Weather Shocks and Structural Transformation: Micro Evidence from Africa

Chunxiao Jing *

October 14, 2026

Abstract

Structural transformation is a central driver of economic development, yet the movement of labor from agriculture into more productive non-agricultural employment has remained slow in Sub-Saharan Africa. In this paper, I examine how household constraints shape this process using panel data from Niger, Malawi, and Ethiopia. I construct a weather-predicted NDVI growth index using a Random Forest and show that favorable weather-induced productivity shocks reduce agricultural labor in all three countries. These reductions are accompanied by manufacturing employment gains in Niger and Malawi but a substantial increase in idle labor—working-age individuals without agricultural or non-agricultural employment—in Ethiopia. Subsistence constraints weaken labor release in Niger and Malawi, while distance, land rights, and regional employment conditions shape adjustment in Ethiopia. I then develop a structural model linking Lewis-type labor release, labor transition under mobility and land-rights frictions, and Harris–Todaro employment absorption. Estimating the model using simulated method of moments indicates that subsistence constraints primarily limit transformation in Niger and Malawi, whereas transition frictions and limited employment absorption are more important in Ethiopia. Finally, counterfactual experiments evaluate individual and combined policies. Interventions that relax subsistence constraints generate the largest non-agricultural employment gains in Niger and Malawi, while land-rights reform and stronger employment absorption are more effective in Ethiopia. These findings highlight the importance of country-specific policies that translate agricultural labor release into productive employment.

1 Introduction

Structural Transformation, the movement of labor from low-productivity agriculture to more productive non-agricultural activities, is a central driver of economic growth in developing countries (Lewis, 1954; Syrquin,

*The Department of Agricultural and Applied Economics, University of Wisconsin Madison. Email: cjing7@wisc.edu. I thank my advisor, Jeremy Foltz, and Gisella Kagy and Vladimir Gimpelson for their guidance and helpful comments. I thank Samson Kastro Dake, Bereket Yaekob, and Eyasu Wansamo Wakayo for their valuable assistance with the fieldwork in Ethiopia. I thank Justin Sandefur for helpful comments on the policy analysis. I also thank participants at the UW–Madison Agricultural and Applied Economics Development Workshop and the Agricultural and Applied Economics Association (AAEA) Annual Meeting for valuable feedback. This research received funding through the Structural Transformation and Economic Growth (STEG) programme, supported by the UK's Foreign, Commonwealth & Development Office (FCDO), under contract STEG_LOA_203003. The views expressed are not necessarily those of FCDO. All errors are my own.

1988). Yet, the pace of structural transformation has been considerably slower in Sub-Saharan Africa than in other developing regions (Morsy et al., 2023; Mensah et al., 2018; Fox et al., 2017). Agriculture continues to employ the majority of workers in many African countries despite its relatively low productivity (Christiaensen and Maertens, 2022; Herrendorf et al., 2022; Barrett et al., 2017; McMillan and Headey, 2014). This persistent concentration of labor in agriculture limits productivity growth and slows economic development. Therefore, understanding why labor does not move successfully into more productive non-agricultural activities is a central question for development economics.

Agricultural productivity is viewed as a key driver of structural transformation in developing countries (Lewis, 1954; Gollin et al., 2002; Ngai and Pissarides, 2007; Gollin et al., 2007). However, the response of labor allocation to productivity improvements remains theoretically ambiguous. Higher agricultural productivity may relax subsistence constraints and release labor from agriculture (Benjamin, 1992; Alvarez-Cuadrado and Poschke, 2011; Bustos et al., 2016; Gollin, 2023), but it may also increase the returns to farming and encourage households to retain labor in agricultural production (Bustos et al., 2016;Karayalcin and Pintea, 2022). Even when productivity improvements release labor from agriculture, workers may still face barriers to entering productive non-agricultural employment. As a result, productivity improvements may generate different patterns of structural transformation, depending on the constraints households face.

Existing theories emphasize different stages of this process. The Lewis framework highlights the release of surplus labor from agriculture, with subsistence constraints potentially limiting households' ability to reduce agricultural labor (Lewis, 1954; Gollin et al., 2007). Studies of labor reallocation emphasize mobility costs (Bryan et al., 2014; Donaldson, 2018), land institutions (De Janvry et al., 2015; Adamopoulos et al., 2024), and other frictions that shape access to non-agricultural activities. The Harris–Todaro framework emphasizes expected earnings, which depend on both wages and the probability of obtaining employment (Harris and Todaro, 1970). Workers may therefore leave agriculture without securing non-agricultural jobs. Together, these perspectives motivate a three-stage framework: labor release, transition into non-agricultural activities, and employment absorption.

Despite these theoretical insights, we have limited empirical understanding of which stage constrains structural transformation and why the relevant constraints differ across countries. Although structural transformation is an aggregate process, aggregate sectoral employment shares can conceal differences in how households reallocate labor and which constraints they face. I therefore examine this macroeconomic question at the household level, distinguishing labor release, labor transition, and labor absorption. This paper asks two related questions. First, do agricultural productivity improvements release labor by relaxing subsistence constraints, or retain labor by increasing the returns to farming? Second, what determines whether labor released from agriculture is absorbed into productive non-agricultural employment? Answering these questions helps explain why similar agricultural productivity improvements can lead to different employment outcomes and which constraints policies must address to promote structural transformation.

To answer these questions, I combine household panel data from three Sub-Saharan African countries (Niger, Malawi, and Ethiopia). To address the potential endogeneity of agricultural productivity in explaining households' labor reallocation decisions, I construct a weather-predicted NDVI growth index using rainfall and temperature, capturing productivity variation driven by weather rather than household production decisions. Positive weather-induced productivity shocks reduce agricultural labor in all three countries, consistent with the labor-release effect dominating the incentive to retain labor in farming. However, the accompanying employment responses differ markedly. In Niger and Malawi, agricultural labor reductions are accompanied by increases in manufacturing employment, whereas in Ethiopia they are accompanied by a substantial increase in idle labor–working-age household members who are not engaged in agricultural or non-agricultural employment. These findings distinguish the release of labor from its absorption into productive non-agricultural employment.

To examine why similar productivity shocks generate different employment responses, I test how these responses vary with baseline constraints at each stage of structural transformation. I first examine subsistence status to assess the Lewis-type labor-release mechanism: whether food requirements limit households' ability to reduce agricultural labor. Subsistence constraints weaken labor release in Niger and Malawi and attenuate the manufacturing response in Malawi. I then examine distance and land rights to assess the transition-friction mechanism: whether barriers to accessing non-agricultural opportunities shape labor reallocation. Land rights are associated with a larger manufacturing response in Malawi and stronger labor release in Ethiopia, while in Ethiopia, households farther from roads, markets, or population centers experience larger reductions in agricultural labor and larger increases in idle labor following favorable weather-induced productivity shocks. Finally, I examine both community absorption and regional labor-market conditions to assess the Harris–Todaro absorption mechanism: whether employment opportunities are sufficient to absorb labor released from agriculture. Higher baseline community absorption is associated with a larger manufacturing response in Niger, while complementary Ethiopian evidence suggests that stronger regional hiring attenuates the idle-labor response. Together, these findings motivate a framework that distinguishes three stages of structural transformation: labor release, labor transition, and labor absorption.

To quantify the relative importance of these constraints and evaluate policies that relax them, I develop a structural model incorporating labor release, labor transition, and labor absorption. The model allows subsistence constraints, transition frictions, and employment opportunities to jointly determine household labor allocation, with the model using simulated method of moments. The estimates suggest that subsistence constraints primarily limit labor release in Niger and Malawi, whereas transition barriers and limited employment absorption are more important in Ethiopia. Counterfactual policy experiments show that the effectiveness of interventions depends on which constraint they relax, highlighting the importance of policies that address the stage limiting structural transformation in each country.

This paper makes three main contributions to the literature on structural transformation and development pol-

icy. First, I distinguish labor release, labor transition, and labor absorption within a common household-level analysis. Studies of agricultural productivity examine whether productivity improvements shift labor out of agriculture (Bustos et al., 2016; Krishna and Meenakshi, 2022; Wang, 2025; Vicente Cateia et al., 2024). Related studies emphasize mobility costs and land institutions that shape access to non-agricultural activities (Bryan et al., 2014; Donaldson, 2018; Adamopoulos and Restuccia, 2020; Deininger and Jin, 2008; Deininger and Jin, 2006), while another literature examines whether non-agricultural sectors generate sufficient productive employment (Harris and Todaro, 1970; McMillan and Headey, 2014; Diao et al., 2019; Diao et al., 2025; Rodrik, 2016). I extend this literature by bringing these different aspects of structural transformation into a common household-level analysis. I examine labor release, labor transition, and labor absorption within a common framework, showing how constraints at each stage shape responses to weather-induced productivity shocks across Niger, Malawi, and Ethiopia. I also explicitly include idle labor as a possible outcome, extending the analysis beyond movement into informal or low-productivity employment to account for labor that leaves agriculture without entering another sector. This approach distinguishes constraints that limit labor release from those that prevent released labor from being absorbed into productive non-agricultural employment.

Second, I develop a structural framework that jointly quantifies these constraints and evaluates policies that relax them. Existing models provide insights into agricultural productivity and resource allocation (Lagakos and Waugh, 2013; Adamopoulos and Restuccia, 2014), mobility frictions (Tombe and Zhu, 2019), and employment opportunities (Harris and Todaro, 1970). Related policy research examines interventions that improve agricultural productivity (Barrett et al., 2019; Suri and Udry, 2022; Byerlee et al., 2009), expand employment opportunities (McKenzie, 2017), or reduce mobility barriers (Donaldson, 2018; Bryan et al., 2014). I build on these insights by allowing subsistence constraints, transition frictions, and employment absorption to jointly determine household labor allocation within the same model. Estimating this framework using simulated method of moments allows me to quantify the relative importance of these constraints across countries under the model's assumptions. Counterfactual experiments compare interventions targeting subsistence needs, mobility costs, land rights, and labor absorption, examining how each performs when other constraints remain. These comparisons help identify which constraints policies should target to increase productive non-agricultural employment and promote structural transformation, highlighting the importance of country-specific interventions that address the constraints households face.

Third, I construct a weather-predicted NDVI growth index to study weather-induced agricultural productivity shocks. Existing studies use rainfall, temperature, and other weather measures to examine labor-market responses (Alfano and Görlach, 2024; Colmer, 2021; Branco and Féres, 2021; Emerick, 2018; Jessoe et al., 2018; Cattaneo and Peri, 2016; Hirvonen, 2016), while another strand exploits technology-induced changes in agricultural productivity (Bustos et al., 2016). I complement these approaches by constructing a weather-predicted NDVI growth index that captures nonlinear relationships between weather and vegetation growth.

The index captures weather-driven variation in agricultural conditions, providing a plausibly exogenous proxy for agricultural productivity shocks under the identifying assumptions. This approach links climatic conditions to household labor-allocation decisions and helps clarify an agricultural productivity channel through which climate change may affect household employment.

The remainder of the paper is organized as follows. Section 2 describes the data sources and the construction of the weather productivity measure. Section 3 outlines the empirical strategy and presents the reduced-form results. Section 4 introduces the theoretical model. Section 5 estimates the structural parameters. Section 6 calibrates the policies and conducts the policy counterfactual analysis. Section 7 concludes and discusses the policy implications.

2 Data and Measurement

Answering the questions requires measuring both agricultural productivity and household labor allocation over time. I combine nationally representative household panel surveys with high-resolution geospatial weather data to construct a panel of agricultural productivity and labor allocation in Niger, Malawi, and Ethiopia. This section first describes the household and geospatial data sources before introducing the measure of weather-induced agricultural productivity used throughout the analysis.

2.1 Country Selection and Survey Design

This paper uses household panel data from the World Bank Living Standards Measurement Study–Integrated Surveys on Agriculture (LSMS-ISA) for Niger, Malawi, and Ethiopia. The selection of these countries is primarily driven by the requirements of the empirical design rather than representativeness. The analysis examines changes in household labor allocation over time and requires continuously tracked households that can be followed across multiple survey waves. Moreover, because the empirical analysis compares changes in labor allocation following weather-induced agricultural productivity shocks, the observation interval should be comparable across countries.

The three selected countries satisfy these requirements. In each country, the LSMS-ISA surveys re-interview the same households after approximately three years, allowing the construction of comparable household panels. The common survey interval ensures that changes in household labor allocation are measured over similar time horizons, improving the comparability of transition dynamics across countries. In addition, all three surveys provide detailed labor allocation modules and georeferenced household locations, making it possible to construct comparable weather-productivity measures using satellite and climate data.

Table 1 summarizes the survey design. Although the survey years differ across countries, each panel spans approximately three years and follows the same households between survey waves. The empirical analysis exploits within-household changes rather than cross-sectional differences, so tracking comparable observa-

tion intervals is more important.

Finally, these countries provide useful variation for studying structural transformation. All three are Sub-

Table 1: Household Panel Surveys Used in the Analysis

Country	Survey Waves	Panel Interval	Household Tracking	GPS Available
Niger	2011, 2014	3 years	Yes	Yes
Malawi	2016, 2019	3 years	Yes	Yes
Ethiopia	2018, 2021	3 years	Yes	Yes

Notes: The table summarizes the household panel surveys used in this paper. All three surveys are part of the World Bank Living Standards Measurement Study–Integrated Surveys on Agriculture (LSMS-ISA). Households are continuously tracked across survey waves, allowing the construction of comparable household panels. Each panel spans approximately three years, ensuring that household labor transitions are measured over similar observation windows across countries. Interview periods differ because surveys are conducted to coincide with each country's agricultural calendar.

