Climate Change and Labor Reallocation: Evidence from Six Decades of the Indian Census

Maggie Y. Liu, Yogita Shamdasani, Vis Taraz*

December 1, 2021

Abstract

How do rising temperatures affect long-term labor reallocation in developing economies? In this paper, we examine how increases in temperature impact structural transformation and urbanization within Indian districts between 1951 and 2011. We find that rising temperatures are associated with lower shares of workers in non-agriculture, with effects intensifying over a longer time frame. Supporting evidence suggests that local demand effects play an important role: declining agricultural productivity under higher temperatures reduces the demand for non-agricultural goods and services, which subsequently lowers non-agricultural labor demand. Our results illustrate that rising temperatures limit sectoral and rural-urban mobility for isolated households.

JEL Classification: O12, O13, Q54, R23

Keywords: climate change, structural transformation, labor reallocation, urbanization, migration

*Liu: Smith College, Department of Economics, Wright Hall, 5 Chapin Drive, Northampton, MA 01063-6317 (email: yliu77@smith.edu). Shamdasani: National University of Singapore, Department of Economics, 1 Arts Link, Singapore 117570 (email: yogita@nus.edu.sg). 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 Scott Fulford for sharing data and Christine Cai for sharing code. We are grateful for feedback from Emily Breza, Jonathan Colmer, Kyle Emerick, John Hoddinott, Deepak Saraswat and numerous seminar and conference participants. We are also grateful to the editor and the anonymous referees who helped us improve the paper significantly.

1 Introduction

Developing economies are commonly characterized by large productivity gaps across the economy — for example, between agricultural and non-agricultural sectors (McMillan et al., 2014; Gollin et al., 2014; Herrendorf and Schoellman, 2018) and between rural and urban areas (Lewis, 1954; Young, 2013). Reallocation of workers both across sectors and space could thus be beneficial for economic development if it allows for a more efficient allocation of human capital.[1] While weather shocks favorable to agricultural productivity have been shown to encourage sectoral reallocation (Emerick, 2018), higher temperatures have adverse impacts on agricultural incomes and productivity (Schlenker and Roberts, 2009; Taraz, 2018; Aragón et al., 2021). Thus, it is possible that rising temperatures under climate change may hinder labor reallocation in developing economies where the vast majority of workers engage in rural agriculture.[2]

In light of this, this paper addresses a critical empirical question: do rising temperatures affect the pace of reallocation of workers within local labor markets, namely through structural transformation and urbanization? A key feature of the empirical exercise we undertake is our focus on responses to medium- to long-term climate variations — which capture adaptation or intensification effects more accurately than short-term weather fluctuations — thus better approximating impacts of future anthropogenic climate change (Dell et al., 2014; Burke and Emerick, 2016; Kolstad and Moore, 2020).

We address this question in the context of India, where rural-urban mo-

[1]Labor reallocation would bring meaningful aggregate productivity gains if productivity gaps are not driven by self selection, i.e., workers with different skills/ability sorting into certain sectors (Roy, 1951). Recent work using panel data and specifications with individual fixed effects has documented that the size of observed sectoral productivity gaps is smaller than previous estimates that have typically relied on cross-sectional data (Hamory et al., 2021; Alvarez, 2020; Lagakos et al., 2020).

[2]For sectoral or spatial arbitrage to take place, there must be both an availability of non-agricultural jobs with higher pay and a sufficiently low cost of switching between sectors or regions. Negative shocks to agricultural incomes could spill over to other sectors, impacting demand for non-agricultural labor. In addition, with declining agricultural incomes, a combination of mobility costs and liquidity constraints could present a larger barrier to labor reallocation.

bility is low (Munshi and Rosenzweig, 2016) and structural transformation, particularly the movement from agriculture to manufacturing, is slow and "stunted" (Binswanger-Mkhize, 2013). We are particularly interested in local labor market responses, so we focus on sectoral and rural-urban movements *within* Indian districts. Local movements within a district likely present those impacted by climate change with low-cost immediate adaptation strategies; our analysis thus captures switching decisions of potential movers on the margin.

We develop a simple general equilibrium model that explores the relationship between agricultural productivity and sectoral allocation of labor. In developing countries, agriculture constitutes a substantial share of the local economy and provides an important source of income for a large share of the population. Productivity shocks to the agricultural sector can thus be important for other sectors of the economy in general equilibrium. Under this scenario, adverse agricultural productivity shocks may reduce the demand for non-agricultural labor through its impact on consumption of local non-agricultural goods and services. We capture this dynamic — which we refer to as "local demand effects" — in the model. The model predicts that declines in agricultural productivity can lead to increases in agricultural labor supply if the output elasticity of agriculture with respect to labor is less than one, and if the income elasticity of demand for the agricultural good is also less than one.

To test the model's predictions, we assemble a district-level panel data set spanning 1951 to 2011, combining measures of worker shares across agricultural and non-agricultural sectors, urbanization rates, and intra-district rural-to-urban migrant shares with decadal temperature and precipitation variables constructed using monthly gridded weather data. Our use of sub-national data allows us to track *local* labor market responses, focusing on movements within a district — a margin that has been shown to be especially important in the context of India (Kone et al., 2018).

We adopt two specifications to quantify the impact of rising temperatures on structural transformation over different temporal scales. First, we estimate a panel fixed effect model that exploits decade-to-decade fluctuations in weather. 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-time fixed effects, decade-to-decade fluctuations in weather are quasi-random. Second, following Burke and Emerick (2016), we estimate a long differences model that exploits variation in long-term temperature trends over the span of our sixty-year data set, conditional on regional trends.

We find two main results. First, we find that rising temperatures inhibit structural transformation in Indian districts. The magnitude of this effect — estimated using a decadal panel specification — is economically meaningful: a 1°C increase in mean decadal temperature in an average Indian district leads to a 17.0% increase in the share of the labor force who are agricultural laborers, and a 8.2% decline in the share of the labor force engaged in non-agriculture. In contrast, we find no detectable impact of rising temperatures on urbanization and intra-district rural-to-urban migration. This latter result is in line with Henderson et al. (2017), who find no average impact of adverse changes in climate on urbanization in Sub-Saharan Africa.

Second, we find that the adverse effects of rising temperatures on structural transformation are amplified when we examine impacts over a longer time frame. The larger estimates from our long differences specification suggest that not only do individuals fail to adapt to sustained, higher temperatures, they also experience intensification of these adverse impacts over time. Our results are consistent with sustained warming making reallocation infeasible for a larger share of households.

We find evidence consistent with our results being driven by local demand effects. In particular, declining farm incomes arising from lower agricultural productivity lead to a contraction in demand for local non-agricultural goods and services, and this in turn leads to a reduction in non-agricultural labor demand. We document a decline in both food and non-food consumption among households in response to rising temperatures, consistent with a contraction in local demand. We also find that the decline in the share of the labor force engaged in non-agriculture is driven by reductions in the service sector, which is by nature non-tradable.

We also investigate several other potential mechanisms through which rising temperatures could affect labor reallocation. Rising temperatures could generate labor productivity losses of differential magnitudes across agricultural and non-agricultural sectors, which in turn could impact relative labor demand. Alternatively, rising temperatures could impact labor reallocation under a scenario where mobility costs are non-zero, as income losses driven by higher temperatures could make it more difficult for workers to afford these costs. Rising temperatures could also affect labor reallocation through consolidation of agricultural land or shifts in cultivation practices. Additional empirical analysis yields weak evidence in support of these alternative mechanisms playing an important role.

The main contribution of this paper is to shed light on sectoral and spatial reallocation of workers in response to slow-onset changes in climate in a developing country. The bulk of the existing literature on labor reallocation and rural-to-urban migration in developing countries uses short-run variation in weather to estimate effects (Mueller et al., 2014; Bohra-Mishra et al., 2014; Maystadt et al., 2016; Emerick, 2018; Jessoe et al., 2018; Colmer, 2021). While exploiting short-run weather variation is advantageous because it allows the econometrician to control for a number of potential confounding factors, the primary drawback of this method is that it provides well-identified estimates of the short-term effects of *weather* shocks, not climate, on outcome variables (Auffhammer, 2018; Kolstad and Moore, 2020).

Long-term responses to climate change may differ fundamentally from short-term responses to weather fluctuations because they account for potential adaptation or intensification driven by permanent changes in weather that may have taken place over time. Exploiting longer-term climate fluctuations thus provides a better estimate of how agents will respond to anthropogenic climate change (Burke and Emerick, 2016), yet to date this approach has not been deployed broadly in studies of developing countries.[3] Impacts of long-

[3]One notable exception is Henderson et al. (2017), who document that long-term in-

term climate shocks could differ substantially in developing settings as agents may have distinct or limited means of adaptation, due to features of under-development or high costs to reallocation. We thus extend this literature by assembling a unique panel data set that spans all districts in India over six decades in order to examine long-term responses to slow-onset changes in climate.

Our results on structural transformation are largely consistent with earlier works mentioned above that exploit short-run variation in weather, but there are some important nuances worth highlighting. Emerick (2018) finds that transitory high rainfall shocks in India increase the non-agricultural labor share due to increased demand for local non-tradables. Although Emerick (2018) analyzes precipitation rather than temperature, his findings are broadly consistent with ours — beneficial weather shocks decrease the share of the labor force engaged in agriculture, while adverse weather shocks increase it. Our results are also complementary to Jessoe et al. (2018), who find that transitory increases in extreme annual temperatures lead to a reduction in local non-agricultural wage employment in villages in rural Mexico, with effects operating through an agricultural channel. In contrast to our long-term results, Colmer (2021) finds that short-term increases in temperature are associated with a *reduction* in the agricultural labor share at the district level from 2004-2008. Using the same data source as Colmer (2021) over a longer time frame, we show that both his short-term result and our long-term result co-occur in the data. Our estimates using decade-to-decade variation in climate thus suggest that leaving agriculture could be a short-term adjustment, yet it may not be a viable option in the medium- to long-term.

We next consider our results on urbanization and migration in the context of the literature on rural-urban movements. Our results are complementary to Henderson et al. (2017), who document that long-term increases in dryness have no average impact on urbanization in Sub-Saharan Africa. We study a

creases in dryness have no average impact on urbanization in 29 countries in Africa. The focus of this paper is to estimate impacts of rising temperatures on structural transformation. For completeness, we also estimate impacts on urbanization and find null results, similar to those documented in Henderson et al. (2017).

similar time period and also find a null effect of rising temperatures on urbanization in India. Further, both our papers document impacts in response to adverse changes in climate that are spatially heterogeneous: Henderson et al. (2017) find that urbanization increases in a subset of regions with manufacturing centers, while we document a decline in urbanization in regions with sparse road networks. Our paper examines impacts on *sectoral* mobility, which allows us to explore how climate change interacts with the pace of structural transformation in a developing economy. We also study a longer time frame — our data spans 288 districts over sixty years — which enables us to quantify impacts over two different temporal scales and to explore whether impacts are attenuated or magnified in the longer-term.

The rest of this paper is organized as follows. In Section 2, we discuss several mechanisms through which rising temperatures could impact labor reallocation. 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, explore underlying mechanisms and present robustness checks. In Section 6, we conclude.

2 Mechanisms

In this section, we outline three potential mechanisms through which rising temperatures could affect labor reallocation in the short- and long-term.

Relative labor productivity loss Rising temperatures adversely impact labor productivity in both the agricultural sector (Schlenker and Roberts, 2009; Dell et al., 2012; Taraz, 2018) and the non-agricultural sector (Hsiang, 2010; Somanathan et al., 2021). In a two-sector economy with perfectly inelastic labor supply and zero costs to moving across sectors, if rising temperatures cause a relatively greater labor productivity loss in the agricultural sector, this would lead to a decline in demand for agricultural labor, and subsequently lower agricultural wages and employment. If, on the other hand, heat-induced labor productivity loss is relatively greater in the non-agricultural sector, then we would expect to find the opposite: rising temperatures would trigger a reduction in the demand for non-agricultural labor.

Local demand effects On the other hand, changes in local demand arising from a general equilibrium mechanism introduce additional effects on labor reallocation. For simplicity, we assume rising temperatures reduce only agricultural productivity and bring no productivity loss to the non-agricultural sector. If the income elasticity of non-agricultural goods is greater than that of agricultural goods, then a decline in farm incomes arising from lower agricultural productivity may result in a subsequent contraction in demand for local non-agricultural goods and services.[4] This in turn would trigger a reduction in demand for non-agricultural labor, and lead to more labor engaged in the agricultural sector.

In Appendix A, we develop a mathematical model that generates local demand effects. We are interested in predicting how a decrease in agricultural productivity induced by rising temperatures will affect labor supply to the non-agricultural sector. We are also interested in understanding how transportation costs might modulate this relationship. Our model includes a rural region, which has an agricultural sector and a service sector, and an urban region, which has a manufacturing sector and a service sector. There are iceberg costs to transporting the agricultural good, and the service good is not tradable across the two regions.

We assume that the elasticity of agricultural output with respect to labor is less than one, and that individuals in both regions have Stone-Geary preferences, with the income elasticity of the agricultural good being less than one. With these assumptions, our model delivers the prediction that a decrease in agricultural productivity will lead to a decline in the non-agricultural labor supply in the rural region. Furthermore, when we explore the role of iceberg transportation costs, we find that these local demand effects are concentrated in areas with higher transportation costs: for example, places with sparse road networks. The mechanisms underlying our theoretical model are similar to the mechanisms in many models of structural transformation (Matsuyama,

[4]Subsistence (or quasi-subsistence) agriculture is one scenario that can generate income elasticities such as these (Gollin and Rogerson, 2014).

1992; Kongsamut et al., 2001; Gollin et al., 2002; Gollin and Rogerson, 2014; Herrendorf et al., 2014; Bustos et al., 2016).

Liquidity and mobility costs Both mechanisms outlined above rely on an important assumption of perfect labor mobility across sectors. However, reallocation from the agricultural to non-agricultural sector cannot take place if liquidity and mobility costs are too high.[5] For example, a worker may incur travel costs (e.g. bus fares) to leave their village for a non-agricultural job in other parts of the district, or may require access to credit or savings to cover the upfront cost of job search or to sustain possible unemployment spells. Assuming rising temperatures disproportionately affect labor productivity in the agricultural sector, the resulting loss in farm income would make it more difficult for a worker to afford these costs, making sectoral movements to the relatively unaffected non-agricultural sectors less likely. These effects can be generated from a two-period Roy-Borjas model (Roy, 1951; Borjas, 1987) in which individuals reallocate their labor subject to incentive-compatibility and feasibility constraints.[6] Furthermore, we would expect these effects to be concentrated in places where liquidity and mobility costs are higher: for example, places with sparse road networks and/or poor access to bank credit.

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) and Migration (D-series) data tables, which provide district-level summaries of demographic and economic indicators. For each census year, we construct four key outcome measures at the district level:

[5] Liquidity costs could arise from credit and savings constraints (Bryan et al., 2014; Angelucci, 2015; Gazeaud et al., 2021), while mobility costs could arise from informational frictions or transportation costs (Gollin and Rogerson, 2014; Gertler et al., 2019; Shamdasani, 2021).

[6] See Cattaneo and Peri (2016) for an example of such a model applied to the migration decision.

the share of the labor force who are agricultural laborers, the share of the labor force who are non-agricultural workers, the share of the total population residing in urban areas, and the share of the male population who are rural-to-urban intra-district migrants.[7] Note that our measures of agricultural labor share and non-agricultural worker share are not perfectly collinear because we focus only on wage workers in agriculture, and thus exclude cultivators from our measure.[8]

We focus on sectoral and spatial movements *within* districts for two reasons. First, Census data limitations preclude us from being able to identify individuals who have switched sectors and moved to another district, nor can we identify the origin districts of individuals who have moved across districts. Second, cross-district migration rates in our empirical setting are low (Kone et al., 2018), therefore our analysis of within-district labor reallocation in response to climate shocks is unlikely to be confounded by spillovers between districts due to cross-district migration.

To account for the fact that districts split and boundaries are adjusted over time, we use concordance tables from Kumar and Somanathan (2017) and Singh et al. (2011) 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.[9] This results in 288 consistent districts, as illustrated in Appendix Figure B1a. The various splits and boundary changes between 1961 and 2011 can be deduced from the grey boundaries that trace out the 2011 Census districts delineation. Appendix Figure B1b

[7]Our migration measure based on last residence captures the most recent permanent migration, and is subject to under-counting seasonal or past moves. It should therefore be considered a lower bound of the true magnitude of migration.

[8]Recent work has also documented that cultivators in India have limited occupational and spatial mobility (Fernando, 2020).

[9]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.

highlights the six regions that span all the districts that form our analysis sample.

3.2 National Sample Survey Data

We supplement our Census data with rich individual- and household-level survey data from the National Sample Survey (NSS) spanning the years 1987-2012 — the longest time period for which publicly available NSS data contains district identifiers. We extract data from two NSS modules — the "employment and unemployment" module and the "consumer expenditure" module.

For each NSS year for which these modules are canvassed, we aggregate the individual/household-level data to generate district-level averages of the following outcomes: the share of the labor force who are agricultural workers, the share of the labor force who are non-agricultural workers, the share of the labor force who are engaged in manufacturing, services, and construction, as well as average total consumption, food consumption, and non-food consumption.[10] We harmonize the district-level panel data to have consistent district boundaries over time, applying the same methodology described above for the Census data.

3.3 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 (Matsuura and Willmott, 2015a,b).[11] We construct district-level weather data 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

[10] Agricultural worker share constructed using the NSS includes both agricultural laborers and cultivators; the NSS does not make a distinction between the two. Non-agricultural worker share constructed using the NSS is analogous to that constructed using Census data. We provide more details on the construction of each outcome variable in Data Appendix C.

