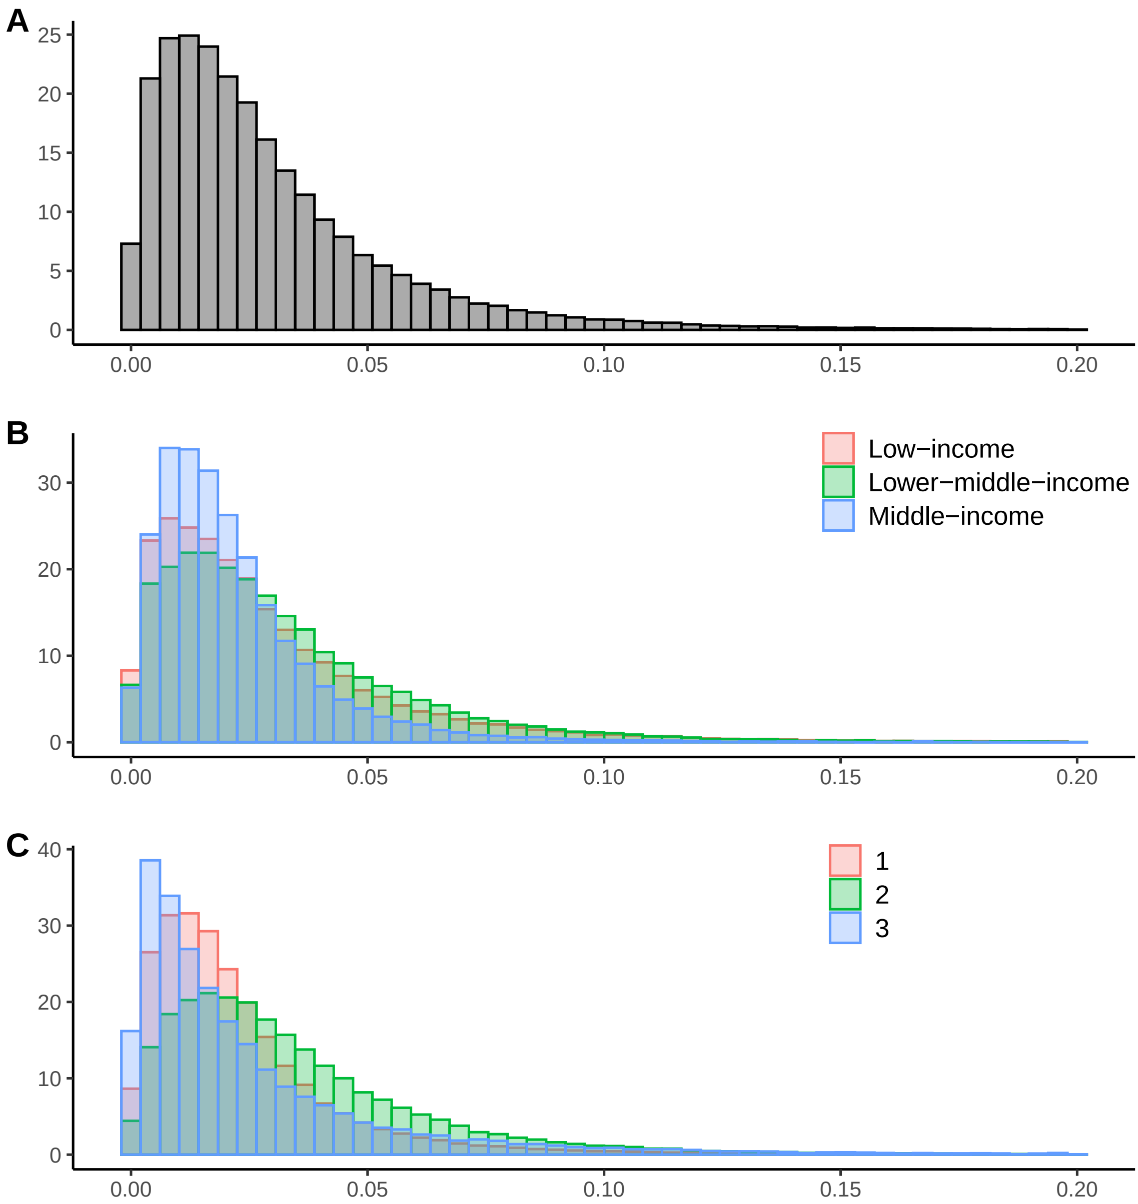

Figure C1.3: The figures show the distribution of city footprint growth rates. Panel A shows the distribution for all observations in the data. Panel B shows the superimposed distributions by income category. Panel C shows the superimposed distributions by agricultural employment, where 1 is the first tercile (lowest share of agricultural employment), and 3 is the third tercile (highest share).

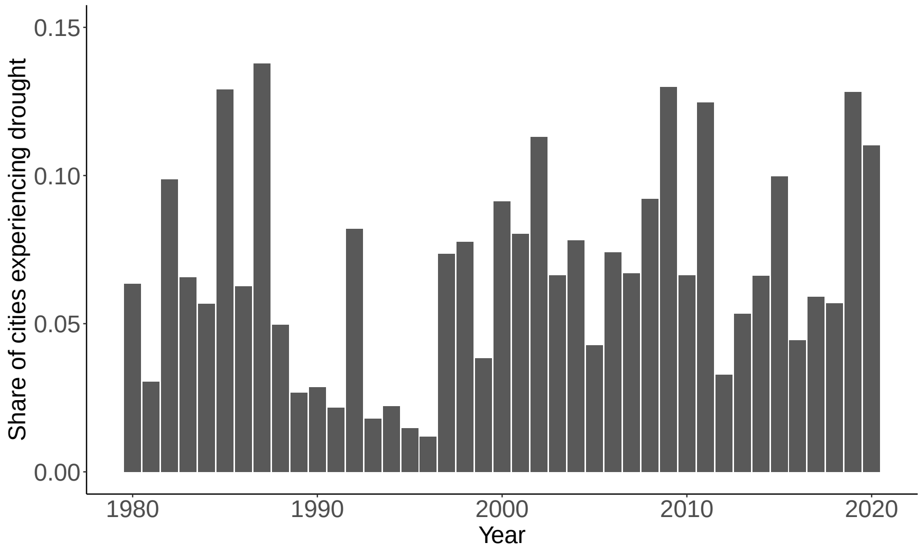

(a)

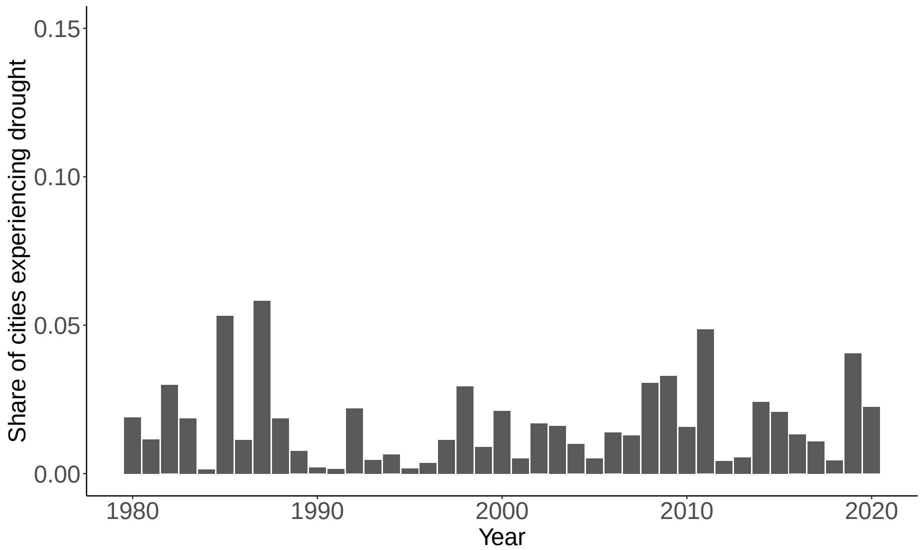

(b)

Figure C1.4: Panel a) shows the share of cities that experience a 1.5 SD drought in a given year. Panel b) shows the share of cities that experience a 2 SD drought in a given year.