Saharan African economies where agriculture remains the dominant source of employment and agricultural production is highly dependent on weather conditions. The purpose of selecting these countries is to exploit comparable household panels while examining whether similar agricultural productivity shocks generate different patterns of labor allocation. This variation allows the paper to identify which stage of structural transformation–labor release, labor reallocation, or labor absorption–constitutes the binding constraint in each setting.

2.2 Background of the Countries Studied

Figure 1 highlights the three countries included in this study: Niger, Malawi, and Ethiopia. These countries provide a useful setting for studying structural transformation because they differ substantially in their economic development, institutional environments, and geographic contexts, while sharing a common feature, that is, agriculture remains the primary source of employment for a large share of the population.

The three countries represent different stages of economic development. In 2024, Ethiopia had the highest GDP per capita (916 USD), followed by Niger (585 USD) and Malawi (552 USD). These differences are important because the severity of subsistence constraints, the availability of non-agricultural employment, and the feasibility of labor reallocation are likely to vary across countries at different income levels.

The sample also spans three major regions of Sub-Saharan Africa: Niger in West Africa, Ethiopia in East Africa, and Malawi in Southern Africa. Beyond geography, the countries differ in their historical and institutional environments, including colonial legacies, market development, and rural livelihood systems. This diversity provides useful variation for examining whether similar productivity shocks lead to different patterns of structural transformation under different institutional settings.

A particularly important institutional difference concerns land tenure. Ethiopia operates under a state-owned land system in which rural households hold use rights that are closely tied to residence and cultivation. As a result, migration into non-agricultural employment may increase the risk of losing access to land, potentially

discouraging labor mobility (Abate and Kiros, 2010; Manyazewal and Shiferaw, 2019). In contrast, land institutions in Niger and Malawi generally provide greater scope for customary land rental and transfer. These institutional differences motivate the paper's focus on the role of household land rights as a key mechanism through which agricultural productivity shocks translate into labor reallocation across countries. Stronger household land rights reduce the opportunity cost of leaving agriculture by allowing households to better preserve the value and security of their land assets when pursuing non-agricultural opportunities.

Overall, these differences make the three countries well-suited for examining whether similar agricultural productivity shocks generate different patterns for labor release, labor reallocation, and labor absorption.



Figure 1: Niger, Malawi, and Ethiopia

2.3 Household Survey Data

The empirical analysis uses household panel data from the World Bank's Living Standards Measurement Study-Integrated Surveys on Agriculture (LSMS-ISA), which provides nationally representative panel surveys for developing countries (Grosh and Glewwe, 1998; Bentze and Wollburg, 2024). The LSMS-ISA surveys collect detailed information on household demographics, labor allocation, agricultural production, consumption, expenditures, and food security, making them particularly well suited for studying structural transformation (Lagakos and Shu, 2023).

To facilitate cross-country comparison, I construct balanced household panels with a common three-year interval between survey waves. The final sample consists of Niger (2011, 2014), Malawi (2016, 2019), and Ethiopia (2018, 2021). Restricting attention to similar observation windows reduces differences arising from survey frequency while improving the comparability of estimated productivity responses across countries. Because the key explanatory variable captures weather-induced agricultural productivity, the main estimation

samples focus on households exposed to agricultural production. Household links account for differences in how the surveys track movers and split households.

Accounting for household splits is important for studying structural transformation. Individuals who obtain employment outside agriculture may leave their original households and establish or join households elsewhere. These departing members may represent precisely the labor reallocation that this study seeks to measure. I therefore preserve links between household branches and their origins where the surveys permit, and use information on departing members where split households are not separately followed. Figures 2 and 3 illustrate these procedures.

In Niger and Malawi, the surveys track movers and split households where possible. I link second-wave households to their original households, preserving connections across related branches, as illustrated in Figure 2. This allows the analysis to retain individuals whose employment transitions involve leaving their original households.

In Ethiopia, split households are not separately followed, but the original household roster records departing members' employment status. Among 192 individuals who left for work, 145 reportedly obtained employment and 47 had not. As illustrated in Figure 3, I assign employed leavers to manufacturing in the main specification and alternatively to services in sensitivity checks. Those without employment enter idle labor.

Using these household links and labor records, I construct household labor allocations across four mutually exclusive categories: agriculture, manufacturing, services, and idle labor. Agricultural labor includes both on-farm self-employment and agricultural wage employment. Manufacturing and services include both wage employment and self-employment in the corresponding sectors, subject to the treatment of Ethiopian leavers described above. Idle labor refers to working-age household members who report neither agricultural nor non-agricultural employment during the survey reference period, regardless of whether they are actively searching for employment. Accordingly, this category should be interpreted as a measure of labor underutilization rather than unemployment in the conventional labor-force sense.

2.4 Climate Data

To measure weather-driven variation in agricultural productivity, I combine the LSMS household surveys with three complementary geospatial datasets capturing vegetation conditions, precipitation, and temperature. These variables provide plausibly exogenous measures of local growing conditions and form the basis for constructing the weather-induced agricultural productivity measure introduced in the next subsection.

Household locations are obtained from the LSMS enumeration-area (EA) GPS coordinates. Following the LSMS disclosure protocol, these coordinates are randomly displaced to protect respondent confidentiality. To account for this displacement, climate variables are extracted by averaging raster values within 2 km buffers

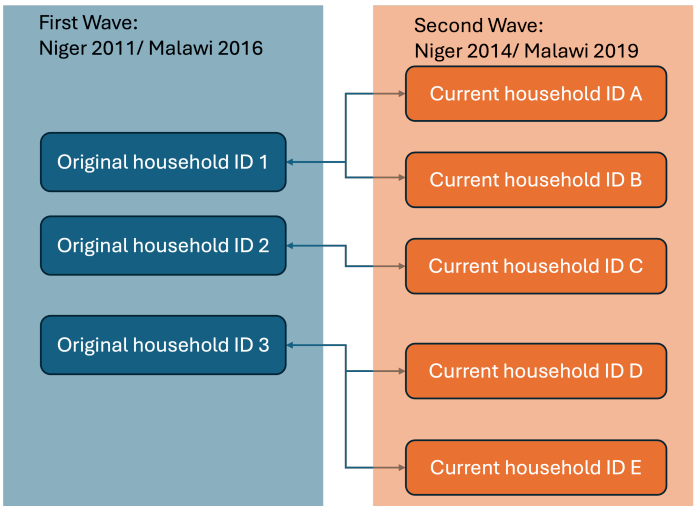


Figure 2: Linking Second-Wave Households to Origin Households in Niger and Malawi

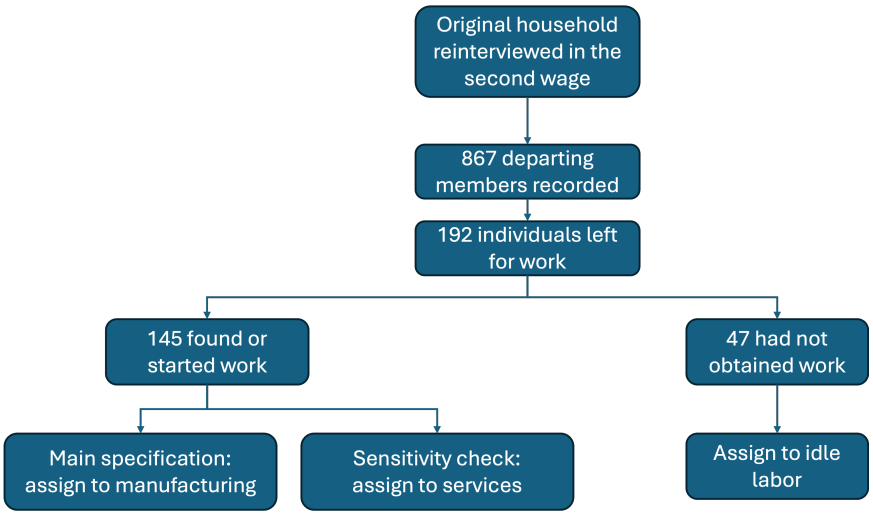


Figure 3: Accounting for Work-Related Household Departures in Ethiopia

for urban enumeration areas and 5 km buffers for rural enumeration areas, consistent with the LSMS-ISA recommendations.

Vegetation conditions are measured using the Normalized Difference Vegetation Index (NDVI) derived from the MODIS MOD13Q1 satellite product at 250-meter spatial resolution.[1] NDVI is one of the most widely used remote-sensing measures of vegetation health and has been extensively employed as a proxy for agricultural productivity in economics, agriculture, and environmental science. Unlike rainfall and temperature, which measure climatic inputs into production, NDVI captures the realized response of vegetation to weather conditions and local growing environments. Because it reflects the cumulative effects of precipitation, temperature, soil conditions, and crop growth, NDVI provides a more direct indicator of agricultural productivity than weather variables alone. Annual NDVI is constructed by averaging all MODIS observations within each year and then averaging across all pixels contained within the household buffer.

Rainfall data are obtained from the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS), which has a spatial resolution of approximately 0.05° (about 5.5 km). The buffered extraction described above is applied to these raster data to reduce measurement error arising from GPS displacement. Rainfall is included because precipitation is one of the primary determinants of agricultural production in the predominantly rain-fed farming systems of Sub-Saharan Africa.

Temperature data are obtained from the ERA5-Land reanalysis dataset at approximately 0.1° (about 10 km) spatial resolution. The annual mean temperature is assigned to each household according to its buffered location. Temperature complements rainfall by capturing thermal conditions that influence crop growth, including heat stress and growing-season conditions. Although the ERA5-Land resolution is coarser than the household buffers, it accurately reflects local climatic conditions relevant for agricultural production.

Because the household surveys are conducted at three-year intervals, climate variables are constructed annually throughout the observation period, including the three years preceding the first survey wave. These annual measures provide continuous information on local growing conditions and allow agricultural productivity to be measured consistently over time. In the next subsection, I combine NDVI, rainfall, and temperature to construct a weather-induced measure of agricultural productivity that serves as the key explanatory variable throughout the empirical analysis.

2.5 Measuring Weather-Induced Agricultural Productivity

2.5.1 Identification Strategy

Estimating the effect of agricultural productivity on labor allocation raises an important identification challenge because observed agricultural productivity is jointly determined with household production decisions.

[1] NDVI is calculated as $(NIR-Red)/(NIR+Red)$, where NIR and Red denote near-infrared and red spectral reflectance, respectively. Healthy vegetation absorbs visible red light for photosynthesis while strongly reflecting near-infrared light, resulting in higher NDVI values.

Households simultaneously choose labor allocation, technology adoption, fertilizer use, and other management practices, while unobserved local economic conditions may influence both agricultural productivity and labor supply decisions (Branco and Féres, 2021; Emerick, 2018; Jessoe et al., 2018). Consequently, directly using observed agricultural output or productivity as an explanatory variable would generally produce biased estimates of the effect of agricultural productivity on structural transformation.

To address this challenge, I construct a measure of weather-induced agricultural productivity, defined as the component of vegetation growth explained solely by weather conditions. Figure 4 illustrates the identification strategy. Observed annual NDVI growth contains two sources of variation. The first is weather-driven variation, arising from rainfall and temperature, which is plausibly exogenous to household production decisions. The second reflects human and other non-weather factors, including labor allocation, technology adoption, fertilizer use, management practices, and other local conditions. The objective is to isolate only the weather-driven component and use it as the treatment variable throughout the empirical analysis.

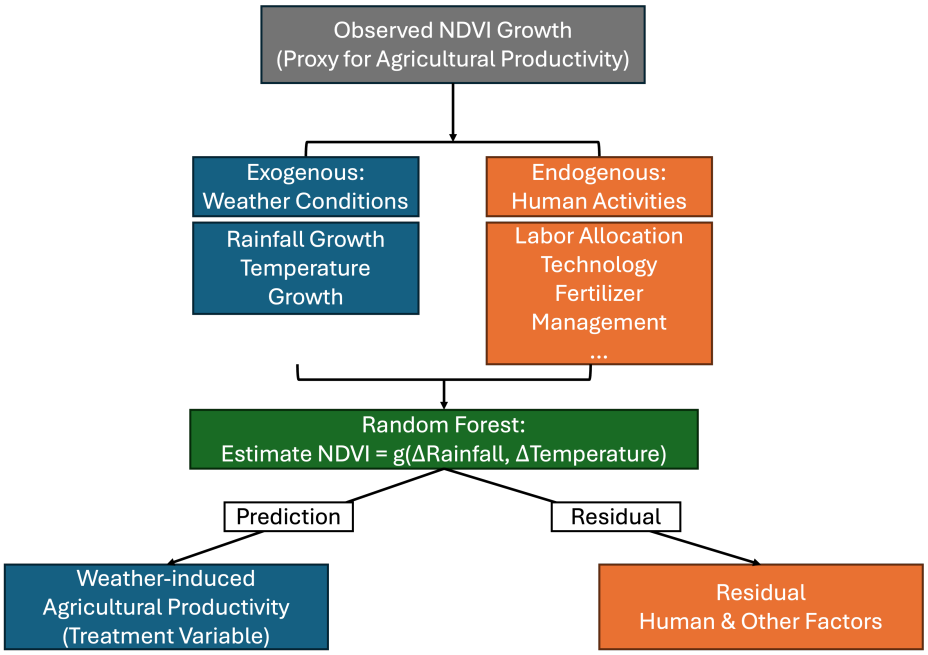


Figure 4: Weather-Induced Agricultural Productivity Identification Strategy

This decomposition can be motivated by a standard agricultural production function in which agricultural output depends on both weather conditions and household production decisions. Appendix A provides the formal derivation. The resulting treatment variable is defined as

$$\widehat{\Delta\theta}_{it} = g(\Delta\text{Rainfall}_{it}, \Delta\text{Temperature}_{it}), \tag{1}$$

where $\widehat{\Delta\theta}_{it}$ denotes the predicted weather-induced agricultural productivity growth, and $g(\cdot)$ is the Random Forest prediction function that maps annual weather changes to annual vegetation growth.