[11] This monthly weather data set has been used in numerous papers on India, including Allcott et al. (2016), Emerick (2018), and Kaur (2019).

the district centroid. We construct measures of average 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 growing season weather variables to ten-year averages.

3.4 Infrastructure and Yields Data

We use data on road infrastructure and crop yields from the Village Dynamics in South Asia (VDSA) Meso dataset. For road infrastructure, 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. Appendix Figure B2a 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 ($0.10km/km^2$, as indicated by the solid vertical line in the figure), and zero otherwise. Appendix Figure B2b plots a heat map of the road density measure across all districts, with shades of red and blue denoting districts with above and below median road density respectively.

The VDSA Meso data set contains annual measures of district-level agricultural yields spanning the years 1966 to 2010. We construct a composite yield measure that aggregates yields across all crops that have non-missing price data, using base-year 1966-1970 crop prices as weights.

3.5 Bank Credit Data

We use data on bank credit from the Basic Statistical Returns, collected by the Reserve Bank of India. In particular, we use the total bank credit in each district in 1972 — the earliest year for which this data is available. We construct a district-level bank credit per capita measure by dividing total bank credit by total population. Appendix Figure B3a summarizes the distribution of the bank credit per capita measure across all districts. We create a binary measure that takes the value 1 if bank credit per capita in a district is above the median level of the full distribution (19 Rupees, as indicated by the solid vertical line in the figure), and zero otherwise. Appendix Figure B3b plots a heat map of the bank credit per capita measure across all districts, with shades of red and blue denoting districts with above and below median road density respectively.

Further details on the various data sources and the construction of all key variables described in this section can be found in Data Appendix C.

3.6 Descriptive Statistics

Table 1 provides summary statistics for decadal weather and Census variables for each census year. We report means and standard deviations of key variables for all balanced districts in the sample.

First, the table summarizes the two weather variables — temperature and precipitation. The first row confirms that temperatures have been rising over time. The growing season average monthly temperature is 0.42 °C higher in 2011, relative to 1961. On the other hand, the second row 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.

Next, the table summarizes the four Census outcome measures. The share of the labor force who are agricultural laborers is 15.1% on average in 1961, and increases to 30% by 2011. We see a similar increase over time in the share of the labor force who are non-agricultural workers — this rises from 27.2% in 1961 to 41.7% in 2011. We also see growth in urbanization — the share of total population residing in urban areas is 15.9% on average in 1961 and increases steadily over time to 27.5% in 2011. This is reflected in intra-district rural-to-urban migration patterns, which have also increased over the decades.[12] To complement Table 1, Figure 1 plots the spatial distribution of changes in the long-run climate and outcome variables from 1961 to 2011, by district.

[12] Appendix Table B1 provides summary statistics for decadal weather and NSS variables for each year the NSS is canvassed. We observe a similar pattern of rising temperatures and increasing shares of non-agricultural workers over time.

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-f) are heterogeneous across space. The share of total population residing in urban areas has increased by more than 18 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. The intra-district migrant share has decreased in one-third of the districts, mostly in the Northern and Eastern regions.

Appendix Figure B4 illustrates the relationship between long-run changes in non-agricultural worker share and long-run changes in decadal temperature for all districts in our sample. We observe that an increase in temperature is associated with a reduction in non-agricultural worker share. Motivated by this pattern in the raw data, we proceed to rigorously test the robustness of this relationship with the full panel data set, using two empirical specifications that we describe in the next section.

4 Empirical Specification

4.1 Panel Approach

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

$$lnY_{jsrt} = \beta T_{jsrt} + \gamma P_{jsrt} + \alpha_j + \alpha_t + \alpha_{srt} + \epsilon_{jsrt}, \tag{1}$$

where Y_{jsrt} represents the outcome of interest, in district j, located in state s and region r, in year t.[13] For our main specification, Y_{jsrt} represents the

[13] For migrant share, $lnY_{jsrt} = ln(Y_{jsrt} + 0.01)$ to account for zeros in the sample.

share of agricultural laborers, the share of non-agricultural workers, the share of the total population residing in urban areas, or the share of rural-to-urban intra-district migrants in the male population. T_{jsrt} is the average temperature measured in degrees Celsius 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 and α_t is a vector of year fixed effects that controls for changes over time. α_{srt} is a term that controls for time-varying region-specific or state-specific effects. Depending on our specification, this is either a vector of region-specific linear time trends, state-specific linear time trends or region-year fixed effects; the goal is to control for unobserved 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).[14]

The identifying assumption is that, conditional on the inclusion of district and year fixed effects, along with the region/state-year controls, any remaining variation in decadal temperature and precipitation is essentially random. This in turn allows for a causal interpretation of the β and γ coefficients as the effect of slow-onset changes in climate on structural transformation and urbanization.

Predictions for the sign of β, the coefficient of interest, are theoretically ambiguous. If rising temperatures lead to relatively greater productivity losses in the agricultural sector, this would lead to a decline in demand for agricultural labor. Under this scenario, we would expect $\beta < 0$ in regressions where the share of agricultural laborers is the dependent variable, and $\beta > 0$ in regressions where the non-agricultural worker share, urbanization, or the mi-

[14] To implement Conley standard errors, we use Stata routines from Hsiang (2010), Colella et al. (2019), and Fetzer (2020).

grant share is the dependent variable. On the other hand, if a decline in farm incomes arising from low agricultural productivity under rising temperatures leads to local demand effects or exacerbates liquidity and mobility costs, we would expect to see a larger share of workers in agriculture over time as higher temperatures persist — we would thus find the opposite results under this scenario.

4.2 Long Differences Approach

Our panel specification with district and year fixed effects identifies the impact of location-specific changes in decadal temperatures on labor reallocation: these estimates can be interpreted as medium-run responses to climate change. In the longer-run, changes in average climate may affect labor reallocation differently if the medium-run effects are mediated through adaptation, or compounded and intensified over time.

To complement our panel estimates of the effect of climate on structural transformation and urbanization, we follow Burke and Emerick (2016) and estimate a long differences regression of the form:

$$\Delta lnY_{jsr} = \beta_{LD}\Delta T_{jsr} + \gamma_{LD}\Delta P_{jsr} + \alpha_r + \epsilon_{jsr}. \tag{2}$$

In this regression, ΔlnY_{jsr} represents the difference (in natural logarithm) of the outcome of interest between two periods: 1961–1981 and 1991–2011.[15] The independent variables ΔT_{jsr} and ΔP_{jsr} are differences in the average decadal growing-season temperature and precipitation over congruent periods.[16] α_r is a region fixed effect that allows for differential time trends across the regions in our sample. ϵ_{jsr} is an idiosyncratic error term.

Because each observation is a district-specific difference of two periods in

[15]The climatology literature refers to a 30-year lagged average as a "climate normal" (Arguez and Vose, 2011; Bento et al., 2021). This motivated our choice of window length for the end points of our long-difference specification.

[16]More specifically, the outcomes and dependent variables in 1961-1981 are calculated as the average of 1961, 1971, and 1981 decadal observations, and those in 1991-2011 are calculated as the average of 1991, 2001, and 2011 decadal observations.

time, any time-invariant district factors are differenced out. Unbiased estimates of β_{LD} and γ_{LD} thus rely on changes in temperature and precipitation between the two periods that are not correlated with time-trends that also affect labor reallocation. We argue that conditional on region fixed effects, long-term changes in district temperature and precipitation are likely exogenous with respect to our outcomes. As with Equation 1, the signs of the coefficients of interest, β_{LD}, are theoretically ambiguous.

In addition to the magnitude and sign of β_{LD}, the magnitude of β_{LD} *relative to* β from Equation 1 is also of interest. As discussed in Dell et al. (2014) and Burke and Emerick (2016), a comparison of panel and long difference estimates can provide insight on the existence of adaptation versus intensification effects. In their seminal paper, Burke and Emerick (2016) find that the adverse impact of high temperatures on crop yields in the U.S. is attenuated slightly when they move from an annual panel specification to a long differences specification that spans 1980-2000. However, in most specifications they find no statistically significant difference between their panel and long difference estimates, suggesting that there is only minimal longer-term adaptation to high temperatures in their study context. We consider the role of adaptation versus intensification effects in our setup. To fix ideas, let us focus on the share of agricultural laborers, and suppose that in the panel specification we find $\beta > 0$: in other words, higher temperatures increase the share of agricultural laborers. Then, moving to the long differences estimate, suppose we find $\beta > \beta_{LD} > 0$. In this case, in the longer time horizon, rising temperatures still have an adverse effect of increasing the share of agricultural laborers, but the magnitude of this effect is diminished over the longer time frame, suggesting that agents are able to adapt in the presence of sustained temperature increases. On the other hand, suppose that we find $\beta_{LD} > \beta > 0$. This would suggest that the effect of rising temperatures on the share of agricultural laborers is intensified as we move from the decadal panel specification to the long differences specification, which implies that adverse climate effects compound and intensify over the longer term.

5 Results

5.1 Panel Approach

Table 2 presents the effects of temperature and precipitation on four key outcomes — share of agricultural laborers, share of non-agricultural workers, urbanization, and share of intra-district migrants — using the panel approach (Equation 1). For each outcome, the first column presents regression estimates using region-specific linear time trends while the second column presents regression estimates with region-year fixed effects. The latter is our preferred specification since region-year fixed effects control for arbitrary unobserved region-specific confounding factors over time. We report results with standard errors clustered at the district level in parentheses and with Conley standard errors in brackets.

In Columns 1 and 2 of Table 2, we document a positive, statistically significant effect of rising average temperatures on the share of agricultural laborers. Using our preferred specification, we find that a 1°C increase in temperature leads to a 17.0% increase in the share of the labor force engaged in agriculture.[17] At the same time, we find that an increase in decadal precipitation may have a negative effect on the share of agricultural laborers, though this effect is not statistically different from zero once we include region-year fixed effects. Further, the magnitude of the impact of precipitation on the share of agricultural laborers is much smaller than that of temperature.[18] 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, 2021).

We find a negative effect of rising average temperatures on the share of non-agricultural workers in Columns 3 and 4 of Table 2: a 1°C increase in temperature is associated with a 8.2% reduction in the share of the labor force

[17] We calculate effect sizes using the formula $\%\Delta y = e^{(\beta \Delta x)} - 1$.

[18] Based on the coefficients in Column 1 of Table 2, where the effect of rainfall is statistically significant, a one standard deviation increase in decadal temperature would have a roughly nine times greater impact on the share of agricultural laborers than a one standard deviation increase in decadal precipitation.

engaged in non-agriculture. At the same time, we find that an increase in decadal precipitation has no detectable impact on the share of non-agricultural workers in the labor force. Next, in Columns 5 and 6 of Table 2, we fail to find any detectable impact of rising average temperatures on urbanization rates. This result is in line with Henderson et al. (2017), who similarly document no average impact of adverse changes in climate on urbanization in Sub-Saharan Africa. Finally, in Columns 7 and 8 of Table 2, we find that rising temperatures have no detectable impact on the share of rural-to-urban intra-district migrants — while the point estimate is negative, it is imprecisely estimated.

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 our other outcomes. We conduct an additional test where we run the same regression in Equation 1 with region-year fixed effects, dropping precipitation. The coefficients on temperature are almost identical: 0.162 (with agricultural labor share as the dependent variable), -0.086 (with non-agricultural worker share as the dependent variable), 0.001 (with urbanization as the dependent variable), and -0.018 (with intra-district migrant share as the dependent variable). Given these patterns, we focus only on interpreting the coefficients on temperature in our analysis moving forward.

Given that employment in agriculture is seasonal, it is common for workers to engage in multiple activities across different sectors within a year (Emerick, 2018; Breza et al., 2021). In Appendix Table B2, we further investigate how labor is reallocated in response to rising temperatures using detailed classification of workers' employment. The Census splits worker counts into two groups: main and marginal workers, which we interpret as proxies for full-time and short-term, seasonal workers respectively. We find that rising average temperatures lead to an increase in both main and marginal agricultural labor shares (Panel A). Next, the NSS data contains information on the usual primary and secondary occupations of employed individuals. We find suggestive evidence that rising average temperatures lead to an increase in both primary and secondary worker shares in agriculture, and a corresponding decrease in primary worker shares in non-agriculture (Panel B). Taken together, this evidence supports that the margin of adjustment comes from workers who primarily engage in non-agriculture switching to agriculture full-time or on a short-term basis in response to rising temperatures.

The existing literature has demonstrated significant non-linear effects of temperature — in which damages from rising temperatures intensify above a certain threshold — for outcomes including crop yields (Schlenker and Roberts, 2009), economic production (Burke and Emerick, 2016), labor supply (Graff Zivin and Neidell, 2014) and migration (Bohra-Mishra et al., 2014). We explore non-linear temperature effects in our empirical setting by allowing the effect of rising temperatures to vary based on whether a district's long-run average growing season temperature is above or below the median across all districts.[19] This specification allows us to explore whether increases in temperature have an intensified effect in districts with a higher baseline level of heat.

Results in Appendix Table B3 provide evidence that there are non-linear effects of temperature on agricultural labor share and non-agricultural worker share. While higher temperatures increase the share of the labor force engaged in agriculture across both hotter and less hot districts, the point estimate for hotter districts is larger in magnitude and more precisely estimated than that for less hot districts. Similarly, higher temperatures reduce the share of the labor force engaged in non-agriculture in both hotter and less hot districts, but the point estimate for hotter districts is larger in magnitude and more precisely estimated than that for less hot districts. These patterns are consistent with previous findings that damages from rising temperatures may intensify above a certain threshold. However, given that we do not have a sufficiently large

[19] As our analysis includes agricultural channels, it would be ideal to use a temperature data set with daily observations to capture non-linear temperature effects. Daily temperature data would enable us to construct daily temperature bins (Schlenker and Roberts, 2009) or degree days (Deschênes and Greenstone, 2007; D'Agostino and Schlenker, 2016), measures which are used widely in the climate change literature. However, the commonly used daily gridded weather data sets such as the Modern-Era Retrospective Analysis for Research and Applications (Rienecker et al., 2011) and ERA-Interim (Dee et al., 2011) only have coverage starting in 1979, corresponding to the era of modern remotely sensed data. As a result, these daily weather data sets are not compatible with our Census data which begins in 1961.

sample to detect a statistically significant difference across the two groups, we acknowledge that these results are suggestive.

Lastly, we carry out an additional empirical exercise that allows us to contrast the long-term effects of rising temperatures documented here with short-term effects that have been previously documented in the literature (Colmer, 2021). Specifically, we run two specifications using data from the National Sample Survey and present results in Appendix Table B4. First, we implement our panel specification (Equation 1) using this alternative data source in Panel A. We document that rising average decadal temperatures lead to a substantial increase in the share of the labor force engaged in agriculture and a corresponding decline in the share of the labor force engaged in non-agriculture. These precisely estimated effects are consistent with our results in Table 2, suggesting that our findings are robust across the two data sets. Next, we split the climate variables into current year averages, similar to that used in Colmer (2021), and decadal averages, similar to that used in our panel approach, in Panel B.[20] This specification allows us to look at both the short-term and long-term effects of temperature simultaneously. We document that rising average current temperatures are associated with a decline in the share of the labor force engaged in agriculture and a corresponding increase in the share of the labor force engaged in non-agriculture — these effects are consistent with the short-term effects documented in Colmer (2021). Concurrently, we find that rising average decadal temperatures have opposite effects that are consistent with the long-term patterns we document in Panel A as well as in Table 2.

5.2 Long Differences Approach

Table 3 presents the effects of temperature and precipitation on our four key outcomes using the long differences approach (Equation 2). Taken together, the signs of the long differences estimates are consistent with the panel estimates in Table 2. Controlling for regional trends, we find that districts

[20] The decadal averages paired with NSS outcomes are averages of one-year to ten-year lags; therefore they do not overlap with the current year climate measures.

that experienced greater increases in decadal temperatures from 1961-1981 to 1991-2011 have higher shares of agricultural laborers and lower shares of non-agricultural workers, with both effects significant at the 1% level.

Importantly, for both of these outcomes, our long differences coefficient estimates are *larger* in absolute value than the corresponding panel coefficient estimates.[21] For example, the coefficient on temperature in the agricultural labor share regression is 0.157 in our panel specification (Table 2), and it increases to 0.3819 in our long differences specification (Table 3). As discussed in Section 4.2, this suggests intensification effects of warming. We thus find adverse impacts of rising temperatures on structural transformation in the medium-term, and the impacts from sustained higher temperatures appear even larger over a longer time frame.[22]

5.3 Mechanisms

In this section, we empirically investigate potential mechanisms through which rising temperatures could affect labor reallocation.

To begin, we test the effect of rising temperatures on agricultural yields — the basis underlying all three mechanisms outlined in Section 2. While the temperature-yield relationship has already been established in the literature (Schlenker and Roberts, 2009; Taraz, 2018), we replicate this result using data from our empirical context in Panel A of Appendix Table B5. We find a strong, negative effect of higher current year average temperatures on yields: a 1°C increase in current growing season temperature is associated with a 5.9% reduction in yields aggregated across all crops. These results provide empirical support that shocks to agricultural productivity may be driving the impacts

[21] For the share of agricultural laborers, the p-value of the difference is 0.0041, while for the share of non-agricultural workers, the p-value is 0.047.

[22] These intensified longer-run effects of rising temperatures are consistent with a scenario where local demand effects described in Section 2 act as the main mechanism. Broadly speaking, as warming negatively affects agricultural productivity, there is a reduced demand for non-agricultural goods and services, leading to a contraction in the size of the non-agricultural labor force. As persistent warming over a longer time frame prolongs the reduction in local demand, the contraction of the non-agricultural sector is intensified over time.

of climate change on labor reallocation.