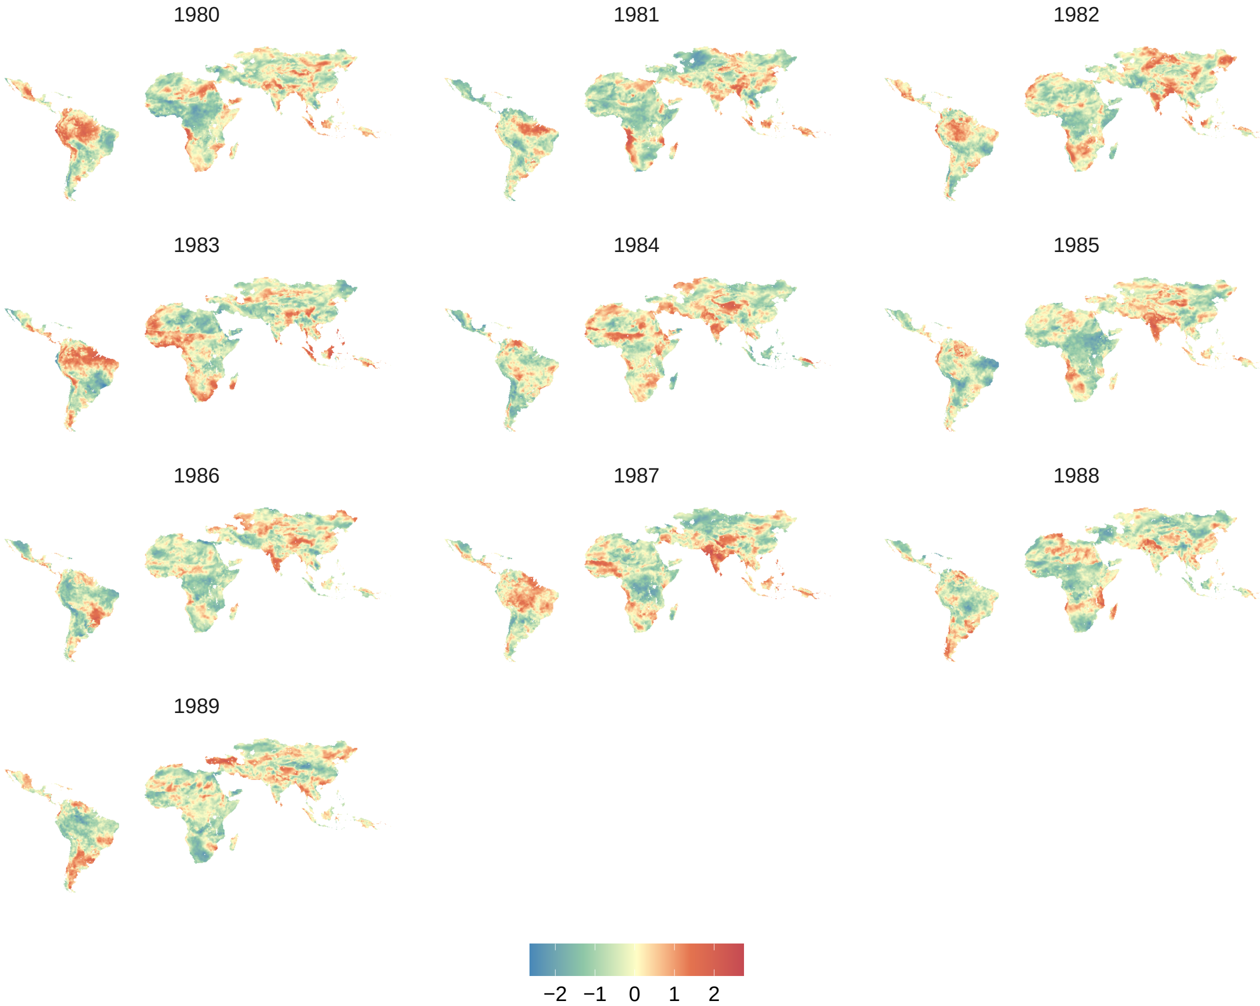

Figure C1.5: Each sub-figure shows the distribution of the yearly average pixel-level SPEI values across the sample, calculated using data from (Vicente-Serrano et al., 2022), over the main growing season, between 1980 and 1989. Only countries in the sample are displayed. Higher values indicate drier conditions relative to a long term average, and lower values indicate wetter conditions.

