Climate Change, Structural Transformation, and Infrastructure: Evidence from India

Maggie Liu, Yogita Shamdasani, Vis Taraz*

December 31, 2019

Click here for most updated version

Abstract

In developing economies with large productivity gaps, reallocation of workers both across space and sectors is crucial for economic development as it allows for a more efficient allocation of human capital. In this paper, we examine how rising temperatures affect the pace of reallocation of workers within local labor markets. Specifically, relying on six decades of district-level census data from India, we explore how decadal changes in temperature have affected urbanization and structural transformation within districts. We find evidence that rising temperatures are associated with lower rates of urbanization, higher shares of workers in agriculture, and lower shares of workers in non-agriculture. These effects are concentrated in districts with sparse road infrastructure networks, suggesting that higher temperatures exacerbate liquidity constraints faced by rural, isolated households, and subsequently limit rural-urban and sectoral mobility. Our findings demonstrate that the impacts of climate change can be unequal even *within* a country, and are aggravated by underdevelopment.

JEL Classification: O12, O13, Q54, R23

Keywords: structural transformation, urbanization, climate change, road infrastructure

*Liu: Smith College, Department of Economics, Wright Hall, 5 Chapin Drive, Northampton, MA 01063-6317 (email: yliu77@smith.edu). Shamdasani: University of Pittsburgh, Posvar Hall, 230 S Bouquet St, Pittsburgh, PA 15260 (email: yshamdasani@pitt.edu). Taraz: Smith College, Department of Economics, Wright Hall, 5 Chapin Drive, Northampton, MA 01063-6317 (email: vtaraz@smith.edu). Julia Bouzaher, Simren Nagrath, and Katya Israel-Garcia provided excellent research assistance. We thank Christine Cai for sharing code. We are grateful for feedback from conference participants at the 2019 Northeast Universities Development Conference (NEUDC) and the 2019 Georgetown Center for Economic Research Biennial Conference.

1 Introduction

It is well-established that developing economies are characterized by large productivity gaps across the economy—for example, between rural and urban areas (Lewis, 1954; Young, 2013) and between agricultural and non-agricultural sectors (McMillan et al., 2014; Gollin et al., 2014; Gollin and Rogerson, 2014). Reallocation of workers both across space and sectors is thus crucial for economic development as it allows for a more efficient allocation of human capital. At the same time, rising temperatures under climate change may potentially impede the pace of labor reallocation in developing economies, where the vast majority of workers engage in rural agriculture. Higher temperatures adversely impact crop yields (Schlenker and Roberts, 2009), and thus, in presence of liquidity constraints and costly switching, it is possible that higher temperatures may slow labor reallocation.

In light of this, this paper addresses two related empirical questions. The first question explores whether rising temperatures under climate change affect the pace of reallocation of workers within local labor markets — namely through urbanization and structural transformation — with important implications for economic growth. The second question investigates the extent to which access to transport infrastructure modulates the relationship between rising temperatures and reallocation. Importantly, for spatial or sectoral arbitrage to take place, there must be both a productivity gap and a sufficiently low cost of mobility.[1] We are thus interested in understanding whether extensive transport infrastructure — well-developed road networks, in particular — can facilitate the reallocation of workers from low to high productivity sectors in the presence of climate change.

We address these questions in the context of India, where rural-urban mobility is low (Munshi and Rosenzweig, 2016) and structural transformation, particularly the movement from agriculture to manufacturing, is slow and "stunted" (Binswanger-Mkhize, 2013). We

[1] A large literature has highlighted various barriers to spatial and/or sectoral reallocation of workers. This includes credit constraints (Banerjee and Newman, 1993), information frictions (Bryan et al., 2014), dependency on insurance networks in rural areas (Munshi and Rosenzweig, 2016), state boundaries (Kone et al., 2018) and poor transport infrastructure (Asher and Novosad, 2019; Shamdasani, 2019).

are particularly interested in local labor market responses, so we focus on rural-urban and sectoral movements *within* a district. We develop a simple model that highlights how rising temperatures can affect local labor reallocation, and how the presence of road infrastructure can modulate this relationship. We then assemble a district-level panel data set spanning six decades, combining measures of urbanization rates and worker shares across agricultural and non-agricultural sectors with decadal temperature and precipitation variables that are constructed using monthly gridded weather data. We further incorporate a baseline district-level measure of road network density in order to explore the role of transport infrastructure in modulating these relationships. Our identification strategy relies on the assumption—in line with the recent climate–economy literature (Dell et al., 2014)—that, conditional on district and region-by-year fixed effects, decade-to-decade fluctuations in weather are quasi-random.

We find two main results. First, we find that rising temperatures inhibit urbanization and structural transformation in Indian districts. The magnitude of these effects are economically meaningful—a 1°C increase in mean decadal temperatures in an average Indian district leads to a 5.6% decline in the share of the total population residing in urban areas, a 5.9% increase in the share of the total population engaged in agriculture and a 7.4% decline in the share of the total population engaged in non-agriculture. Second, we find that transport infrastructure plays an important role in mitigating these effects. While higher temperatures lead to an increase in the agricultural labor share and a reduction in urbanization and non-agricultural worker share in districts with sparse road networks, they lead to a decline in the agricultural labor share and no change in the urbanization rate and non-agricultural worker share in districts with dense road networks.

Our results are consistent with lower agricultural incomes under higher temperatures leading to binding liquidity constraints for rural, isolated households in sparsely connected districts. This, in turn, reduces the ability of these workers to move from rural to urban areas and/or to move out of agriculture into non-agriculture. On the other hand, workers in densely connected districts are able to move out of agriculture when they experience lower agricultural incomes under higher temperatures. Taken together, these results suggest that the impacts of rising temperatures can be unequal even *within* a country, and are exacerbated by underdevelopment.

Our paper contributes to two strands of the economics literature. First, we contribute to the literature that studies the response of rural-urban and sectoral movements to changes in weather and climate. Regarding rural-urban movements, Henderson et al. (2017) document that long-term increases in dryness have no average impact on urbanization in Sub-Saharan Africa, but that they do increase urbanization in the subset of regions with manufacturing centers. Cattaneo and Peri (2016) document that higher temperatures reduce rural-to-urban migration rates in low-income countries, consistent with binding liquidity constraints. Peri and Sasahara (2019) find that higher temperatures lead to a reduction in rural-to-urban migration in low-income countries, and an increase in rural-to-urban migration in middle-income countries.[2] Regarding sectoral reallocation, Emerick (2018) finds that transitory high rainfall shocks in India increase the non-agricultural labor share, due to an increased demand for local non-tradables. Also looking at India, Colmer (2019) finds that short-term increases in temperature are associated with a reduction in the agricultural labor share at the district level. The discrepancy between our results and Colmer (2019) may be due to the difference in the sample period used (1961–2011 vs. 2003–2008), but is also in line with existing literature that suggests the long-run impacts of higher temperatures may differ from short-run impacts (Hornbeck, 2012; Auffhammer and Schlenker, 2014; Dell et al., 2014; Moore and Lobell, 2014; Burke and Emerick, 2016).

Second, we contribute to the expansive literature on the effects of transport infrastructure. The papers most relevant to our current study examine the interactions between transport infrastructure and either sectoral reallocation or weather shocks. Regarding sectoral reallocation, Asher and Novosad (2019) show that the construction of paved roads in rural India

[2] Our findings on urbanization are consistent with Peri and Sasahara (2019), who classify India as a low-income country.

leads to reallocation of workers from agriculture to non-agriculture. Shamdasani (2019) finds that this sectoral reallocation is heterogeneous across space—households within close proximity to town centers move out of agriculture, while remote, isolated households still remain in agriculture. Regarding infrastructure and weather shocks, Burgess and Donaldson (2010) show that railroad access in colonial era India mitigates the impact of short-run adverse weather shocks on famines by enabling openness to trade.

Our paper makes several contributions to the literature. First, we provide evidence of sectoral labor reallocation in response to slow onset changes in climate. These findings complement the existing literature, which has to date focused either on the response of rural-urban movements to slow onset changes in climate (Cattaneo and Peri, 2016; Henderson et al., 2017; Peri and Sasahara, 2019) or of sectoral reallocation to short-term (annual) weather fluctuations (Emerick, 2018; Colmer, 2019). Second, our use of sub-national data allows us to track *local* labor market responses, focusing on movements within a district. This margin has been shown to be especially important in the context of India (Kone et al., 2018) and hence our results complement earlier work that has studied movements over greater distances. Third, we provide evidence that, in the modern era, access to well-developed transport infrastructure networks can modulate local labor market responses to higher temperatures, complementing existing evidence on the role of infrastructure in modulating rainfall shocks during the colonial period (Burgess and Donaldson, 2010). Lastly, our data allows us to examine whether regions within a country respond differently to rising temperatures. We find that the effects of higher temperatures are spatially heterogeneous, which suggests that there are important within-country dynamics to consider when thinking about the implications of climate policy.

Our findings demonstrate that underdevelopment—measured here by sparse road infrastructure networks—exacerbates the negative impacts of climate change by hindering potential reallocation of workers from low to high productivity sectors. Recent work has found that trade (movement of goods) is going to be a significant margin of adaptation to climate change (Gouel and Laborde, 2018). We argue that, in the context of India, worker mobility within local labor markets—movement of human capital between rural and urban areas, as well as between agricultural and non-agricultural sectors—could serve as an important form of adaptation to climate change. Furthermore, given that earlier work has documented that mobility costs such as state boundaries act as significant barriers to labor mobility (Kone et al., 2018), our work emphasizes the importance to reconsider and adjust mobility barriers, such as sparse road networks, in light of climate change.

The rest of this paper is organized as follows. In Section 2, we describe a two-period Roy-Borjas model that relates agricultural productivity to rural-urban and sectoral movements at varying levels of switching costs. In Section 3, we detail our data sources and present descriptive statistics. In Section 4 we describe our empirical specification. In Section 5, we discuss our results and present robustness checks. In Section 6, we conclude.

2 Model

2.1 Set-up

We develop a simple, two-period Roy-Borjas model (Roy, 1951; Borjas, 1987) that explores the costs and benefits to individuals of choosing sectors within an economy and captures the potential existence of binding liquidity constraints. More specifically, individuals compare their potential incomes (dependent on weather conditions) in different sectors, and make the switching decision based on income differentials net of switching costs. Cattaneo and Peri (2016) present a similar model that explains the hump-shaped relationship between migration rates and income across countries. Our model, however, focuses on the optimizing choices of workers selecting between sectors within a district (e.g. rural vs. urban, agricultural vs. non-agricultural). In addition, we integrate heterogeneity in switching costs and explore how it affects workers' sectoral switching decisions.[3]

To begin, consider a district with two sectors, "agriculture" and "non-agriculture," indicated by the superscripts A and N, respectively.[4] Here we use A and N for notation purposes. The two sectors can be generalized to any "low" and "high" productivity sectors such as urban and rural sectors; therefore, in addition to sector switching between agriculture and non-agriculture, the model can also be used for binary choices such as within-district rural-urban migration. In the first period, each individual works either in the agricultural or non-agricultural sector and earns the wage for that sector. Individuals differ in their skills, and their wages in either sector depend on their skills. At the start of the second period, individuals in the agricultural sector choose between remaining in their initial sector or switching to the other sector, depending on the potential wages they would earn in each sector in the second period. Workers in the agricultural sector A earn w_i^A, while workers in the non-agricultural sector N earn w_i^N. To simplify the model, we assume that individuals have a zero discount rate. We also assume that the wages in a given sector are identical across the two periods. The wage of individual i in sector J $(J = A, N)$ in the first and second periods is given by:

$$w_i^J = \mu^J + \beta^J \epsilon_i \tag{1}$$

where μ^J is the baseline wage in sector J earned by a typical worker (with median skills), and $\beta^J \epsilon_i$ captures the portion of wage attributed to individual i's de-meaned value of skills.