2.5.2 Construction

To estimate the weather-driven component, I use annual NDVI growth as a proxy for changes in agricultural productivity. NDVI has been shown to closely track agricultural production and crop yields in both economics and remote-sensing studies (Donaldson and Storeygard, 2016; Lobell et al., 2015). Because NDVI reflects the realized response of vegetation to local growing conditions, it provides an objective measure of agricultural productivity that is observed independently from household survey responses.

LSMS-ISA randomly displaces rural enumeration-area coordinates by up to 5 km to protect respondent confidentiality. Following the World Bank disclosure protocol, the remote-sensing variables (NDVI, temperature, and rainfall) used for weather-productivity construction are extracted using a 5km buffer around each rural GPS point, ensuring that the actual enumeration area remains within the extraction window despite the displacement. Because the agricultural production is measured for households residing within the sampled enumeration area, and plot locations are generally referenced relative to the household homestead in LSMS-ISA surveys, this buffering strategy substantially reduces measurement error arising from the displacement while preserving the relevant local production environment. (Ali et al. (2019))

The relationship between weather conditions (temperature and rainfall) and vegetation growth (NDVI) is inherently nonlinear. Crop growth depends on interactions between rainfall and temperature rather than monotonic changes in either variable alone. For example, additional rainfall may improve agricultural productivity under moderate temperatures but reduce productivity when temperatures are already low or when excessive rainfall causes flooding. Linear regression, therefore, provides an inadequate representation of the weather-productivity relationship.

To address the problem, we estimate the following relationship:

$$\Delta NDVI_{it} = g(\Delta Rainfall_{it}, \Delta Temperature_{it}) + \varepsilon_{it} \tag{2}$$

where $g(\cdot)$ is a function estimated using a Random Forest algorithm. The predicted component,

$$g(\Delta Rainfall_{it}, \Delta Temperature_{it})$$

, represents the variation in vegetation growth systematically explained by weather conditions and is interpreted as the weather-induced agricultural productivity shock used in the subsequent analysis.[2] The residual component, ε_{it}, captures the remaining variation in vegetation growth arising from non-weather factors, such as labor allocation, technology, fertilizer, land management practices, or other human activities. By using only the predicted component, the treatment variable isolates exogenous weather-induced productivity varia-

[2]The weather-productivity shock is a generated regressor obtained from out-of-sample Random Forest predictions. The reported second-stage standard errors are conditional on the estimated prediction function and therefore do not account for first-stage prediction uncertainty.

tion rather than changes in vegetation associated with household production decisions or other local factors.

The relationship between weather and vegetation is highly nonlinear and involves complex interactions between rainfall and temperature that are difficult to capture using conventional linear specifications. Random Forest is well-suited for estimating $g(\cdot)$ because it flexibly models nonlinearities and high-order interactions without imposing a parametric functional form, while its ensemble-learning framework reduces prediction variance by averaging across a large number of decision trees. We use growth rates rather than levels because changes in agricultural productivity, rather than productivity levels, drive households' labor reallocation decisions in our theoretical model.

The prediction model is estimated using 5-fold cross-validation. Specifically, the remote sensing dataset is randomly partitioned into 5 mutually exclusive folds. For each fold, the Random Forest is trained on the remaining 4 folds and used to predict NDVI growth for the held-out fold. Because each predicted value is generated by a model that excludes the corresponding observation from the training sample, the resulting weather-productivity measure is not affected by in-sample overfitting. All prediction statistics reported below are calculated from these out-of-sample predictions.

The out-of-sample predictions isolate the component of annual vegetation growth that is systematically explained by weather conditions and are used as the weather-induced agricultural productivity shock in the subsequent empirical analysis. The out-of-sample (R^2) equals 0.29 for Niger, 0.60 for Malawi, and 0.30 for Ethiopia, while the mean absolute prediction error remains low across all three countries. These results suggest that rainfall and temperature explain a meaningful share of annual vegetation dynamics while leaving the remaining variation to household production decisions and other non-weather factors.

2.5.3 Distribution of Weather-Induced Agricultural Productivity

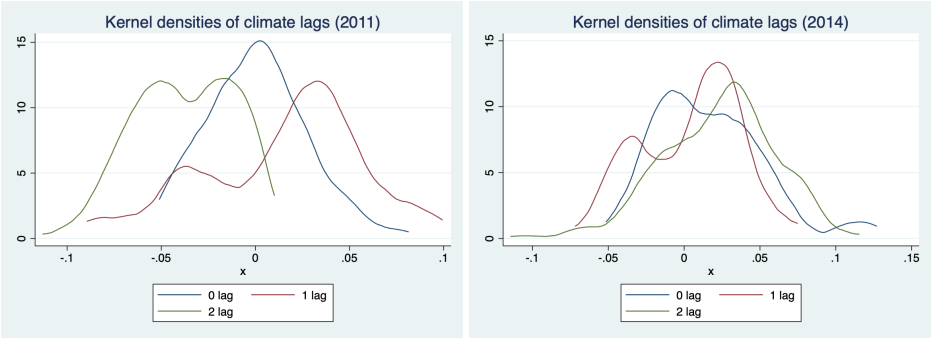
Figure 5 presents the distribution of weather-induced agricultural productivity across enumeration areas (EAs) for each country and year. Positive values indicate weather conditions that improve agricultural productivity relative to the previous year, whereas negative values correspond to less favorable growing conditions. Different colors represent different crop years within each country.

The distributions exhibit substantial variation across both space and time. Some years display distributions centered above zero, indicating broadly favorable weather conditions (for example, 2010 in Niger and 2019 in Malawi), while others are centered below zero, reflecting widespread adverse growing conditions. In several years, the distribution is concentrated around zero, suggesting relatively normal agricultural conditions. These patterns demonstrate that weather-induced agricultural productivity varies considerably across enumeration areas even within the same year, providing the cross-sectional and temporal variation required for identification.

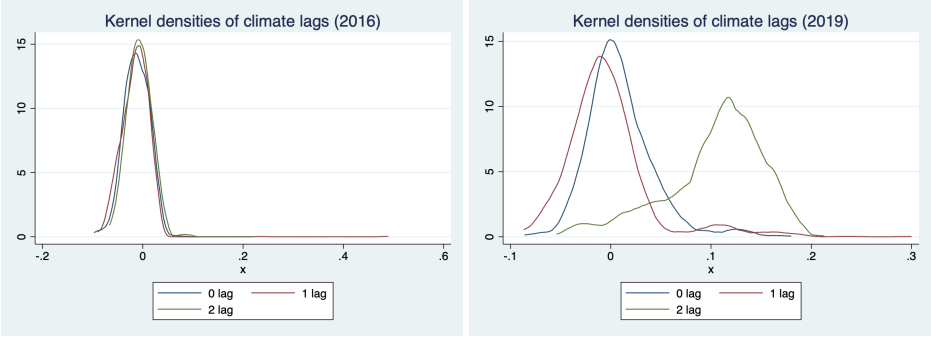
Importantly, the weather-productivity shock represents a temporary, exogenous change in agricultural pro-

ductivity rather than a permanent difference in agricultural potential across locations. This feature makes it well suited for identifying households' short-run labor reallocation responses to productivity changes while avoiding confounding long-run differences in economic development, technology changes, or infrastructure. The estimated responses capture how households adjust labor allocation following an exogenous productivity shock, providing the behavioral relationships that underpin the subsequent structural analysis. Although the weather shock is transitory, the resulting labor reallocation can be long-term if transitions across sectors involve fixed costs or other adjustment frictions. We will return to this distinction when motivating the structural model in Section 5.

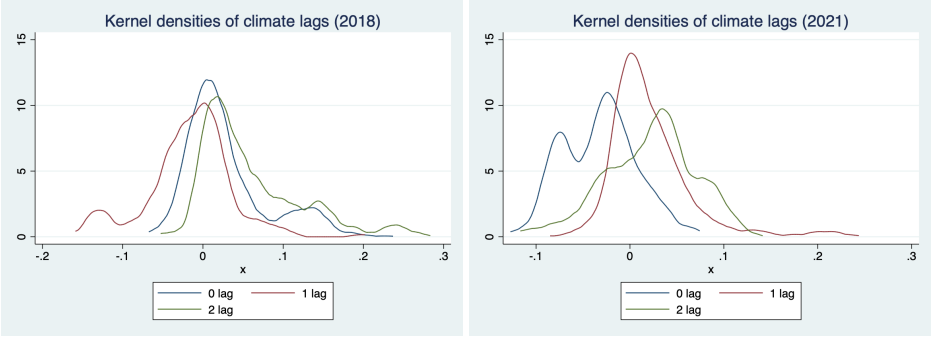
Figure 5: Distribution of Weather Productivity



(a) Niger in 2011 **(b)** Niger in 2014



(c) Malawi in 2016 **(d)** Malawi in 2019



(e) Ethiopia in 2018 **(f)** Ethiopia in 2021

2.6 Summary Statistics

Table 2 reports summary statistics for the variables used in both the reduced-form and structural analyses. The dependent variables are the share of working-age household members (ages 15-85) allocated to agriculture,

manufacturing, services, and underutilized (idle) labor. These shares are calculated using all working-age household members as the denominator. Consequently, the four reported shares do not necessarily sum to one because the remaining share consists of working-age individuals outside the labor force, such as students, retirees, homemakers, or individuals not participating in economic activities.

Agriculture remains the dominant source of employment in all three countries, accounting for 63.4 percent of household labor in Niger, 77.6 percent in Malawi, and 66.0 percent in Ethiopia. Non-agricultural employment is concentrated primarily in services, while manufacturing accounts for only a small share of household labor. A notable difference across countries is the prevalence of idle labor. Idle labor (underutilized labor) refers to working-age individuals who are unemployed or only marginally attached to productive employment through unpaid household agricultural work or temporary and casual activities (see Appendix A for the harmonized construction). Ethiopia exhibits a substantially larger idle labor share (18.2 percent) than Niger (2.8 percent) and Malawi (0.4 percent). Throughout the paper, idle labor refers to working-age household members who are not engaged in either agricultural or non-agricultural employment and therefore measures the extent to which labor is not absorbed into productive sectors. The high incidence of idle labor in Ethiopia provides preliminary evidence that labor released from agriculture may not always translate into productive non-agricultural employment.

The table also reports the key variables used to examine the mechanisms underlying structural transformation. Food subsistence serves as a proxy for subsistence constraints, minimum distance to the nearest market captures mobility barriers, and land type classifies households according to differences in household land-rights index. These variables exhibit substantial variation across countries, providing useful heterogeneity for the empirical and structural analyses. Food subsistence is most prevalent in Malawi (70.3 percent of households), followed by Niger (48.4 percent), while Ethiopia has the lowest incidence (15.4 percent). Average distance to markets ranges from 3.7 kilometers in Malawi to approximately 19 kilometers in Niger, indicating considerable cross-country differences in market accessibility. The land-type variable captures households' reported inheritance and collateral rights over agricultural land. It equals 0 for neither right and 1 for at least one right. About 90% of households in Malawi report at least one of these rights, compared with only 6.7% in Niger.

The estimation sample consists of 1,890 households from 139 communities in Niger, 955 households from 81 communities in Malawi, and 1,120 households from 213 communities in Ethiopia, with each household observed in both survey waves. Together, these descriptive statistics reveal substantial cross-country heterogeneity in labor allocation, household constraints, and market access despite a common dependence on agriculture. The following section examines whether these differences translate into heterogeneous responses of household labor allocation to weather-induced agricultural productivity shocks.

Table 2: Summary Table

	(1)	(2)	(3)	(4)	(5)	(6)
	Niger		Malawi		Ethiopia	
	Mean	SD	Mean	SD	Mean	SD
HH ag share	0.634	0.406	0.776	0.357	0.660	0.462
HH manuf share	0.118	0.276	0.052	0.188	0.023	0.134
HH serv share	0.208	0.333	0.154	0.301	0.094	0.265
HH idle share	0.028	0.135	0.004	0.054	0.182	0.371
Weather Productivity Shock (1-lag)	0.016	0.037	-0.011	0.027	0.004	0.048
Food Subsistence	0.484	0.500	0.703	0.457	0.154	0.361
Min-distance	18.993	24.166	3.724	3.196	16.786	18.284
Household Land-rights Index	0.067	0.250	0.900	0.299	0.627	0.4837
Household Obs	1890		955		1120	
Community Obs	139		81		213	

Labor shares are calculated as the proportion of all working-age household members (ages 15–85) allocated to agriculture, manufacturing, services, and underutilized labor (idle). The shares do not necessarily sum to one because working-age individuals outside the labor force are not included in these four categories.

3 Empirical Evidence

3.1 Empirical Strategy

In this section, I examine two central structural transformation questions. First, does higher agricultural productivity release labor from agriculture, as predicted by the Lewis-type mechanism (Lewis, 1954)? Second, is this release accompanied by increased non-agricultural employment or greater labor underutilization, consistent with the labor-absorption mechanism emphasized by Harris and Todaro (Harris and Todaro, 1970)?