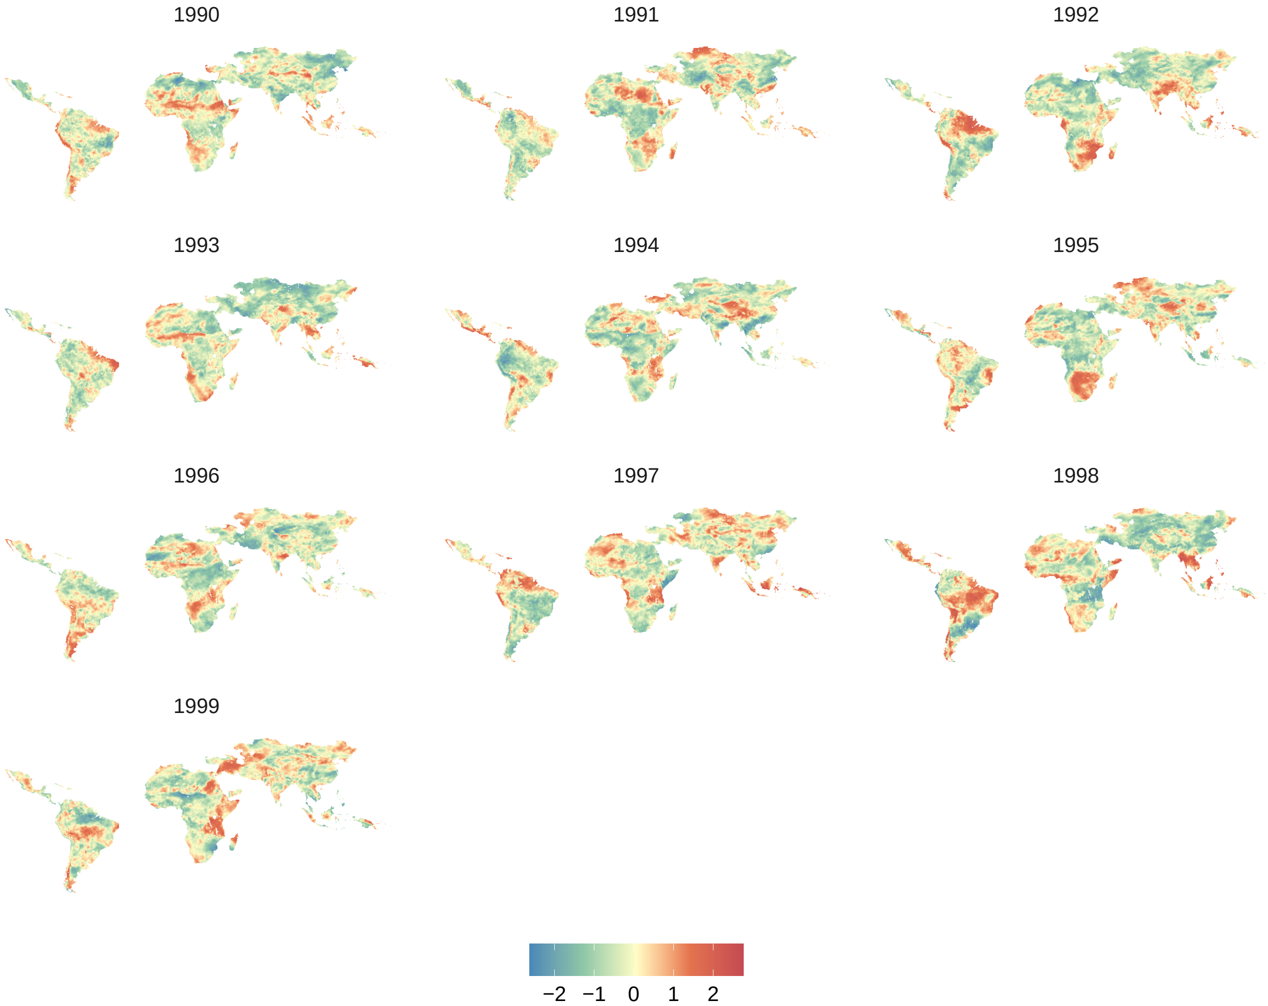

Figure C1.6: Each sub-figure shows the distribution of the yearly average pixel-level SPEI values across the sample, calculated using data from (Vicente-Serrano et al., 2022), over the main growing season, between 1990 and 1999. Only countries in the sample are displayed. Higher values indicate drier conditions relative to a long term average, and lower values indicate wetter conditions.

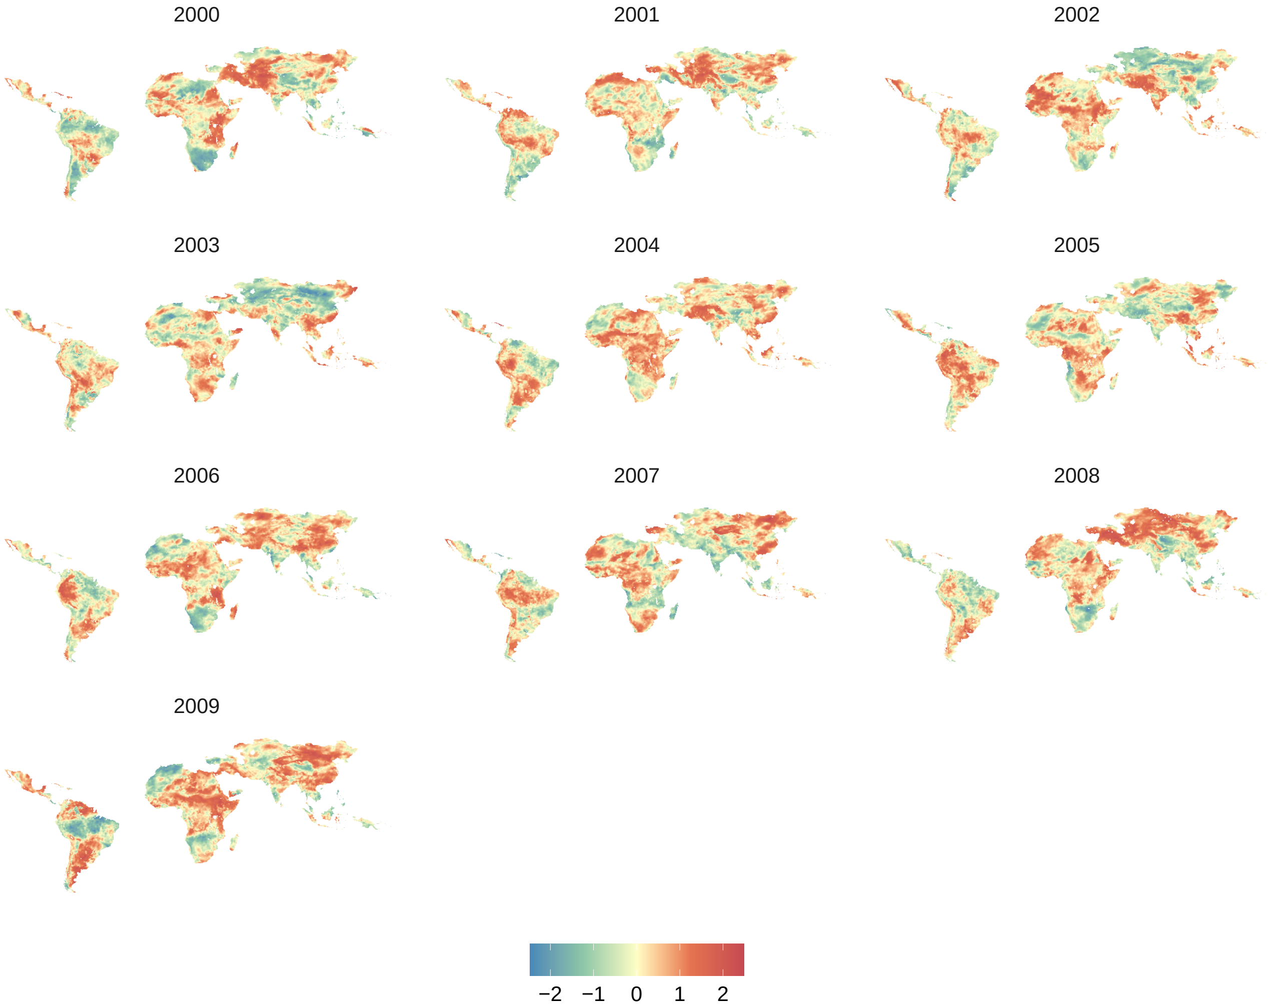

Figure C1.7: Each sub-figure shows the distribution of the yearly average pixel-level SPEI values across the sample, calculated using data from (Vicente-Serrano et al., 2022), over the main growing season, between 2000 and 2009. Only countries in the sample are displayed. Higher values indicate drier conditions relative to a long term average, and lower values indicate wetter conditions.