Next, we gauge the validity of local demand effects and liquidity and mobility costs as possible mechanisms by exploring heterogeneous effects. Specifically, we examine whether certain features of development in a given district — road connectivity and access to credit, in particular — play a role in modulating the effect of climate change on labor reallocation. If local demand effects drive our results, we would expect to see intensified effects in districts with stronger local price effects (sparse road networks) and in places where it is harder for individuals to smooth consumption across periods (poor access to bank credit).[23] If liquidity and mobility costs drive our results, we would expect to see intensified effects in areas where these costs are effectively higher, specifically districts with sparse road networks and poor access to bank credit. Both mechanisms predict the same set of heterogeneous patterns; we turn to the data to rigorously test if the empirical results are consistent with these two mechanisms.

In Table 4, we augment our specification in Equation 1 to allow for the net effects of temperature and precipitation to vary based on two distinct features of development: road connectivity in Panel A and access to formal credit in Panel B.[24] We find that the effects of rising average temperatures on the share

[23]In Appendix A, we formally show that local demand effects are intensified in areas with higher transportation costs. While our set-up does not allow for multiple periods in order to fully model the impact of credit access, the parsimonious specification suggests that credit access could have additional effects on labor reallocation through *consumption*. In particular, we note that agricultural productivity shocks change the equilibrium non-agricultural labor share through a reduction in consumption of non-agricultural goods and services. Individuals in districts with access to credit are better able to smooth consumption across periods relative to those in districts with limited credit access. Therefore, the responsiveness of non-agricultural labor to agricultural productivity shocks would be attenuated in places with better access to bank credit.

[24]We interact the weather variables with binary measures and estimate:

$$lnY_{jsrt} = \beta T_{jsrt} + \gamma P_{jsrt} + \beta_D T_{jsrt} * D_j + \gamma_D P_{jsrt} * D_j + \alpha_j + \alpha_t + \alpha_{\mathbf{D}t} + \alpha_{srt} + \epsilon_{jsrt}$$

, where D_j is a binary variable that takes the value 1 if the baseline road density/bank credit per capita in district j is above the median of the distribution across all districts, and 0 otherwise. We also include a heterogeneous-group-by-year fixed effect $\alpha_{\mathbf{D}t}$, which allows for each subgroup to have different unobserved shocks over time. All other terms are as defined above in Equation 1. In this specification, β and γ capture the effects of decadal changes in of agricultural laborers and the share of non-agricultural workers documented in Table 2 are driven entirely by districts with below median road density and below median bank credit per capita at baseline (Panels A and B, Columns 1-4). While we document an overall null effect of rising average temperatures on urbanization rates and the share of rural-to-urban intra-district migrants in Table 2, heterogeneity analysis in Table 4 reveals interesting patterns that are complementary to our results on sectoral reallocation — rising average temperatures in districts with below median road density at baseline are associated with a reduction in the share of the total population residing in urban areas (Panel A, Columns 5 and 6), and this reduction is entirely attenuated in districts with well-developed road networks. Impacts of rising temperatures on rural-to-urban migrant share are complementary to our results on urbanization (Panel A, Columns 7 and 8).[25]

Across all our outcome measures, the average effects of rising temperatures in districts with sparse road networks and limited access to formal credit are larger than the average effects in the full sample, which suggests that these features of underdevelopment amplify the impacts of rising temperatures on the degree of structural transformation and urbanization in Indian districts. These heterogeneous results are consistent with *both* local demand effects and liquidity and mobility costs as possible mechanisms.

To disentangle these two mechanisms, we turn to additional tests in Tables 5 and 6. First, we examine the effect of rising average temperatures on household consumption across three categories — total consumption, food

temperature and precipitation respectively in districts with limited road connectivity/credit access, while $\beta + \beta_D$ and $\gamma + \gamma_D$ capture the effects of decadal changes in temperature and precipitation respectively in districts with extensive road connectivity/credit access.

[25]One important concern is that districts with more roads or commercial banks might also be districts in which agricultural productivity is less sensitive to higher temperatures — for example, due to differences in geography or in access to technological innovations such as heat tolerant seeds. In Panels B and C of Appendix Table B5, we estimate the temperature-yield relationship for each subgroup of districts. Across both panels, we find that the coefficients on the interaction terms are small and statistically insignificant, which suggests that there are no differential impacts of current year average temperature on yields across districts with below/above median road density (panel B) and across districts with below/above median bank credit per capita (panel C).

consumption and non-food consumption. We document significant declines in household consumption across all three categories in Panel A of Table 5, which suggests a reduction in demand for both farm and non-farm output in response to rising temperatures. Next, we decompose the effect of rising average temperatures on non-agricultural worker shares across three sectors — services, construction and manufacturing. We find significant reductions in the share of the labor force engaged in services, which is by nature non-tradable (Panel B, Table 5). In contrast, there is no detectable impact on the share of the labor force engaged in construction and manufacturing — while the estimated coefficients are also negative, they are imprecisely estimated and not statistically different from zero at standard levels of significance. These patterns are consistent with local demand effects driving a reduction in demand for local non-agricultural goods and services, which subsequently leads to a decline in the share of the labor force engaged in these sectors.

Further, we explore the heterogeneous effect of rising average temperatures across different segments of the population. In particular, we allow for the net effects of temperature and precipitation to vary by social grouping (caste) and by educational attainment. If liquidity and mobility costs dominate, we would expect effects to be concentrated among workers for whom these constraints are more likely to bind — we proxy for this group of workers using lower-caste and lower-educated individuals. In Table 6, we find that the patterns of labor reallocation documented in Table 2 are similar for lower-caste and higher-caste households (Panel A), as well as for lower-educated and higher-educated households (Panel B).[26] Given that these different segments of the population appear to be responding similarly, it is unlikely that these effects are driven primarily by exacerbated liquidity and mobility costs.

Taken together, the above analyses support a local demand effects mechanism since we find that the impacts of rising temperatures on labor reallocation are i) larger in less developed districts with sparse roads and limited access to

[26] For agricultural worker share, the coefficients on decadal average temperatures are statistically indistinguishable at standard levels across the education groups and the social groups.

banks, ii) larger in the services (non-tradable) sector, iii) transmitted through reductions in demand for both farm and non-farm goods, and iv) similar across segments of the population with varying capacities to fund liquidity and mobility costs.

In the remainder of this section, we consider several other mechanisms that could explain how rising temperatures affect labor reallocation, and we present these findings in Appendix Tables B6, B7, and B8. First, we examine whether warming has a direct impact on non-agricultural sectors. As described in Section 2, productivity in the non-agricultural sectors could be negatively affected by temperature, for example, through heat stress on workers (Hsiang, 2010; Somanathan et al., 2021). Under this scenario, we would expect to see a reduction in demand for non-agricultural labor, and importantly, the resulting decline in non-agricultural worker shares should be present in *both* rural and urban areas. In Appendix Table B6, to test whether the above pattern appears in our context, we examine the effect of rising average temperatures on non-agricultural worker shares *separately* for rural and urban areas within a district. We find that the negative effect of rising temperatures on the share of non-agricultural workers is concentrated in rural areas (Column 2), while the share of non-agricultural workers in urban areas remains unaffected by rising temperatures (Column 4). The lack of effects in urban areas suggests that our findings cannot be explained by a direct effect of rising temperatures on productivity in the non-agricultural sector.[27]

Second, if reductions in agricultural productivity under climate change lead to land consolidation, this could mechanically result in an increase in the share of agricultural laborers. Under this scenario, we would expect to find a corresponding reduction in the share of *cultivators*. In Panel A of Appendix Table B7, we examine the impacts of rising temperatures on cultivator share — defined as main cultivators divided by total workers — using our panel

[27] The lack of effects in urban areas also helps us reconcile our findings on sectoral reallocation and urbanization. Recall that in Table 2, we find a negative impact of rising temperatures on sectoral reallocation, but do not observe a similar impact on urbanization. This can be attributed to the fact that the decline in sectoral reallocation from agriculture to non-agriculture is driven by workers in rural areas only.

approach. We find a positive and marginally significant effect of rising average temperatures on the share of cultivators (Column 2). Turning to heterogeneous impacts, the effect of rising average temperatures is positive in underdeveloped districts with below median density road networks/bank credit per capita at baseline (Columns 3-6). The coefficients on the interactions terms are statistically indistinguishable from zero, suggesting that there is no differential impact of rising temperatures on the share of cultivators in districts with above median density road networks/bank credit per capita at baseline. This suggests that cultivators are more likely to be characterized by limited occupational and spatial mobility, even in districts with access to well-developed road networks and formal credit. This could be driven by distortions in land markets (Foster and Rosenzweig, 2021; Bolhuis et al., 2021), insecure property rights (Gottlieb and Grobovšek, 2019) or cultural norms (Fernando, 2020).

Third, rising temperatures could induce a shift in cultivation practices. If farmers respond to reductions in agricultural productivity under climate change by cultivating dry-season crops or by switching to more labor-intensive crops, this would result in an increase in demand for agricultural labor. In Appendix Table B8, we examine impacts of rising temperatures on the share of total land cultivated with dry-season and labor-intensive crops using our panel approach. We find that rising temperatures lead to a slight reduction in cultivation across these two categories of crops, which suggests that our findings are not driven by shifts in farmers' cultivation practices.

The bulk of the empirical evidence presented in this section is consistent with a local demand mechanism driving our labor reallocation results. To summarize, rising temperatures reduce agricultural productivity, which subsequently lowers consumption of local non-agricultural goods and services, leading to a contraction in the non-agricultural labor share. Although these effects are felt commonly by individuals across skill and social groups, they are concentrated in underdeveloped areas with sparse road connectivity and/or poor access to credit.

On a final note, we now revisit contrasting effects of short-term versus long-term temperature shocks on labor reallocation as previously discussed in

Section 5.1. One explanation could be that the degree to which individuals can smooth consumption in response to income shocks depends on the time-scale of these shocks. In particular, we posit that individuals are likely more able to smooth their consumption in response to brief, transitory income shocks, however, they have limited capacity to do so when facing long-term productivity shocks such as persistent warming. In our context, higher temperatures bring minimal local demand effects as individuals successfully consumption smooth in the short-term. Therefore, the relative labor productivity loss mechanism dominates, and year-to-year increase in temperatures draws labor to the non-agricultural sector on a short time-scale. However, individuals are less able to consumption smooth in response to sustained shocks to agricultural productivity that persist over longer time-scales. Thus, when we evaluate the effects of rising temperatures using a longer time-frame, we see that local demand effects intensify and dwarf the labor productivity loss mechanism. This can explain why higher decadal average temperatures lead to a reduction in the amount of labor supplied to the non-agricultural sector on longer time-scales.

5.4 Robustness Checks

In this section, we explore the robustness of our results. First, we show that our panel results are robust to a broader definition of the share of the labor force who are agricultural workers. In our main analysis, we focus on agricultural wage workers only. We now expand our definition to include cultivators as well. In Panel B of Appendix Table B7, we examine the impacts of rising temperatures on the share of agricultural laborers and cultivators — defined as total agricultural laborers and cultivators divided by total workers — and we find a similar positive, statistically significant effect of rising average temperatures on the share of agricultural laborers and cultivators.[28]

Second, we test the robustness of our panel results to changes in the sam-

[28]Further, we find that the coefficient on decadal temperature in Column 2 of Table 2 (with agricultural labor share as the outcome) is not statistically different from that in Column 2 of Appendix Table B7 Panel B (with agricultural labor and cultivator share as the outcome) — the p-value is 0.202.

ple. In our main analysis, we restrict our sample to districts for which the dependent variable is non-missing in all years. We report results using the full unbalanced sample in Panel A of Appendix Table B9. The estimated coefficients and significance levels are largely unchanged under the inclusion of these unbalanced districts.

Third, we show that our panel results are robust to using alternative definitions of the temperature and precipitation measures. In defining long-run changes in climate, we use decadal averages of temperature and precipitation, weighted by the inverse distance of weather grids to district centroids. In Panel B of Appendix Table B9, 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 fall within a district's boundary. In Panel C of Appendix Table B9, we use logs of temperature and precipitation instead of levels, and we find that this specification strengthens our results.

Fourth, we explore the robustness of our results to a distributed lag model. Our use of decadal averages is motivated by our interest in the impact of slow-onset changes in climate. In Appendix Table B10, we present a distributed lagged average model to explore whether our decadal results are driven mostly by recent shocks, or whether there are persistent long-term effects of shocks several years prior. We break down our decadal temperature measure into three smaller averages: current temperature to two-year lagged temperature; three-year lagged temperature to six-year lagged temperature; and seven-year lagged temperature to nine-year lagged temperature. The results are mixed, depending on the outcome variable. For agricultural labor share, the effects appear to be driven by slightly older shocks (3 to 6 years), whereas for non-agricultural worker share, the effects appear to be driven by shocks in more recent years. Although no clear patterns emerge, Appendix Table B10 does suggest that our decadal results are not driven *solely* by recent temperature shocks.

Fifth, we show that our panel results are robust to controlling for trends in other variables that may influence labor reallocation. More specifically, we re-estimate our panel specification including time-varying controls for area cultivated by high-yielding varieties (a proxy for access to the Green Revolution), a labor regulation strictness index (a proxy for flexibility in hiring workers), a continuous road density measure, the number of markets (proxies for infrastructure in the district), and the number of banks (a proxy for financial development of the district). Appendix Table B11 illustrates that our results are robust to the inclusion of these controls — the coefficients on agricultural labor share and non-agricultural worker share are precisely estimated and are of similar magnitude to the coefficients in Table 2. We choose not to include these time-varying controls in our main specification in order to avoid the potential bias of "bad controls" — control variables that could themselves be affected by temperature (Angrist and Pischke, 2009; Hsiang et al., 2013).

Sixth, we test the robustness of our long difference results to alternative end points. Our long difference specification takes the difference in outcomes between two periods: 1961–1981 and 1991–2011. Each period is composed of three Census observations and three decadal observations for the independent weather variables. In Appendix Table B12, we use an alternate window, taking the difference between the periods 1961–1971 and 2001–2011 — we now average together two observations instead of three for both the Census and weather variables. The signs and precision of the coefficients on temperature for agricultural labor share and non-agricultural worker share are robust to these alternative end points. Moreover, the point estimates continue to be larger in magnitude than the point estimates in the panel specification, ruling out potential adaptation to warming over the longer-term.

Seventh, we demonstrate that the heterogeneous impacts summarized in Table 4 are robust to alternate thresholds for our heterogeneity dummies. In our main specification, we define high road density and high bank credit districts to be those whose baseline values of road density and bank credit per capita respectively are above the median of the distribution across all districts at baseline. We test the sensitivity of our results to two alternate thresholds — the 40th and 60th percentiles of the distribution across all districts at baseline. In Appendix Table B13, we find that our results are largely robust to these two alternate thresholds. We thus conclude that our heterogeneous impacts are not sensitive to the choice of threshold used.

Eighth, we show that the heterogeneous impacts are also robust to the inclusion of other baseline district-level controls interacted with temperature and precipitation. In Appendix Tables B14 and B15, we sequentially include controls such as baseline wages and irrigated land interacted with weather variables to the heterogeneous specification. The stability of our coefficients of interest as we cumulatively add these controls across the columns suggest that the heterogeneous impacts summarized in Table 4 are not driven by factors other than access to roads or bank credit.

6 Conclusion

As temperatures rise, agricultural productivity will decline and this may impact the spatial and sectoral allocation of workers in the economy. 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, 2021).

In this paper, we add to this base of knowledge by exploring responses to slow-onset changes in temperature, measured using decadal averages in a panel specification, and using changes over our entire 60-year sample in a long differences specification. Under both specifications, we find that higher temperatures inhibit structural transformation in Indian districts. This finding has important policy implications for India and other low- and middle-income countries. The existing climate-agronomy literature has demonstrated that higher temperatures have severe adverse effects on crop yields and agricultural incomes (Schlenker and Roberts, 2009; Taraz, 2018). We document that sustained higher temperatures not only dramatically lower yields, but also inhibit the movement of labor *out* of agriculture, potentially magnifying the human welfare impacts of climate change.

Furthermore, we find suggestive evidence that these impacts are driven by local demand effects and that they are concentrated in districts with sparse road networks or low access to formal bank credit. These results suggest that individuals in areas with these features of underdevelopment are more susceptible to the adverse effects of higher temperatures. Earlier research demonstrates that social safety net programs can cushion the impact of adverse environmental events (Deryugina, 2017; Garg et al., 2020) — this may in turn alleviate the negative impacts on structural transformation. In a similar vein, interventions that make agricultural incomes less susceptible to high temperatures may also alleviate negative impacts.

Finally, our results add to the broader climate–economy literature by demonstrating the importance of analyzing slow-onset changes in climate, as estimates based on long-term variations in climate give a better approximation of how agents will respond to anthropogenic climate change (Burke and Emerick, 2016). Relative to our decadal specification, our long differences specification detects magnified adverse effects of rising temperatures on structural transformation in the long-term. This suggests that, not only do individuals fail to adapt to higher temperatures in the medium-term, but they also endure accumulating impacts from sustained warming that renders sectoral reallocation more challenging in the long-term. A promising avenue for future research is to continue exploring climate impacts using slow-onset changes and long panels of data, particularly in low- and middle-income countries that are especially vulnerable to climate change.

References

Allcott, H., Collard-Wexler, A., and O'Connell, S. D. (2016). How do electricity shortages affect industry? Evidence from India. *American Economic Review*, 106(3):587–624.

Alvarez, J. A. (2020). The agricultural wage gap: Evidence from Brazilian micro-data. *American Economic Journal: Macroeconomics*, 12(1):153–73.