To answer these questions, I combine annually constructed weather-induced productivity shocks with household panel data observed at three-year intervals. Then I regress the weather productivity shock on households' sectoral labor share, including agricultural, manufacturing, service, and idle labor share, and explore how weather productivity shocks lead to different responses across sectors and countries:

$$s_{ijt} = \beta_j^{OLS} \Delta NDVI_{g,t-1} + \delta_i + \lambda_t + \varepsilon_{ijt}. \tag{3}$$

where s_{ijt} denotes household i's labor share in sector j at survey wave t, where j includes agriculture, manufacturing, services, and idle labor. The treatment variable $\Delta NDVI_{g,t-1}$ measures observed NDVI growth in community g over the lagged treatment window and serves as a proxy for changes in agricultural productivity shock. To control for confounding factors, I include two sets of fixed effects: original-household fixed effects δ_i, which control for persistent characteristics shared by households linked to the same baseline household[3], and year fixed effects λ_t, which control for national shocks and common changes across survey waves. The

[3]In all three countries, I link split-off households and individuals back to their original households.

OLS coefficient β_j^{OLS} measures the within-panel association between observed NDVI growth and the household labor share in sector j.

Observed NDVI growth may reflect agricultural labor inputs, crop choices, and land management, rather than productivity changes alone. Although the treatment window precedes the observed labor outcome, persistent household decisions and local shocks may affect both NDVI and subsequent labor allocation, making OLS estimates endogenous. To address this concern, I instrument observed NDVI growth with weather-predicted NDVI growth, following the weather-based identification logic of Allen and Atkin (2022). Specifically, I use a Random Forest model to predict NDVI growth from rainfall and temperature, averaging bootstrap predictions generated out of fold, with all observations from the same community assigned to the same fold. I then use these predictions as an instrument in a linear first-stage regression, with the original household, community, and year fixed effects in the outcome equation:

$$\Delta NDVI_{g,t-1} = \pi \widehat{Z}_{g,t-1} + \delta_i^F + \lambda_t^F + v_{igt}, \tag{4}$$

where $\widehat{Z}_{g,t-1}$ denotes the out-of-fold, bootstrap-averaged weather prediction of NDVI growth. I estimate Equation (4) by two-stage least squares, using this prediction as an instrument for observed NDVI growth and including the same fixed effects in both stages. The resulting IV coefficient measures the labor-allocation response to weather-induced variation in observed NDVI growth, under the identifying assumptions.

3.2 Does Agricultural Productivity Release Labor?

Table 3 examines whether higher agricultural productivity releases labor from agriculture. The OLS estimates in columns (1), (4), and (7) are negative in all three countries, although only the estimate for Malawi is statistically significant. After instrumenting observed NDVI growth with its weather-predicted component, the coefficients become more negative: the IV estimates in columns (2), (5), and (8) are -0.592 for Niger, -1.794 for Malawi, and -1.171 for Ethiopia. The IV estimate for Niger is significant at the 10 percent level, while those for Malawi and Ethiopia are significant at the 5 percent level. The reduced-form estimates also indicate that favorable weather-predicted NDVI growth reduces agricultural labor shares. These results are consistent with the Lewis-type labor-release mechanism: agricultural productivity improvements can allow households to meet their food and income needs with less labor devoted to farming. The estimates capture the overall labor-release response.

To interpret the economic magnitudes, I evaluate the IV estimates at a one-standard-deviation increase in observed NDVI growth, which scales the response by the variation observed in each country. A one-standard-deviation increase in NDVI growth is 0.068 in Niger, 0.039 in Malawi, and 0.083 in Ethiopia. Applying the IV coefficients, these increases imply reductions in agricultural labor shares of approximately 4.0, 7.0, and 9.7 percentage points, respectively. These reductions are equivalent to approximately 6.3, 9.0, and 15.0

percent of the respective mean agricultural labor shares.

Thus, productivity improvements are associated with reductions in households' allocation of labor to agriculture. However, labor release alone does not ensure structural transformation: workers must also reach non-agricultural opportunities and find employment. The next analysis examines whether lower agricultural labor shares are accompanied by increased non-agricultural employment or greater labor underutilization.

Table 3: NDVI Growth and Household Agricultural Labor Allocation

	Household agricultural labor share								
	Niger			Malawi			Ethiopia		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Estimator	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Observed lagged NDVI growth	-0.135	-0.592*		-0.593**	-1.794**		-0.047	-1.171**	
	(0.138)	(0.307)		(0.247)	(0.834)		(0.172)	(0.575)	
Weather-predicted NDVI growth			-0.640**			-0.824**			-0.827**
			(0.311)			(0.392)			(0.368)
Mean outcome	0.634	0.634	0.634	0.776	0.776	0.776	0.649	0.649	0.649
SD of observed lagged NDVI growth	0.068	0.068	0.068	0.039	0.039	0.039	0.083	0.083	0.083
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,780	3,780	3,780	1,910	1,910	1,910	2,240	2,240	2,240
Households	1,890	1,890	1,890	955	955	955	1,120	1,120	1,120
Communities	139	139	139	81	81	81	213	213	213
First-stage coefficient		1.081***			0.459***			0.706***	
		(0.152)			(0.114)			(0.100)	
Kleibergen–Paap F-statistic		50.88			16.28			50.15	
First-stage partial R^2		0.275			0.121			0.193	
Anderson–Rubin F-test p-value		0.042			0.039			0.026	

Notes: The dependent variable is the household agricultural labor share. Columns (1), (4), and (7) report OLS estimates; columns (2), (5), and (8) report IV estimates; and columns (3), (6), and (9) report reduced-form estimates. Within each country, all specifications use the agricultural outcome's IV estimation sample. The IV specifications instrument observed lagged NDVI growth with community-grouped out-of-fold weather prediction. The reduced-form specifications regress agricultural labor shares directly on this instrument. First-stage coefficients relate observed lagged NDVI growth to the excluded instrument. The SD of observed lagged NDVI growth is calculated in the corresponding estimation sample. All specifications include household and year fixed effects. Standard errors clustered by household and community are in parentheses. Stars indicate conventional coefficient tests; Anderson–Rubin F-test p-values are reported separately. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.3 Where Does Released Labor Go?

Table 4 examines whether agricultural labor release is accompanied by increased non-agricultural employment or greater labor underutilization. In Niger, the IV estimate for manufacturing is positive, with a coefficient of 0.541. The service-sector response is close to zero and not statistically significant, while there is little evidence of an idle-labor response. These results suggest that manufacturing activities provide an outlet for labor released from agriculture. In Malawi, the clearest response also occurs in manufacturing, and the IV coefficient is 1.281. There is also weaker evidence of increased idle labor ($p = 0.068$), suggesting that non-agricultural employment does not absorb all of the labor released. The service-sector response is not statistically significant. The reduced-form estimates show corresponding increases in manufacturing in both Niger and Malawi.

Ethiopia exhibits a different pattern. The IV manufacturing coefficient is close to zero and not statistically significant, whereas the service-sector coefficient is negative at -0.354 and significant at the 10 percent

level. The idle-labor coefficient is positive at 1.274 and significant at the 5 percent level. The reduced-form estimates likewise indicate an increase in idle labor and weaker evidence of a decline in services. Thus, agricultural productivity improvements appear to release labor without a corresponding expansion of non-agricultural employment. This pattern is consistent with the distinction between labor release and labor absorption emphasized by the Harris–Todaro mechanism: improved agricultural conditions can free labor, but productive employment also requires access to opportunities and sufficient labor demand.

Table 4: NDVI Growth and Household Non-Agricultural Labor Allocation

	Niger			Malawi			Ethiopia		
	(1) Manufacturing	(2) Services	(3) idle labor	(4) Manufacturing	(5) Services	(6) idle labor	(7) Manufacturing	(8) Services	(9) idle labor
Panel A: OLS									
Observed lagged NDVI growth	0.231**	-0.030	-0.030	0.218	0.285	0.039	-0.022	0.045	0.029
	(0.100)	(0.125)	(0.051)	(0.145)	(0.188)	(0.052)	(0.043)	(0.087)	(0.152)
Panel B: IV (2SLS)									
Instrumented lagged NDVI growth	0.541***	0.027	0.076	1.281**	-0.381	0.427*	0.062	-0.354*	1.274**
	(0.182)	(0.236)	(0.104)	(0.599)	(0.611)	(0.230)	(0.121)	(0.207)	(0.523)
Panel C: Reduced form									
Weather-predicted NDVI growth	0.585***	0.030	0.082	0.589**	-0.175	0.196*	0.044	-0.250*	0.900***
	(0.173)	(0.255)	(0.109)	(0.232)	(0.262)	(0.106)	(0.085)	(0.137)	(0.316)
Mean outcome	0.118	0.209	0.028	0.052	0.154	0.005	0.033	0.094	0.180
SD of observed lagged NDVI growth	0.068	0.068	0.068	0.039	0.039	0.039	0.083	0.083	0.083
Observations	3,780	3,780	3,780	1,910	1,910	1,910	2,240	2,240	2,240
Households	1,890	1,890	1,890	955	955	955	1,120	1,120	1,120
Communities	139	139	139	81	81	81	213	213	213
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variables are household manufacturing, services, and idle-labor shares. Standard errors clustered by household and community are in parentheses. Within each outcome and country, OLS, IV, and reduced-form specifications use the corresponding IV estimation sample. The excluded instrument is the community-grouped out-of-fold weather prediction of lagged NDVI growth. All specifications include household and year fixed effects. The SD of observed lagged NDVI growth is calculated in the corresponding estimation sample. Stars indicate conventional coefficient tests. The dash denotes the Ethiopia services reduced-form standard error, which was unavailable in the supplied output. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.4 Mechanisms of Structural Transformation

3.4.1 Empirical Strategy

The preceding results show that agricultural productivity improvements release labor, but the accompanying changes in non-agricultural employment differ across countries. I next examine the constraints that may explain these differences. Following the three-stage framework, I consider subsistence constraints that limit labor release from agriculture, mobility and land-rights frictions that shape access to non-agricultural opportunities, and labor-market conditions that determine whether released labor can be absorbed into employment. These analyses connect the observed labor-allocation responses to the Lewis-type release mechanism, transition frictions, and the Harris–Todaro absorption mechanism.

In this section, I examine the mechanisms by testing whether labor-allocation responses to weather-induced agricultural productivity shocks differ with these constraints. For subsistence status, distance, land rights, and labor-market absorption, I use baseline characteristics, measured in the first survey wave and held fixed across waves, to reduce concerns that the moderators themselves respond to productivity shocks over the survey interval. Following the baseline reduced-form approach, I interact the lagged weather-predicted NDVI

growth index with each baseline constraint and estimate the responses separately by country and sector:

$$s_{ijt} = \beta_j Z_{g,t-1} + \gamma_j (Z_{g,t-1} \times C_{i0}) + \alpha_{ij} + \lambda_{jt} + \varepsilon_{ijt}, \tag{5}$$

where $Z_{g,t-1}$ denotes the lagged weather-predicted NDVI growth index and C_{i0} denotes the baseline constraint under consideration. All constraints are measured in the baseline survey wave—2011 in Niger, 2016 in Malawi, and 2018 in Ethiopia—and held fixed across waves. Baseline absorption is constructed from the community leave-one-household-out ratio of non-agricultural employment to the sum of non-agricultural employment and idle labor. In all three countries, the absorption indicator equals one when the baseline ratio is at or above the baseline household mean and zero otherwise. The main effect of C_{i0} is absorbed by household fixed effects. Standard errors are clustered by household and community, using original household and community identifiers. This approach follows Basu (2025), who interacts lagged drought exposure with baseline land ownership and education in a household fixed-effects model. Related work by Colmer (2021) examines how labor-market flexibility shapes responses to temperature-driven agricultural shocks.

This specification directly estimates how labor allocation responds to the weather-predicted NDVI growth index and how that response differs with baseline constraints. The coefficients capture responses to the weather-induced component of agricultural conditions, rather than effects per unit of observed NDVI growth. The reduced-form approach also avoids scaling these responses by first-stage relationships that are weak in several of the interacted IV specifications. Identification relies on the weather index and its interaction with the baseline constraint being uncorrelated with unobserved determinants of changes in labor allocation, conditional on the included fixed effects.

The coefficient γ_j measures how the response to the weather index varies with the constraint. For a binary moderator, β_j is the response for households with $C_{i0} = 0$, while $\beta_j + \gamma_j$ is the response for households with $C_{i0} = 1$. For a continuous moderator, the estimated response is $\beta_j + \gamma_j C_{i0}$. I report tests of the interaction coefficient and combined responses for binary moderators. To express magnitudes in comparable units within each country, responses can be scaled by the standard deviation of the weather index. The response to a one-standard-deviation increase in the index at $C_{i0} = c$ is $(\beta_j + \gamma_j c) \times SD(Z)$.

These analyses ask which constraints shape the labor-allocation responses to agricultural productivity improvements. Do subsistence constraints affect how much labor is released from agriculture? Do distance and land rights affect whether households move into non-agricultural activities? Do initial labor-market conditions affect whether released labor enters employment or remains underutilized? By examining how the responses differ with these baseline characteristics, the mechanism tests connect the observed changes in labor allocation to the release, transition, and absorption stages. Together, they help explain why similar improvements in agricultural productivity can lead to different patterns of structural transformation across countries.

3.4.2 Subsistence Constraints and Labor Release

The first stage of the framework asks whether subsistence constraints limit the release of labor from agriculture. Households that must meet a minimum food requirement may retain more labor in agriculture even when productivity improves. I explore how this constraint affects labor allocation by allowing baseline subsistence status to affect the responsiveness of household labor shares to weather-induced agricultural productivity shocks. Specifically, I interact the weather-predicted NDVI growth index with baseline subsistence status in a reduced-form specification.