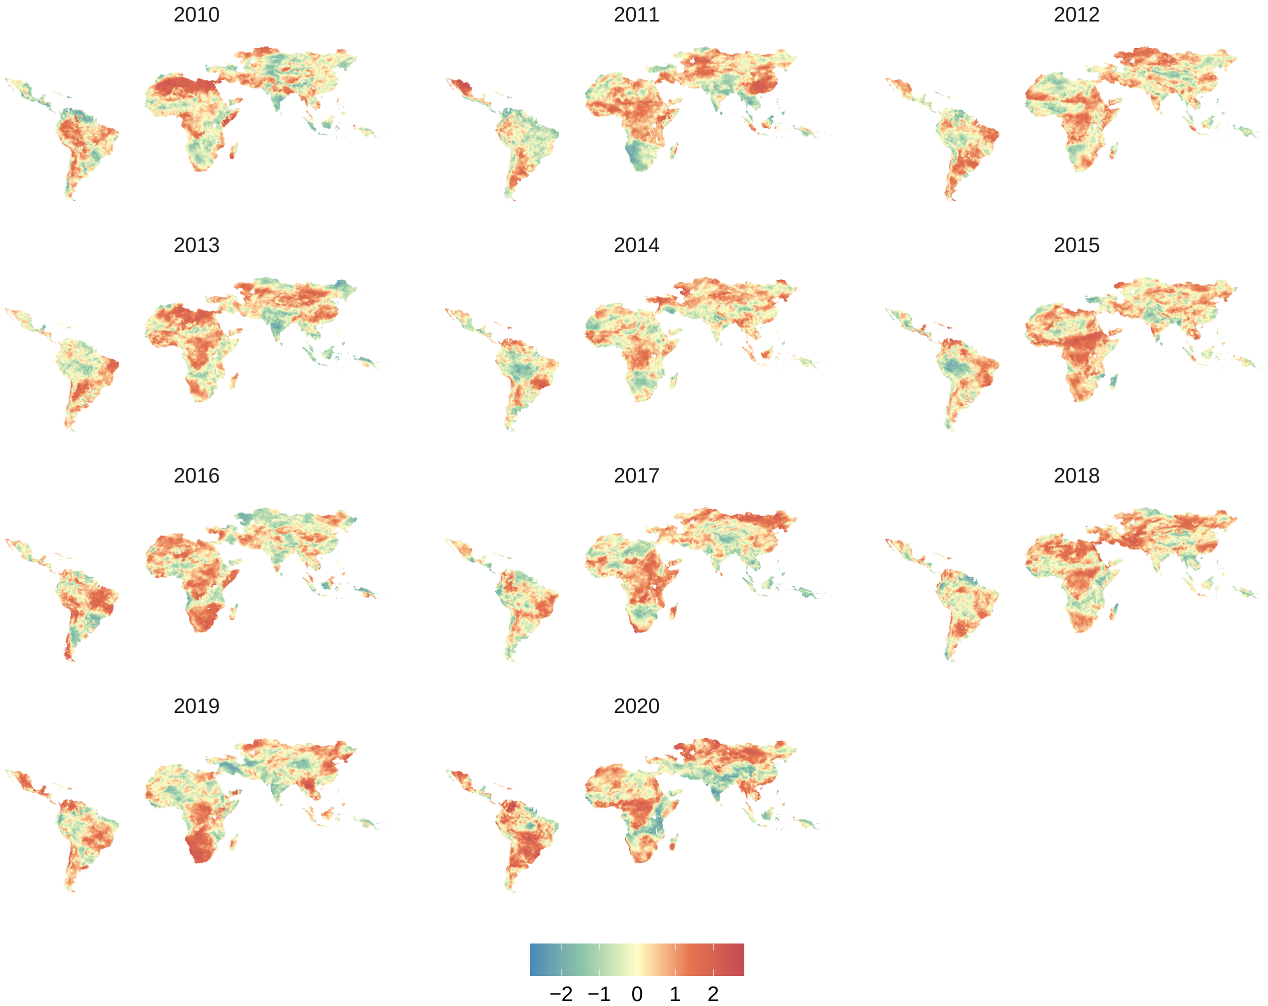

Figure C1.8: Each sub-figure shows the distribution of the yearly average pixel-level SPEI values across the sample, calculated using data from (Vicente-Serrano et al., 2022), over the main growing season, between 2010 and 2020. Only countries in the sample are displayed. Higher values indicate drier conditions relative to a long term average, and lower values indicate wetter conditions.

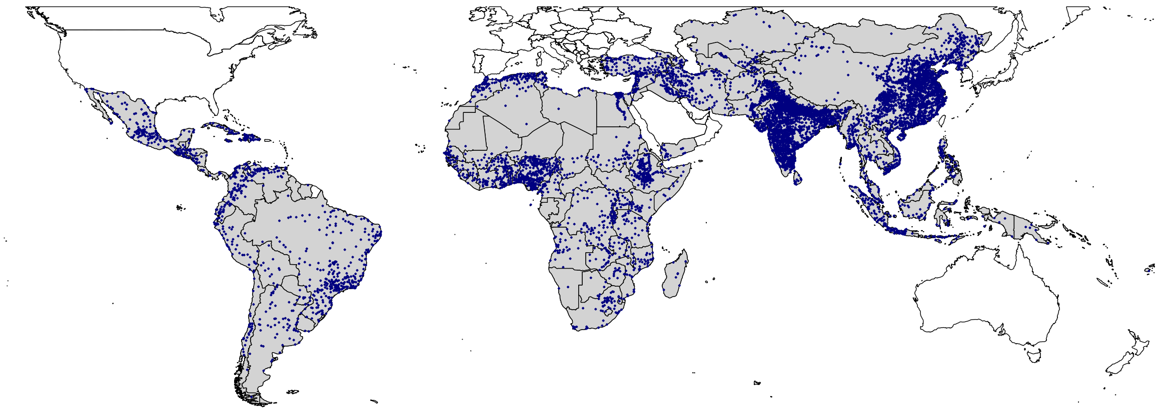

(a)

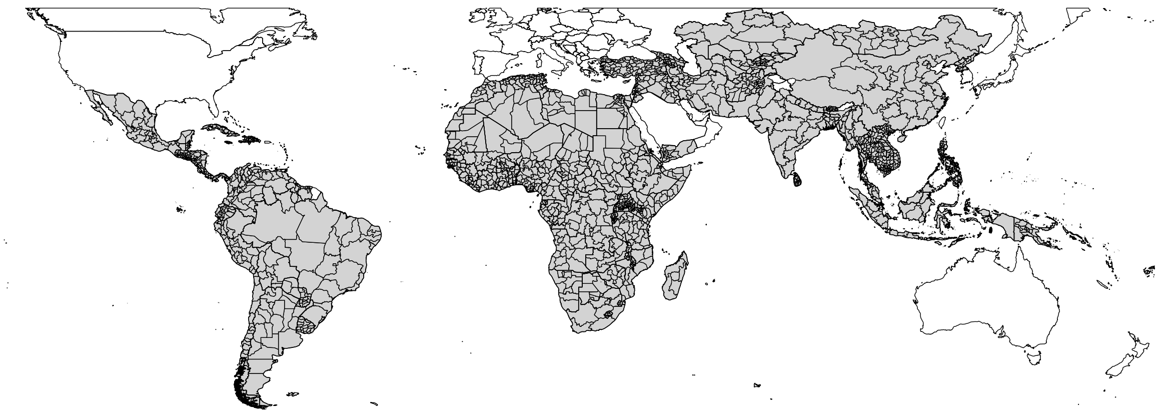

(b)

Figure C1.9: Panel a) shows the countries which are included in the sample, colored in gray, and the points show each city. Panel b) shows the ADM1 regions of the countries included in the sample, again colored gray.