Angelucci, M. (2015). Migration and financial constraints: Evidence from Mexico. *Review of Economics and Statistics*, 97(1):224–228.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics*. Princeton University Press.

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

Arguez, A. and Vose, R. S. (2011). The definition of the standard WMO climate normal: The key to deriving alternative climate normals. *Bulletin of the American Meteorological Society*, 92(6):699–704.

Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives*, 32(4):33–52.

Bento, A., Miller, N. S., Mookerjee, M., and Severnini, E. R. (2021). A unifying approach to measuring climate change impacts and adaptation. *Working Paper*.

Besley, T. and Burgess, R. (2004). Can labor regulation hinder economic performance? Evidence from India. *The Quarterly Journal of Economics*, 119(1):91–134.

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*, 48(26/27):5–13.

Bohra-Mishra, P., Oppenheimer, M., and 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.

Bolhuis, M., Rachapalli, S., and Restuccia, D. (2021). Misallocation in Indian agriculture. *Working Paper*.

Borjas, G. J. (1987). Self-selection and the earnings of immigrants. *American Economic Review*, 77(4):531–553.

Breza, E., Kaur, S., and Shamdasani, Y. (2021). Labor rationing. *American Economic Review*, 111(10):3184–3224.

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. *Working Paper*.

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

Bustos, P., Caprettini, B., and Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from Brazil. *American Economic Review*, 106(6):1320–65.

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

Colella, F., Lalive, R., Sakalli, S. O., and Thoenig, M. (2019). Inference with arbitrary clustering. *Working Paper*.

Colmer, J. (2021). Temperature, labor reallocation, and industrial production: Evidence from India. *American Economic Journal: Applied Economics*, 13(4):101–24.

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

D'Agostino, A. L. and Schlenker, W. (2016). Recent weather fluctuations and agricultural yields: Implications for climate change. *Agricultural Economics*, 47(S1):159–171.

Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., ..., and Vitart, F. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, 137(656):553–597.

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.

Deryugina, T. (2017). The fiscal cost of hurricanes: Disaster aid versus social insurance. *American Economic Journal: Economic Policy*, 9(3):168–98.

Deschênes, O. and 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.

Emerick, K. (2018). Agricultural productivity and the sectoral reallocation of labor in rural India. *Journal of Development Economics*, 135:488–503.

Fernando, A. N. (2020). Shackled to the soil? Inherited land, birth order, and labor mobility. *Journal of Human Resources*, pages 0219–10014R2.

Fetzer, T. (2020). Can workfare programs moderate conflict? Evidence from India. *Journal of the European Economics Association*, 18(6):3337–3375.

FICCI (2015). Labour in India: A growing challenge. Retrieved from `http://ficci.in/spdocument/20550/FICCI-agri-Report%2009-03-2015.pdf`.

Fishman, R., Carrillo, P., and Russ, J. (2019). Long-term impacts of exposure to high temperatures on human capital and economic productivity. *Journal of Environmental Economics and Management*, 93:221–238.

Foster, A. D. and Rosenzweig, M. R. (2021). Are there too many farms in the world? Labor-market transaction costs, machine capacities and optimal farm size. *Working Paper*.

Fulford, S. L. (2013). The effects of financial development in the short and long run: Theory and evidence from India. *Journal of Development Economics*, 104:56–72.

Garg, T., Jagnani, M., and Taraz, V. (2020). Temperature and human capital in India. *Journal of the Association of Environmental and Resource Economists*, 7(6):1113–1150.

Gazeaud, J., Mvukiyehe, E., and Sterck, O. (2021). Cash transfers and migration: Theory and evidence from a randomized controlled trial. *Review of Economics and Statistics.* publication. `https://doi.org/10.1162/rest_a_01041`.

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. *Working Paper*.

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

Gollin, D., Parente, S., and Rogerson, R. (2002). The role of agriculture in development. *American Economic Review*, 92(2):160–164.

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

Gottlieb, C. and Grobovšek, J. (2019). Communal land and agricultural productivity. *Journal of Development Economics*, 138:135–152.

Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26.

Hamory, J., Kleemans, M., Li, N. Y., and Miguel, E. (2021). Reevaluating agricultural productivity gaps with longitudinal microdata. *Journal of the European Economic Association*, 19(3):1522–1555.

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

Herrendorf, B., Rogerson, R., and Valentinyi, Á. (2014). Growth and structural transformation. In *Handbook of Economic Growth*, volume 2, pages 855–941. Elsevier.

Herrendorf, B. and Schoellman, T. (2018). Wages, human capital, and barriers to structural transformation. *American Economic Journal: Macroeconomics*, 10(2):1–23.

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.

Hsiang, S. M., Burke, M., and Miguel, E. (2013). Quantifying the influence of climate on human conflict. *Science*, 341(6151).

Hu, Z. and Li, T. (2019). Too hot to handle: The effects of high temperatures during pregnancy on adult welfare outcomes. *Journal of Environmental Economics and Management*, 94:236–253.

ICRISAT (2015). Meso level data for India, 1966-2011: Collected and compiled under the project on Village Dynamics in South Asia. Retrieved from `http://vdsa.icrisat.ac.in/vdsa-mesodoc.aspx`.

Imbert, C. and Papp, J. (2020). Short-term migration, rural public works, and urban labor markets: Evidence from India. *Journal of the European Economic Association*, 18(2):927–963.

Jain, A., O'Sullivan, R., and Taraz, V. (2020). Temperature and economic activity: Evidence from India. *Journal of Environmental Economics and Policy*, 9(4):430–446.

Jessoe, K., Manning, D. T., and Taylor, J. E. (2018). Climate change and labour allocation in rural Mexico: Evidence from annual fluctuations in weather. *The Economic Journal*, 128(608):230–261.

Kaur, S. (2019). Nominal wage rigidity in village labor markets. *American Economic Review*, 109(10):3585–3616.

Kolstad, C. D. and Moore, F. C. (2020). Estimating the economic impacts of climate change using weather observations. *Review of Environmental Economics and Policy*, 14(1):1–24.

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.

Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *Review of Economic Studies*, 68(4):869–882.

Kumar, H. and Somanathan, R. (2017). Creating long panels using Census data (1961–2001). *Economic and Political Weekly*, 52(29):105–109.

Lagakos, D., Marshall, S., Mobarak, A. M., Vernot, C., and Waugh, M. E. (2020). Migration costs and observational returns to migration in the developing world. *Journal of Monetary Economics*.

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

Matsuura, K. and Willmott, C. J. (2015a). Terrestrial air temperature: 1900–2014 gridded monthly time series (v 4.01). Retrieved from http://climate.geog.udel.edu/~climate/html_pages/download.html#ghcn_T_P_clim3.

Matsuura, K. and Willmott, C. J. (2015b). Terrestrial precipitation: 1900–2014 gridded monthly time series (v 4.01). Retrieved from http://climate.geog.udel.edu/~climate/html_pages/download.html#ghcn_T_P_clim3.

Matsuyama, K. (1992). Agricultural productivity, comparative advantage, and economic growth. *Journal of Economic Theory*, 58(2):317–334.

Maystadt, J.-F., Mueller, V., and Sebastian, A. (2016). Environmental migration and labor markets in Nepal. *Journal of the Association of Environmental and Resource Economists*, 3(2):417–452.

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.

Mueller, V., Gray, C., and Kosec, K. (2014). Heat stress increases long-term human migration in rural Pakistan. *Nature Climate Change*, 4(3):182–185.

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

Rienecker, M. M., Suarez, M. J., Gelaro, R., Todling, R., Bacmeister, J., Liu, E., ..., and Kim, G.-K. (2011). MERRA: NASA's modern-era retrospective analysis for research and applications. *Journal of Climate*, 24(14):3624–3648.

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. (2021). Rural road infrastructure and agricultural production: Evidence from India. *Journal of Development Economics*, 152:102686.

Singh, A., Sethi, S. R., Chandramouli, C., et al. (2011). *Administrative Atlas of India*. Office of the Registrar General and Census Commissioner, India.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827.

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.

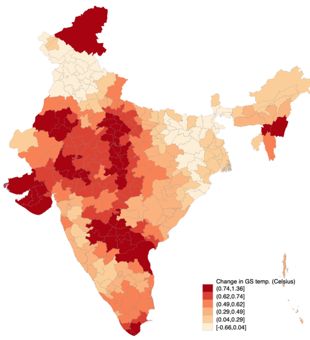

(a) Temperature

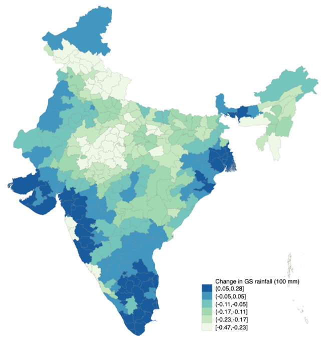

(b) Precipitation

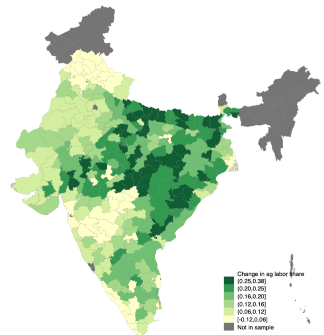

(c) Agricultural Labor Share

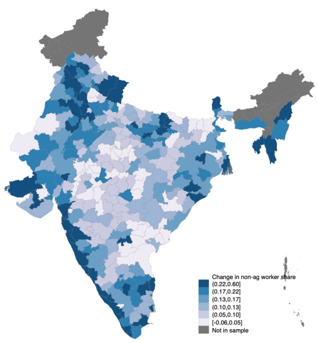

(d) Non-Agricultural Worker Share

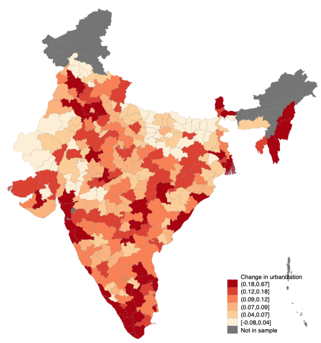

(e) Urbanization

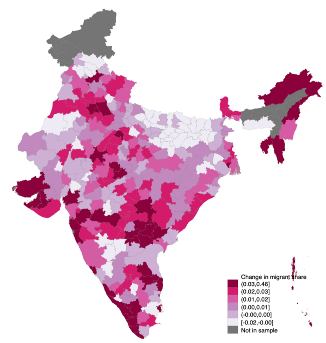

(f) Migrant Share

Figure 1: Figure illustrates long-run changes in the ten-year average of growing season temperature (panel a), growing season precipitation (panel b), share of agricultural laborers (panel c), share of non-agricultural workers (panel d), urbanization (panel e), and share of intra-district migrants (panel f) across all balanced districts. These changes are computed by subtracting the value of each variable in 1961 from the corresponding value in 2011. Data are district-level panel data constructed from the Indian Census.

Table 1: Summary Statistics by Year

Year	1961	1971	1981	1991	2001	2011	Total
10-Year Avg. GS Temperature (Celsius)	23.78	23.87	23.94	23.95	24.00	24.20	23.96
	(3.192)	(3.268)	(3.296)	(3.323)	(3.338)	(3.312)	(3.286)
10-Year Avg. GS Rainfall (100 mm)	1.283	1.164	1.314	1.192	1.154	1.189	1.216
	(0.657)	(0.615)	(0.640)	(0.633)	(0.598)	(0.650)	(0.634)
Agricultural Labor Share	0.151	0.257	0.238	0.224	0.257	0.300	0.238
	(0.103)	(0.130)	(0.128)	(0.117)	(0.131)	(0.148)	(0.135)
Non-Agricultural Worker Share	0.272	0.257	0.284	0.296	0.385	0.417	0.318
	(0.149)	(0.152)	(0.157)	(0.164)	(0.183)	(0.188)	(0.176)
Urbanization	0.159	0.176	0.207	0.228	0.245	0.275	0.215
	(0.139)	(0.148)	(0.152)	(0.158)	(0.168)	(0.182)	(0.163)
Migrant Share	0.0198	.	0.0241	0.0227	0.0221	0.0385	0.0254
	(0.0163)	(.)	(0.0187)	(0.0163)	(0.0190)	(0.0484)	(0.0275)

Note: Table presents summary statistics for the weather variables and Census outcome variables over time for the sample of districts for which non-agricultural worker share is non-missing for all years (N=270). Migrant share data is unavailable for 1971.

Table 2: Effect of Rising Temperatures Using Panel Specification

	Ag Labor Share		Non-Ag Worker Share		Urbanization		Migrant Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.181	0.157	-0.077	-0.086	-0.021	0.001	-0.013	-0.018
	(0.059)***	(0.062)**	(0.033)**	(0.031)***	(0.042)	(0.045)	(0.059)	(0.064)
	[0.077]**	[0.076]**	[0.038]**	[0.035]**	[0.046]	[0.046]	[0.062]	[0.068]
P	-0.153	-0.081	-0.026	-0.001	0.025	0.000	-0.016	-0.000
	(0.059)**	(0.060)	(0.030)	(0.030)	(0.042)	(0.044)	(0.050)	(0.058)
	[0.098]	[0.093]	[0.036]	[0.032]	[0.044]	[0.045]	[0.058]	[0.064]
Region-year trends	Y	N	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y	N	Y
Observations	1,548	1,548	1,620	1,620	1,596	1,596	1,350	1,350

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Temperature and precipitation are decadal averages of the past ten growing seasons. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which the dependent variable is non-missing in all years. 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.

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

Table 3: Effect of Rising Temperatures Using Long Difference Specification

	Ag Labor Share (1)	Non-Ag Worker Share (2)	Urbanization (3)	Migrant Share (4)
T	0.3819	-0.1491	-0.0434	-0.1716
	(0.0995)***	(0.0531)***	(0.0827)	(0.1619)
	[0.2292]*	[0.0621]**	[0.1569]	[0.2092]
P	0.3256	0.0202	-0.2617	0.0350
	(0.2236)	(0.1080)	(0.1830)	(0.4016)
	[0.4429]	[0.0777]	[0.1494]*	[0.5837]
Region FE	Y	Y	Y	Y
Observations	258	270	266	267

Note: The dependent variable in each column is the difference (in natural logarithm) of an outcome between two 30-year periods, 1961-1981 and 1991-2011. The outcomes are the share of agricultural laborers in Column (1), the share of non-agricultural workers in Column (2), urbanization rates in Column (3), and the share of intra-district migrants in Column (4). The independent variables are differences in average growing-season temperature and precipitation over the same time periods. Data are district-level data constructed from the Indian Census. All columns include region fixed effects. We present standard errors 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: Heterogeneous Effects of Rising Temperatures by Road Network Density & Bank Credit per Capita

	Ag Labor Share		Non-Ag Worker Share		Urbanization		Migrant Share	
Panel A: Road Network Density	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.322	0.355	-0.104	-0.137	-0.193	-0.190	-0.102	-0.149
	(0.091)***	(0.099)***	(0.055)*	(0.058)**	(0.076)**	(0.083)**	(0.072)	(0.078)*
	[0.108]***	[0.118]***	[0.054]*	[0.052]***	[0.067]***	[0.068]***	[0.074]	[0.075]**
T x High Road Density	-0.320	-0.362	0.094	0.119	0.222	0.231	0.188	0.227
	(0.110)***	(0.114)***	(0.070)	(0.068)*	(0.092)**	(0.098)**	(0.104)*	(0.108)**
	[0.125]**	[0.128]***	[0.068]	[0.062]*	[0.081]***	[0.084]***	[0.098]*	[0.100]**
Region-year trends	Y	N	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y	N	Y
P-val of sum, cluster	0.974	0.919	0.827	0.633	0.613	0.511	0.269	0.347
P-val of sum, Conley	0.978	0.929	0.860	0.665	0.632	0.519	0.285	0.340
Observations	1,458	1,458	1,458	1,458	1,452	1,452	1,210	1,210
Panel B: Bank Credit Per Capita	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.288	0.271	-0.147	-0.157	-0.064	-0.042	0.012	-0.013
	(0.067)***	(0.070)***	(0.054)***	(0.048)***	(0.067)	(0.069)	(0.066)	(0.069)
	[0.083]***	[0.082]***	[0.066]**	[0.062]**	[0.071]	[0.068]	[0.086]	[0.084]
T x High Bank Credit	-0.210	-0.217	0.108	0.106	0.032	0.030	-0.041	-0.029
	(0.109)*	(0.110)**	(0.065)*	(0.059)*	(0.086)	(0.087)	(0.099)	(0.099)
	[0.108]*	[0.111]*	[0.072]	[0.066]	[0.080]	[0.078]	[0.116]	[0.114]
Region-year trends	Y	N	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y	N	Y
P-val of sum, cluster	0.389	0.558	0.328	0.181	0.556	0.826	0.726	0.632
P-val of sum, Conley	0.418	0.580	0.298	0.124	0.515	0.808	0.738	0.652
Observations	1,548	1,548	1,620	1,620	1,596	1,596	1,350	1,350

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Temperature and precipitation are decadal averages of the past ten growing seasons. Data are district-level panel data constructed from the Indian Census. Panel A explores heterogeneous effects by road network density; *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, year and high road density-by-year fixed effects, as well as controls for decadal precipitation and decadal precipitation interacted with the road density dummy. Panel B explores heterogeneous effects by bank credit per capita; *High Bank Credit* is a binary variable that takes the value 1 if the district has above median bank credit per capita at baseline. All columns include district, year and high bank credit-by-year fixed effects, as well as controls for decadal precipitation and decadal precipitation interacted with the bank credit dummy. We restrict our sample to districts for which road density/bank credit data is non-missing and 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 5: Effect of Rising Temperatures on Consumption and Sectoral Worker Shares