The estimates are shown in Table 5. In Niger, the interaction coefficient for agriculture is positive and significantly different from zero at the 5% level. The estimated agricultural response is -0.929 for unconstrained households, while the interaction coefficient is 0.590. Consistent with subsistence constraints limiting labor release, constrained households reduce agricultural labor less in response to favorable weather-induced productivity shocks. In terms of magnitudes, the estimated reduction in agricultural labor is approximately 64% smaller for constrained households, with a combined coefficient of -0.339. The service-sector interaction is negative but not statistically significant, so the estimates do not establish a difference in the service-sector response. The manufacturing interaction is close to zero, suggesting little difference in this response across the two groups.

Malawi shows a similar pattern in agricultural labor allocation. The interaction coefficient is positive and significant at the 1% level. The agricultural response is -1.942 for unconstrained households, while the interaction coefficient is 1.676, giving a combined response of -0.266 for constrained households. The estimated reduction in agricultural labor is therefore approximately 86% smaller among constrained households. The manufacturing response is positive and significant for unconstrained households, but the interaction is negative and significant at the 5% level, indicating a smaller manufacturing response among constrained households. This finding reinforces the main regression results, which show that labor released from agriculture in Malawi shifts toward manufacturing. It further suggests that subsistence constraints limit this transition by reducing both labor release from agriculture and the accompanying increase in manufacturing employment.

In Ethiopia, the agricultural interaction is also positive, but it is not statistically significant. The estimated agricultural response is -1.148 for unconstrained households and -0.443 for constrained households. Although these point estimates suggest less labor release among constrained households, the estimates do not establish a difference between the two groups. None of the interactions for manufacturing, services, or idle labor is statistically significant. For unconstrained households, favorable weather-induced productivity shocks reduce agricultural labor and increase both manufacturing employment and idle labor, with the larger estimated increase occurring in idle labor. This pattern raises the question of whether labor released from agriculture can reach and be absorbed into non-agricultural employment.

Together, the estimates show that baseline subsistence status shapes the agricultural labor response in Niger and Malawi. In both countries, constrained households exhibit smaller reductions in agricultural labor following favorable weather-induced productivity shocks, consistent with the Lewis-type release mechanism. In Ethiopia, the responses do not differ significantly by baseline subsistence status, while the increase in idle labor points to the importance of examining the transition and absorption stages.

Table 5: Subsistence Constraints and Labor Release

	(1) Agriculture	(2) Manufacturing	(3) Services	(4) Idle labor
Panel A: Niger				
Weather-predicted NDVI growth index	-0.929***	0.610***	0.212	0.062
	(0.343)	(0.206)	(0.284)	(0.118)
Weather index × baseline subsistence constraint	0.590**	-0.051	-0.372	0.041
	(0.289)	(0.235)	(0.246)	(0.142)
Observations	3,780	3,780	3,780	3,780
Households	1,890	1,890	1,890	1,890
Communities	139	139	139	139
Panel B: Malawi				
Weather-predicted NDVI growth index	-1.942***	1.126***	0.080	0.235
	(0.460)	(0.363)	(0.280)	(0.223)
Weather index × baseline subsistence constraint	1.676***	-0.841**	-0.300	-0.056
	(0.535)	(0.422)	(0.436)	(0.265)
Observations	1,884	1,884	1,884	1,884
Households	942	942	942	942
Communities	81	81	81	81
Panel C: Ethiopia				
Weather-predicted NDVI growth index	-1.148***	0.202***	-0.178	0.931***
	(0.368)	(0.071)	(0.140)	(0.328)
Weather index × baseline subsistence constraint	0.706	-0.187	-0.382	0.021
	(0.910)	(0.129)	(0.311)	(0.753)
Observations	2,240	2,240	2,240	2,240
Households	1,120	1,120	1,120	1,120
Communities	213	213	213	213

Notes: Standard errors are in parentheses. Subsistence status is measured in 2011 in Niger, 2016 in Malawi, and 2018 in Ethiopia and held fixed across waves. The baseline subsistence indicator equals one for constrained households. The main weather-index coefficient is the response for unconstrained households; the response for constrained households is the sum of the main and interaction coefficients. All specifications are reduced-form regressions of household labor shares on the weather-predicted NDVI growth index and its interaction with baseline subsistence status. Coefficients are reported in the original index units. All specifications include household and year fixed effects, using original-household fixed effects in Malawi. Standard errors are clustered by household and community (original household and community in Malawi). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.4.3 Mobility Frictions and Labor Reallocation

The transition stage asks whether access to non-agricultural opportunities shapes the reallocation of labor released from agriculture. Greater distance may increase the costs of searching for work, commuting, and accessing markets. I explore this mechanism by allowing baseline distance to affect the responsiveness of household labor shares to weather-induced agricultural productivity shocks. Distance is measured as the minimum distance to a road, market, or population center.

The estimates are shown in Table 6. None of the distance interactions is statistically significant at the 10% level in Niger or Malawi. In Niger, the agricultural interaction is negative, while the manufacturing and service-sector interactions are positive. These point estimates do not follow the pattern in which greater distance limits movement out of agriculture, and the estimates do not establish differences in labor-allocation

responses by baseline distance. In Malawi, the positive agricultural interaction and negative manufacturing interaction suggest less labor release and a smaller manufacturing response at greater distances. However, both interactions are imprecisely estimated, and the service-sector interaction is close to zero.

In Ethiopia, the agricultural interaction is negative and significant at the 10% level, while the idle-labor interaction is positive and significant at the 1% level. The corresponding interaction coefficients are -0.049 and 0.064. This pattern indicates that households farther from roads, markets, or population centers experience a larger reduction in agricultural labor and a larger increase in idle labor following favorable weather-induced productivity shocks. The manufacturing and service-sector interactions are negative but not statistically significant. These findings are consistent with a separation between labor release and employment absorption: greater distance does not prevent labor from leaving agriculture, but the accompanying increase in idle labor suggests that released workers may face barriers to reaching non-agricultural employment.

Related studies distinguish improvements in transport access from the employment opportunities available to workers. In rural India, Asher and Novosad (2020) find that new roads facilitate movement out of agriculture toward work outside the village, while employment in village firms expands only slightly. In South Africa, Banerjee and Sequeira (2023) find that transport subsidies increase job-search intensity without increasing the probability of employment. In Ethiopia, Abebe et al. (2021) find that transport subsidies improve formal employment in the short run, but these effects dissipate after four years, while assistance with signaling skills generates more persistent gains. These findings suggest that reducing travel costs can improve access and search without necessarily producing sustained employment gains when other barriers remain.

In the present analysis, baseline distance does not identify statistically significant differences in labor-allocation responses in Niger or Malawi. In Ethiopia, greater distance is associated with a larger agricultural labor reduction and a larger idle-labor increase, consistent with mobility frictions limiting access to employment after labor is released. However, the estimates do not establish differences in manufacturing or service-sector responses. This motivates examining land rights separately as another friction that may shape households' ability to move into non-agricultural activities.

Land rights may also shape households' ability to move into non-agricultural activities. When retaining land rights depends on continued cultivation, households may keep labor in agriculture to avoid losing access to their land. More secure rights can reduce this risk, allowing household members to work elsewhere while retaining their land. Where rights permit renting or selling, rental income or sale proceeds may also help cover the costs of searching for work, commuting, or relocating. Consistent with these channels, De Janvry et al. (2015) find that land certification in Mexico increases migration by removing the link between land rights and continued land use. In Ethiopia, Ayele and Siba (2017) find that certification increases participation in unskilled non-agricultural work. I explore this mechanism by interacting the weather-predicted NDVI growth index with baseline land rights in a reduced-form specification. Land rights are measured in the first survey

Table 6: Mobility Frictions and Labor Reallocation: Baseline Distance

	(1) Agriculture	(2) Manufacturing	(3) Services	(4) Idle labor
Panel A: Niger				
Weather-predicted NDVI growth index	-0.272	0.500**	-0.165	-0.017
	(0.368)	(0.213)	(0.314)	(0.104)
Weather index × baseline distance	-0.016	0.004	0.008	0.004
	(0.010)	(0.007)	(0.005)	(0.004)
Observations	3,780	3,780	3,780	3,780
Households	1,890	1,890	1,890	1,890
Communities	139	139	139	139
Panel B: Malawi				
Weather-predicted NDVI growth index	-1.079*	0.596*	-0.163	0.299*
	(0.600)	(0.336)	(0.376)	(0.154)
Weather index × baseline distance	0.088	-0.003	-0.004	-0.036
	(0.115)	(0.089)	(0.086)	(0.026)
Observations	1,910	1,910	1,910	1,910
Households	955	955	955	955
Communities	81	81	81	81
Panel C: Ethiopia				
Weather-predicted NDVI growth index	-0.254	0.244**	-0.135	-0.051
	(0.464)	(0.106)	(0.183)	(0.405)
Weather index × baseline distance	-0.049*	-0.005	-0.007	0.064***
	(0.025)	(0.003)	(0.007)	(0.022)
Observations	2,240	2,240	2,240	2,240
Households	1,120	1,120	1,120	1,120
Communities	213	213	213	213

Notes: Standard errors are in parentheses. Distance is measured in 2011 in Niger, 2016 in Malawi, and 2018 in Ethiopia and held fixed across waves. It is the minimum distance to a road, market, or population center in all three countries. The main weather-index coefficient is the response at baseline distance zero; the response at distance d is $\beta_j + \gamma_j d$. All specifications are reduced-form regressions of household labor shares on the weather-predicted NDVI growth index and its interaction with baseline distance. Coefficients are reported in the original index units. All specifications include household and year fixed effects, using original-household fixed effects. Standard errors are clustered by household and community (original household and community). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

wave and held fixed across waves.

The estimates are shown in Table 7. None of the land-rights interactions is statistically significant at the 10% level in Niger. The estimates therefore do not establish differences in labor-allocation responses by baseline land rights.

In Malawi, the manufacturing interaction is positive and significant at the 5% level, with a coefficient of 0.992. The estimated manufacturing response is -0.299 for households without baseline land rights and 0.693 for households with rights. The combined response for households with rights is also significant at the 5% level. This finding reinforces the main regression results showing movement toward manufacturing in Malawi and is consistent with land rights facilitating access to non-agricultural employment. The agricultural and service-sector interactions are not statistically significant. Idle-labor responses cannot be assessed because this outcome has no variation in the estimation sample.

In Ethiopia, the agricultural interaction is negative and significant at the 1% level, with a coefficient of -1.874. The estimated agricultural response is 0.090 for households without baseline land rights and -1.784 for households with rights. Households with rights therefore experience a larger reduction in agricultural labor following favorable weather-induced productivity shocks. This pattern is consistent with secure land rights allowing households to reduce their reliance on farming without losing access to land. The manufacturing interaction is positive and significant at the 5% level, with a coefficient of 0.269, indicating a larger manufacturing response among households with rights. However, the idle-labor interaction is also positive and significant at the 5% level, with a coefficient of 1.107. These findings suggest that stronger labor release among households with rights is accompanied by increases in both manufacturing employment and idle labor. The service-sector interaction is not statistically significant.

Together, the estimates suggest that land rights shape different margins of adjustment in Malawi and Ethiopia. In Malawi, households with baseline land rights experience a larger manufacturing response, while in Ethiopia they experience a larger reduction in agricultural labor and larger increases in both manufacturing employment and idle labor. These findings are consistent with land rights easing transition frictions, although the analysis does not separately identify the tenure-security and land-income channels. The Ethiopian results also highlight that stronger labor release can generate some non-agricultural employment gains without being fully absorbed into employment, motivating an examination of labor-market absorption.

3.4.4 Labor-Market Conditions and Absorption

The final stage of the framework asks whether labor released from agriculture can be absorbed into non-agricultural employment. Following the Harris–Todaro mechanism, leaving agriculture does not guarantee employment, because workers may face limited job opportunities in the receiving sector. Favorable agricultural productivity shocks may therefore release labor without generating a corresponding increase in produc-

Table 7: Land Rights and Labor Reallocation

	(1) Agriculture	(2) Manufacturing	(3) Services	(4) Idle labor
Panel A: Niger				
Weather-predicted NDVI growth index	-0.845**	0.664***	0.207	0.003
	(0.381)	(0.197)	(0.321)	(0.046)
Weather index × baseline land rights	1.205	-1.113	0.282	-0.495
	(1.213)	(0.743)	(0.781)	(0.376)
Observations	2,648	2,648	2,648	2,648
Households	1,324	1,324	1,324	1,324
Communities	125	125	125	125
Panel B: Malawi				
Weather-predicted NDVI growth index	0.062	-0.299	0.443	—
	(0.773)	(0.323)	(0.716)	—
Weather index × baseline maximum land rights	-0.651	0.992**	-0.501	—
	(0.904)	(0.444)	(0.762)	—
Observations	1,582	1,582	1,582	1,582
Households	791	791	791	791
Communities	80	80	80	80
Panel C: Ethiopia				
Weather-predicted NDVI growth index	0.090	0.008	-0.227	0.282**
	(0.302)	(0.095)	(0.163)	(0.351)
Weather index × baseline land rights	-1.874***	0.269**	-0.038	1.107**
	(0.633)	(0.126)	(0.237)	(0.527)
Observations	2240	2240	2240	2240
Households	1120	1120	1120	1120
Communities	213	213	213	213

Notes: Standard errors are in parentheses. Land rights are measured in 2011 in Niger, 2016 in Malawi, and 2018 in Ethiopia and held fixed across waves. The baseline land-rights moderator is binary. The main weather-index coefficient is the response for households with the moderator equal to zero; the response for households with the moderator equal to one is the sum of the main and interaction coefficients. All specifications are reduced-form regressions of household labor shares on the weather-predicted NDVI growth index and its interaction with baseline land rights. Coefficients are reported in the original index units. All specifications include household and year fixed effects, using original-household fixed effects. Missing component rights indicators are coded as zero, while unclassified land types remain missing. In Malawi, idle labor is zero throughout the estimation sample; its coefficients and standard errors are therefore shown as dashes. The reported Malawi sample counts also apply to this outcome. Standard errors are clustered by household and community (original household and community in Malawi). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

tive employment, leaving some workers idle. I examine this mechanism by testing whether stronger baseline community absorption increases the non-agricultural employment response and attenuates the increase in idle labor. Specifically, I interact the weather-predicted NDVI growth index with an indicator that equals one when the baseline community leave-one-household-out absorption ratio is at or above the baseline household mean and zero otherwise. The indicator is measured in the first survey wave and held fixed across waves.