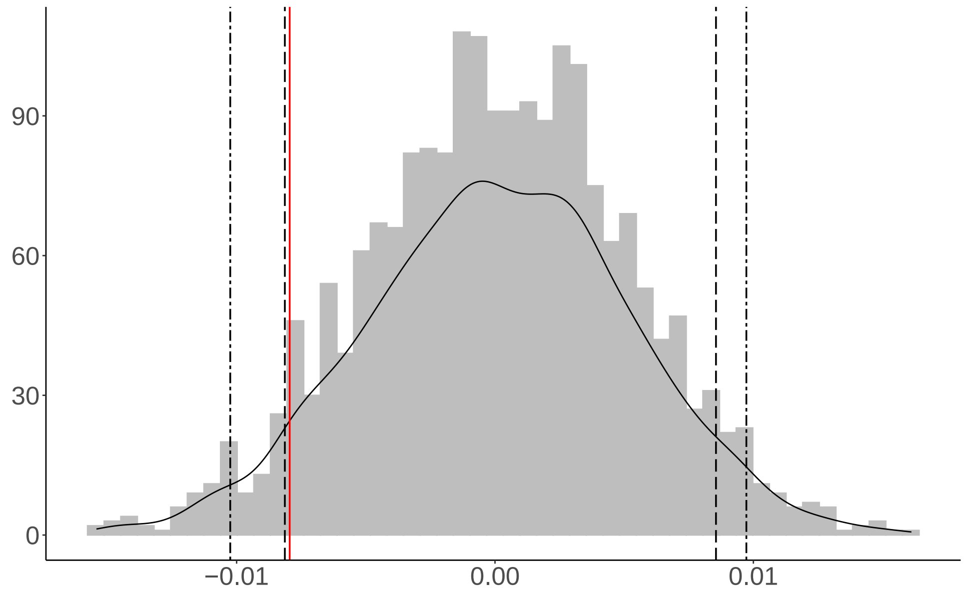

(a) 1.5 SD.

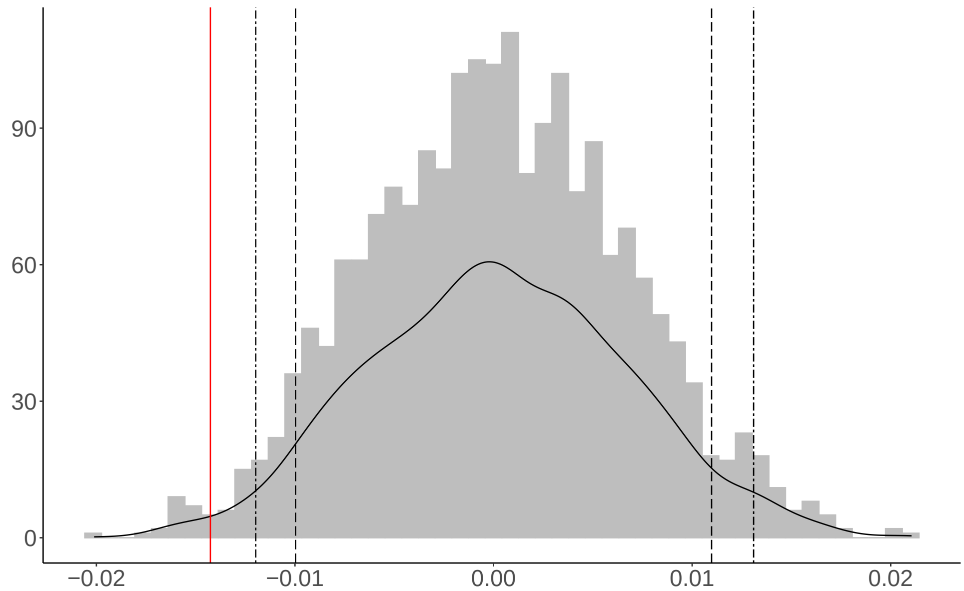

(b) 2 SD.

Figure C1.10: The histograms show the distributions of the estimated effect of a drought on city growth after 7 years (β_7) from estimating Equation 2 2,000 times using simulated data, where drought conditions are randomly assigned across years. Panel a) shows the results for 1.5 SD droughts, panel b) for 2 SD droughts. The solid red lines show the estimated effect using real data. The dot-dashed, and dashed lines reflect the critical values for a two-sided hypothesis test that the effect of a drought is zero at a 10 %, and 5 % level of significance, respectively.

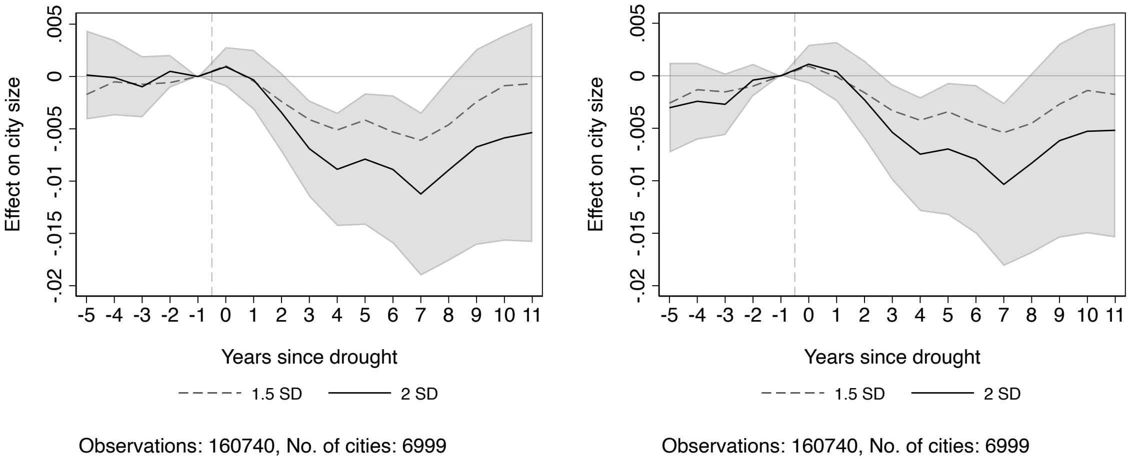

(a) Intermediate region FE.

(b) Winsorized.

Figure C1.11: Estimates show the dynamic treatment effects of drought on city size. Panel a) shows the effects when replacing subregion-by-year fixed effects with intermediate-by-year fixed effects. Panel b) shows the effects when city growth is winsorized at 1st/99th percentiles using the specification in Equation 2. City fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Dashed lines show the point estimates for each lag (or lead) of a 1.5 SD drought, and solid lines the estimates for a 2 SD drought. The shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

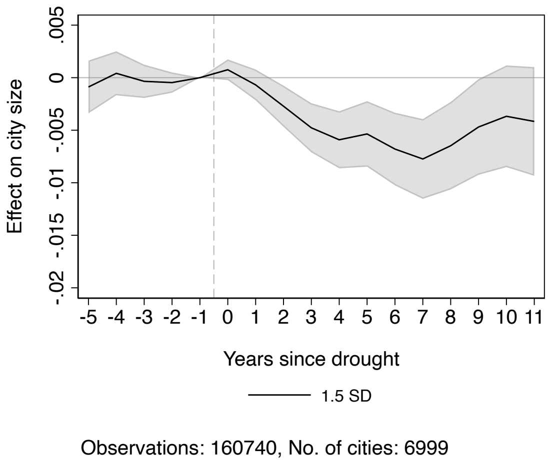

(a) 250 kilometer cut-off.