	Total		Food		Non-Food	
Panel A: Consumption	(1)	(2)	(3)	(4)	(5)	(6)
T	-0.096	-0.091	-0.069	-0.071	-0.113	-0.122
	(0.031)***	(0.033)***	(0.025)***	(0.027)***	(0.046)**	(0.050)**
	[0.039]**	[0.035]***	[0.031]**	[0.031]**	[0.053]**	[0.053]**
P	0.064	0.063	0.004	0.001	0.112	0.127
	(0.048)	(0.052)	(0.039)	(0.045)	(0.073)	(0.078)
	[0.052]	[0.056]	[0.042]	[0.047]	[0.083]	[0.092]
Region-year trends	Y	N	Y	N	Y	N
State-year trends	N	Y	N	Y	N	Y
Observations	1,590	1,590	1,590	1,590	1,590	1,590
	Services		Construction		Manufacturing	
Panel B: Non-ag sectoral share	(1)	(2)	(3)	(4)	(5)	(6)
T	-0.139	-0.260	0.040	-0.132	-0.006	-0.151
	(0.070)**	(0.073)***	(0.181)	(0.198)	(0.102)	(0.117)
	[0.091]	[0.090]***	[0.231]	[0.196]	[0.122]	[0.126]
P	0.071	0.180	0.024	0.383	0.077	0.235
	(0.095)	(0.102)*	(0.381)	(0.418)	(0.160)	(0.183)
	[0.129]	[0.145]	[0.379]	[0.411]	[0.181]	[0.225]
Region-year trends	Y	N	Y	N	Y	N
State-year trends	N	Y	N	Y	N	Y
Observations	2,120	2,120	2,120	2,120	2,120	2,120

Note: The dependent variable in Panel A is the natural logarithm of the district average per capita annual total consumption in Columns (1) and (2), of annual food consumption in Columns (3) and (4), and of annual non-food consumption in Columns (5) and (6). The dependent variable in Panel B is the natural logarithm of the share of workers engaged in services in Columns (1) and (2), of the share of workers engaged in construction in Columns (3) and (4), and of the share of workers engaged in manufacturing in Columns (5) and (6). Temperature and precipitation are decadal averages of the past ten growing seasons. Data are district-level panel data aggregated from the National Sample Survey. For both panels, we restrict our sample to districts for which the dependent variable is non-missing in all years. In addition, we restrict our sample to districts with non-missing observations of non-agricultural shares across all years in the PCA data. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Heterogeneous Effects of Rising Temperatures by Social Grouping and Education

	Non-scheduled caste/tribe				Scheduled caste/tribe			
	Ag Worker Share		Non-Ag Worker Share		Ag Worker Share		Non-Ag Worker Share	
Panel A: By social grouping	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.305	0.252	-0.087	-0.206	0.405	0.397	-0.161	-0.327
	(0.075)***	(0.085)***	(0.065)	(0.065)***	(0.129)***	(0.108)***	(0.097)*	(0.090)***
	[0.090]***	[0.088]***	[0.088]	[0.081]**	[0.149]***	[0.121]***	[0.107]	[0.090]***
P	-0.056	0.032	-0.097	0.027	0.026	-0.024	-0.110	0.082
	(0.232)	(0.165)	(0.083)	(0.090)	(0.154)	(0.137)	(0.136)	(0.144)
	[0.238]	[0.174]	[0.113]	[0.128]	[0.159]	[0.150]	[0.142]	[0.155]
Region-year trends	Y	N	Y	N	Y	N	Y	N
State-year trends	N	Y	N	Y	N	Y	N	Y
Observations	2,120	2,120	2,120	2,120	2,120	2,120	2,120	2,120

	Primary school or above				Below primary school			
	Ag Worker Share		Non-Ag Worker Share		Ag Worker Share		Non-Ag Worker Share	
Panel B: By education	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.262	0.297	-0.028	-0.155	0.238	0.224	-0.194	-0.311
	(0.072)***	(0.073)***	(0.052)	(0.050)***	(0.068)***	(0.054)***	(0.078)**	(0.078)***
	[0.091]***	[0.087]***	[0.069]	[0.068]**	[0.080]***	[0.068]***	[0.104]*	[0.087]***
P	0.028	0.019	-0.010	0.105	0.018	-0.005	-0.032	0.125
	(0.147)	(0.144)	(0.073)	(0.076)	(0.113)	(0.131)	(0.126)	(0.122)
	[0.179]	[0.179]	[0.096]	[0.107]	[0.122]	[0.135]	[0.147]	[0.157]
Region-year trends	Y	N	Y	N	Y	N	Y	N
State-year trends	N	Y	N	Y	N	Y	N	Y
Observations	2,120	2,120	2,120	2,120	2,120	2,120	2,120	2,120

Note: The dependent variable is the natural logarithm of the share of agricultural workers in Columns (1), (2), (5) and (6), and of the share of non-agricultural workers in Columns (3), (4), (7) and (8), within the specified social or education group. Temperature and precipitation are decadal averages of the past ten growing seasons. Data are district-level panel data aggregated from the National Sample Survey. We restrict our sample to districts for which the dependent variable is non-missing in all years. In addition, we restrict our sample to districts with non-missing observations of non-agricultural shares in all years in the PCA data. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A Theoretical Framework

We develop a simple general equilibrium model of sectoral allocation with three sectors and two regions. We have two goals in developing this model. First, we want to demonstrate that adverse shocks to agricultural productivity can decrease the allocation of labor to the non-agricultural sector via local demand effects. Second, we want to show that high transportation costs can intensify this effect. We develop a model in the spirit of the model developed in Gollin and Rogerson (2014).[29] More broadly, the mechanisms underlying our theoretical model are similar to the mechanisms in many models of structural transformation (Matsuyama, 1992; Kongsamut et al., 2001; Gollin et al., 2002; Gollin and Rogerson, 2014; Herrendorf et al., 2014; Bustos et al., 2016).

Our model features three sectors: the agricultural sector (denoted by A), the manufacturing sector (M), and the service sector (S). There are two regions: the rural region (R) which produces agriculture and services, and the urban region (U) which produces manufacturing and services. Services are not tradable between the regions, but agricultural and manufacturing goods can be traded. We assume iceberg transportation costs of size q for the agricultural good only. Specifically, we assume that if one unit of the agricultural good is transported to the urban region, there will be fractional losses q so that only $(1-q)$ units will be received.

Individuals in both regions have Stone-Geary preferences over the three goods, which, for tractability, we assume are of the form:

$$u(c_A, c_M, c_S) = ln(c_A - \bar{a}) + ln(c_M - \bar{m}) + ln(c_S)$$

We assume that $\bar{a} > 0$, so that the income elasticity of the agricultural good is less than one. Labor is the only input for all three sectors. Each region has one unit of available labor, and the labor in region J devoted to sector I is

[29] Our model differs in three important ways from the model developed in Gollin and Rogerson (2014): we model one rural region instead of two; we shut down migration between the regions; and we model a (non-tradable) service sector in addition to the agricultural and manufacturing sectors.

denoted by L_I^J. Urban labor is divided between manufacturing and services so that $L_M^U + L_S^U = 1$, and all agricultural goods consumed in the urban region are imported from the rural region. In the rural region, labor is divided between the agricultural and services sectors, so that $L_A^R + L_S^R = 1$ and all manufacturing goods consumed in the rural region are imported from the urban region. We abstract away from the possibility of migration between regions. The production function of the agricultural sector in the rural region is given by:

$$Y_A^R = \theta_A^R (L_A^R)^\beta = \theta_A^R (1 - L_S^R)^\beta$$

where θ_A^R is rural agricultural total factor productivity and β is the elasticity of output with respect to labor. We assume that $\beta < 1$.

In the context of our study, shocks to θ_A^R will be driven by higher temperatures. The dependence of agricultural productivity on temperature is well-documented in the empirical literature (Schlenker and Roberts, 2009; Dell et al., 2012; Taraz, 2018). We note that in the Indian context, there is also empirical evidence that higher temperatures reduce non-agricultural productivity (Somanathan et al., 2021), but these reductions are smaller in magnitude relative to impacts on agricultural productivity (Jain et al., 2020).[30]

The output of the manufacturing sector in the urban region is given by:

$$Y_M^U = \theta_M^U L_M^U = \theta_M^U (1 - L_S^U)$$

The output of the service sector in each region is given by:

$$Y_S^J = \theta_S^J L_S^J$$

where θ_S^J is the service sector total factor productivity in region J. We further assume that $\theta_S^R = 1$.

To solve for the competitive equilibrium of our model, we note that our model contains no externalities and hence we can apply the First Welfare

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

Theorem. In other words, the competitive equilibrium will be the same as the solution to the Social Planner's problem, and that is what we will solve for, applying equal weights to each region. The maximization problem for the social planner is given by:

$$\max_{c_A^R,c_A^U,c_M^R,c_M^U,c_S^R,c_S^U,L_S^R,L_S^U} ln(c_A^R-\bar{a})+ln(c_M^R-\bar{m})+ln(c_S^R)+ln(c_A^U-\bar{a})+ln(c_M^U-\bar{m})+ln(c_S^U)$$

where c_I^J be the consumption of good I in region J. This maximization problem is subject to the following four feasibility constraints:

$$\frac{c_A^U}{(1-q)}+c_A^R = \theta_A^R(1-L_S^R)^\beta \tag{A1}$$

$$c_M^U+c_M^R = \theta_M^U(1-L_S^U) \tag{A2}$$

$$c_S^R = L_S^R \tag{A3}$$

$$c_S^U = \theta_S^U L_S^U \tag{A4}$$

Equation (A1) states that the sum of agricultural consumption in the urban region (factoring in transportation costs) and rural region must equal the output of the agricultural sector in the rural region. Equation (A2) states that the total manufacturing consumption across both regions must be equal to the manufacturing output of the urban region. Equations (A3) and (A4) state that service consumption in each region must equal service production in that same region, since services are not tradable.

Substituting Equations (A3) and (A4) into the maximization problem, we get a simplified maximization that is subject to Equations (A3) and (A4) only:

$$\max_{c_A^R,c_A^U,c_M^R,c_M^U,L_S^R,L_S^U} ln(c_A^R-\bar{a})+ln(c_M^R-\bar{m})+ln(\theta_S^R L_S^R)+ln(c_A^U-\bar{a})+ln(c_M^U-\bar{m})+ln(\theta_S^U L_S^U)$$

Manipulating the first order constraints of this maximization problem, we are able to solve for the following expressions for the rural and urban consumption bundles:

$$
\begin{aligned}
c_A^R =& \frac{\theta_A^R(1-L_S^R)^\beta(1-q)-\bar{a}q}{2(1-q)} \\
c_A^U =& \frac{\theta_A^R(1-L_S^R)^\beta(1-q)+\bar{a}q}{2} \\
c_M^U =& c_M^R = \frac{\theta_M^U(1-L_S^U)}{2}
\end{aligned}
$$

We can then write a further simplified version of the Social Planner's Problem, by substituting in these expressions for c_A^M, c_A^S, c_A^U, and c_S^U:

$$
\begin{aligned}
\max_{L_S^R, L_S^U} ln&\left(\frac{\theta_A^R(1-L_S^R)^\beta(1-q)-(2-q)\bar{a}}{2(1-q)}\right)+ln\left(\frac{\theta_M^U-\theta_M^U L_S^U-2\bar{m}}{2}\right)+ln\left(L_S^R\right) \\
&+ln\left(\frac{\theta_A^R(1-L_S^R)^\beta(1-q)-(2-q)\bar{a}}{2}\right)+ln\left(\frac{\theta_M^U-\theta_M^U L_S^U-2\bar{m}}{2}\right)+ln\left(\theta_S^U L_S^U\right)
\end{aligned}
$$

Using the rule for a logarithm of a quotient, we can simplify this maximization problem to:

$$
\begin{aligned}
\max_{L_S^R, L_S^U} 2ln&\left(\theta_A^R(1-L_S^R)^\beta(1-q)-(2-q)\bar{a}\right)+2ln\left(\theta_M^U-\theta_M^U L_S^U-2\bar{m}\right)+ln\left(L_S^R\right) \\
&+ln\left(\theta_S^U L_S^U\right)-ln(1-q)
\end{aligned}
$$

Next, we take for the first order conditions for L_S^R and we get:

$$
\frac{-2\beta(1-q)\theta_A^R(1-L_S^R)^{\beta-1}}{\theta_A^R(1-L_S^R)^\beta(1-q)-(2-q)\bar{a}}+\frac{1}{L_S^R}=0
$$

Manipulating this expression, we get

$$
(1-L_S^R)^\beta(1-q)-2\beta(1-q)(1-L_S^R)^{\beta-1}L_S^R=\frac{\bar{a}(2-q)}{\theta_A^R} \tag{A5}
$$

Now, we can solve for the comparative statics of interest. First, we can solve for the impact of a change in agricultural productivity on rural labor supply to the service sector. Taking the implicit derivative of Equation (A5) with respect to agricultural productivity θ_A^R and rearranging terms, we get:

$$\frac{\delta L_S^R}{\delta \theta_A^R} = \frac{\bar{a}(2-q)}{(\theta_A^R)^2(1-q)} \times \frac{1}{3\beta(1-L_S^R)^{\beta-1} + 2\beta(1-\beta)(1-L_S^R)^{\beta-2} L_S^R} \tag{A6}$$

We can now determine the sign of the expression on the right-hand side of Equation (A6). We have assumed that $\bar{a} > 0$ and $q < 1$, so the numerator of the fraction is positive. We have also assumed that $\beta < 1$, and $L_S^R < 1$, so the denominator is also positive. Therefore we have shown that $\frac{\delta L_S^R}{\delta \theta_A^R} > 0$, which means that an increase in agricultural productivity triggers an increase of labor allocated to the service sector (and hence decrease labor supply to the agricultural sector). Conversely, an adverse shock to agricultural productivity will decrease rural service sector employment and increase rural agricultural labor supply.

In addition to the direct effect of agricultural productivity shocks on non-agricultural labor supply, we are also interested in the role of transportation costs in modulating this relationship. Looking at Equation (A6), we note that increasing q (in the range $0 < q < 1$) will increase the right-hand side of Equation (A6), and hence $\frac{\delta^2 L_S^R}{\delta \theta_A^R \delta q} > 0$, which means that places with higher transportation costs will face intensified local demand effects.

It is worth noting two important limitations of our model. First, our model does not allow for migration between rural and urban regions. This assumption may be reasonable in short-term, since the costs of seasonal migration in India is high and workers prefer local public works to migration (Imbert and Papp, 2020). In the medium- to long-term, the level of within-district migration in India is also very low (Kone et al., 2018). Nevertheless, the implications of our model should be caveated when applied to other contexts with higher levels of cross-region migration. On the one hand, local demand effects could be dampened as cross-region migration arbitrages away the difference in agricultural productivity shocks. On the other hand, spatial linkages created by migration could expose regions unaffected by rising temperatures to agricultural risks elsewhere.

Second, our set-up does not allow for multiple periods, and therefore it cannot provide comparisons of the short- and long-term effects. Without formally modeling multiple periods, the predictions from the above model can be taken as short-term dynamics. To gauge possible intensification or adaptation effects in the longer-term, we posit two possible extensions to the model. The first is to introduce liquidity and mobility costs, which may intensify the effects of rising temperatures over time. If farmers' agricultural incomes are stochastic, it follows that a short duration of high temperatures will reduce farm income, which renders the costs of switching sectors infeasible for a relatively small fraction of farmers. However, a longer period of sustained high temperatures will lead to long-lasting farm income reductions, leaving a much greater fraction of the population not able to afford the liquidity and mobility costs. A second possible extension is to consider costly investment in human capital. Under this set-up, warming affects not only the affordability of switching sectors for *current* workers, but also the human capital investments of *future* workers. A growing literature documents that high temperatures have negative and persistent impacts on human capital.[31] It may be the case that the adverse impacts of warming on structural transformation in the short-term can be compounded by the dampening of human capital accumulation in the longer-term.

[31]For example, Garg et al. (2020) find that higher temperatures in India reduce contemporaneous human capital due to an agricultural income channel. Fishman et al. (2019) find that high temperatures in Ecuador around the time at birth have long-term effects on human capital and earnings productivity that persist into adulthood; Hu and Li (2019) find similar effects looking at China.

B Appendix Figures and Tables

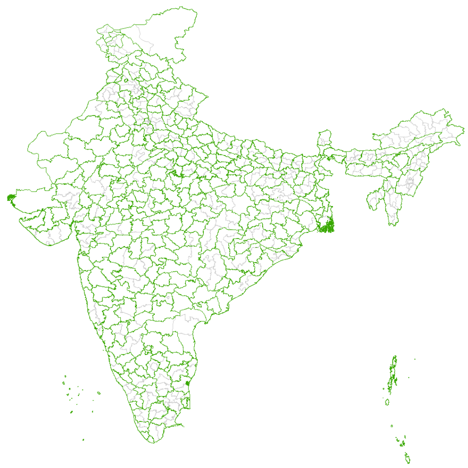

(a) Consistent district boundaries, 1961-2011

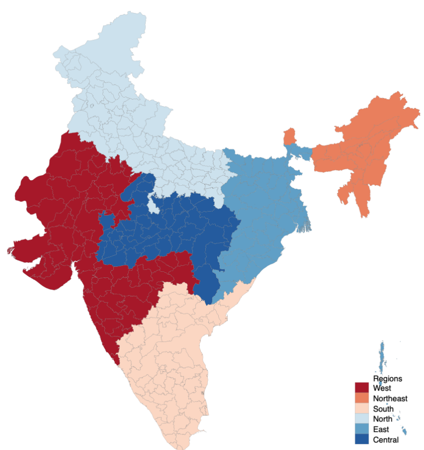

(b) Districts, by region

Figure B1: Figure illustrates the 288 consistent district boundaries over 1961-2011 (panel a), and the 287 districts by region used in the analysis (panel b). Lakshadweep is dropped due to lack of weather records. We classify districts into six regions based on the Government of India's administrative regional classification.