The estimates are shown in Table 8. In Niger, the manufacturing interaction is positive and significantly different from zero at the 10% level, with a coefficient of 0.557. The estimated manufacturing response is 0.145 for households in the below-average absorption group and approximately 0.702 for households in the at-or-above-average group. Consistent with stronger local labor markets facilitating employment absorption, households with higher baseline absorption experience a larger manufacturing response. The interactions for agriculture, services, and idle labor are not statistically significant.

None of the absorption interactions is statistically significant at the 10% level in Malawi or Ethiopia. The estimates therefore do not establish differences in labor-allocation responses between the two absorption groups in either country. These findings provide no clear evidence that stronger baseline absorption increases non-agricultural employment gains or reduces the idle-labor response in these countries. The baseline ratio describes initial employment conditions, but it may not fully capture the availability of additional jobs when agricultural labor is released.

I therefore examine complementary measures of regional labor-market conditions in Ethiopia to further assess the Harris–Todaro mechanism. In this mechanism, workers respond to expected earnings, which depend on both wages and the probability of obtaining employment. Limited job availability can therefore leave some workers without employment even when they leave agriculture. Specifically, I test whether stronger regional labor demand attenuates the increase in idle labor following favorable weather-induced productivity shocks. These measures capture employment opportunities beyond the household's immediate community. The analysis uses the combined rural and urban LSMS sample to examine absorption across the broader regional labor market. Because the measures and estimation sample differ from the baseline analysis, these results provide complementary evidence on the absorption mechanism.

The first measure is constructed from the World Bank Enterprise Surveys as the share of firms within each region that expand employment. Because the Enterprise Survey waves do not coincide with the LSMS waves, I match the 2014 Enterprise Survey to the 2018 LSMS wave and the 2024 Enterprise Survey to the 2021 LSMS wave. The second measure uses changes in regional unemployment rates constructed from Ethiopia's Urban Employment and Unemployment Survey (UEUS). The Enterprise Survey captures firms' employment expansion, while the UEUS provides an alternative indicator of regional employment conditions based on urban workers. A negative interaction with the share of firms expanding employment is consistent with stronger hiring attenuating the idle-labor response. A positive interaction with changes in unemployment is consistent

with worsening employment conditions increasing the idle-labor response.

The results are shown in Table 9. Using the Enterprise Survey measure, the interaction between the weather-induced productivity shock and the share of firms expanding employment is negative and significant at the 10% level in both specifications, with coefficients of -10.671 and -11.524. Consistent with the absorption mechanism, regions with a larger share of expanding firms experience a smaller increase in idle labor following favorable productivity shocks. In column (2), the additional urban interaction is also negative and significant at the 10% level, suggesting stronger attenuation among urban households.

The unemployment measure provides further, though less precise, evidence. In column (3), the interaction is positive and significant at the 5% level, with a coefficient of 0.181. Regions experiencing larger increases in unemployment therefore exhibit a larger idle-labor response following favorable productivity shocks. The corresponding interaction remains positive but is not statistically significant in column (4), and the additional urban interaction is also insignificant. The evidence is thus consistent with worsening regional employment conditions limiting absorption, although it is not statistically significant in every specification.

Although the reduced-form evidence provides patterns consistent with the proposed mechanisms, it cannot separately quantify their relative importance or predict the effects of relaxing each constraint. The next section develops a structural model of labor release, reallocation, and absorption. Under the model's behavioral and identification assumptions, the estimated parameters allow me to compare policy counterfactuals that relax different constraints.

3.5 Summary of Reduced-Form Evidence

Table 10 summarizes the main findings from the reduced-form analysis. Positive weather-induced productivity shocks consistently release labor from agriculture across all three countries, providing evidence consistent with the labor-release mechanism emphasized by Lewis (1954). However, successful structural transformation requires more than labor release alone. In Niger and Malawi, the main estimates show movement toward manufacturing employment, while in Ethiopia labor release is accompanied by a substantial increase in idle labor. These findings suggest that agricultural productivity improvements do not necessarily translate into productive non-agricultural employment.

The mechanism analysis further suggests that different constraints operate at different stages of structural transformation. In Niger and Malawi, baseline subsistence constraints weaken labor release from agriculture. In Malawi, they also attenuate the manufacturing response, while baseline land rights are associated with a larger manufacturing response. In Ethiopia, greater baseline distance is associated with larger reductions in agricultural labor and larger increases in idle labor. Baseline land rights are associated with stronger labor release and larger increases in both manufacturing employment and idle labor. These findings distinguish the ability to leave agriculture from the ability to obtain non-agricultural employment, showing that stronger

Table 8: Labor-Market Absorption and Labor Reallocation

	(1) Agriculture	(2) Manufacturing	(3) Services	(4) Idle labor
Panel A: Niger				
Weather-predicted NDVI growth index	0.242	0.145	-0.324	0.111
	(0.600)	(0.269)	(0.231)	(0.559)
Weather index × above-average baseline absorption	-1.012	0.557*	0.399	-0.085
	(0.660)	(0.308)	(0.342)	(0.562)
Observations	3,748	3,748	3,748	3,748
Households	1,874	1,874	1,874	1,874
Communities	138	138	138	138
Panel B: Malawi				
Weather-predicted NDVI growth index	-0.574	0.559	-0.644	0.293
	(1.042)	(0.491)	(0.569)	(0.357)
Weather index × above-average baseline absorption	-0.497	0.125	0.717	-0.119
	(1.139)	(0.570)	(0.672)	(0.394)
Observations	1,702	1,702	1,702	1,702
Households	851	851	851	851
Communities	73	73	73	73
Panel C: Ethiopia				
Weather-predicted NDVI growth index	-1.163*	0.148**	-0.388**	1.215*
	(0.662)	(0.063)	(0.175)	(0.673)
Weather index × above-average baseline absorption	0.116	0.074	-0.039	-0.260
	(1.015)	(0.147)	(0.350)	(0.896)
Observations	1,472	1,472	1,472	1,472
Households	736	736	736	736
Communities	133	133	133	133

Notes: Standard errors are in parentheses. Baseline absorption is constructed from the community leave-one-household-out ratio of non-agricultural employment to the sum of non-agricultural employment and idle labor. It is measured in 2011 in Niger, 2016 in Malawi, and 2018 in Ethiopia and held fixed across waves. The moderator equals one when the baseline ratio is at or above the baseline household mean and zero otherwise. The main coefficient is the response for the below-average group; the response for the at-or-above-average group is the sum of the main and interaction coefficients. All specifications are reduced-form regressions of household labor shares on the weather-predicted NDVI growth index and its interaction with baseline absorption. Coefficients are reported in the original index units. All specifications include household and year fixed effects, using linked-household fixed effects in Malawi. No moderator-by-year interactions are included. The Ethiopia sample is rural. Standard errors are clustered by household and community (linked household and community in Malawi). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Regional Labor-Market Conditions and Idle Labor in Ethiopia

Dependent variable: household idle-labor share	(1) WBES	(2) WBES	(3) UEUS	(4) UEUS
Weather-induced productivity shock	7.585**	8.980**	0.669**	0.670*
	(1.97)	(2.10)	(1.97)	(1.96)
Weather shock × share of firms expanding employment	-10.671*	-11.524*		
	(-1.73)	(-1.69)		
Weather shock × share of firms expanding employment × urban indicator		-2.504*		
		(-1.75)		
Weather shock × change in unemployment rate			0.181**	0.161
			(2.49)	(1.52)
Weather shock × change in unemployment rate × urban indicator				0.045
				(0.31)
Household fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	4,456	4,456	5,152	5,152
Communities	290	290	360	360
Within R^2	0.041	0.051	0.045	0.046

Notes: t-statistics are in parentheses. Columns (1) and (2) use the regional share of firms expanding employment from the World Bank Enterprise Surveys (WBES). The 2014 Enterprise Survey is matched to the 2018 LSMS wave, and the 2024 Enterprise Survey to the 2021 LSMS wave. Columns (3) and (4) use changes in regional unemployment rates constructed from the 2016, 2018, and 2020 Urban Employment and Unemployment Surveys (UEUS). Columns (2) and (4) additionally report interactions with an urban indicator. The sample includes rural and urban LSMS households. The regional measures and sample differ from those in Table 8; these results provide complementary, descriptive evidence. Standard errors are clustered at the LSMS community level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

labor release can generate employment gains while also leaving some workers idle.

The absorption analysis provides additional evidence on this distinction. In Niger, higher baseline community absorption is associated with a larger manufacturing response. The corresponding interactions are not statistically significant in Malawi or Ethiopia. Complementary evidence from the combined rural and urban Ethiopian sample shows that stronger regional hiring is associated with a smaller idle-labor response, while rising unemployment is associated with a larger response in one specification. These patterns are consistent with the Harris–Todaro absorption mechanism, although the regional interactions are descriptive and the evidence varies across specifications. Together, the findings suggest that a common policy prescription is unlikely to address the constraints observed in all three countries.

Although the reduced-form evidence provides patterns consistent with the proposed mechanisms, it does not establish the causal effects of the baseline constraints themselves or quantify their relative importance. The next section therefore develops a structural model incorporating labor release, labor reallocation, and labor absorption, and uses the estimated model to evaluate the effects of relaxing these constraints through alternative policy interventions.

Table 10: Summary of Reduced-Form Evidence Across the Three Stages of Structural Transformation

Stage	Finding	Niger	Malawi	Ethiopia
Stage I: Labor Release	Favorable productivity shocks reduce agricultural labor	✓	✓	✓
	Subsistence constraints weaken agricultural labor release	✓	✓	—
Stage II: Labor Reallocation	Main destination of labor released from agriculture	Manufacturing	Manufacturing	Idle labor
	Baseline distance moderates labor-allocation responses	—	—	✓
	Baseline land rights moderate labor-allocation responses	—	✓	✓
Stage III: Labor Absorption	Higher baseline community absorption strengthens employment responses	✓	—	—
	Stronger regional hiring attenuates the idle-labor response	N/A	N/A	✓
	Rising regional unemployment amplifies the idle-labor response	N/A	N/A	✓

Notes: The table summarizes the main reduced-form findings. ✓ indicates a statistically significant response in the stated direction at the 10% level or better. — indicates that the relevant interactions are not statistically significant at the 10% level; it does not establish the absence of the mechanism. The destination row summarizes sectoral labor-share responses rather than directly tracked worker transitions. Baseline moderators are measured in the first survey wave and held fixed across waves. Regional hiring and unemployment results use the combined rural and urban Ethiopian sample and provide complementary, descriptive evidence. N/A indicates that comparable regional labor-market measures were not evaluated for that country.

4 Structural Model

4.1 Overview

The reduced-form evidence provides patterns consistent with constraints on labor release, labor transition, and labor absorption, but it does not quantify how these constraints individually and jointly shape household labor allocation or evaluate the effects of relaxing them. To address these questions, I develop a structural model that incorporates all three stages within a common framework. The model allows subsistence needs, transition frictions, and employment opportunities to jointly determine household labor allocation. I

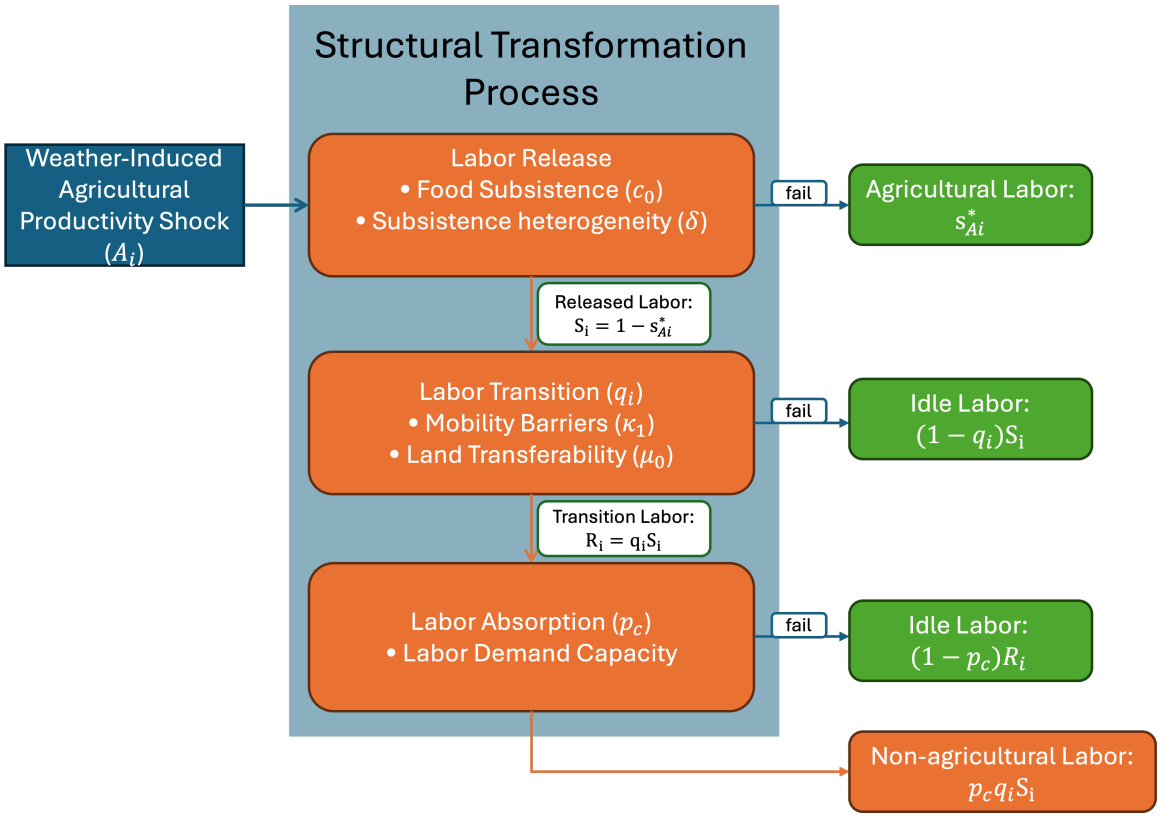


Figure 6: Sequential Structure of the Structural Transformation Model

use the estimated model to evaluate both individual policies and combined interventions, examining which constraints policies should target and whether addressing multiple constraints together generates additional employment gains.