Median wages in both sectors depend on temperature T, which we write as $\mu^J(T)$. The dependence of wages on temperature is motivated by empirical evidence that higher temperatures reduce agricultural productivity (Schlenker and Roberts, 2009; Taraz, 2018) and also agricultural wages (Colmer, 2019). The productivity of each sector varies from the other and the non-agricultural sector is more productive than the agricultural sector (Gollin et al.,

[3]In our empirical analysis, we will exploit spatial variation in road network densities as a source of heterogeneity in switching costs.

[4]We abstract away from sectoral switching across different districts, as cross-district migration rates in India are relatively low (Kone et al., 2018).

2014), which is reflected in the fact that: $\mu^N(T) > \mu^A(T), \forall T$, or, in other words, for any given temperature the median wage in the non-agricultural sector will always be higher than the median wage in the agricultural sector at that same temperature.[5] In the Indian context, there is also empirical evidence that higher temperatures reduce non-agricultural productivity (Somanathan et al., 2015), but that these reductions are smaller in magnitude than the impacts on the agricultural sector (Jain et al., 2019).[6] Mathematically, we can summarize these assumptions as:

$$\frac{\partial \mu^A}{\partial T} < \frac{\partial \mu^N}{\partial T} < 0 \tag{2}$$

In equation 1, the term β^J captures the return to skills in sector J and the term ϵ_i represents the skills of individual i. Without loss of generality, we assume that ϵ_i is normally distributed with an average of 0 and a standard deviation of 1. For simplicity, we assume that the skills of an individual are perfectly transferable from the agricultural sector to the non-agricultural sector. However, the returns to skills in the non-agricultural sector are different than that in the agricultural sector. Following evidence from the existing literature (Kijima, 2006; Azam, 2012), we assume that the returns to skills are higher in the non-agricultural sector than the agricultural sector: $\beta^N > \beta^A$.

Additionally, we assume that there is a cost C to switching between agriculture and non-agriculture. To simplify the analysis, we express switching costs C in "time-equivalent" terms, c; that is, $c = C/w_i^A$. This switching cost could capture informational frictions and search costs more broadly (Bryan et al., 2014), or could capture specifically the cost to switch sectors due to poor transport infrastructure. Based on existing empirical literature (Asher and Novosad, 2019; Gertler et al., 2019), we assume that increasing the quality of road networks (r) reduces sectoral switching costs. Thus we can write switching costs as a decreasing function of road networks: $c(r)$, and we have $\frac{\partial c}{\partial r} < 0$. As in Grogger and Hanson

[5] Similarly we can assume that the urban sector is more productive than the rural sector (Henderson et al., 2017)

[6] This India-specific evidence is consistent with cross-country evidence that, in poor countries, the agricultural sector is more sensitive to higher temperatures than the non-agricultural sector, but that both respond negatively to higher temperatures (Dell et al., 2012).

(2011) and Cattaneo and Peri (2016), we assume that individuals' preferences are linear in their net wages (wages net of sectoral switching costs). An individual will choose to switch sectors if and only if the income earned in the non-agricultural sector (net of switching costs) exceeds the net income earned in the agricultural sector.[7] In particular, an individual i in the agricultural sector A will choose to switch to the non-agricultural sector N if:

$$\mu^N(T) + \beta^N \epsilon_i - c(r) > \mu^A(T) + \beta^A \epsilon_i. \tag{3}$$

Rearranging, we can solve this inequality to see that an individual will choose to switch to the non-agricultural sector if:

$$\epsilon_i > \frac{\mu^A(T) - \mu^N(T) + c(r)}{\beta^N - \beta^A} \tag{4}$$

Following Cattaneo and Peri (2016), we can think of equation 4 as an "incentive-compatibility" constraint: individuals from the agricultural sector will switch to the non-agricultural sector only if their income differentials gained from the switch (non-agricultural wages minus their agricultural wages) exceed the switching cost. In light of the Borjas (1987) selection model, equation 4 predicts positive selection, since the returns to skills are higher in the non-agricultural sector ($\beta^N - \beta^A > 0$). In other words, all the individuals that choose to switch sectors will have skills above some threshold determined by sector productivity differentials, switching costs and the difference in returns to skills. Ceteris paribus, if the non-agricultural sector is more productive than the agricultural sector by a bigger margin, if switching costs are lower, or if the returns to skills are higher in non-agriculture by a bigger margin, then a bigger fraction of the agricultural workers will switch sectors.

An individual's decision to switch sectors must also satisfy a second constraint: a "feasi-

[7] In general, an individual will choose to switch sectors if and only if the potential income differential exceeds the switching costs. However, in our specified model, no individual will switch from the non-agricultural sector to the agricultural sector since $w_i^N > w_i^A$; therefore the (negative) income differential from sector N to A will never afford the switching costs.

bility" constraint. Let us assume that sectoral switching happens at the start of the second period, that individuals cannot borrow due to a liquidity constraint, and that sectoral switching costs must be paid up front. Then, an individual can switch sectors only if the costs of switching do not exceed their total savings at the end of period 1. Given our simplified model, with no income sources except for labor, the total savings can be no greater than w_i^A. Thus, we can write the "feasibility constraint" as:

$$\mu^A(T) + \beta^A \epsilon_i > c(r), \tag{5}$$

which we can rearrange to write as:

$$\epsilon_i > \frac{c(r) - \mu^A(T)}{\beta^A} \tag{6}$$

Equation 6 is also a positive selection equation where only individuals in the agricultural sector with skills above the threshold determined by their wages, switching costs and returns to skills in agriculture will be able to afford the switch. Ceteris paribus, if the agricultural productivity is higher, if switching costs are lower, or if the returns to skills in agriculture are higher, then a bigger fraction of the agricultural workers can overcome the liquidity constraint and switch sectors if they choose to do so.

2.2 Propositions

The fraction of individuals who will switch from the agricultural sector to the non-agricultural sector is equal to one minus the cumulative density of a normal distribution at the highest of the two thresholds defined in Equations 4 and 6. For any given district, only one of the two thresholds can be binding, and thus we derive the following proposition:

Proposition 1. *The overall effect of a temperature increase on the rate of sectoral switching is ambiguous. If the incentive-compatibility constraint binds, then an increase in average* *temperature will be associated with an increase in the rate of sectoral switching to the non-agricultural sector. If, on the other hand, the feasibility constraint binds, then an increase in average temperature will be associated with a decrease in the rate of sectoral switching to the non-agricultural sector.*

Proof. For districts where the feasibility-compatibility constraint binds, the share of people switching to the non-agricultural sector is equal to the share of the population with skills above that threshold, given by the expression:

$$1 - \Phi\left(\frac{c(r) - \mu^A(T)}{\beta^A}\right) \tag{7}$$

where Φ is the CDF of a standard normal distribution. The expression is increasing in $\mu^A(T)$, because the CDF Φ is a monotonically increasing function. Given our earlier assumption that increases in temperature T decrease $\mu^A(T)$, we have that the expression is decreasing in T.

On the other hand, for districts where the incentive-compatibility constraint binds, the share of people switching to the non-agricultural sector is equal to the share of the population with skills above that threshold, given by the expression:

$$1 - \Phi\left(\frac{\mu^A(T) - \mu^N(T) + c(r)}{\beta^N - \beta^A}\right) \tag{8}$$

where Φ is the CDF of a standard normal distribution. The expression is decreasing in $\mu^A(T) - \mu^N(T)$, because the CDF Φ is a monotonically increasing function. Given our earlier assumption that increases in temperature T decrease $\mu^A(T)$ more than $\mu^N(T)$, we have that the expression is increasing in T. □

Given our interest in road network quality, we also present some propositions related to road network heterogeneity. If the returns to skill in the non-agricultural sector are sufficiently higher than the returns to skill in the agricultural sector then, for any given initial temperature T, there is a value of road quality $r*$, such that the incentive compatible and feasibility thresholds are equal. We consider this value as marking the difference between high quality (H) and low quality road network districts (L).[8] Given the distinction between these two types of districts, we generate the following two predictions:

Proposition 2. *For districts with a sufficiently low quality road network, an increase in average temperature is associated with a decrease in the rate of sectoral switching.*

Proof. For districts whose road quality is lower than $r*$, defined as low quality road network districts, L, only the feasibility threshold (6) is binding. Hence the share of people switching to the non-agricultural sector is equal to the share of individuals with skills above that threshold, given by:

$$1 - \Phi\left(\frac{c(r) - \mu^A(T)}{\beta^A}\right) \tag{9}$$

As derived above, we have that this expression is decreasing in T. Hence, for districts with a sufficiently low quality road network, an increase in average temperature is associated with a decrease in the rate of sectoral switching. $\square$

Proposition 3. *For districts with a sufficiently high quality road network, an increase in average temperature is associated with an increase in the rate of sectoral switching.*

Proof. For districts whose road quality is higher than $r*$, defined as high quality road network districts, L, only the incentive-compatibility threshold (6) is binding. Hence the share of people switching to the non-agricultural sector is equal to the share of individuals with skills above that threshold, given by:

$$1 - \Phi\left(\frac{\mu^A(T) - \mu^N(T) + c(r)}{\beta^N - \beta^A}\right) \tag{10}$$

As derived above, we have that this expression is increasing in T. Hence, for districts with a

[8]The specific value is given by $c(r*) = \dfrac{\mu^A(T)\beta^N - \mu^N(T)\beta^A}{\beta^N - 2\beta^A}$.

sufficiently high quality road network, an increase in average temperature is associated with a increase in the rate of sectoral switching. □

3 Data

3.1 Census Data

We use data from the decadal national Census of India, spanning the years 1961 to 2011. Specifically, we extract outcome measures from the Primary Census Abstract (PCA) Data Tables, which provide summaries of district-level demographic and economic indicators.[9] For each census year, we construct three key outcome measures at the district level: the share of the total population residing in urban areas, the share of the total population who are agricultural laborers, and the share of the total population who are non-agricultural workers.[10]

To account for the fact that districts split and boundaries are adjusted over time, we use concordance tables from Kumar and Somanathan (2017) to construct consistent district boundaries that span the same area between 1961 and 2011. Specifically, we map every district in each census year to its parent district in 1961.[11] This results in 288 consistent districts, as illustrated in Figure 1a. The various splits and boundary changes between 1961

[9]For the years 1961-1991, we use data from Vanneman and Barnes (2000). For the years 2001 and 2011, we use data from the Census website. The data can be accessed at `http://www.censusindia.gov.in/DigitalLibrary/Tables.aspx`.