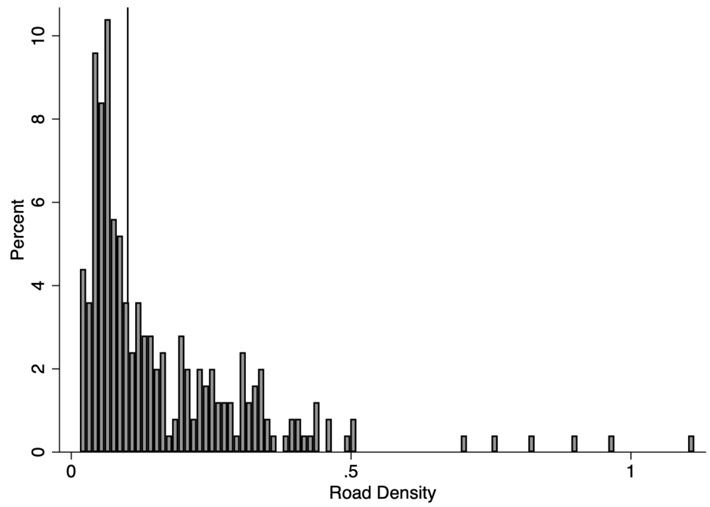

(a) Histogram

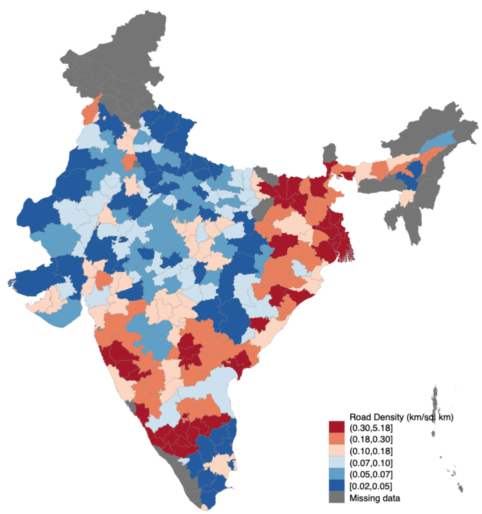

(b) Spatial Distribution

Figure B2: 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 of the distribution (0.10 km/km^2).

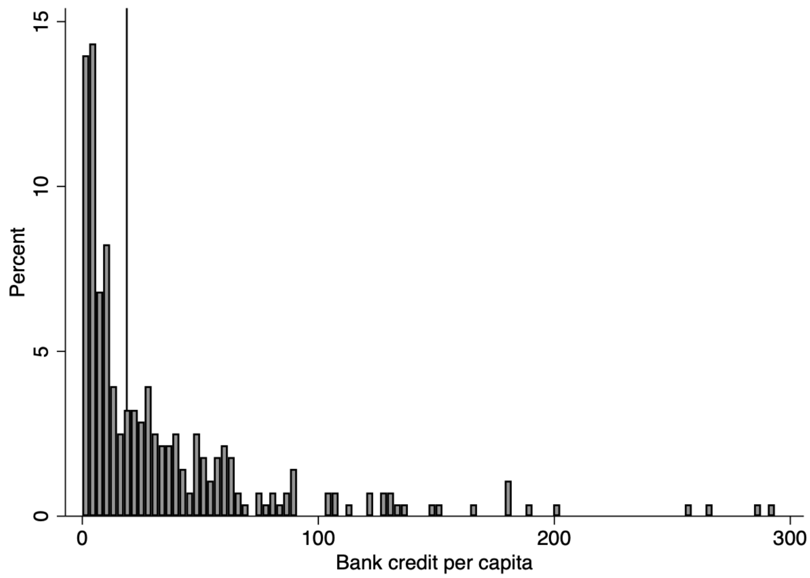

(a) Histogram

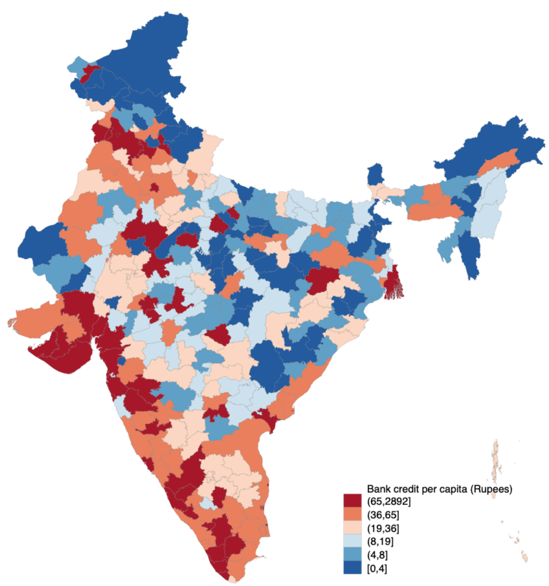

(b) Spatial Distribution

Figure B3: Figure plots the bank credit per capita (Rupees) 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 of the distribution (19 Rupees).

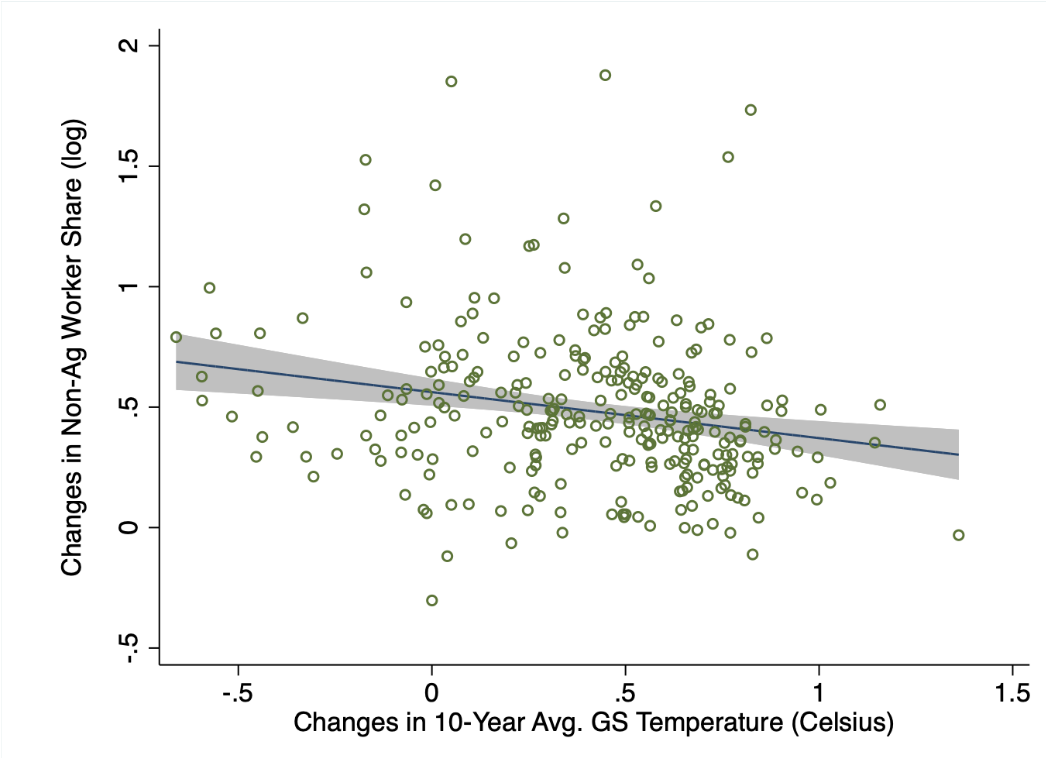

Figure B4: Figure plots the relationship between long-run changes in non-agricultural worker share and changes in 10-year average growing season temperature between 1961 and 2011. Data comes from district-level panel data constructed from the Indian Census. Each dot represents a district in our sample. For each district, we take the difference in the ten-year average growing season temperature between 1961 and 2011, as well as the difference in the natural log of the non-agricultural worker share between 1961 and 2011, and plot them against one another. Fitted linear regression lines and 95% confidence intervals are presented along with the scatter plots.

Table B1: Summary Statistics by Year

Year	1987-88	1993-94	1999-2000	2004-05	2005-06	2007-08	2009-10	2011-12	
NSS Round	43	50	55	61	62	64	66	68	Total
10-Year Avg. GS Temperature (Celsius)	23.95	23.95	23.92	24.04	24.09	24.16	24.17	24.22	24.06
	(3.311)	(3.322)	(3.352)	(3.337)	(3.347)	(3.320)	(3.303)	(3.309)	(3.322)
10-Year Avg. GS Rainfall (100 mm)	1.306	1.167	1.171	1.141	1.114	1.135	1.144	1.170	1.168
	(0.637)	(0.634)	(0.603)	(0.597)	(0.602)	(0.621)	(0.634)	(0.651)	(0.624)
Agricultural Worker Share	0.543	0.650	0.545	0.472	0.382	0.505	0.390	0.363	0.481
	(0.172)	(0.218)	(0.177)	(0.124)	(0.150)	(0.153)	(0.122)	(0.120)	(0.183)
Non-Agricultural Worker Share	0.371	0.311	0.421	0.491	0.572	0.455	0.580	0.606	0.476
	(0.161)	(0.209)	(0.163)	(0.112)	(0.146)	(0.145)	(0.117)	(0.113)	(0.179)
Manufacturing Worker Share	0.0993	0.0818	0.0998	0.119	0.144	0.106	0.114	0.121	0.111
	(0.0644)	(0.0710)	(0.0673)	(0.0634)	(0.0845)	(0.0698)	(0.0674)	(0.0661)	(0.0716)
Services Worker Share	0.220	0.192	0.273	0.307	0.357	0.264	0.347	0.336	0.287
	(0.105)	(0.145)	(0.113)	(0.0808)	(0.0992)	(0.0975)	(0.0865)	(0.0879)	(0.118)
Construction Worker Share	0.0448	0.0294	0.0419	0.0574	0.0611	0.0767	0.110	0.112	0.0667
	(0.0562)	(0.0262)	(0.0280)	(0.0289)	(0.0320)	(0.0458)	(0.0596)	(0.0506)	(0.0516)

Note: Table presents summary statistics for the weather variables and National Sample Survey outcome variables over time for the sample of districts used in regression analysis — this include districts with non-missing observations of non-agricultural shares in all years in the PCA data (N=270). The 50th round in 1993-1994 has incomplete coverage of the urban population — only a quarter of the districts have their urban households represented in the survey. Therefore, the district-level summary statistics in 1993-1994 have relatively higher shares of agricultural workers, and lower shares of non-agricultural workers, compared to those in other rounds.

Table B2: Effect of Rising Temperatures on Primary and Secondary Occupations

	Agricultural Labor Share					
	Total		Main		Marginal	
Panel A: Main and Marginal Employment in Census	(1)	(2)	(3)	(4)	(5)	(6)
T	0.198	0.176	0.155	0.143	0.518	0.473
	(0.069)***	(0.073)**	(0.069)**	(0.072)**	(0.128)***	(0.136)***
	[0.088]**	[0.089]**	[0.097]	[0.092]	[0.219]**	[0.214]**
P	-0.086	-0.051	0.011	0.060	-0.088	-0.016
	(0.070)	(0.072)	(0.067)	(0.065)	(0.125)	(0.124)
	[0.103]	[0.106]	[0.101]	[0.097]	[0.243]	[0.231]
Region-year trends	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y
Observations	1,290	1,290	1,290	1,290	1,290	1,290

	Ag Worker Share				Non-Ag Worker Share			
	Primary Occupation		Secondary Occupation		Primary Occupation		Secondary Occupation	
Panel B: Primary and Secondary Employment in NSS	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.191	0.213	0.211	0.132	-0.076	-0.202	-0.335	-0.223
	(0.061)***	(0.054)***	(0.079)***	(0.080)*	(0.067)	(0.062)***	(0.166)**	(0.181)
	[0.074]***	[0.070]***	[0.082]**	[0.082]	[0.085]	[0.071]***	[0.177]*	[0.187]
P	-0.086	-0.110	0.079	0.042	-0.058	0.073	-0.341	-0.184
	(0.119)	(0.096)	(0.086)	(0.101)	(0.093)	(0.098)	(0.210)	(0.230)
	[0.137]	[0.126]	[0.095]	[0.099]	[0.125]	[0.140]	[0.221]	[0.238]
Region-year trends	Y	N	Y	N	Y	N	Y	N
State-year trends	N	Y	N	Y	N	Y	N	Y
Observations	2,120	2,120	2,072	2,072	2,120	2,120	2,072	2,072

Note: The dependent variables in Panel A are the natural logarithm of the shares of total agricultural laborers (columns 1-2), of main agricultural laborers (columns 3-4), and of marginal agricultural laborers (columns 5-6). Temperature and precipitation are decadal averages of the past ten growing seasons. The sample in Panel A comes from district-level data constructed from the Indian Census. The sample is restricted to districts for which the dependent variable is non-missing in all years, and excludes data from 1991 as counts of main and marginal agricultural workers are not available in that year. The dependent variables in Panel B are the natural logarithm of the share of workers engaged in agriculture as a primary occupation (columns 1-2), engaged in agriculture as a secondary occupation (columns 3-4), engaged in non-agriculture as a primary occupation (columns 5-6), and engaged in non-agriculture as a secondary occupation (columns 7-8). The sample in Panel B comes from district-level data aggregated from the National Sample Survey. The sample is restricted to districts for which i) the dependent variables are non-missing in all years in the NSS data, and ii) the non-agricultural share variable is non-missing across all years in the PCA data. All columns in both panels include district and year fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B3: Heterogeneous Effect of Rising Temperatures by Long-Run Temperature

	Ag Labor Share		Non-Ag Worker Share		Urbanization		Migrant Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T x Less Hot District	0.156	0.170	-0.049	-0.071	0.036	0.050	-0.078	-0.103
	(0.088)*	(0.090)*	(0.047)	(0.044)	(0.058)	(0.060)	(0.070)	(0.074)
	[0.114]	[0.110]	[0.050]	[0.048]	[0.062]	[0.058]	[0.066]	[0.066]
T x Hot District	0.239	0.203	-0.097	-0.079	-0.047	-0.043	0.082	0.103
	(0.087)***	(0.089)**	(0.044)**	(0.042)*	(0.064)	(0.069)	(0.101)	(0.103)
	[0.088]***	[0.087]**	[0.048]**	[0.041]*	[0.060]	[0.063]	[0.102]	[0.104]
Region-year trends	Y	N	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y	N	Y
Observations	1,548	1,548	1,620	1,620	1,596	1,596	1,350	1,350

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Temperature is the decadal average of the past ten growing seasons. "Hot District" is a binary variable that takes the value 1 if a district's average growing season temperature for the period 1901-2014 is above the median for that period; "Less Hot District" takes the value 1 if a district's average temperature for 1901-2014 is below the median. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which the dependent variable is non-missing in all years. All columns include district fixed effects and year fixed effects and control for decadal precipitation interacted with the hot and less hot district dummies. 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 B4: Effect of Rising Temperatures using National Sample Survey

	Ag Worker Share		Non-Ag Worker Share	
Panel A: Panel specification	(1)	(2)	(3)	(4)
T (Decadal Average)	0.283	0.281	-0.122	-0.248
	(0.063)***	(0.059)***	(0.065)*	(0.061)***
	[0.083]***	[0.077]***	[0.087]	[0.074]***
P (Decadal Average)	-0.014	-0.023	-0.048	0.103
	(0.133)	(0.146)	(0.094)	(0.096)
	[0.154]	[0.161]	[0.122]	[0.135]
Region-year trends	Y	N	Y	N
State-year trends	N	Y	N	Y
Observations	2,120	2,120	2,120	2,120
Panel B: Short and long-term effects	(1)	(2)	(3)	(4)
Current Year T	-0.086	-0.061	0.122	0.099
	(0.025)***	(0.030)**	(0.023)***	(0.024)***
	[0.032]***	[0.037]*	[0.033]***	[0.032]***
Decadal Average T	0.285	0.270	-0.123	-0.239
	(0.063)***	(0.061)***	(0.063)*	(0.061)***
	[0.080]***	[0.075]***	[0.078]	[0.069]***
Region-year trends	Y	N	Y	N
State-year trends	N	Y	N	Y
Observations	2,120	2,120	2,120	2,120

Note: The dependent variable is the natural logarithm of the share of individuals engaged in agriculture in Columns (1) and (2), and of the share of individuals engaged in non-agricultural sectors in Columns (3) and (4). Data are district-level panel data aggregated from the National Sample Survey. In panel A, temperature and precipitation measures are decadal averages of the past ten growing seasons, and in Panel B, they are current growing season monthly averages and decadal averages of the past ten growing seasons. All columns include district and year fixed effects. We restrict our sample to districts for which the dependent variable is non-missing in all years. In addition, we restrict our sample to districts with non-missing observations of non-agricultural shares across all years in the PCA data. 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 B5: Effect of Rising Temperatures on Agricultural Yields