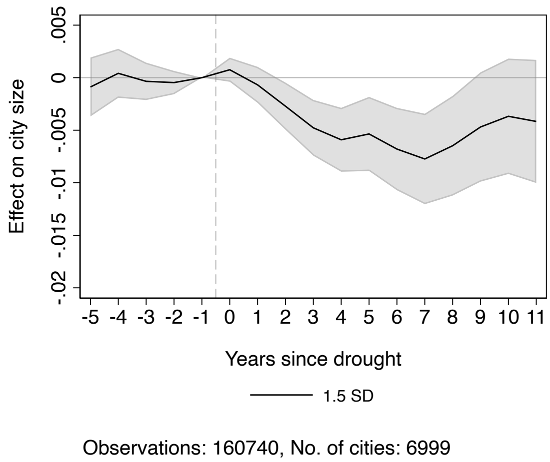

(b) 500 kilometer cut-off.

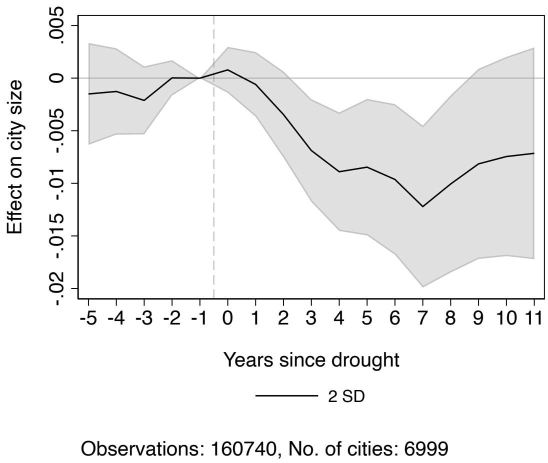

(c) 250 kilometer cut-off.

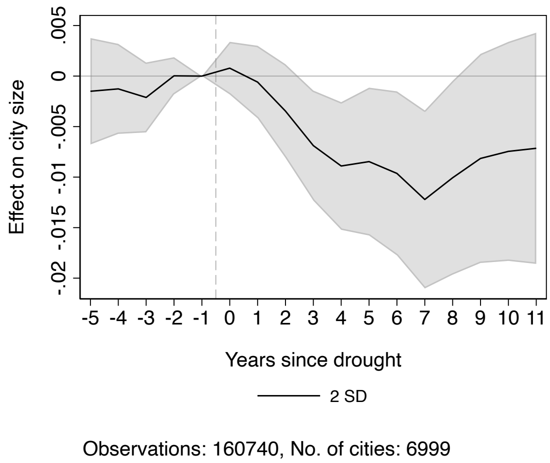

(d) 500 kilometer cut-off.

Figure C1.12: Estimates show the dynamic treatment effects of drought on city size. Standard errors are estimated following Conley (1999), allowing for serial correlation up to 20 years. Panels a) and b) show the results for a 1.5 SD drought, using a 250 and 500-kilometer cut-off respectively. Panels c) and d) show the results for a 2 SD drought, using a 250 and 500-kilometer cut-off respectively. City and subregion-by-year fixed effects are included in all estimations. Solid lines show point estimates at each lag (or lead) and the shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

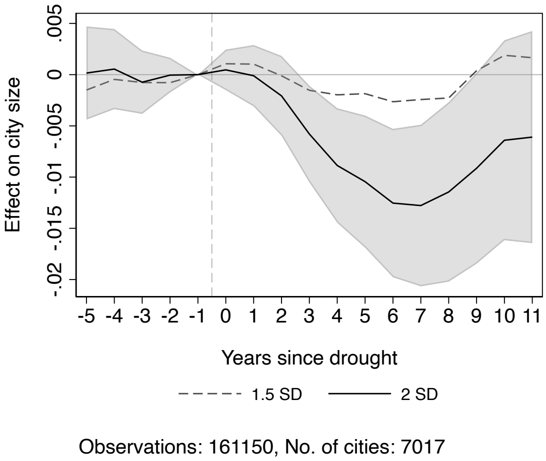

(a) Calendar year.

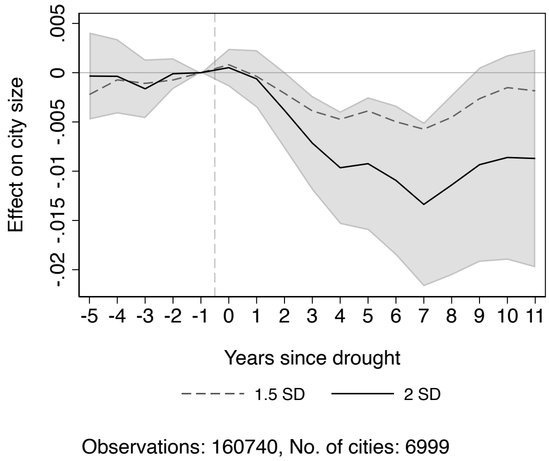

(b) SPEI-6.

Figure C1.13: Estimates show the dynamic treatment effects of drought on city size. Panel a) shows the effects when using the entire calendar year to define drought exposure. Panel b) shows the results of using SPEI-6 instead of SPEI-3. City and subregion-by-year fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Dashed lines show the point estimates for each lag (or lead) of a 1.5 SD drought, and solid lines the estimates for a 2 SD drought. The shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

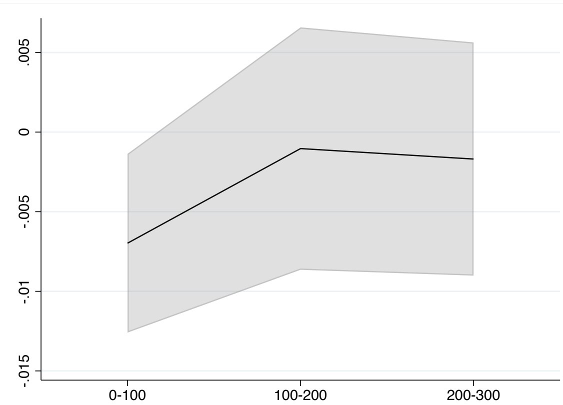

(a) 1.5 SD.

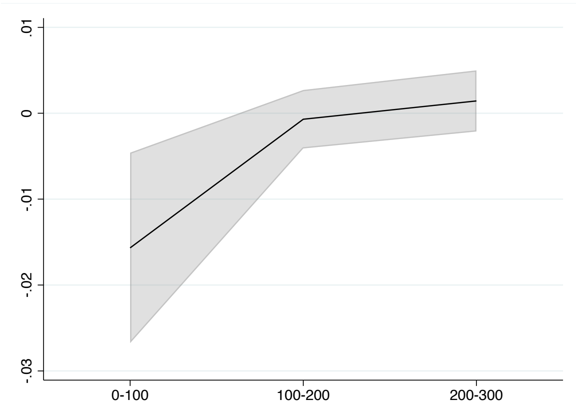

(b) 2 SD.