[10]We define non-agricultural workers as the sum of workers across two census categories: household industry workers and other workers. The 2011 Census defines household industry as "an industry conducted by one or more members of the household at home or within the village in rural areas and only within the precincts of the house where the household lives in urban areas", and other workers as "workers other than cultivators, agricultural laborers or workers in Household Industry". Examples given for the latter category include workers engaged in the public sector, manufacturing, construction, trade, business etc.

[11]For example, Kancheepuram and Thiruvallur districts in Tamil Nadu were formed when Chengalpattu district split in 2001. In this case, we designate Chengalpattu as the consistent district from 1961 to 2011. There are also instances where a district does not have a unique parent district— this happens when a district is carved out of two or more original districts. For these cases, we create a "greater" parent district which is the superset of all parent districts. As an example, Narmada district in Gujarat was carved out of two districts— Vadodara and Bharuch— in 2001, therefore we designate "Vadodara and Bharuch" as the consistent district boundary from 1961 to 2011.

and 2011 can be deduced from the grey boundaries that trace out the 2011 Census districts delineation. Figure 1b highlights the 287 districts across six regions that form our analysis sample.[12]

3.2 Weather Data

We use gridded monthly data on temperature and precipitation from the Terrestrial Precipitation: Monthly Time Series (1900–2014), version 4.01, and the companion Terrestrial Air Temperature data set.[13] We match the grid points to each of our districts by taking the weighted average of all grid points within 100 kilometers of each district's centroid, using weights that are the inverse of the squared distance between the grid point and the district centroid. We focus on temperature and precipitation during the main agricultural growing season months (June through February) as these have the greatest impacts on agriculture. Given our interest in responses to slow-onset changes in climate, we aggregate the monthly weather variables to ten-year averages.

3.3 Yield Data

We use data on yields from the Village Dynamics in South Asia Meso Dataset (VDSA), compiled by researchers at the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT, 2015), spanning the years 1996 to 2010. We construct an aggregate yield measure that quantifies yields across all the crops in VDSA that have non-missing price data,[14] using 1966-1970 crop prices as weights. We also construct yield measures separately for rice and wheat.

[12] We classify districts into six regions based on the Government of India's administrative regional classification. Lakshadweep is dropped from the analysis sample due to lack of weather records.

[13] Willmott, C. J. and K. Matsuura (2015). Terrestrial Precipitation: 1900-2014 Gridded Monthly Time Series (1900 - 2014) (V 4.01), `http://climate.geog.udel.edu/~climate/html_pages/Global2014/README.GlobalTsP2014.html`, Willmott, C. J. and K. Matsuura (2015). Terrestrial Air Temperature: 1900-2014 Gridded Monthly Time Series (1900 - 2014) (V 4.01), `http://climate.geog.udel.edu/~climate/html_pages/Global2014/README.GobalTsT2014.html`

[14] These crops are rice, wheat, sugarcane, cotton, groundnut, sorghum, maize, pearl millet, finger millet, barley, chickpeas, pigeon pea, sesame, rapeseed and mustard, castor, and linseed.

3.4 Infrastructure Data

We use data on road infrastructure from the Village Dynamics in South Asia Meso Dataset (VDSA), compiled by researchers at the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT, 2015). In particular, we use the total length of roads in kilometers in each district in 1970— the earliest year for which this data is available. We construct a district-level road density measure by dividing the total length of roads by the total surface area, computed in ArcGIS from the consistent district boundaries outlined in Figure 1a.

Figure 2a summarizes the distribution of the road density measure across all districts. We create a binary measure that takes the value one if road density in a district is above the median level of the full distribution ($0.11km/km^2$, as indicated by the solid vertical line in the figure), and zero otherwise. Figure 2b plots a heat map of the road density measure across all districts, with shades of red denoting districts with above median road density, and shades of blue denoting districts with below median road density. Districts with above median road density are scattered across the country, with a greater concentration in two regions: the East and the South.

3.5 Descriptive Statistics

Table 1 provides summary statistics across the three data sets, reported for each census year. We report means and standard deviations of key variables for all districts in the sample, and separately for districts with below and above median road density.

First, the table describes the road density measure. The first row in Panel A shows that road density is $0.183km/km^2$ on average across all districts in 1970. There is significant heterogeneity in this measure— the mean road density in districts above the median level is $0.316km/km^2$ (Panel C), which is more than five times higher than $0.0552km/km^2$, the mean road density in districts below the median level (Panel B).

Next, the table summarizes the weather variables. The second row in Panel A confirms that temperatures have been rising over time. The growing season average monthly temperature is 0.43 °C higher in 2011, relative to 1961. On the other hand, the third row in Panel A suggests that growing season average monthly precipitation has not changed monotonically over time— in fact, we see a decline of 9mm over the same time period. These patterns are similar for both subgroups in Panels B and C.

Finally, the table summarizes the three outcome measures from the Census. The share of total population residing in urban areas is 15.5% on average in 1961, and increases steadily over time. By 2011, this share has increased to 26.3%. Agricultural labor shares and non-agricultural worker shares have also increased over the decades. At each point of time, the three measures are lower in districts with below median road density in Panel B, relative to districts with above median road density in Panel C. It is interesting to note, however, that agricultural labor shares increased at a faster pace in districts with below median road density (107% growth between 1961 and 2011, compared to 67% growth in districts with above median road density during the same period), while non-agricultural worker shares increased at a faster pace in districts with above median road density, reaching 18% in 2011 (15.2% in districts with below median road density).

Figure 3 plots the spatial distribution of changes in the long-run climate and outcome variables from 1961 to 2011, by district.[15] Inland India has experienced much larger increases in temperature, relative to the coastal areas (panel a). Perhaps unsurprisingly, inland India has also experienced larger declines in precipitation compared to areas closer to the coast (panel b). It is evident that temporal changes in the Census outcomes (panels c-e) are heterogeneous across space. The share of total population residing in urban areas has increased by more than 17 percentage points in one-sixth of the districts, while another one-sixth of the districts have experienced less than a 4 percentage points increase. The biggest gains in agricultural labor shares over the decades appear to be concentrated among districts in the Eastern, Central, and Southern regions, while the lower half of the Western and Northern regions have experienced the smallest gains.

[15] Detailed plots of each of these measures in each decade can be found in Appendix B.

4 Empirical Specification

To estimate the effect of climate on urbanization and structural transformation, we estimate a regression of the form:

$$lnY_{jsrt} = \beta lnT_{jsrt} + \gamma lnP_{jsrt} + \alpha_j + \alpha_{rt} + \epsilon_{jsrt}, \tag{11}$$

where Y_{jsrt} represents the outcome of interest — share of the total population residing in urban areas, share of agricultural laborers or share of non-agricultural workers — in district j, located in state s and region r, in year t (= 1961, 1971, 1981, 1991, 2001, 2011). T_{jsrt} is the average temperature measured in °C over the growing season months (June through February) in the past decade ending in year t, and P_{jsrt} is the average precipitation measured in millimeters over the growing season months in the past decade ending in year t. α_j is a vector of district fixed effects that controls for any time-invariant district-specific factors that may be correlated with climate or local economic patterns. α_{rt} is a vector of region-decade fixed effects that controls for any unobserved region-specific factors that may be correlated with climate or local economic patterns over time. Lastly, ϵ_{jsrt} is an idiosyncratic error term. We cluster our errors at the district level to allow for potential serial correlation over time within each district. We also report Conley standard errors that allow for spatial correlation up to 500km and arbitrary serial correlation in the error term (Conley, 1999).[16] Additionally, we present results where we include state-specific linear time trends that control for any smoothly varying changes over time that may be occurring at the state level.

Following Cattaneo and Peri (2016), the dependent variable in the regression is the natural logarithm of Y_{jsrt}, and it depends on the logarithm of T_{jsrt} and P_{jsrt}. The identifying assumption is that, conditional on the inclusion of district and region-year fixed effects, along with state-specific linear time trends, any remaining variation in decadal temperature and precipitation is essentially random. This in turn allows for a causal interpretation of the β

[16] To implement Conley standard errors, we use Stata routines from Hsiang (2010) and Fetzer (2014).

and γ coefficients as the effect of slow-onset changes in climate on urbanization and structural transformation.

As described in Section 2, predictions for the sign of the coefficient of interest β are theoretically ambiguous. If higher temperatures reduce agricultural productivity and subsequently, the demand for agricultural labor, this should result in an outflow of workers from agriculture to non-agriculture, as well as prompt rural-to-urban migration. Under this scenario, we would expect $\beta > 0$ in regressions where urbanization or the non-agricultural worker share is the dependent variable, and $\beta < 0$ in regressions where the share of agricultural laborers is the dependent variable. On the other hand, lower agricultural incomes under higher temperatures can lead to binding liquidity constraints, which in turn, reduce the ability of workers to move across sectors and/or from rural to urban areas. (Cattaneo and Peri, 2016; Peri and Sasahara, 2019). Under this scenario, we would find the opposite results: $\beta < 0$ in regressions where urbanization or the non-agricultural worker share is the dependent variable, and $\beta > 0$ in regressions where the share of agricultural laborers is the dependent variable.

To examine whether transport infrastructure plays a role in modulating the effect of climate change on urbanization and structural transformation, we allow for the net effects of temperature and precipitation to vary based on the extent of road connectivity in a given district. We interact the weather variables with a binary measure for road density and estimate a regression of the form:

$$\begin{aligned} lnY_{jsrt} =& \beta lnT_{jsrt} + \gamma lnP_{jsrt} + \beta_D lnT_{jsrt} * D_j + \gamma_D lnP_{jsrt} * D_j \\ &+ \alpha_j + \alpha_{\mathbf{D}t} + \alpha_{rt} + \epsilon_{jsrt}, \end{aligned} \tag{12}$$

where D_j is a binary variable that takes the value 1 if the road density in district j is above the median of the distribution across all districts in 1970, and 0 otherwise. We also include $\alpha_{\mathbf{D}t}$, heterogeneous-group-by-year fixed effect, which allows for each subgroup (districts with below-median road density and districts with above-median road density) to have different unobserved shocks over time. All other terms are as defined above. In this specification, β and γ capture the effects of decadal changes in temperature and precipitation respectively in districts with sparse road networks, while $\beta + \beta_D$ and $\gamma + \gamma_D$ capture the effects of decadal changes in temperature and precipitation respectively in districts with dense road networks.

5 Results

5.1 Main Effects

Table 2 presents the main effects of temperature and precipitation on three key outcomes — urbanization, share of agricultural laborers, and share of non-agricultural workers. For each outcome, the first column presents regression estimates using Equation 11, while the second column presents regression estimates with the inclusion of state-specific linear time trends. We report results with standard errors clustered at the district level in parentheses and with Conley standard errors that allow for spatial clustering up to 500km and arbitrary serial correlation in brackets.