Figure 6 illustrates the three-stage framework. In the first stage, agricultural productivity and household subsistence needs determine how much labor can be released from agriculture. Higher productivity allows households to meet their food requirements with less agricultural labor, capturing the Lewis-type labor-release mechanism. In the second stage, spatial mobility costs and land institutions shape whether released labor can transition into non-agricultural activities. In the third stage, employment opportunities determine whether workers entering the non-agricultural labor market obtain jobs or remain idle. This stage captures the employment-probability channel emphasized by Harris and Todaro (1970). Together, the three stages distinguish the ability to leave agriculture from the ability to enter and obtain non-agricultural employment.

4.2 Environment

The economy consists of households indexed by i. Each household is endowed with one unit of labor, which can be allocated across three activities: agricultural production, non-agricultural employment, and idle labor. Let

$$s_{Ai} + s_{Mi} + s_{Ni} = 1, \tag{6}$$

where s_{Ai} denotes the share of household labor allocated to agriculture, s_{Mi} denotes the share allocated to non-agricultural employment, and s_{Ni} denotes the share that remains idle. Idle labor represents household members who are not engaged in either agricultural or non-agricultural economic activities.

Households are heterogeneous in their agricultural productivity, household size, land institutions, and geographic location. Agricultural productivity is determined by a weather-induced productivity shock,

$$A_i = A(\theta_i), \tag{7}$$

where θ_i denotes the weather-induced agricultural productivity measure constructed in Section 2. A higher value of θ_i increases agricultural productivity and relaxes the household's production constraint.

The timing of the model follows the agricultural production cycle. At the beginning of each period, households observe the realized weather-induced productivity shock. They then determine the amount of labor required for agricultural production, decide whether released labor can access non-agricultural opportunities, and finally participate in the local labor market. Consequently, the final allocation of labor across agriculture, non-agricultural employment, and idle activities is determined sequentially through the three stages illustrated in Figure 6.

4.3 Stage I: Labor Release

The first stage of the model determines how much labor households can release from agriculture following a weather-induced productivity shock. Following Lewis (1954), households must first allocate sufficient labor to agriculture to satisfy their food subsistence requirement before supplying labor to non-agricultural activities. Improvements in agricultural productivity increase output for a given labor input, allowing households to meet their food needs with fewer agricultural workers and thereby releasing labor for non-agricultural activities.

Agricultural production is given by

$$y_{Ai} = A_i s_{Ai}^{\gamma}, \tag{8}$$

where A_i denotes weather-induced agricultural productivity and $\gamma \in (0,1)$ captures diminishing marginal returns to agricultural labor.

Each household faces a minimum food subsistence requirement,

$$\bar{c}_i = c_0(1+\delta Z_i)h_i, \tag{9}$$

where h_i denotes household size, Z_i is an indicator for whether the household is food-subsistence constrained, c_0 captures the baseline subsistence requirement, and δ measures the additional food requirement faced by

subsistence-constrained households.[4] When $\delta = 0$, all households face the same subsistence requirement regardless of food-security status.

Households choose agricultural labor to maximize utility from agricultural consumption and labor available for non-agricultural activities:

$$\max_{s_{Ai} \in [0,1)} \left\{ \log\left(A_i s_{Ai}^{\gamma} - \bar{c}_i\right) + \beta \log(1 - s_{Ai}) \right\}, \tag{10}$$

subject to

$$A_i s_{Ai}^{\gamma} \geq \bar{c}_i. \tag{11}$$

The first term represents agricultural consumption net of the household's subsistence requirement, while the second term captures the value of labor available for non-agricultural activities. The parameter $\beta > 0$ governs the relative value households place on labor released from agriculture.

The subsistence requirement implies a minimum feasible agricultural labor allocation,

$$\underline{s}_{Ai} = \left(\frac{\bar{c}_i}{A_i}\right)^{1/\gamma}. \tag{12}$$

For an interior solution, the optimal allocation satisfies the first-order condition

$$\frac{A_i \gamma s_{Ai}^{\gamma-1}}{A_i s_{Ai}^{\gamma} - \bar{c}_i} = \frac{\beta}{1 - s_{Ai}}. \tag{13}$$

Let s_{Ai}^{int} denote the corresponding interior solution. The household's optimal agricultural labor allocation is therefore

$$s_{Ai}^{*} = \max\left\{1, \max\left\{\underline{s}_{Ai}, s_{Ai}^{int}\right\}\right\}. \tag{14}$$

If the subsistence requirement cannot be met even when the household allocates its entire labor endowment to agriculture, the household reaches a corner solution with $s_{Ai}^{*} = 1$. In this case, no labor is released from agriculture ($S_i = 0$)

The amount of labor released from agriculture is

$$S_i = 1 - s_{Ai}^{*}, \tag{15}$$

where S_i denotes the labor available for non-agricultural activities after satisfying the household's food subsistence requirement.

Equation (15) summarizes the labor-release stage of the model. A positive weather-induced productivity

[4]The parameter δ allows subsistence requirements to differ across households by capturing the additional food requirement faced by subsistence-constrained households. This heterogeneity reflects that households experiencing more severe food insecurity must allocate more labor to agriculture before participating in non-agricultural activities.

shock raises agricultural productivity A_i, reducing the minimum labor required to satisfy subsistence needs and increasing the amount of labor released from agriculture. Conversely, adverse weather shocks reduce agricultural productivity, requiring households to retain more labor in agriculture and leaving less labor available for structural transformation.

4.4 Stage II: Labor Reallocation

The second stage of the model determines whether labor released from agriculture can successfully access non-agricultural opportunities. Releasing labor from agriculture does not automatically imply successful structural transformation. Even after agricultural productivity has increased and households have labor available to leave farming, workers may face barriers that prevent them from entering non-agricultural activities. These barriers arise from both geographic frictions and institutional constraints, which jointly determine the opportunity cost of leaving agriculture.

Let $q_i \in [0,1]$ denote the probability that labor released from agriculture successfully reaches the non-agricultural labor market. This probability captures households' ability to overcome reallocation barriers before employment opportunities can be realized. The reallocation probability is specified as

$$q_i = \Lambda\left[\mu_0^{NT}(1-T_i) + \mu_0^{T} T_i + \kappa_1 dist_i\right], \tag{16}$$

where $\Lambda(\cdot)$ denotes the logistic cumulative distribution function, T_i is an indicator equal to one if the household processes stronger land rights and zero otherwise. For the structural estimation, the three-category land-rights measure used in the reduced-form analysis is collapsed into this binary indicator, with land Type 1 and 2 classified as strong land rights ($T_i = 1$) and land Type 0 classified as weak land rights ($T_i = 0$). The variable $dist_i$ measures the household's mobility barrier, proxied by the distance to the nearest economic center.

The parameters μ_0^{NT} and μ_0^{T} capture the baseline probability of accessing non-agricultural opportunities under weak land-rights ($T_i = 0$) and strong land-rights ($T_i = 1$) institutions, respectively. The parameter κ_1 measures the effect of spatial mobility barriers on labor reallocation. A larger value of μ_0^{T} relative to μ_0^{NT} implies that stronger land rights lowers the opportunity cost of leaving agriculture and facilitates labor reallocation, while a more negative value of κ_1 implies that greater mobility barriers reduce access to non-agricultural opportunities.

Conditional on labor being released from agriculture, the amount of labor successfully reallocated to the non-agricultural labor market is

$$R_i = q_i S_i \tag{17}$$

where S_i denotes labor released from agriculture in Stage I. The variable R_i summarizes the amount of labor that has successfully overcome reallocation barriers and therefore becomes available for non-agricultural employment. Whether this reallocated labor is ultimately employed or remains idle is determined by local labor demand in the next stage of the model.

4.5 Stage III: Labor Absorption

The final stage of the model determines whether workers who have successfully accessed non-agricultural opportunities obtain productive employment. Access to the labor market does not guarantee employment. Following Harris and Todaro (1970), workers compete for a limited number of non-agricultural jobs, and the availability of these jobs depends on local labor demand. Consequently, even after overcoming reallocation barriers, some workers may remain outside productive employment if labor demand is insufficient.

Let $p_i \in [0,1]$ denote the probability that workers entering the non-agricultural labor market are successfully absorbed into productive employment. The parameter p_i summarizes the labor absorption capacity of the local economy.

Conditional on reaching the non-agricultural labor market, the amount of labor employed outside agriculture is

$$s_{Mi} = p_i R_i \tag{18}$$

where $R_i = q_i S_i$ is the amount of labor successfully reallocated in Stage II.

Idle labor consists of two groups of workers: those who fail to access non-agricultural opportunities and those who access the labor market but fail to obtain employment. Therefore,

$$s_{Ni} = (S_i - R_i) + (1 - p_i)R_i \tag{19}$$

The model interprets idle labor as the component of labor allocation generated by unsuccessful transitions during the structural transformation process. It captures workers who remain idle because they are unable to reach productive non-agricultural opportunities or fail to obtain employment after entering the labor market. The model is not intended to explain all sources of observed idle labor, such as persistent baseline non-participation unrelated to agricultural productivity shocks.

which can be equivalently written as

$$s_{Ni} = (1 - p_i q_i)S_i \tag{20}$$

For notational convenience, define

$$\pi_i = p_i q_i \tag{21}$$

where π_i denotes the overall probability that labor released from agriculture ultimately enters productive non-agricultural employment. The employment and idle labor allocations can therefore be written as

$$s_{Mi} = \pi_i S_i \tag{22}$$

And

$$s_{Ni} = (1 - \pi_i) S_i \tag{23}$$

Equations (18)–(21) summarize the labor-absorption stage of the model. A higher labor-absorption probability increases the share of reallocated workers entering productive non-agricultural employment, while a lower absorption probability results in a larger share remaining idle following labor release. Thus, successful structural transformation depends not only on releasing labor from agriculture or enabling labor reallocation, but also on the capacity of local labor markets to absorb workers into productive employment.

In the theoretical framework, the absorption probability p_i is allowed to vary across households and locations to reflect heterogeneity in local labor demand. For empirical tractability, however, the structural estimation assumes a common within-country absorption probability ($p_i = p$), so that the estimated parameter captures the average labor absorption capacity of the national economy while abstracting from within-country regional variation.

4.6 Model Implication

The model characterizes structural transformation as a sequential process in which households move through three distinct stages. Weather-induced agricultural productivity first determines the amount of labor released from agriculture through the household's subsistence constraint. Released labor must then overcome reallocation barriers to access non-agricultural opportunities. Finally, workers who successfully reach the non-agricultural labor market are either absorbed into productive employment or remain idle, depending on local labor demand.

Taken together, the three stages determine the household labor allocation implied by the model:

$$S_i = 1 - s_{Ai}^*, \tag{24}$$

$$R_i = q_i S_i, \tag{25}$$

$$s_{Mi} = p_i R_i, \tag{26}$$

$$s_{Ni} = (S_i - R_i) + (1 - p_i) R_i. \tag{27}$$

The model contains four groups of structural parameters corresponding to the three stages of structural transformation. The first-stage parameters, c_0 and δ, govern households' subsistence requirements and therefore determine the amount of labor released from agriculture following a productivity shock. The second-stage

parameters, μ_0^{NT}, μ_0^T, and κ_1, determine the probability that released labor can access non-agricultural opportunities by capturing the effects of land institutions and spatial mobility barriers. The third-stage parameter, p, captures the probability that workers who reach the non-agricultural labor market are successfully absorbed into productive employment.

These structural parameters are not directly observed in the data. Instead, they are inferred by matching the model-implied labor allocation generated by the three-stage framework to the corresponding labor allocation observed in the household survey data. The next section estimates these parameters using the simulated method of moments and then uses the estimated model to evaluate the effects of alternative policy interventions on structural transformation.

5 Structural Estimation

5.1 Parameterization

The structural model developed in the previous section contains two types of parameters: calibrated parameters that are fixed prior to estimation and structural parameters that are estimated from the data. The calibrated parameters determine the agricultural production technology and the mapping from weather-induced productivity shocks to agricultural productivity. Specifically, agricultural productivity is constructed as

$$A_i = \exp(a_0 + a_1 \Delta\theta_{gt}), \tag{28}$$

where the calibration parameters a_0 and a_1 are obtained from the reduced-form relationship between household agricultural income and weather-induced productivity growth reported in Appendix. The remaining production parameters, including the agricultural labor elasticity $\gamma = 0.5$, and the utility parameter $\beta = 0.95$ appearing in the household optimization problem, are calibrated prior to estimation.