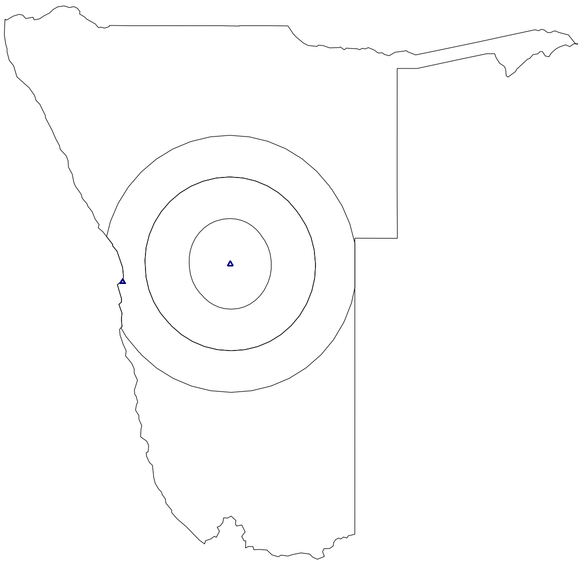

(c) Annuli.

Figure C1.14: Panels a) and b) show the results of estimating a spatial lag model as specified in Equation 7, illustrating the potential spatial spillover effects of drought exposure. Panels a) and b) show the effects of a drought on city i's size. 0-100 kilometers is the drought exposure measures used in the main results, e.g. 1. 100-200 is the effect on city i's size from droughts in cities within a 100-200 kilometer annuli, and 200-300 within a 200-300 kilometer annuli. Panel a) shows the results for 1.5 SD droughts, and panel b) for 2 SD droughts. Panel c) shows an example of the annuli used to construct the spatial lags, around the city of Windhoek, Namibia, where the outer annuli represents the 200-300 kilometer distance. The shaded area is 95 % confidence intervals.

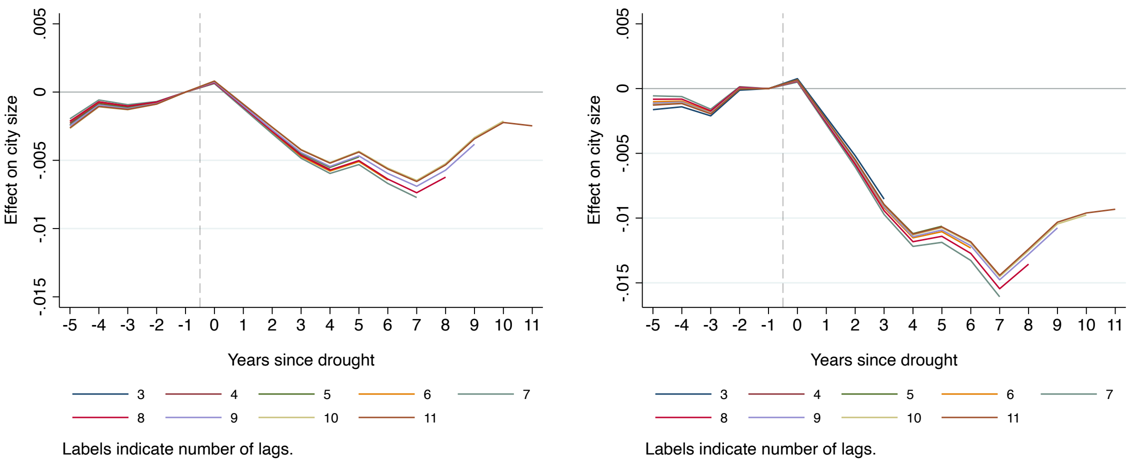

(a) 1.5 SD.

(b) 2 SD.

Figure C1.15: Estimates show the dynamic treatment effects of drought on city size. Each line shows the point estimates when including a different number of lags, from 3 to 11. I include 5 leads in each estimation. Panel a) shows the results for a 1.5 SD drought, panel b) a 2 SD drought. City and subregion-by-year fixed effects are included in all estimations. Confidence intervals have been omitted. Normalized to the year before drought ($\beta_{t=-1} = 0$).

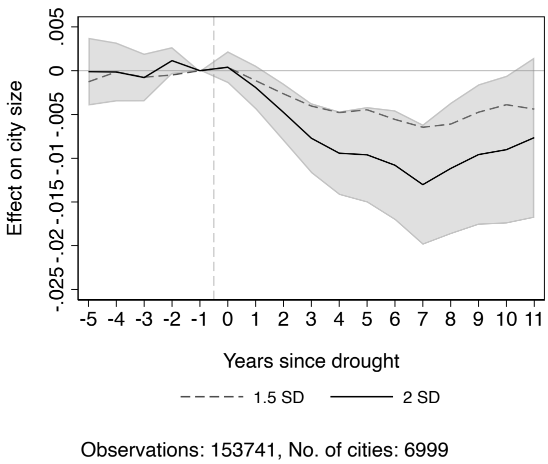

(a) AR(1).

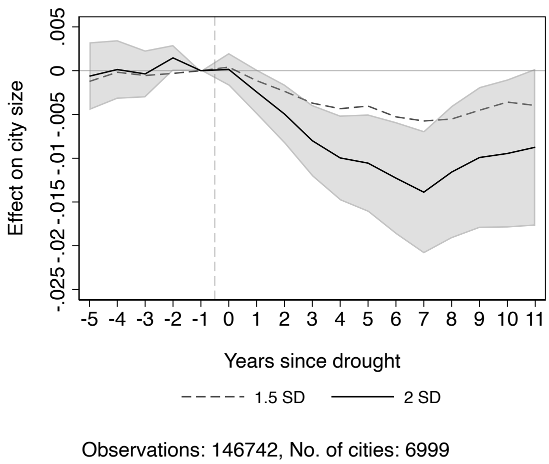

(b) AR(2).

Figure C1.16: Estimates show the dynamic treatment effects of drought on city size. and 5. Estimates are the dynamic treatment effects of a drought on city size. Panel a) shows the results of the AR(1) model, specified in Equation 4. Panel b) shows the results of the AR(2) model, specified in Equation 5. City and subregion-by-year fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Solid lines show point estimates at each lag (or lead) and the shaded region show 95% confidence intervals for a 2 SD drought. Normalized to the year before drought ($\beta_{t=-1} = 0$).

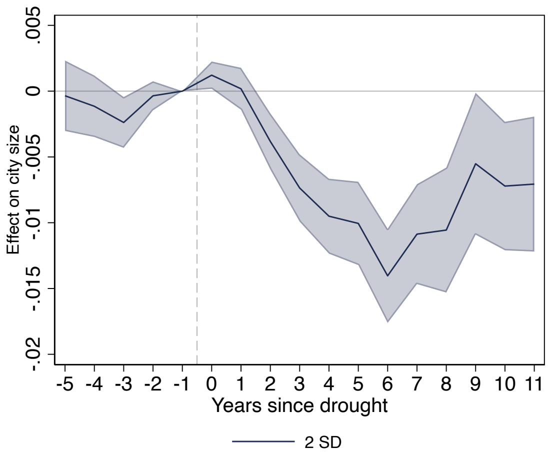

(a) Local projections.

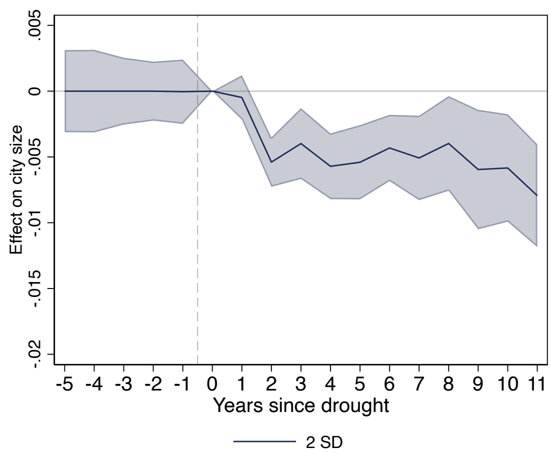

(b) DC-DH estimator.