In Columns 1 and 2 of Table 2, we find a strong, negative effect of rising average temperatures on urbanization rates. The estimated coefficient in Column 2 indicates that a 1% increase in decadal temperature is associated with a 1.4% reduction in the share of the total population residing in urban areas, and this effect is statistically significant at the 1% level. This translates to a 5.6% decline in the urbanization rate for an average district in our sample, if the mean decadal temperature were to increase by 1°C. At the same time, we find that changes in decadal precipitation has no detectable impact on urbanization rates. Further, the estimated coefficient is very small in magnitude — a 1% increase in decadal precipitation is associated with a 0.01% reduction in the share of the total population residing in urban areas. This is in line with recent studies that have documented significantly larger impacts of temperature on agricultural production relative to rainfall in the Indian context (Burgess et al., 2017; Colmer, 2019).

Next, in Columns 3 and 4 of Table 2, we find a positive, statistically significant effect of rising average temperatures on the share of agricultural laborers — a 1% increase in temperature leads to a 1.4% increase in the share of the population engaged in agriculture. Finally, in Columns 5 and 6 of Table 2, we find a strong, negative effect of rising average temperatures on the share of non-agricultural workers that is statistically significant at the 1% level. The estimated coefficient indicates that a 1% increase in temperature is associated with a 1.8% reduction in the share of the total population engaged in non-agriculture. Thus, if the mean decadal temperature were to increase by 1°C, this would result in a 5.9% increase in the share of agricultural laborers and a 7.4% decline in the share of non-agricultural workers for an average district in our sample. Consistent with the idea that temperature has larger impacts on agricultural production relative to rainfall, we find that rising average precipitation has a small, negative effect on the share of agricultural workers, and no detectable effect on the share of non-agricultural workers.[17]

Taken together, these results are consistent with higher temperatures negatively impacting agricultural yields and incomes. This, in turn, causes liquidity constraints to bind and reduces the ability of workers to move from rural to urban areas and/or to move out of agriculture into non-agriculture. Rising temperatures thus appear to inhibit urbanization and structural transformation in Indian districts. While the temperature-yield relationship has been established previously in the literature, we replicate this result using data from our empirical context in Appendix Table A1. Across three different yield measures, we consistently find a strong, negative effect of rising average temperatures on yields — for example, a 1% increase in temperature is associated with a 0.586% reduction in yields aggregated across all crops in VDSA.

[17]We run an additional test where we run the same regression in Equation 11, dropping precipitation. The coefficients on temperature are very similar: -1.342 (with urbanization as the dependent variable), 1.515 (with agricultural labor share as the dependent variable), and -1.750 (with non-agricultural worker share as the dependent variable).

5.2 Heterogeneous Effects

Table 3 presents heterogeneous effects of temperature by baseline road network density on three key outcomes—urbanization, share of agricultural laborers, and share of non-agricultural workers. For each outcome, the first column presents regression estimates using Equation 12, while the second column presents regression estimates with the inclusion of state-specific linear time trends. As before, we report results with standard errors clustered at the district level in parentheses and with Conley standard errors that allow for spatial clustering up to 500km and arbitrary serial correlation in brackets.

In Columns 1 and 2 of Table 3, we find that the strong, negative effect of rising average temperatures on urbanization rates documented in Table 2 is concentrated among districts with below median road density at baseline. The estimated coefficient indicates that a 1% increase in temperature in districts with sparse road networks is associated with a 2% reduction in the share of the total population residing in urban areas, significant at the 1% level. This would translate to a 8.3% decline in the urbanization rate for an average below-median district in our sample, if the mean decadal temperature were to increase by 1°C. On the other hand, a 1% increase in temperature in districts with dense road networks has no detectable impact on the urbanization rate—the p-value of the sum of coefficients is 0.6539.

Next, in Columns 3 and 4 of Table 3, we find contrasting effects of rising average temperatures on the share of agricultural laborers across the two subgroups—we find a positive effect in districts with below median road density at baseline, and a negative effect in districts with above median road density at baseline. The estimated coefficient indicates that a 1% increase in temperature in districts with sparse road networks is associated with a 2.2% increase in the share of the total population engaged in agriculture, significant at the 5% level. On the other hand, the estimated coefficient indicates that a 1% increase in temperature in districts with dense road networks is associated with a 3.3% *decline* in the share of the total population engaged in agriculture, significant at the 5% level.

Finally, in Columns 5 and 6 of Table 3, we find that the strong, negative effect of rising average temperatures on the share of non-agricultural workers documented in Table 2 is driven entirely by districts with below median road density at baseline. The estimated coefficient indicates that a 1% increase in temperature in districts with sparse road networks is associated with a 2.2% reduction in the share of the total population engaged in non-agriculture, significant at the 1% level. This would translate to a 9.2% decline in the share of non-agricultural workers for an average below-median district in our sample, if the mean decadal temperature were to increase by 1°C. On the other hand, a 1% increase in temperature in districts with dense road networks has no detectable impact on the share of non-agricultural workers—the p-value of the sum of coefficients is 0.8855.

Notice that the average effects of rising temperatures in districts with sparse road networks are larger than the average effects in the full sample. This suggests that underdevelopment *amplifies* the impacts of rising temperatures on the degree of urbanization and structural transformation in Indian districts. These results also demonstrate that transport infrastructure networks within a district play an important role in mitigating the impacts of rising temperatures.

5.3 Robustness Checks

Tables 4 and 5 present results from a series of robustness tests, outlined below.

In our main analysis, we restrict our sample to districts for which the dependent variable is non-missing in all years. In Panel A of Table 4, we report results using the full unbalanced sample. The estimated coefficients and significance levels are largely unchanged under the inclusion of these unbalanced districts.

Next, in order to maximize power, our main empirical specification uses a binary variable that takes the value one if road density in the district is above the median of our sample distribution. In Panel B of Table 4, we use an alternate cutoff—we define *High Road Density* as 1 if the road density measure is within the top 30th percentile of our sample distribution. We find that the results are robust to this alternate cutoff. Finally, we show in Panel C of

Table 4 that the results are also robust under a more conservative clustering of standard errors at the state level, though we do lose precision on some of the estimates—for example, the coefficient on the share of agricultural workers is now significant at the 5% level (compared to 1% level when clustering at the district level).

In defining long-run changes in climate, we use decadal averages of temperature and precipitation. In Panel A of Table 5, we show that our results are robust to using averages defined over a five-year window instead — under this scenario, we compute long-run average temperature and precipitation for outcomes in year t using monthly data from years $t-4$ to t.[18] Finally, in Panel B of Table 5, we show that the results also hold under an alternate construction of the temperature and precipitation variables — here, we take the average weather across all grid points that overlap with a district's boundary.

6 Conclusion

As temperatures rise, productivity in the agricultural sector will drop relative to productivity in other sectors, and it may be beneficial for individuals to move from agricultural to non-agricultural sectors and/or from rural to urban areas. Earlier work on India has demonstrated that individuals do switch sectors in response to short-term weather shocks, and that such switching has important economic benefits (Emerick, 2018; Colmer, 2019).

In this paper, we add to this base of knowledge by exploring responses to slow-onset changes in temperature, measured using decadal averages. Exploiting panel data spanning a 60-year period, we find that higher temperatures inhibit urbanization and structural transformation in Indian districts, and that effects are heterogeneous across space. Specifically, in districts with low road density, higher temperatures lead to an increase in the share of the population engaged in agriculture and a reduction in urbanization rates and the share of the population engaged in non-agriculture. On the other hand, in districts with high road density, higher temperatures lead to a decline in the share of the population engaged in agri-

[18] Our results are also robust to using alternate (e.g. six-, seven-, eight-year etc. long) windows.

culture, while urbanization rates and the share of the population engaged in non-agriculture remain unchanged.

Our results add to a growing climate–economy literature, and support the idea that public infrastructure can be a key element to facilitate private individual-level adaptation to climate change. Future work should explore the interaction of higher temperatures, structural transformation and transport infrastructure in other low- and middle-income countries. Another beneficial avenue for future work would be to investigate the role of other public infrastructure and publicly funded assets in facilitating individual adaptation to climate change.

References

Asher, S. and Novosad, P. (2019). Rural roads and local economic development. Retrieved from `http://www.dartmouth.edu/~novosad/asher-novosad-roads.pdf`.

Auffhammer, M. and Schlenker, W. (2014). Empirical studies on agricultural impacts and adaptation. *Energy Economics*, 46(C):555–561.

Azam, M. (2012). Changes in wage structure in urban india, 1983–2004: A quantile regression decomposition. *World Development*, 40(6):1135–1150.

Banerjee, A. V. and Newman, A. F. (1993). Occupational choice and the process of development. *Journal of Political Economy*, 101(2):274–298.

Binswanger-Mkhize, H. P. (2013). The stunted structural transformation of the indian economy: Agriculture, manufacturing and the rural non-farm sector. *Economic and Political Weekly*, pages 5–13.

Borjas, G. J. (1987). Self-selection and the earnings of immigrants. *The American Economic Review*, pages 531–553.

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

Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2017). Weather, climate change and death in India. Retrieved from `https://pdfs.semanticscholar.org/a6d6/dfe69800c66ae6bcd28fb631bf266b048c1d.pdf`.

Burgess, R. and Donaldson, D. (2010). Can openness mitigate the effects of weather shocks? Evidence from India's famine era. *American Economic Review*, 100(2):449–453.

Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140.

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

Colmer, J. (2019). Temperature, labor reallocation, and industrial production: Evidence from India. Retrieved from `https://drive.google.com/file/d/0B-BakBtoHwF8UjNtU3ZONmdjOTA/view`.

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

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature*, 52(3):740–798.

Emerick, K. (2018). Agricultural productivity and the sectoral reallocation of labor in rural India. *Journal of Development Economics*, pages 1–64.

Fetzer, T. (2014). Social insurance and conflict: Evidence from India. Retrieved from `https://www.cemfi.es/ftp/pdf/papers/Seminar/fetzer.pdf`.

Gertler, P. J., Gonzalez-Navarro, M., Gracner, T., and Rothenberg, A. D. (2019). Road quality, local economic activity, and welfare: Evidence from Indonesia's highways. Retrieved from `https://are.berkeley.edu/sites/are.berkeley.edu/files/user/profile2/main/publications/Indonesia_Roads.pdf`.

Gollin, D., Lagakos, D., and Waugh, M. E. (2014). The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993.

Gollin, D. and Rogerson, R. (2014). Productivity, transport costs and subsistence agriculture. *Journal of Development Economics*, 107(C):38–48.

Gouel, C. and Laborde, D. (2018). The crucial role of international trade in adaptation to climate change. Retrieved from `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3278521`.

Grogger, J. and Hanson, G. H. (2011). Income maximization and the selection and sorting of international migrants. *Journal of Development Economics*, 95(1):42–57.