	Agricultural Yields	
Panel A: Main Effects	(1)	(2)
T	-0.0840	-0.0606
	(0.0173)***	(0.0115)***
	[0.0205]***	[0.0203]***
P	0.0906	0.1506
	(0.0115)***	(0.0133)***
	[0.0145]***	[0.0142]***
Region-time trends	Y	N
Region-decade FE	N	Y
Observations	11,860	11,860
Panel B: By Road Density	(1)	(2)
T	-0.094	-0.088
	(0.022)***	(0.022)***
	[0.025]***	[0.024]***
T x High Road Density	0.010	0.006
	(0.025)	(0.025)
	[0.028]	[0.027]
Region-year trends	Y	N
Region-decade FE	N	Y
P-val of sum, cluster	0.0001	0.0001
P-val of sum, Conley	0.0009	0.0012
Observations	11,860	11,860
Panel C: By Bank Credit per Capita	(1)	(2)
T	-0.101	-0.096
	(0.017)***	(0.017)***
	[0.021]***	[0.021]***
T x High Bank Credit	0.032	0.031
	(0.025)	(0.024)
	[0.025]	[0.024]
Region-year trends	Y	N
Region-decade FE	N	Y
P-val of sum, cluster	0.0047	0.0084
P-val of sum, Conley	0.0084	0.0129
Observations	11,860	11,860

Note: The dependent variable is the natural logarithm of aggregate yields. Temperature and precipitation are annual averages over the growing season months. We use annual data from VDSA spanning 1966 to 2010. All columns include district and year fixed effects. Panels B and C control for precipitation and precipitation interacted with the heterogeneity measure. 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 B6: Effect of Rising Temperatures on Rural & Urban Non-Agricultural Worker Share

	Non-Agricultural Worker Share			
	Rural		Urban	
	(1)	(2)	(3)	(4)
T	-0.084	-0.106	-0.007	-0.016
	(0.040)**	(0.038)***	(0.011)	(0.012)
	[0.060]	[0.055]*	[0.017]	[0.016]
P	-0.015	0.014	-0.031	-0.009
	(0.042)	(0.043)	(0.014)**	(0.014)
	[0.061]	[0.055]	[0.014]**	[0.014]
Region-year trends	Y	N	Y	N
Region-year FE	N	Y	N	Y
Observations	1,608	1,608	1,596	1,596

Note: The dependent variable is the natural logarithm of the share of rural non-agricultural workers in Columns (1) and (2), and the natural logarithm of the share of urban non-agricultural workers in Columns (3) and (4). Temperature and precipitation are decadal averages of the past ten growing seasons. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which the dependent variable is non-missing in all years. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B7: Effect of Rising Temperatures on Cultivator & Agricultural Worker Share

Panel A: Cultivator Share	(1)	(2)	(3)	(4)	(5)	(6)
T	0.090	0.116	0.036	0.075	0.077	0.105
	(0.071)	(0.074)	(0.055)	(0.057)	(0.039)**	(0.039)***
	[0.060]	[0.059]**	[0.053]	[0.051]	[0.043]*	[0.037]***
T x High Road Density			0.006	-0.013		
			(0.074)	(0.074)		
			[0.066]	[0.063]		
T x High Bank Credit					-0.008	-0.018
					(0.113)	(0.114)
					[0.099]	[0.100]
P-val of sum, cluster			0.445	0.264	0.546	0.444
P-val of sum, Conley			0.438	0.226	0.451	0.343
Observations	1,620	1,620	1,620	1,620	1,620	1,620
Panel B: Ag Labor + Cultivator Share						
T	0.072	0.093	0.083	0.121	0.115	0.141
	(0.036)**	(0.038)**	(0.045)*	(0.048)**	(0.039)***	(0.039)***
	[0.043]*	[0.045]**	[0.038]**	[0.038]***	[0.049]**	[0.049]***
T x High Road Density			-0.037	-0.066		
			(0.061)	(0.062)		
			[0.048]	[0.048]		
T x High Bank Credit					-0.094	-0.103
					(0.062)	(0.060)*
					[0.062]	[0.058]*
Region-year trends	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y
P-val of sum, cluster			0.303	0.231	0.690	0.481
P-val of sum, Conley			0.271	0.170	0.703	0.498
Observations	1,548	1,548	1,548	1,548	1,548	1,548

Note: The dependent variable is the natural logarithm of the share of cultivators in Panel A, and of the share of agricultural workers (agricultural laborers and cultivators) in Panel B. Temperature is the decadal average of the past ten growing seasons. All columns include district and year fixed effects. We restrict our sample to districts for which the dependent variable is non-missing in all years. *High Road Density* is a binary variable that takes the value 1 if the district has above median road density at baseline. *High Bank Credit* is a binary variable that takes the value 1 if the district has above median bank credit per capita at baseline. Data are district-level panel data constructed from Indian Census. Columns (1) and (2) control for decadal precipitation. Columns (3) and (4) control for decadal precipitation interacted with the high road density dummy and include high road density-by-year fixed effects. Columns (5) and (6) control for decadal precipitation interacted with the high bank credit dummy and include high bank credit-by-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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B8: Effect of Rising Temperatures on Crop Area Shares

	Dry Season Crop Area		Labor-Intensive Crop Area	
	(1)	(2)	(3)	(4)
T	-0.0038	-0.0040	-0.0094	-0.0079
	(0.0027)	(0.0017)**	(0.0030)***	(0.0019)***
	[0.0035]	[0.0035]	[0.0041]**	[0.0041]***
P	0.0078	0.0074	0.0104	0.0061
	(0.0021)***	(0.0019)***	(0.0027)***	(0.0020)***
	[0.0028]***	[0.0028]**	[0.0041]**	[0.0042]**
Region-year trends	Y	N	Y	N
Region-decade FE	N	Y	N	Y
Observations	11,705	11,705	11,705	11,705

Note: The dependent variable is the share of crop area planted with dry season crops in Columns (1) and (2), and the share of crop area planted with labor-intensive crops in Columns (3) and (4). Temperature and precipitation are decadal averages of the past ten growing seasons. We use annual data from VDSA spanning 1966 to 2010. The dry season (*rabi*) crops are wheat, pearl millet, barley, chickpea, pigeon pea, rapeseed and mustard seed, linseed, and sunflower. The labor-intensive crops are defined to be those that require 700 or more average person-hours per hectare, which are rice, groundnut, cotton, and sugarcane (FICCI, 2015). 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B9: Robustness to Alternate Samples and Variable Definitions

	Ag Labor Share			Non-Ag Worker Share			Urbanization			Migrant Share		
Panel A: Unbalanced Panel	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
10-Year Avg. GS Temperature (Celsius)	0.206**	0.337***	0.160*	-0.069*	-0.140**	-0.135**	0.017	-0.185**	-0.006	0.006	-0.152**	0.012
	(0.094)	(0.097)	(0.095)	(0.037)	(0.056)	(0.057)	(0.045)	(0.082)	(0.068)	(0.061)	(0.077)	(0.066)
T x High Road Density		-0.375***			0.125*			0.229**			0.223**	
		(0.114)			(0.066)			(0.095)			(0.107)	
T x High Bank Credit			0.046			0.095			-0.006			-0.034
			(0.164)			(0.062)			(0.082)			(0.095)
P-val: T + T x High Road D.		0.601			0.706			0.458			0.380	
P-val: T + T x High Bank C.			0.146			0.306			0.820			0.790
Observations	1691	1498	1691	1705	1498	1705	1689	1494	1689	1411	1242	1411
Panel B: Grid Point Average T & P												
10-Year Avg. GS Temperature (Celsius), geoinpoly	0.138**	0.310***	0.271***	-0.0723**	-0.118**	-0.151***	0.0229	-0.185**	-0.00590	0.00941	-0.133*	0.0286
	(0.0601)	(0.0903)	(0.0671)	(0.0312)	(0.0572)	(0.0488)	(0.0447)	(0.0831)	(0.0687)	(0.0640)	(0.0743)	(0.0672)
T x High Road Density		-0.314***			0.119*			0.248**			0.238**	
		(0.108)			(0.0682)			(0.0992)			(0.108)	
T x High Bank Credit			-0.258**			0.123**			0.00858			-0.0505
			(0.106)			(0.0578)			(0.0863)			(0.101)
P-val: T + T x High Road D.		0.954			0.975			0.319			0.225	
P-val: T + T x High Bank C.			0.882			0.446			0.962			0.810
Observations	1524	1434	1524	1578	1434	1578	1560	1428	1560	1325	1195	1325
Panel C: Log T & P												
Log of 10-Year Avg. GS Rainfall (100 mm)	2.535***	3.320***	2.587***	-1.803***	-1.902***	-2.285***	-0.732	-2.187***	-1.122*	-0.915	-1.517***	-0.615
	(0.642)	(1.174)	(0.707)	(0.338)	(0.358)	(0.470)	(0.755)	(0.499)	(0.590)	(0.572)	(0.443)	(0.444)
ln T x High Road Density		-4.252**			1.731*			3.428**			3.450*	
		(2.008)			(1.035)			(1.542)			(1.942)	
ln T x High Bank Credit			-1.709			1.398			0.899			-0.948
			(2.249)			(0.997)			(1.436)			(2.062)
P-val: ln T + ln T x High Road D.		0.581			0.859			0.403			0.313	
P-val: ln T + ln T x High Bank C.			0.686			0.325			0.865			0.453
Observations	1542	1458	1542	1614	1458	1614	1596	1452	1596	1345	1210	1345

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Temperature is the decadal average of the past ten growing seasons. *High Road Density* is a binary variable that takes the value 1 if the district has above median road density at baseline. Data are district-level panel data constructed from the Indian Census. The samples in Panels B and C are restricted to districts for which the dependent variable is non-missing in all years. All columns control for decadal precipitation and include district and region-by-year fixed effects. Columns (2), (5), (8), and (11) also control for decadal precipitation interacted with the high road density dummy and include high road density-by-year fixed effects. *High Bank Credit* is a binary variable that takes the value 1 if the district has above median bank credit per capita at baseline. Columns (3), (6), (9), and (12) control for decadal precipitation interacted with the high bank credit dummy and include high bank credit-by-year fixed effects. We present standard errors clustered by district in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B10: Effect of Rising Temperatures: Distributed Lagged Averages

	Ag Labor Share		Non-Ag Worker Share		Urbanization		Migrant Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Average (T_0 to T_{-2})	0.023	0.068	-0.048	-0.067	0.064	0.061	0.029	-0.036
	(0.043)	(0.044)	(0.024)**	(0.026)***	(0.028)**	(0.032)*	(0.054)	(0.063)
	[0.076]	[0.074]	[0.034]	[0.033]**	[0.033]**	[0.036]*	[0.058]	[0.059]
Average (T_{-3} to T_{-6})	0.088	0.094	0.009	-0.024	-0.058	-0.055	-0.035	0.034
	(0.031)***	(0.037)**	(0.015)	(0.018)	(0.022)***	(0.029)*	(0.042)	(0.052)
	[0.053]*	[0.056]*	[0.022]	[0.022]	[0.028]**	[0.033]*	[0.052]	[0.057]
Average (T_{-7} to T_{-9})	0.076	-0.020	-0.054	0.006	-0.019	0.005	0.000	-0.022
	(0.035)**	(0.044)	(0.023)**	(0.025)	(0.023)	(0.028)	(0.039)	(0.047)
	[0.060]	[0.067]	[0.028]*	[0.030]	[0.027]	[0.031]	[0.051]	[0.057]
Region-year trends	Y	N	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y	N	Y
Observations	1,548	1,548	1,620	1,620	1,596	1,596	1,350	1,350

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which the dependent variable is non-missing in all years. All columns include district and year fixed effects. All regressions control for the corresponding distributed lagged averages for precipitation. 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 B11: Robustness to Controlling for Time-Varying Covariates

	Ag Labor Share		Non-Ag Worker Share		Urbanization		Migrant Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.159	0.129	-0.054	-0.068	-0.007	0.013	0.014	0.015
	(0.056)***	(0.058)**	(0.032)*	(0.031)**	(0.044)	(0.047)	(0.056)	(0.062)
	[0.077]**	[0.074]*	[0.040]	[0.034]**	[0.049]	[0.049]	[0.061]	[0.068]
P	-0.125	-0.053	-0.026	-0.000	0.044	0.020	-0.038	-0.004
	(0.054)**	(0.053)	(0.030)	(0.031)	(0.042)	(0.042)	(0.045)	(0.050)
	[0.096]	[0.087]	[0.038]	[0.033]	[0.048]	[0.046]	[0.061]	[0.064]
High-yielding-variety area	0.155	0.147	0.087	0.106	0.074	0.071	-0.214	-0.211
	(0.066)**	(0.067)**	(0.048)*	(0.049)**	(0.059)	(0.062)	(0.058)***	(0.062)***
	[0.077]**	[0.077]*	[0.056]	[0.055]*	[0.054]	[0.055]	[0.061]***	[0.064]***
Labor regulation strictness index	-0.077	-0.071	0.002	0.008	0.003	-0.001	0.044	0.045
	(0.021)***	(0.023)***	(0.012)	(0.013)	(0.013)	(0.015)	(0.016)***	(0.018)**
	[0.030]**	[0.033]**	[0.015]	[0.014]	[0.015]	[0.016]	[0.021]**	[0.022]**
Road density	-0.126	-0.143	0.022	0.021	0.030	0.044	0.027	-0.005
	(0.042)***	(0.031)***	(0.012)*	(0.014)	(0.042)	(0.046)	(0.109)	(0.122)
	[0.053]**	[0.050]***	[0.018]	[0.018]	[0.039]	[0.041]	[0.091]	[0.099]
Number of markets	0.003	0.002	-0.002	-0.002	0.000	0.000	-0.003	-0.002
	(0.006)	(0.005)	(0.001)*	(0.001)	(0.004)	(0.003)	(0.008)	(0.008)
	[0.004]	[0.004]	[0.001]*	[0.001]	[0.003]	[0.002]	[0.006]	[0.006]
Number of banks	-0.049	-0.053	-0.001	-0.002	-0.008	-0.005	-0.007	-0.006
	(0.012)***	(0.013)***	(0.006)	(0.006)	(0.008)	(0.008)	(0.013)	(0.013)
	[0.011]***	[0.011]***	[0.005]	[0.005]	[0.008]	[0.008]	[0.012]	[0.012]
Region-year trends	Y	N	Y	N	Y	N	Y	N
Region-year FE	N	Y	N	Y	N	Y	N	Y
Observations	1,546	1,546	1,546	1,546	1,529	1,529	1,283	1,283

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Temperature and precipitation are decadal averages of the past ten growing seasons. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which the dependent variable is non-missing in all years. All columns include district and year fixed effects. All time-varying covariates are decadal averages. District-level high-yielding variety area, road density, and number of markets are from the VDSA data set. The state-level labor market strictness index, from Besley and Burgess (2004), ranges from 3 to -3; positive values denote states that are more rigid (pro-worker); negative values denote states that are more flexible (pro-employer). The number of banks per district is measured in 100's and is from Fulford (2013). See Appendix C for more details. 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 B12: Effect of Rising Temperatures using Long-Difference Specification with Alternative End Points

	Ag Labor Share	Non-Ag Worker Share	Urbanization	Migrant Share
	(1)	(2)	(3)	(4)
T	0.2768	-0.1376	0.0177	-0.3828
	(0.1011)***	(0.0551)**	(0.0876)	(0.1593)**
	[0.2243]	[0.0587]**	[0.1298]	[0.1878]**
P	-0.9920	0.1561	-0.3520	0.6279
	(0.3688)***	(0.1987)	(0.2972)	(0.5536)
	[0.4429]**	[0.1828]	[0.2864]	[0.8081]
Region FE	Y	Y	Y	Y
Observations	258	270	266	264

Note: The dependent variable is the share of agricultural laborers in Column (1), the share of non-agricultural workers in Column (2), urbanization rates in Column (3), and the share of intra-district migrants in Column (4). The dependent variable in each column is the difference (in natural logarithm) of an outcome between two 20-year periods, 1961-1971 and 2001-2011. The outcome in 1961-1971 are calculated as the average of 1961 and 1971 decadal observations, and that in 2001-2011 are calculated as the average of 2001 and 2011. The independent variables are differences in average growing-season temperature and precipitation over the same periods. Data are district-level data constructed from the Indian Census. All columns include region fixed effects. We present standard errors 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 B13: Heterogeneous Effect of Rising Temperatures with Alternate Thresholds

	Ag Labor Share		Non-Ag Worker Share		Urbanization		Migrant Share	
Panel A: Road Network Density	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.329***	0.291***	-0.154**	-0.156***	-0.210**	-0.159**	-0.233***	-0.063
	(0.112)	(0.099)	(0.067)	(0.052)	(0.092)	(0.080)	(0.087)	(0.095)
T x (Road Density > 40th pct)	-0.301**		0.113		0.229**		0.305***	
	(0.121)		(0.075)		(0.105)		(0.107)	
T x (Road Density > 60th pct)		-0.290**		0.155**		0.191*		0.112
		(0.118)		(0.066)		(0.097)		(0.119)
P-val: T + T x High Road D.	0.677	0.996	0.266	0.994	0.739	0.583	0.324	0.458
Observations	1458	1458	1458	1458	1452	1452	1210	1210
Panel B: Bank Credit Per Capita	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T	0.298***	0.212***	-0.181***	-0.156***	-0.037	-0.079	-0.045	-0.008
	(0.077)	(0.064)	(0.051)	(0.041)	(0.077)	(0.062)	(0.074)	(0.072)
T x (Bank Credit > 40th pct)	-0.224**		0.143**		0.036		0.026	
	(0.106)		(0.059)		(0.091)		(0.098)	
T x (Bank Credit > 60th pct)		-0.149		0.126**		0.118		-0.026
		(0.119)		(0.056)		(0.084)		(0.112)
P-val: T + T x High Bank C.	0.373	0.552	0.308	0.484	0.988	0.533	0.812	0.728
Observations	1548	1548	1620	1620	1596	1596	1350	1350