Figure C1.17: Estimates show the dynamic treatment effects of a 2 SD drought on city size. Panel a) shows the results using the local projections specifications as defined in Equation 6. Standard errors are clustered by ADM1-by-year and city. Solid lines show point estimates at each lag (or lead) and the shaded region show 95% confidence intervals. Normalized to the year before a drought ($\beta_{t=-1} = 0$). Panel b) shows the results using the de Chaisemartin and d'Haultfoeuille (2022) estimator. Solid lines show point estimates at each lag (or lead) and the shaded region show 95% confidence intervals. Normalized to the year before treatment changes ($\beta_{t=0} = 0$).

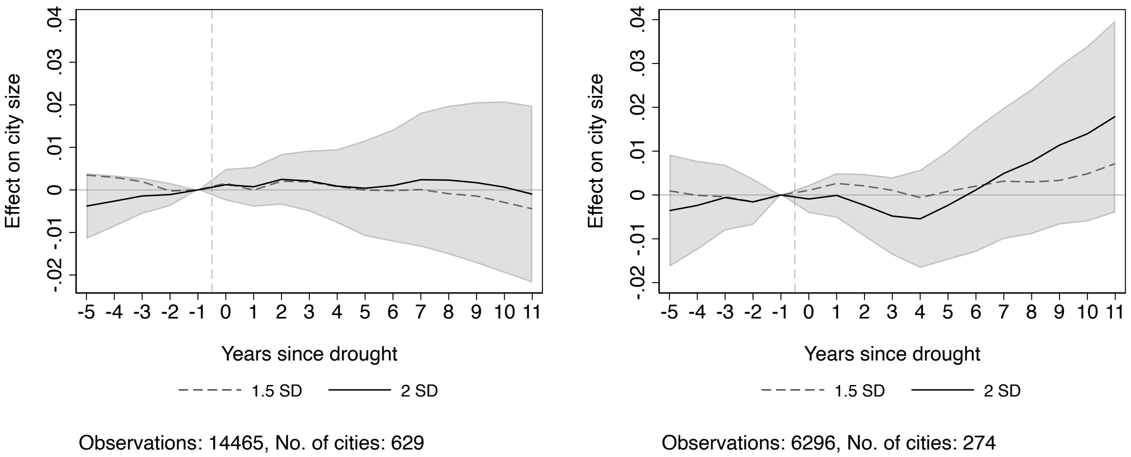

(a) South America.

(b) North America.

Figure C1.18: Estimates based on model in Equation 2. Estimates are the dynamic treatment effects of a drought on city size. Panel a) shows the effects on cities in South America. Panel b) shows the effects on cities in North America. City and subregion-by-year fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Dashed lines show the point estimates for each lag (or lead) of a 1.5 SD drought, and solid lines the estimates for a 2 SD drought. The shaded regions show 95% confidence intervals. Normalized to the year before drought ($\beta_{t=-1} = 0$).

D Additional tables

	Mean	Median	SD
Entire sample			
City pop. 2015	394,127	138,913	1,294,501
GDP per capita	2207	1097	2665
Percent urbanized	42.6	35.9	18.26
Agricultural employment share	52.8	58.5	17.0
City growth (percent)	0.03	0.02	0.05
Africa			
City pop. 2015	293,870	117,559	878,728
GDP per capita	1259	627	1535
Percent urbanized	36.6	36.5	14.9
Agricultural employment share	58.3	55.2	18.6
City growth (percent)	0.03	0.02	0.07
Asia			
City pop. 2015	425,490	150,028	1,403,754
GDP per capita	1863	1023	2200
Percent urbanized	38.4	32.8	12.8
Agricultural employment share	56.6	58.5	10.4
City growth (percent)	0.04	0.03	0.04
North America			
City pop. 2015	451,203	143,241	1,421,678
GDP per capita	5258	5276	3308
Percent urbanized	65.8	73.7	13.7
Agricultural employment share	29.5	25.9	7.9
City growth (percent)	0.03	0.02	0.04
South America			
City pop. 2015	418,154	125,225	135,0774
GDP per capita	5969	4981	3647
Percent urbanized	80.1	82.2	8.4
Agricultural employment share	20.1	19.2	7.2
City growth (percent)	0.03	0.02	0.03
Oceania			
City pop. 2015	132,205	98,267	71,281
GDP per capita	1501	1236	897
Percent urbanized	17.0	13.2	10.4
Agricultural employment share	37.0	37.8	7.0
City growth (percent)	0.01	0.01	0.01

Table D1.1: Descriptive statistics of the entire sample and by continents.

	1 SD	1.5 SD	2 SD
	City growth (%)	City growth (%)	City growth (%)
Year -5	-0.00218**	-0.00262*	-0.00127
	(-2.88)	(-2.22)	(-0.57)
Year -4	-0.00105	-0.00105	-0.00115
	(-1.76)	(-1.06)	(-0.60)
Year -3	-0.000953*	-0.00127	-0.00195
	(-1.99)	(-1.66)	(-1.28)
Year -2	-0.000712*	-0.000877	-0.0000230
	(-2.11)	(-1.77)	(-0.03)
Year -1	0	0	0
	(.)	(.)	(.)
Year 0	0.00116***	0.000801	0.000672
	(3.40)	(1.72)	(0.68)
Year 1	0.000818	-0.000416	-0.000389
	(1.64)	(-0.59)	(-0.26)
Year 2	0.000435	-0.00211*	-0.00347
	(0.68)	(-2.24)	(-1.74)
Year 3	0.000243	-0.00377**	-0.00694**
	(0.31)	(-3.26)	(-2.84)
Year 4	0.000249	-0.00473***	-0.00921**
	(0.27)	(-3.47)	(-3.18)
Year 5	0.00104	-0.00392*	-0.00865*
	(1.00)	(-2.44)	(-2.55)
Year 6	0.000352	-0.00516**	-0.00981*
	(0.30)	(-2.84)	(-2.56)
Year 7	0.000198	-0.00607**	-0.0124**
	(0.15)	(-3.02)	(-2.96)
Year 8	0.000975	-0.00486*	-0.0104*
	(0.70)	(-2.17)	(-2.24)
Year 9	0.00195	-0.00295	-0.00829
	(1.27)	(-1.20)	(-1.67)
Year 10	0.00218	-0.00175	-0.00758
	(1.32)	(-0.66)	(-1.45)
Year 11	0.00168	-0.00200	-0.00729
	(0.94)	(-0.71)	(-1.32)
Observations	160,740	160,740	160,740
Cities	6,999	6,999	6,999
Adjusted R-squared	0.3460	0.3464	0.3459

t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table D1.2: Event study estimates based on model in Equation 2. Estimates shown are the dynamic treatment effects of drought exposure, by levels of intensity, on city growth. City and subregion-by-year fixed effects are included in all estimations. Standard errors are clustered by ADM1-by-year and city. Normalized to the year before drought shocks, $\beta_{-1} = 0$.