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

Hornbeck, R. (2012). The enduring impact of the American Dust Bowl: Short-and long-run adjustments to environmental catastrophe. *American Economic Review*, 102(4):1477–1507.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372.

ICRISAT (2015). Meso level data for India, 1966-2011: Collected and compiled under the project on Village Dynamics in South Asia.

Jain, A., O'Sullivan, R., and Taraz, V. (2019). Temperature and economic activity: Evidence from India. Retrieved from `https://www.vistaraz.com/s/JainOSullivanTaraz_TemperatureActivity_Nov2019.pdf`.

Kijima, Y. (2006). Why did wage inequality increase? Evidence from urban India 1983–99. *Journal of Development Economics*, 81(1):97–117.

Kone, Z. L., Liu, M. Y., Mattoo, A., Özden, Ç., and Sharma, S. (2018). Internal borders and migration in India. *Journal of Economic Geography*, 125:F49–31.

Kumar, H. and Somanathan, R. (2017). Creating long panels using Census data (1961–2001). *Economic and Political Weekly*, 52(29).

Lewis, W. A. (1954). Economic development with unlimited supplies of labour. *The Manchester School*, 22(2):139–191.

McMillan, M. S., Rodrik, D., and Verduzco-Gallo, I. (2014). Globalization, structural change and productivity growth, with an update on Africa. *World Development*, 63:11–32.

Moore, F. C. and Lobell, D. B. (2014). Adaptation potential of European agriculture in response to climate change. *Nature Climate Change*, 4(7):610–614.

Munshi, K. and Rosenzweig, M. (2016). Networks and misallocation: Insurance, migration, and the rural-urban wage gap. *American Economic Review*, 106(1):46–98.

Peri, G. and Sasahara, A. (2019). The impact of global warming on rural-urban migrations: Evidence from global big data. Retrieved from `https://www.nber.org/papers/w25728.pdf`.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2):135–146.

Schlenker, W. and 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.

Shamdasani, Y. (2019). Rural Road Infrastructure and Agricultural Production: Evidence from India. Retrieved from `https://drive.google.com/file/d/1e35zXIFMl2bFGydyG7OJcW3FAChmN9U8/view`.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2015). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. Retrieved from `https://www.isid.ac.in/~pu/dispapers/dp14-10.pdf`.

Taraz, V. (2018). Can farmers adapt to higher temperatures? Evidence from India. *World Development*, 112:205–219.

Vanneman, R. and Barnes, D. (2000). Indian District Data, 1961-1991: Machine-readable data file and code book. Retrieved from `http://vanneman.umd.edu/districts/index.html`.

Young, A. (2013). Inequality, the urban-rural gap, and migration. *The Quarterly Journal of Economics*, 128(4):1727–1785.

(a) Consistent district boundaries, 1961-2011

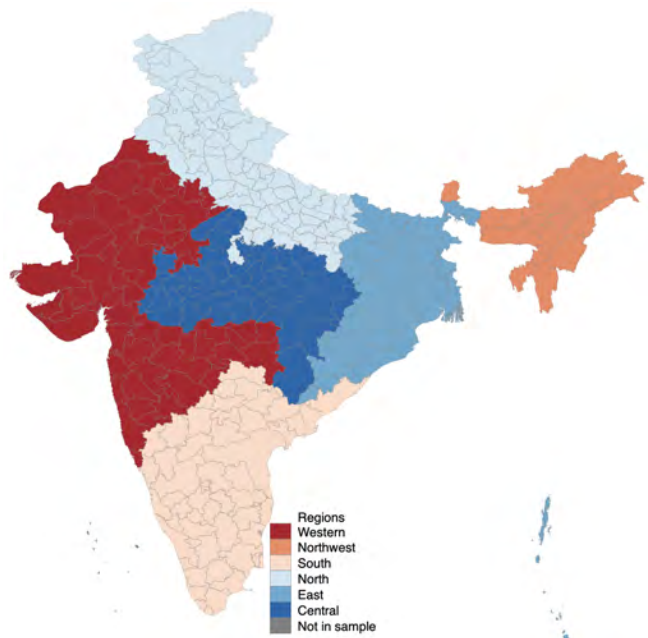

(b) Districts in agricultural states, by region

Figure 1: Figure illustrates the 288 consistent district boundaries over 1961-2011 in panel a, and the 287 districts by region used in the analysis in panel b.

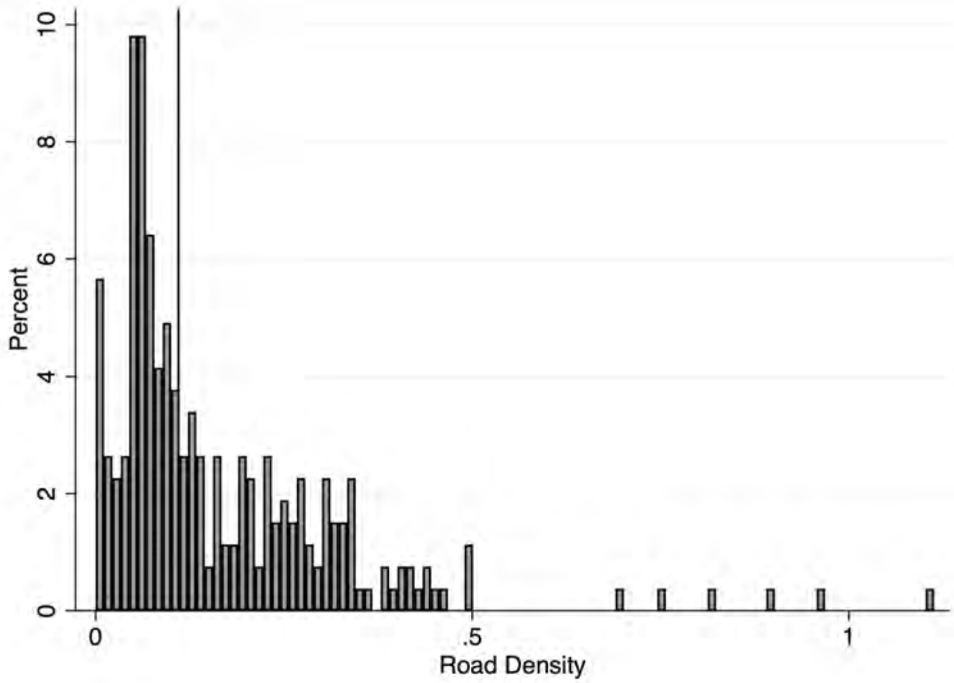

(a) Histogram

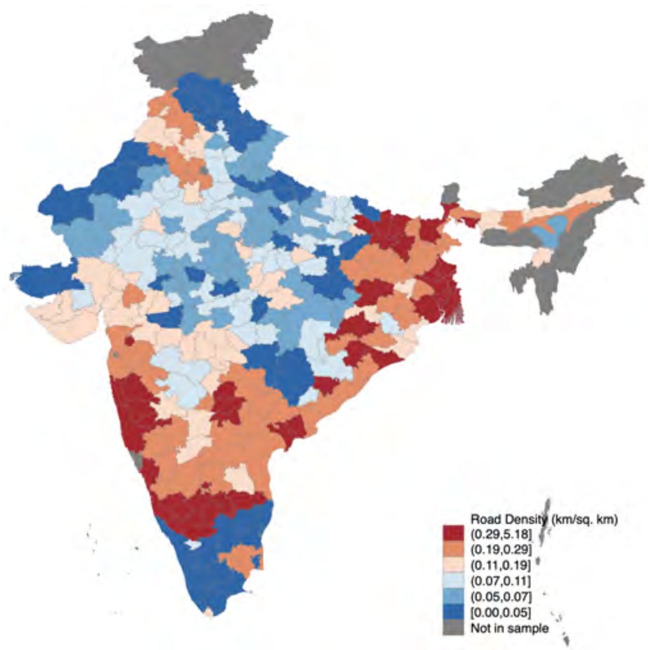

(b) Spatial Distribution

Figure 2: Figure plots the road density (km/km^2) measure across all districts in panel a, and illustrates the distribution of the same measure across space in panel b. The solid vertical line in panel a denotes the median in the distribution ($0.11 km/km^2$).

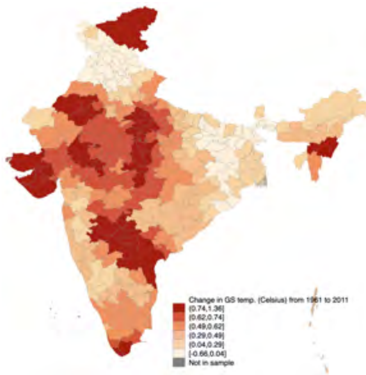

(a) Temperature

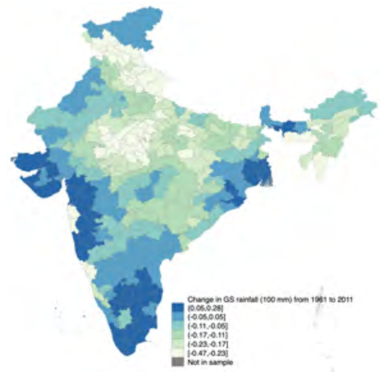

(b) Precipitation

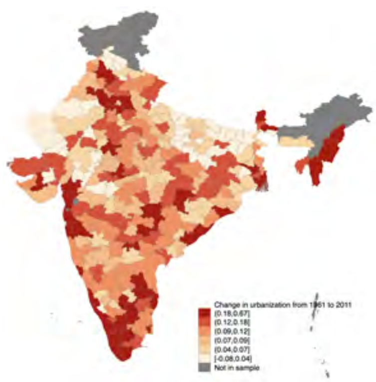

(c) Urbanization

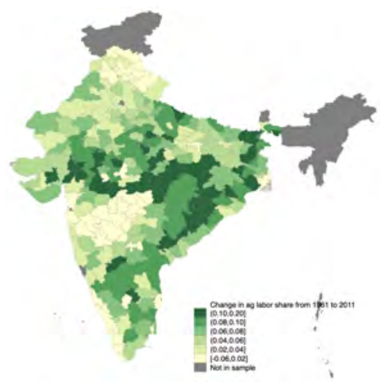

(d) Agricultural Labor Share

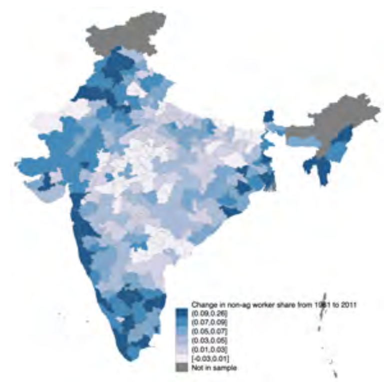

(e) Non-Agricultural Worker Share

Figure 3: Figure illustrates long-run changes in the ten-year average growing season temperature (panel a) and precipitation (panel b), urbanization (panel c), share of agricultural laborers (panel d) and share of non-agricultural workers (panel e) across all balanced districts. These changes are computed by subtracting the value of each variable in 1961 from the corresponding value in 2011.