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Columns (1) and (2), of the share of non-agricultural workers in Columns (3) and (4), of urbanization rates in Columns (5) and (6), and of the share of intra-district migrants in Columns (7) and (8). Temperature is the decadal average of the past ten growing seasons. In Panel A, we use alternate heterogeneity dummies which take the value of 1 if the district's baseline road density is above the 40th or 60th percentile, depending on the column. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which road density/bank credit data is non-missing and the dependent variable is non-missing in all years. All columns include district fixed effects, region-by-year fixed effects, and high road density-by-year fixed effects. We control for decadal precipitation and decadal precipitation interacted with the road density dummy. In Panel B, we use alternate heterogeneity dummies, which take the value of 1 if the district's baseline bank credit per capita is above the 40th or 60th percentile, depending on the column. All columns include district fixed effects, region-by-year fixed effects, and high bank credit-by-year fixed effects. We control for decadal precipitation and decadal precipitation interacted with the high bank credit dummy. We present standard errors clustered by district in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B14: Heterogeneity by Road Density: Interacting Other Baseline Characteristics with Temperature

Panel A: Ag Labor Share	(1)	(2)	(3)	(4)	(5)
T	0.355	0.557	0.420	0.389	0.352
	(0.099)***	(0.103)***	(0.116)***	(0.134)***	(0.153)**
	[0.118]***	[0.133]***	[0.141]***	[0.166]**	[0.179]**
T x High Road Density	-0.362	-0.387	-0.345	-0.353	-0.391
	(0.114)***	(0.116)***	(0.115)***	(0.113)***	(0.123)***
	[0.128]***	[0.128]***	[0.127]***	[0.124]***	[0.137]***
Controls	None	I	II	III	IV
Observations	1,458	1,458	1,428	1,428	1,344
Panel B: Non-Ag Worker Share					
T	-0.137	-0.115	-0.069	-0.007	-0.048
	(0.058)**	(0.063)*	(0.060)	(0.062)	(0.070)
	[0.052]***	[0.063]*	[0.063]	[0.066]	[0.066]
T x High Road Density	0.119	0.117	0.083	0.100	0.127
	(0.068)*	(0.067)*	(0.062)	(0.061)	(0.064)**
	[0.062]*	[0.060]*	[0.059]	[0.057]*	[0.058]**
Controls	None	I	II	III	IV
Observations	1,458	1,458	1,428	1,428	1,344
Panel C: Urbanization					
T	-0.190	-0.121	-0.066	-0.037	-0.225
	(0.083)**	(0.090)	(0.097)	(0.114)	(0.125)*
	[0.068]***	[0.075]	[0.084]	[0.100]	[0.107]**
T x High Road Density	0.231	0.224	0.185	0.191	0.208
	(0.098)**	(0.098)**	(0.097)*	(0.095)**	(0.089)**
	[0.084]***	[0.084]***	[0.086]**	[0.085]**	[0.074]***
Controls	None	I	II	III	IV
Observations	1,452	1,452	1,422	1,422	1,344
Panel D: Migrant Share					
T	-0.149	-0.135	-0.195	-0.243	-0.348
	(0.078)*	(0.086)	(0.110)*	(0.127)*	(0.161)**
	[0.075]**	[0.097]	[0.122]	[0.130]*	[0.159]**
T x High Road Density	0.227	0.222	0.221	0.208	0.168
	(0.108)**	(0.110)**	(0.116)*	(0.114)*	(0.120)
	[0.100]**	[0.102]**	[0.109]**	[0.109]*	[0.111]
Controls	None	I	II	III	IV
Observations	1,210	1,210	1,190	1,190	1,120

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Panel A, of the share of non-agricultural workers in Panel B, of urbanization rates in Panel C, and of the share of intra-district migrants in Panel D. Temperature is the decadal average of the past ten growing seasons. *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 and high road density-by-year fixed effects. We control for decadal precipitation and decadal precipitation interacted with the road density dummy. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which road density data is non-missing and the dependent variable is non-missing in all years. In Columns (2) through (5) we cumulatively add other dummy controls interacted with decadal temperature and with decadal precipitation, to test the stability of our road density heterogeneity coefficient. These controls are I: above median bank credit per capita at baseline; II: I and above median male agricultural wages at baseline; III: II and above median long-run temperature; IV: III and above median proportion irrigated land at baseline. 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 B15: Heterogeneity by Bank Credit: Interacting Other Baseline Characteristics with Temperature

Panel A: Ag Labor Share	(1)	(2)	(3)	(4)	(5)
T	0.271	0.562	0.402	0.354	0.285
	(0.070)***	(0.097)***	(0.111)***	(0.134)***	(0.155)*
	[0.082]***	[0.118]***	[0.125]***	[0.149]**	[0.169]*
T x High Bank Credit	-0.217	-0.219	-0.206	-0.208	-0.243
	(0.110)**	(0.106)**	(0.098)**	(0.097)**	(0.114)**
	[0.111]*	[0.112]*	[0.109]*	[0.109]*	[0.123]**
Controls	None	I	II	III	IV
Observations	1,548	1,458	1,428	1,428	1,344
Panel B: Non-Ag Worker Share					
T	-0.157	-0.194	-0.144	-0.079	-0.146
	(0.048)***	(0.066)***	(0.064)**	(0.068)	(0.078)*
	[0.062]**	[0.070]***	[0.061]**	[0.063]	[0.070]**
T x High Bank Credit	0.106	0.104	0.058	0.061	0.084
	(0.059)*	(0.066)	(0.064)	(0.062)	(0.067)
	[0.066]	[0.075]	[0.068]	[0.067]	[0.074]
Controls	None	I	II	III	IV
Observations	1,620	1,458	1,428	1,428	1,344
Panel C: Urbanization					
T	-0.042	-0.190	-0.146	-0.127	-0.356
	(0.069)	(0.080)**	(0.096)	(0.115)	(0.127)***
	[0.068]	[0.076]**	[0.083]*	[0.098]	[0.108]***
T x High Bank Credit	0.030	0.099	0.070	0.079	0.092
	(0.087)	(0.090)	(0.090)	(0.091)	(0.079)
	[0.078]	[0.095]	[0.088]	[0.088]	[0.079]
Controls	None	I	II	III	IV
Observations	1,596	1,452	1,422	1,422	1,344
Panel D: Migrant Share					
T	-0.013	-0.035	-0.105	-0.160	-0.250
	(0.069)	(0.095)	(0.120)	(0.143)	(0.176)
	[0.084]	[0.117]	[0.134]	[0.145]	[0.166]
T x High Bank Credit	-0.029	-0.094	-0.134	-0.134	-0.192
	(0.099)	(0.109)	(0.105)	(0.105)	(0.100)*
	[0.114]	[0.124]	[0.124]	[0.124]	[0.120]
Controls	None	I	II	III	IV
Observations	1,350	1,210	1,190	1,190	1,120

Note: The dependent variable is the natural logarithm of the share of agricultural laborers in Panel A, of the share of non-agricultural workers in Panel B, of urbanization rates in Panel C, and of the share of intra-district migrants in Panel D. Temperature is the decadal average of the past ten growing seasons. *High Bank Credit* 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 and high bank credit-by-year fixed effects. We control for decadal precipitation and decadal precipitation interacted with the bank credit dummy. Data are district-level panel data constructed from the Indian Census. We restrict our sample to districts for which bank credit data is non-missing and the dependent variable is non-missing in all years. In Columns (2) through (5) we cumulatively add other dummy controls interacted with decadal temperature and with decadal precipitation, to test the stability of our bank credit heterogeneity coefficient. These controls are I: above median road density at baseline; II: I and above median male agricultural wages at baseline; III: II and above median long-run temperature; IV: III and above median proportion irrigated land at baseline. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Data Appendix

Section 3 summarizes the data used in the analysis. In this data appendix, we provide additional information on the various data sources as well as the construction of key variables.

C.1 Census Data

The Primary Census Abstract (PCA) and the Migration (D-series) data tables in the Indian Population Census are the sources of decadal district-level data on demographic and economic indicators (population, worker counts by categories, migrant counts) from 1961 to 2011. For the years 1961-1991, we use data from Vanneman and Barnes (2000), and for the years 2001 and 2011, we use data from the Census website.[32]

The Census classifies workers into four categories: cultivators, agricultural laborers, workers in household industry and other workers. A cultivator is defined as a worker who is "engaged in cultivation of land owned or held from Government or held from private persons or institutions for payments in money, kind or share." An agricultural laborer is defined as a worker who "works on another person's land for wages in money or kind or share. She or he has no risk in the cultivation, but merely works in another person's land for wages." Household industry is defined 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." Other worker is defined as a "worker other than cultivator, agricultural laborer or worker in household industry." Examples of other workers include work in the public sector, manufacturing, construction, trade, business etc.[33]

We construct agricultural labor share as the total count of agricultural laborers divided by total workers. We construct non-agricultural worker share

[32] `https://censusindia.gov.in/pca/`

[33] Details on these categories can be found in the codebook of Vanneman and Barnes (2000) available at `http://vanneman.umd.edu/districts/codebook/laborforce.html`, as well as in the 2011 Census meta data available at `https://www.censusindia.gov.in/2011census/HLO/Metadata_Census_2011.pdf`.

as the sum of workers across two categories – household industry workers and other workers – divided by total workers. Note that the agricultural labor share is not a perfect complement to the non-agricultural worker share as we do not include cultivators when constructing the agricultural labor share. We construct cultivator share as the total count of main cultivators divided by total workers. Note that our measure of cultivator share excludes marginal cultivators (those who "worked for less than six months in the reference period") because we do not have this data for 1961-1991.

The Census further splits agricultural laborers into two categories: main and marginal agricultural laborers. A main agricultural laborer is defined as a worker who "worked for more than six months in the reference period." A marginal agricultural laborer is defined as a worker who "worked for less than six months in the reference period." We construct main (marginal) agricultural labor share as the main (marginal) count of agricultural laborers divided by total workers.

Turning to urbanization and migration, we construct urbanization share as the total urban population divided by total population. The Census definition of urban areas has stayed largely consistent since the 1961 census. Urban areas constitutes (a) Statutory Towns: all places with a municipality, corporation, cantonment board or notified town area committee, etc., (b) Census Towns: all places which satisfied the following criteria: i) A minimum population of 5,000; ii) At least 75 per cent of the male main working population engaged in non-agricultural pursuits; and iii) A density of population of at least 400 persons per sq. km., and (c) Adjoining Outgrowths: a viable unit such as a village or a hamlet (part of a village) contiguous to a town and posses urban features in terms of infrastructure and amenities.

For migration, the Census classifies an individual as an intra-district migrant "if the place in which he is enumerated during the census is other than his place of immediate last residence," and if the last residence is within the same district of his/her current residence. We construct migrant share as the total count of intra-district rural-to-urban male migrants divided by total male population. We consider male migration only as a majority of female migration in India is for marriage, which is outside the scope of our study. We do not have this measure for 1971 due to missing data on migrant counts in Vanneman and Barnes (2000).

C.2 National Sample Survey Data

The Consumer Expenditure (Schedule 1) and Employment and Unemployment (Schedule 10) modules of the National Sample Survey, collected by the National Sample Survey Office (NSSO), are the sources of nationally representative data on demographic and social characteristics, consumption patterns as well as labor market behavior at the individual- and household-level in India.

We use eight rounds of the Employment and Unemployment schedule, spanning the years 1987 to 2012.[34] The time period covered in each round corresponds to the agricultural year from July to the following June. More specifically, the data covers the following time periods (with round number reported in parentheses): July 1987- June 1988 (43rd), July 1993 - June 1994 (50th), July 1999 - June 2000 (55th), July 2004 - June 2005 (61st), July 2005 - June 2006 (62nd), July 2007 - June 2008 (64th), July 2009 - June 2010 (66th), July 2011 - June 2012 (68th).

We restrict the sample to include individuals aged 14 – 65 who participate in the labor force. We use a series of questions regarding individual-level employment activities — this includes employment status and industry of the main activity — during a seven-day reference period. We use industry information from the reference week to classify individuals as working in agricultural and non-agricultural sectors.[35] We aggregate individual-level data to construct the following employment shares at the district level: the share of the labor force who are engaged in agriculture, the share of the labor force who are engaged in non-agriculture, and the share of the labor force who are engaged in manufacturing, services, and construction. We also use informa-

[34] They can be downloaded from http://microdata.gov.in/nada43/index.php/catalog/EUE.

[35] The agricultural sector includes sub-sectors such as crop and animal production, hunting and related service activities, forestry and logging, and fishing and aquaculture. The non-agricultural sectors include mining and quarrying, manufacturing, construction, and services.

tion on each individual's principal and subsidiary employment during the year to construct the following employment shares at the district level: shares of the labor force engaged in agriculture as a primary and secondary occupation, and shares of the labor force engaged in non-agriculture as a primary and secondary occupation.

There is one caveat with the NSS sample described above. The 50th round in 1993-1994 has incomplete coverage of the urban population — most of the districts have their rural households represented, but only a quarter of the districts have their urban households represented in this survey round. We present results using the complete set of NSS rounds in the paper; these results are robust to the exclusion of the 50th round.[36]

In addition, we use six rounds of the Consumption schedule, spanning the years 1993 to 2012.[37] More specifically, the data covers the following time periods (with round number reported in parentheses): July 1993 - June 1994 (50th), July 1999 - June 2000 (55th), July 2004 - June 2005 (61st), July 2007 - June 2008 (64th), July 2009 - June 2010 (66th), July 2011 - June 2012 (68th). We aggregate household-level data to construct the following annual consumption per capita measures at the district level: total consumption, food consumption, and non-food consumption. The consumption variables are adjusted to 2005 base prices using separate purchasing-power-parity conversion rates for urban and rural areas at the state level — this teases out temporal and spatial price level differences such that our consumption variables are comparable across all states, across urban and rural areas, and over all time periods.

C.3 Weather Data

The Terrestrial Precipitation: Monthly Time Series (1900–2014), version 4.01, and the companion Terrestrial Air Temperature data set (Matsuura and Willmott, 2015a,b) is the source of gridded monthly-level data on temperature and precipitation from 1951-2011.

[36] These results are available upon request.

[37] They can be downloaded from http://microdata.gov.in/nada43/index.php/catalog/CEXP.

We first construct district-level weather data 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. This inverse distance weighting method is also used in Burgess et al. (2017) and Taraz (2018). We calculate average temperature and precipitation during the main agricultural growing season (June through February) as these have the greatest impacts on agriculture. Next, we aggregate the growing season weather variables to ten-year averages.

C.4 Infrastructure and Yields Data

The Village Dynamics in South Asia (VDSA) Meso dataset, compiled by researchers at the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT, 2015), consists of a large set of demographic, socioeconomic and agro-ecological variables at the district level. It covers 19 major agricultural states in India at an annual frequency from 1966 to 2010.[38]

The VDSA Meso dataset is the source of district-level data on total length of roads in kilometers. The underlying sources of the VDSA roads data are the annual State Statistical Abstracts. We construct a baseline road infrastructure density measure as the total length of roads in kilometers in each district in 1970 – the earliest year for which this data is available – divided by the total surface area, computed in ArcGIS using the consistent district boundaries illustrated in Appendix Figure B1a. Data on length of roads is missing for 15 districts in 1970. Furthermore, we are unable to construct a road density measure for 22 additional districts since coverage in the VDSA Meso dataset is limited to nineteen states in India.

The VDSA Meso dataset is also the source of annual district-level data on crop yields. The underlying sources of the VDSA data on yields are state-level agricultural agencies such as the Directorate of Agriculture and the Directorate

[38]The states covered in the data base are Andhra Pradesh, Assam, Bihar, Chhattisgarh, Gujarat, Haryana, Himachal Pradesh, Jharkhand, Karnataka, Kerala, Madhya Pradesh, Maharashtra, Orissa, Punjab, Rajasthan, Tamil Nadu, Uttar Pradesh, Uttarakhand, and West Bengal.

of Agriculture and Food Production. We construct an annual yield measure that aggregates yields across all the crops in VDSA that have non-missing price data, using 1966-1970 crop prices as weights. The crops included are rice, wheat, sugarcane, cotton, groundnut, sorghum, maize, pearl millet, finger millet, barley, chickpeas, pigeon pea, sesame, rapeseed and mustard, castor, and linseed.

In one of our robustness tests, we use data on high yielding variety areas from VDSA. Specifically, we control for the fraction of area that is grown with high yield varieties, as a fraction of the total cultivated area in that district in that year.

C.5 Bank Credit Data

The Basic Statistical Returns (BSR) reports, collected by the Reserve Bank of India, is the source of district-level bank credit.[39] The Basic Statistical Returns System was launched in 1971 with the goal of creating a database of scheduled commercial banks. We use the 1972 BSR reports (the earliest available), and digitize Table 2.2, which contains district-wise statistics on the number of functioning offices of scheduled commercial banks, aggregate deposits, and total credit (advances) from all offices as of the last Friday in December 1972. The coverage of data in the 1972 BSR report is 98.7% of aggregate deposits and 98.6% of total credit. We construct a baseline bank credit per capita measure as the total bank credit in a district divided by its total population in 1971.

C.6 Other Data

We draw on two other data sources in our robustness tests. First, we use state-level data on labor regulation strictness from Besley and Burgess (2004). The index ranges from 3 to -3; positive values denote states that are more rigid (pro-worker); negative values denote states that are more flexible (pro-employer). It

[39]These reports are available at `https://dbie.rbi.org.in/DBIE/dbie.rbi?site=publications`.

is based on a tabulation of state-level amendments to the Industrial Disputes Act of 1947, which regulates trade unions, arbitration, and procedures to be followed in the case of an industrial disputes. Because different states passed different amendments to this Act at different points in time, the index from Besley and Burgess (2004) provides a measure of labor regulation with both spatial and temporal variation.

Second, we use data on the number of banks in each district in each year, as a proxy for financial development of the district. The data on the number of banks is based on bank opening data from the Reserve Bank of India, as compiled by Fulford (2013).