Table 1: Summary Statistics by Year

Panel A: All districts	1961	1971	1981	1991	2001	2011	Total
Road Density (continuous)	-	-	-	-	-	-	0.183
							(0.352)
10-Year Avg. GS Temperature (Celsius)	23.88	23.97	24.04	24.05	24.11	24.31	24.06
	(2.992)	(3.056)	(3.094)	(3.121)	(3.138)	(3.117)	(3.085)
10-Year Avg. GS Rainfall (100 mm)	1.246	1.123	1.279	1.153	1.123	1.153	1.180
	(0.608)	(0.560)	(0.601)	(0.575)	(0.566)	(0.608)	(0.589)
Urbanization	0.155	0.168	0.198	0.219	0.235	0.263	0.206
	(0.123)	(0.126)	(0.130)	(0.137)	(0.150)	(0.164)	(0.144)
Agricultural Labor Share	0.0708	0.0985	0.0952	0.0879	0.108	0.130	0.0983
	(0.0511)	(0.0528)	(0.0562)	(0.0514)	(0.0575)	(0.0661)	(0.0588)
Non-Agricultural Worker Share	0.115	0.0907	0.102	0.108	0.149	0.165	0.122
	(0.0503)	(0.0475)	(0.0497)	(0.0523)	(0.0646)	(0.0692)	(0.0621)
Panel B: Districts with below median road density	1961	1971	1981	1991	2001	2011	Total
Road Density (continuous)	-	-	-	-	-	-	0.0552
							(0.0287)
10-Year Avg. GS Temperature (Celsius)	23.29	23.42	23.45	23.49	23.55	23.80	23.50
	(3.920)	(4.009)	(4.038)	(4.073)	(4.099)	(4.082)	(4.027)
10-Year Avg. GS Rainfall (100 mm)	1.201	1.070	1.271	1.100	1.063	1.058	1.127
	(0.464)	(0.430)	(0.523)	(0.435)	(0.457)	(0.478)	(0.471)
Urbanization	0.142	0.156	0.186	0.208	0.221	0.244	0.193
	(0.114)	(0.118)	(0.120)	(0.127)	(0.136)	(0.155)	(0.133)
Agricultural Labor Share	0.0570	0.0833	0.0749	0.0703	0.0923	0.118	0.0826
	(0.0407)	(0.0471)	(0.0498)	(0.0425)	(0.0528)	(0.0642)	(0.0535)
Non-Agricultural Worker Share	0.110	0.0853	0.0951	0.101	0.134	0.152	0.113
	(0.0499)	(0.0470)	(0.0493)	(0.0507)	(0.0600)	(0.0657)	(0.0588)
Panel C: Districts with above median road density	1961	1971	1981	1991	2001	2011	Total
Road Density (continuous)	-	-	-	-	-	-	0.316
							(0.467)
10-Year Avg. GS Temperature (Celsius)	24.49	24.55	24.66	24.64	24.69	24.85	24.65
	(1.247)	(1.303)	(1.359)	(1.400)	(1.404)	(1.393)	(1.353)
10-Year Avg. GS Rainfall (100 mm)	1.294	1.179	1.287	1.208	1.186	1.252	1.234
	(0.727)	(0.667)	(0.675)	(0.691)	(0.657)	(0.708)	(0.687)
Urbanization	0.168	0.181	0.211	0.230	0.250	0.283	0.220
	(0.131)	(0.134)	(0.140)	(0.147)	(0.162)	(0.172)	(0.153)
Agricultural Labor Share	0.0852	0.114	0.116	0.106	0.125	0.142	0.115
	(0.0569)	(0.0539)	(0.0548)	(0.0536)	(0.0577)	(0.0661)	(0.0597)
Non-Agricultural Worker Share	0.120	0.0963	0.109	0.117	0.165	0.179	0.131
	(0.0504)	(0.0477)	(0.0494)	(0.0528)	(0.0656)	(0.0705)	(0.0641)

Note: Table presents summary statistics for the road density measure, weather variables and Census outcome variables over time for the full sample of balanced districts (Panel A, N=270), for below median road density districts (Panel B, N=132), and for above median road density districts (Panel C, N=126). Road density data is missing for 12 districts in the sample.

Table 2: Effect of Rising Temperatures on Urbanization and Structural Transformation

	Urbanization		Ag Labor Share		Non-Ag Worker Share	
	(1)	(2)	(3)	(4)	(5)	(6)
ln T	-0.7318	-1.3525	2.2330	1.4282	-2.0964	-1.7737
	(0.7552)	(0.4411)***	(0.6109)***	(0.5722)**	(0.4905)***	(0.5108)***
	[0.7602]	[0.4741]***	[0.8501]***	[0.7950]*	[0.4715]***	[0.4957]***
ln P	0.0130	-0.0148	-0.2711	-0.1891	-0.0103	-0.0484
	(0.0509)	(0.0505)	(0.0692)***	(0.0677)***	(0.0350)	(0.0356)
	[0.0533]	[0.0498]	[0.1179]**	[0.1023]*	[0.0377]	[0.0333]
State-specific linear time trends	N	Y	N	Y	N	Y
Observations	1,596	1,596	1,542	1,542	1,614	1,614

Note: The dependent variable is the natural logarithm of urbanization rates in Columns (1) and (2), of the share of agricultural laborers in Columns (3) and (4), and of the share of non-agricultural workers in Columns (5) and (6). All columns include district and region-by-year fixed effects. We restrict our sample to districts for which the dependent variable is non-missing in all years. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effect of Rising Temperatures on Urbanization and Structural Transformation: Heterogeneity by Road Network Density

	Urbanization		Ag Labor Share		Non-Ag Worker Share	
	(1)	(2)	(3)	(4)	(5)	(6)
ln T	-1.3010	-2.0088	2.7877	2.2242	-2.1518	-2.2030
	(0.6473)**	(0.4403)***	(0.7926)***	(0.5915)***	(0.5984)***	(0.6564)***
	[0.7887]*	[0.5647]***	[1.0664]***	[0.8645]**	[0.5646]***	[0.6140]***
ln T x High Road Density	2.0247	2.5960	-4.6168	-5.5046	0.6268	2.3359
	(1.4815)	(1.3755)*	(1.8081)**	(1.8139)***	(1.0633)	(1.0973)**
	[1.5052]	[1.3326]*	[2.5192]*	[2.1091]***	[1.3125]	[1.0994]**
ln P	-0.0552	-0.1076	-0.3108	-0.1860	0.0230	-0.0142
	(0.0788)	(0.0803)	(0.0995)***	(0.0966)*	(0.0486)	(0.0484)
	[0.0730]	[0.0717]	[0.1626]*	[0.1457]	[0.0455]	[0.0437]
ln P x High Road Density	0.1721	0.1870	0.1860	0.0603	-0.0372	-0.0157
	(0.1081)	(0.1108)*	(0.1439)	(0.1315)	(0.0750)	(0.0728)
	[0.1003]*	[0.0932]**	[0.1918]	[0.1645]	[0.0770]	[0.0686]
State-specific linear time trends	N	Y	N	Y	N	Y
P-val, cluster: ln T + ln T x Road	0.5933	0.6539	0.2683	0.0566	0.0946	0.8855
P-val, Conley: ln T + ln T x Road	0.5884	0.6391	0.3804	0.0831	0.1966	0.8884
Observations	1,530	1,530	1,542	1,542	1,542	1,542

Note: The dependent variable is the natural logarithm of urbanization rates in Columns (1) and (2), of the share of agricultural laborers in Columns (3) and (4), and of the share of non-agricultural workers in Columns (5) and (6). *High Road Density* is a binary variable that takes the value 1 if the district has above median road density at baseline. All columns include district and region-by-year fixed effects. We restrict our sample to districts for which the dependent variable is non-missing in all years. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Robustness Checks — Panel Specification

	Urbanization		Ag Labor Share		Non-Ag Worker Share	
Panel A: Unbalanced Panel	(1)	(2)	(3)	(4)	(5)	(6)
ln T	-1.352***	-2.232***	1.570***	2.221***	-1.840***	-2.193***
	(0.477)	(0.553)	(0.573)	(0.590)	(0.499)	(0.655)
ln T x High Road Density		3.018**		-6.246***		2.214**
		(1.400)		(1.912)		(1.055)
P-val: ln T + ln T x Road		0.545		0.0281		0.980
Observations	1684	1574	1680	1582	1694	1582
Panel B: Road Density Measure — 30th Percentile Cutoff						
ln T	-1.353***	-1.683***	1.428**	1.982***	-1.774***	-2.044***
	(0.441)	(0.384)	(0.572)	(0.523)	(0.511)	(0.602)
ln T x High Road Density		1.433		-5.245**		2.351*
		(1.515)		(2.235)		(1.301)
P-val: ln T + ln T x Road		0.865		0.132		0.796
Observations	1596	1530	1542	1542	1614	1542
Panel C: Clustering SE at State Level						
ln T	-1.353***	-2.009***	1.428	2.224**	-1.774***	-2.203***
	(0.463)	(0.676)	(0.900)	(0.834)	(0.524)	(0.659)
ln T x High Road Density		2.596*		-5.505*		2.336*
		(1.413)		(2.933)		(1.152)
P-val: ln T + ln T x Road		0.647		0.249		0.895
Observations	1596	1530	1542	1542	1614	1542

Note: The dependent variable is the natural logarithm of urbanization rates in Columns (1) and (2), of the share of agricultural laborers in Columns (3) and (4), and of the share of non-agricultural workers in Columns (5) and (6). *High Road Density* is a binary variable that takes the value 1 if the district has above median road density at baseline in Panels A and C. All columns include district, region-by-year fixed effects and state-specific linear time trends. We restrict our sample to districts for which the dependent variable is non-missing in all years in Panels B and C. We present standard errors clustered by district in Panels A and B, and by state in Panel C. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Robustness Checks — Alternative Definitions

	Urbanization		Ag Labor Share		Non-Ag Worker Share	
Panel A: 5-Year Average T & P	(1)	(2)	(3)	(4)	(5)	(6)
ln T	-0.881*	-1.364***	1.282***	2.048***	-1.822***	-2.309***
	(0.455)	(0.409)	(0.449)	(0.467)	(0.382)	(0.504)
ln T x High Road Density		1.487		-3.772***		2.031**
		(1.017)		(1.325)		(0.860)
P-val: ln T + ln T x Road		0.896		0.157		0.700
Observations	1596	1530	1542	1542	1614	1542
Panel B: Grid Point Average T & P						
ln T	-1.414*	-2.617***	0.715	1.124	-1.633***	-1.920***
	(0.743)	(0.690)	(0.529)	(0.759)	(0.254)	(0.171)
ln T x High Road Density		3.395**		-4.460**		2.428***
		(1.493)		(1.870)		(0.935)
P-val: ln T + ln T x Road		0.559		0.0535		0.585
Observations	1560	1506	1518	1518	1572	1518

Note: The dependent variable is the natural logarithm of urbanization rates in Columns (1) and (2), of the share of agricultural laborers in Columns (3) and (4), and of the share of non-agricultural workers in Columns (5) and (6). *High Road Density* is a binary variable that takes the value 1 if the district has above median road density at baseline. All columns include district, region-by-year fixed effects and state-specific linear time trends. We restrict our sample to districts for which the dependent variable is non-missing in all years. We present standard errors clustered by district in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A Appendix: Supplementary Figures

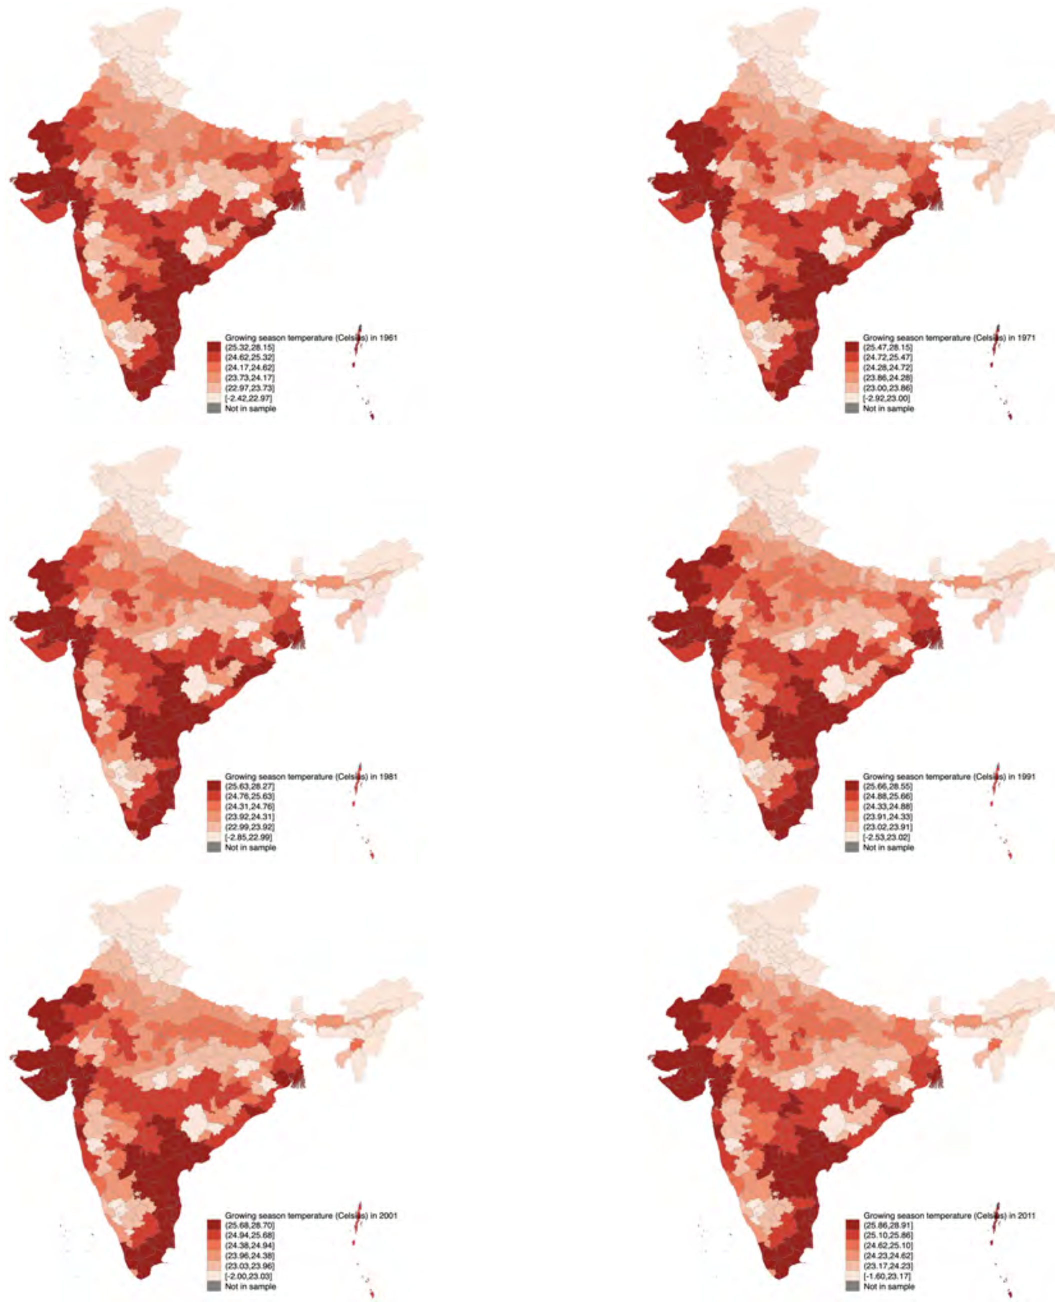

Figure A1: Average growing season (June through February) temperature in Celsius (10-year lagged)

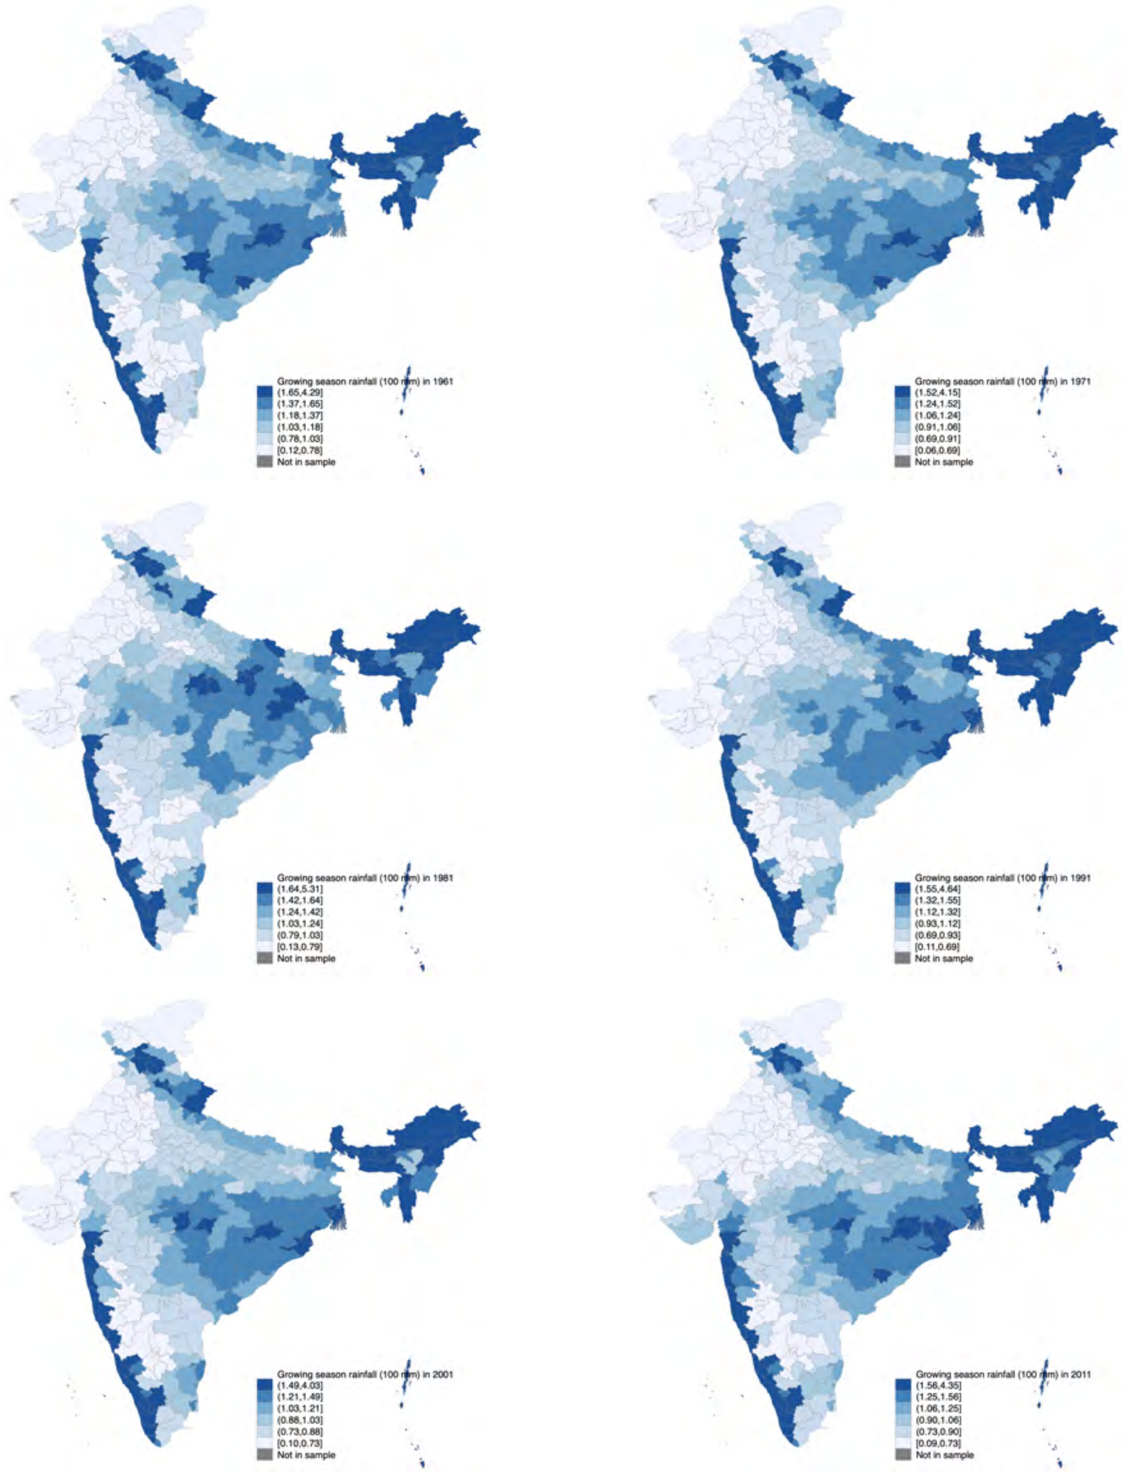

Figure A2: Average growing season (June through February) precipitation in 100 mm (10-year lagged)

Figure A3: Urbanization

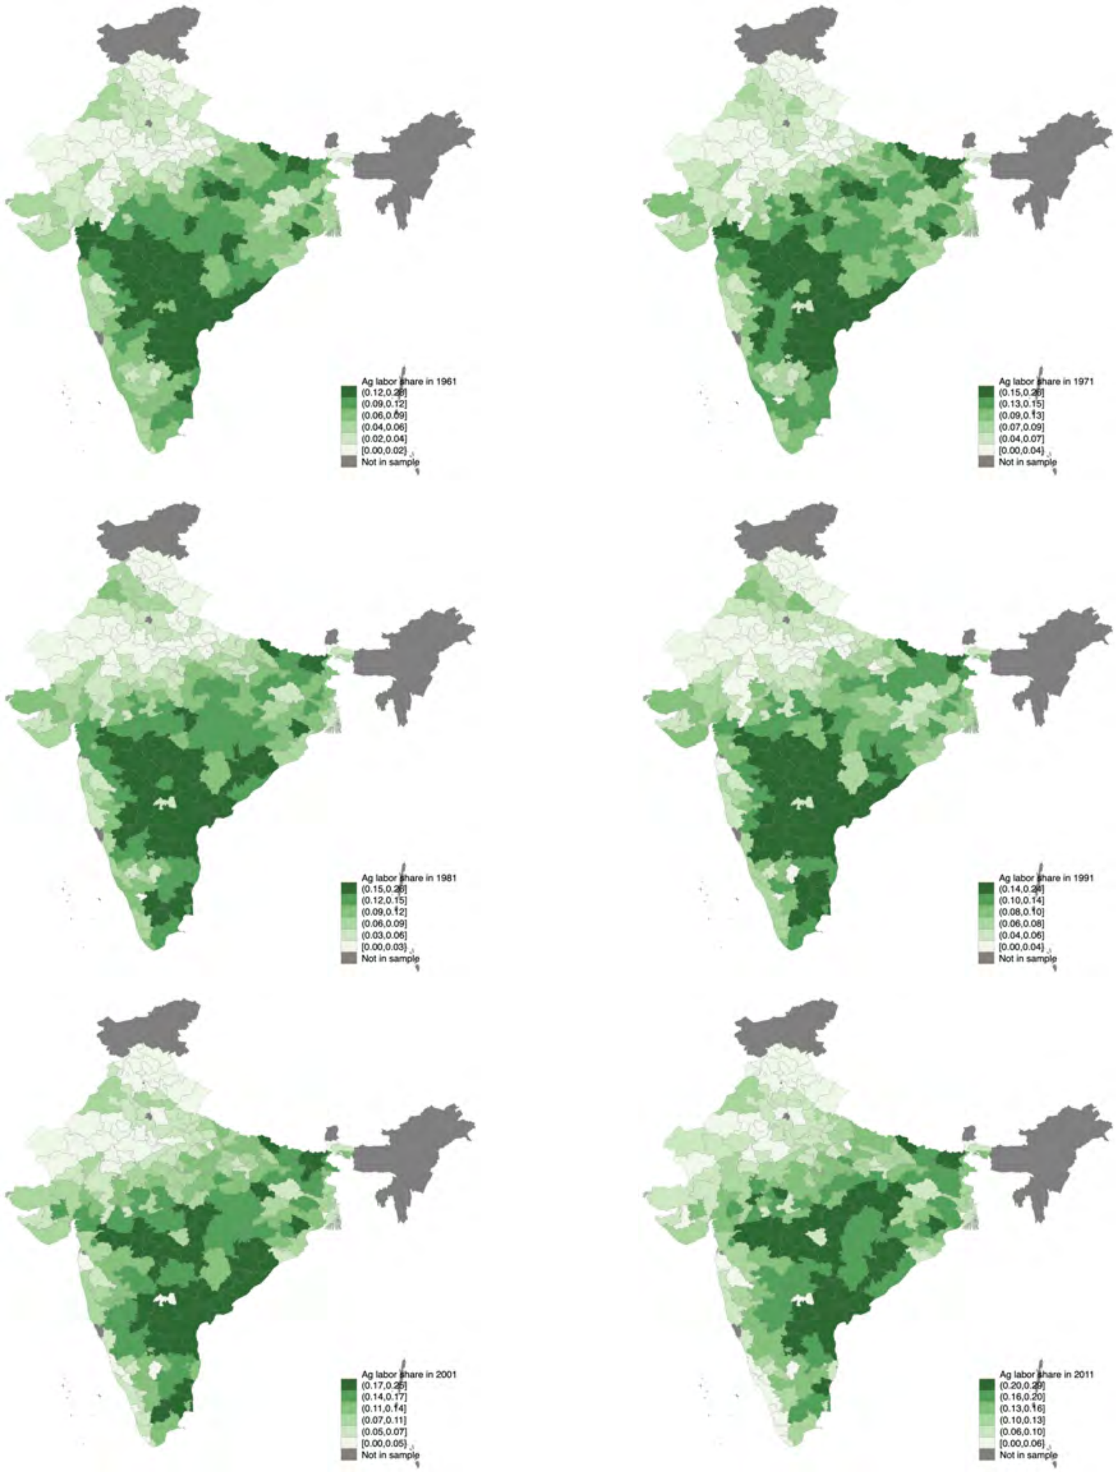

Figure A4: Agricultural Labor Share

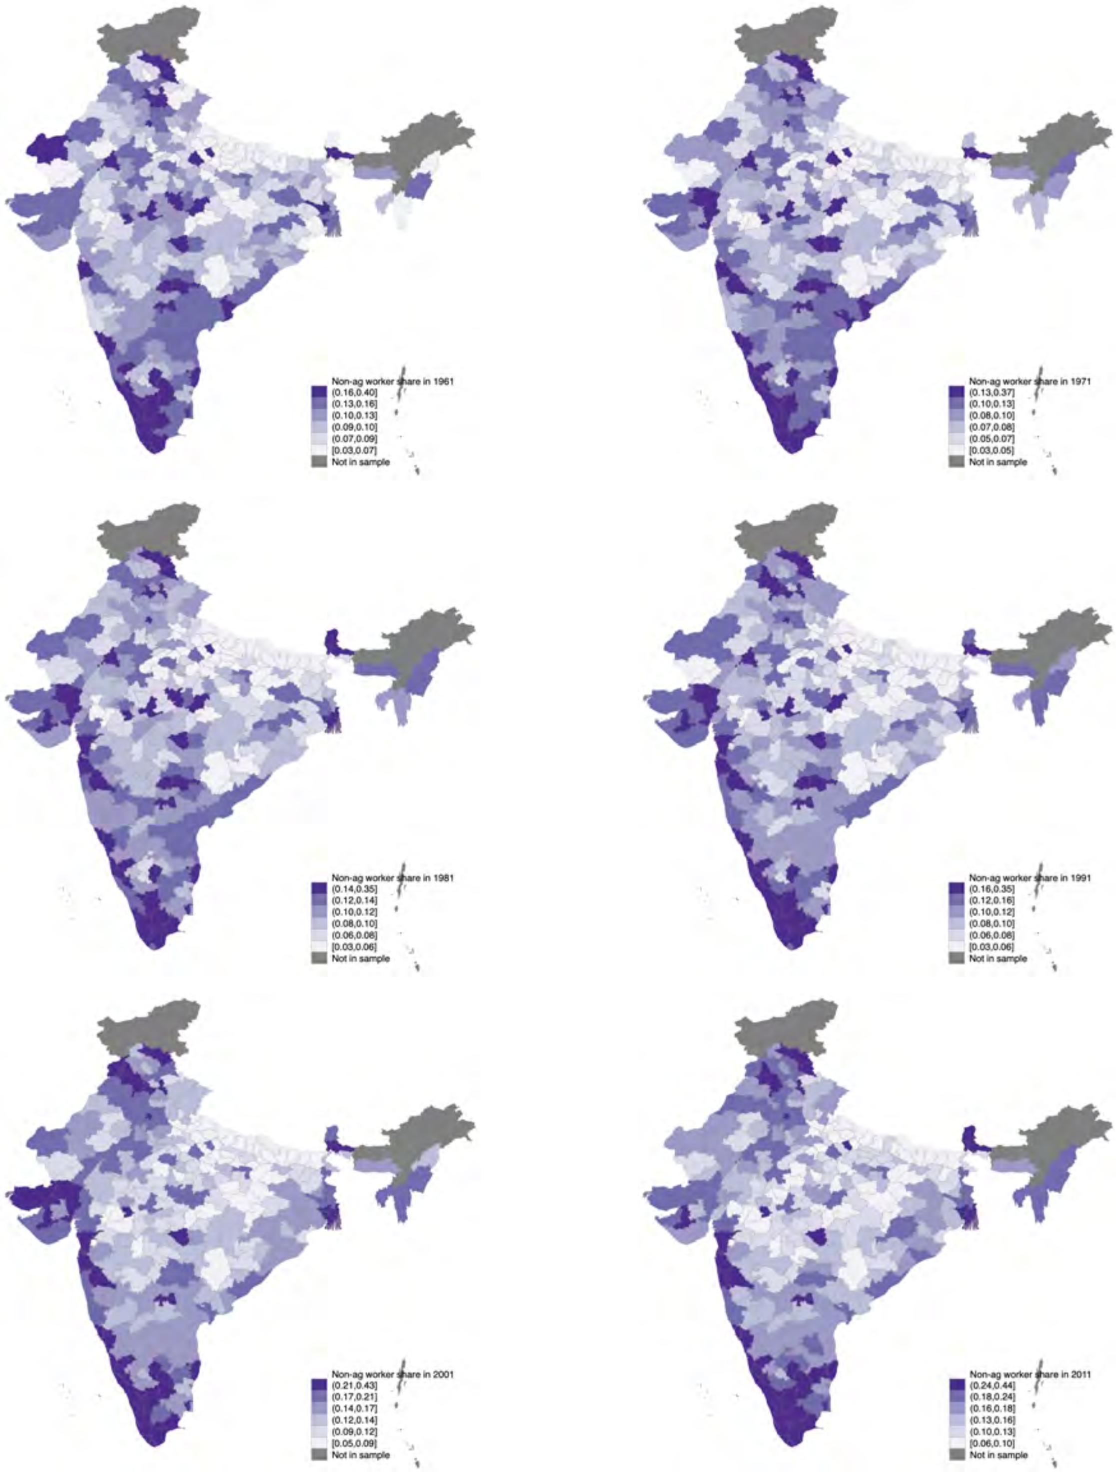

Figure A5: Non-Agricultural Labor Share

B Appendix: Supplementary Tables

Table A1: Effect of Rising Temperatures on Yields

	Aggregate		Rice		Wheat	
	(1)	(2)	(3)	(4)	(5)	(6)
ln T	-0.4807	-0.5860	-0.9120	-0.9029	-0.3329	-0.5082
	(0.1673)***	(0.2350)**	(0.4690)*	(0.4765)*	(0.1361)**	(0.2069)**
	[0.1830]***	[0.2279]**	[0.4563]**	[0.4581]**	[0.1605]**	[0.2062]**
ln P	0.1909	0.2083	0.2027	0.2240	0.0508	0.0784
	(0.0229)***	(0.0231)***	(0.0244)***	(0.0234)***	(0.0142)***	(0.0147)***
	[0.0264]***	[0.0258]***	[0.0326]***	[0.0307]***	[0.0196]***	[0.0187]***
State-specific linear time trends	N	Y	N	Y	N	Y
Observations	11,820	11,820	11,117	11,117	10,343	10,343

Note: The dependent variable is the natural logarithm of aggregate yields in Columns (1) and (2), of rice yields in Columns (3) and (4), and of wheat yields in Columns (5) and (6). All columns include district and year fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$