Conditional on these calibrated parameters, the structural estimation focuses on recovering the behavioral parameters governing the three stages of structural transformation. These parameters summarize households' subsistence requirements, the barriers to labor reallocation, and the ability of local labor markets to absorb workers into productive employment. While the model specifies how these parameters jointly determine household labor allocation, they are not directly observed in the data and must be estimated.

Table 11 summarizes the structural parameters estimated in the model and their corresponding economic interpretations. Consistent with the three-stage framework, each group of parameters governs a distinct stage of structural transformation.

Table 11: Structural Parameters

Stage	Parameter	Economic Interpretation
Labor Release	c_0	Baseline subsistence requirement
	δ	Additional subsistence requirement for food-constrained households
Labor Reallocation	μ_0^{NT}	Baseline access to non-agricultural opportunities under weaker land rights
	μ_0^{T}	Baseline access to non-agricultural opportunities under stronger land rights
	κ_1	Effect of spatial mobility barriers on labor reallocation
Labor Absorption	p	Probability that reallocated workers obtain productive employment

5.2 Estimation Procedure

I estimate the structural parameters using simulated method of moments (SMM). For a candidate parameter vector, the model is solved for each household using its observed characteristics to predict labor allocation across agriculture, productive non-agricultural employment, and idle labor. These household-level predictions are then aggregated into a set of simulated moments that are directly comparable to the corresponding moments computed from the household survey. The structural parameters are chosen to minimize the distance between the simulated and empirical moments, thereby selecting the parameter values that best reproduce the observed labor-allocation patterns.

The moment conditions are organized according to the three stages of the model (see Table 12), allowing each stage to be disciplined by moments directly related to its underlying economic mechanism. Conditional on the fixed preference and technology parameters β and γ, the labor-release parameters, c_0 and δ, are identified using agricultural labor-share moments, including the overall agricultural labor share, agricultural labor shares across four climate-productivity-shock bins, and agricultural labor shares by subsistence status. These moments identify the amount of labor released from agriculture in response to productivity changes.

The labor-reallocation parameters, μ_0^{NT}, μ_0^{T}, and κ_1, are identified from transition-probability moments across land-transferability regimes and distance bins. These moments exploit systematic variation in households' ability to access non-agricultural opportunities arising from land institutions and spatial mobility barriers, thereby identifying heterogeneity in the household-specific reallocation probability, q_i.

Finally, the labor-absorption parameter, p, is identified from the aggregate allocation of released labor between productive non-agricultural employment and underutilized labor. Conditional on the estimated distribution of household-specific reallocation probabilities, the remaining split between successful non-agricultural employment and underutilized labor disciplines the common labor-absorption probability. Thus, labor reallocation is identified from cross-household heterogeneity in structural barriers, whereas labor absorption is identified from the aggregate employment allocation.

The moment conditions are organized according to the three stages of the model (see Table 12). The labor-release parameters, c_0 and δ, are identified using agricultural labor-share moments, including the overall agricultural labor share, agricultural labor shares across four climate-productivity-shock bins, and agricul-

tural labor shares by subsistence status. The labor-reallocation parameters, μ_0^{NT}, μ_0^{T}, and κ_1, are identified using transition-probability moments, which compare the model-implied probability that released labor enters productive non-agricultural employment, $\pi_i = p_i q_i$, with the corresponding transition probabilities observed across land regimes, four distance bins, and their interactions. Finally, the labor-absorption parameter, p, is identified using moments based on the allocation of released workers between productive non-agricultural employment and idle labor.

Let $\Theta = \left(c_0, \delta, \mu_0^{NT}, \mu_0^{T}, \kappa_1, p\right)$ denote the vector of structural parameters to be estimated. For each candidate parameter vector Θ, the model generates a vector of simulated moments $m(\Theta)$. Let $\hat{m}$ denote the corresponding empirical moments computed from the household survey. The SMM estimator solves

$$\hat{\Theta} = \arg\min_{\Theta} \left[m(\Theta) - \hat{m}\right]' W \left[m(\Theta) - \hat{m}\right],$$

where W is a weighting matrix. In the baseline specification, I use an identity weighting matrix so that the estimator minimizes the unweighted squared distance between the model-implied and empirical moments.

Although productive non-agricultural employment depends jointly on labor reallocation and labor absorption, the two mechanisms are identified using different sources of variation. Cross-household heterogeneity in distance and household land-rights identifies the household-specific reallocation probability through observed differences in transition behavior. Conditional on this estimated heterogeneity, the remaining aggregate allocation between productive non-agricultural employment and idle labor identifies the common labor-absorption probability. This separation therefore relies jointly on the transition moments, the labor-allocation moments, and the maintained structural assumptions governing the reallocation and absorption processes.

The estimation proceeds in two steps. In the first step, I estimate the labor-release parameters, c_0 and δ, using only the agricultural labor-share moments because these parameters govern households' decisions to release labor from agriculture. Holding these estimates fixed, I jointly estimate the labor-reallocation and labor-absorption parameters, μ_0^{NT}, μ_0^{T}, κ_1, and p, by matching the transition-probability moments together with the observed allocation of released labor between productive non-agricultural employment and idle labor. This sequential procedure follows the economic structure of the model, allowing the labor-release mechanism to be identified independently before estimating the subsequent reallocation and absorption stages.

Inference is based on a household-cluster bootstrap. For each bootstrap replication, households are resampled with replacement, the complete estimation procedure is repeated, and the resulting parameter estimates are recorded. Standard errors are calculated from the bootstrap distribution, and confidence intervals are constructed using bootstrap percentiles. The baseline results are based on 100 bootstrap replications.

Table 12: Moments Used in Simulated Method of Moments Estimation

Model Stage	Parameters	Primary Identifying Moments	Identification
Labor Release	$c_0,\ \delta$	$E[s_A]$, $E[s_A \mid \text{climate bin}]$, $E[s_A \mid D]$	Variation in agricultural labor shares identifies the amount of labor released from agriculture under different productivity and subsistence conditions.
Labor Reallocation	$\mu_0^{NT},\ \mu_0^{T},\ \kappa_1$	$E[\pi^{pred} - \pi^{obs} \mid \text{distance bin}]$, $E[\pi^{pred} - \pi^{obs} \mid T]$, $E[\pi^{pred} - \pi^{obs} \mid T, \text{distance bin}]$	Cross-household variation in distance and household land-rights identifies heterogeneity in the household-specific reallocation probability q_i.
Labor Absorption	p	$E[s_M^{pred} - s_M^{obs}]$, $E[s_N^{pred} - s_N^{obs}]$	Conditional on the estimated distribution of q_i, the aggregate allocation between productive non-agricultural employment and underutilized labor identifies the common labor-absorption probability.

a. s_A, s_M, and s_N denote agricultural labor, productive non-agricultural employment, and underutilized labor shares, respectively.
b. Climate-productivity bins and distance bins are constructed using quartiles of the corresponding household-level distributions.
c. D denotes household subsistence status. $T = 1$ indicates stronger land rights and $T = 0$ indicates weaker land rights.
d. $\pi_i = pq_i$ denotes the probability that labor released from agriculture successfully enters productive non-agricultural employment, where q_i is household-specific and varies with household land-rights and distance, while p is a country-specific labor-absorption probability.
e. Aggregate labor-allocation moments are primarily informative about π_i. Variation across land-transferability regimes and distance bins disciplines the household-specific reallocation probability q_i, while the remaining aggregate split between productive non-agricultural employment and underutilized labor identifies p.

5.3 Structural Parameter Estimates

Table 13 reports the estimated structural parameters for Niger, Malawi, and Ethiopia. The estimates quantify the mechanisms governing the three stages of structural transformation and complement the reduced-form evidence presented in Section 3 by jointly estimating the relative importance of multiple frictions within a unified structural framework.

The estimated labor-release parameters indicate substantial cross-country differences in the importance of subsistence constraints. Malawi exhibits the largest baseline subsistence requirement ($c_0 = 1.125$), followed by Niger ($c_0 = 0.895$), while Ethiopia has the lowest value ($c_0 = 0.343$). The estimated food subsistence heterogeneity parameter, δ, is positive and statistically significant in Niger and Malawi, implying that households below the food subsistence threshold require substantially more agricultural labor satisfy minimum consumption needs. By contrast, the estimate of δ is statistically insignificant in Ethiopia, suggesting that food subsistence constraints play a much smaller role in determining labor release.

The estimated labor-reallocation parameters also differ markedly across countries. The migration-cost parameter, κ_1, is negative and statistically significant in Niger and Ethiopia, indicating that greater distance from economic centers reduces households' access to non-agricultural opportunities. However, as shown in Panel B, the implied probability effect is substantially larger in Ethiopia than in Niger, indicating that mobility costs constitute a quantitatively more important barrier to labor reallocation in Ethiopia. In contrast, the migration-cost parameter is close to zero and statistically insignificant in Malawi, implying that spatial

mobility barriers are relatively weak.

The latent land effect is positive and statistically significant in Malawi and Ethiopia but insignificant in Niger, suggesting that household land rights substantially improves households' ability to access non-agricultural opportunities only in the former two countries. This pattern is consistent with institutional differences across countries. Ethiopia's state ownership of rural land and restrictions on stronger land rights increase the importance of secure and stronger land rights for households considering non-agricultural employment. Malawi also exhibits a statistically significant household land-rights effect, although its magnitude is considerably smaller than in Ethiopia. By contrast, the insignificant land effect in Niger suggests that the household land-rights measure used here does not substantially explain heterogeneity in labor allocation. This does not necessarily imply that broader land institutions are unimportant, but rather that variation in inheritance and collateral rights appears less consequential for household labor responses in Niger sample.

Because the latent land effect is expressed on the logistic index, Panel B reports the corresponding average effects on economically meaningful probabilities. Stronger household land rights increases the probability of accessing non-agricultural opportunities by approximately 2.3 percentage points in Malawi and 29.6 percentage points in Ethiopia, while having essentially no effect in Niger. Although the structural model identifies statistically significant household land-rights effects in both Malawi and Ethiopia, the much larger probability effect in Ethiopia is consistent with the reduced-form evidence that downstream institutional barriers are considerably more important there.

Finally, the estimated labor-absorption probability highlights substantial differences in local labor demand across countries. Niger and Malawi exhibit nearly complete labor absorption, with estimated probabilities exceeding 0.97, indicating that workers who successfully access non-agricultural opportunities are almost always absorbed into productive employment. Ethiopia, on the other hand, has an estimated absorption probability of only 0.735, implying that roughly one-quarter of workers who successfully reach the non-agricultural labor market remain idle rather than obtaining productive employment.

Overall, the structural estimates closely mirror the reduced-form evidence. Structural transformation in Niger and Malawi is primarily constrained by labor released from agricultural through subsistence requirements, whereas structural transformation in Ethiopia is constrained by both labor reallocation and labor absorption. These estimated mechanisms provide the foundation for the policy counterfactual analysis in the next section.

6 Policy Counterfactual Analysis

The structural estimates quantify the key mechanisms governing labor release, labor reallocation, and labor absorption. However, parameter estimates alone do not indicate which policies are most effective at accelerating structural transformation. The structural model provides a natural framework for answering this question

Table 13: Structural Parameter Estimates and Model-Implied Land Effects

	Niger	Malawi	Ethiopia
Panel A. Structural Parameters			
Subsistence parameter c_0	0.895***	1.125***	0.343***
	(0.030)	(0.094)	(0.019)
Subsistence heterogeneity δ	0.128**	1.230***	-0.105
	(0.055)	(0.217)	(0.171)
Migration cost κ_1	-1.879***	0.010	-0.774***
	(0.036)	(0.185)	(0.139)
Latent land effect $(\mu_0^T - \mu_0^{NT})$	-0.349	2.487***	1.642***
	(0.206)	(0.449)	(0.335)
Absorption probability p	0.973***	0.999***	0.735***
	(0.003)	(0.0004)	(0.034)
Panel B. Model-Implied Land Effects			
$\Delta q = E[q\|T=1] - E[q\|T=0]$	-0.028	0.023	0.296
$\Delta\pi = E[\pi\|T=1] - E[\pi\|T=0]$	-0.027	0.023	0.218

Bootstrap standard errors (100 household-cluster bootstrap replications) are reported in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.
The latent land effect $(\mu_0^T - \mu_0^{NT})$ measures the difference in the transition index between strong and weak land rights. Panel B reports the corresponding average effects on the probability of accessing non-agricultural opportunities (q) and the overall transition probability into productive non-agricultural employment ($\pi = pq$), which are more directly interpretable.
$T = 1$ denotes stronger land rights and $T = 0$ denotes weaker land rights.

because it allows individual mechanisms to be varied while holding all other features of the economy fixed.

The objective of the counterfactual analysis is to quantify the relative importance of the structural frictions identified by the model. Accordingly, each policy experiment perturbs one structural parameter while holding the remaining mechanisms fixed. Therefore, these counterfactuals should be interpreted as mechanism-specific comparative statics rather than comprehensive policy evaluation.

Using the estimated parameters, I conduct a series of counterfactual policy experiments corresponding to the three stages of structural transformation. Specifically, I examine policies that relax subsistence constraints, improve agricultural productivity, reduce mobility barriers, strengthen household land-rights, and increase absorption in the non-agricultural labor market. Comparing the simulated labor allocations across these counterfactual scenarios reveals both the relative importance of each mechanism and the policy interventions most likely to promote structural transformation in each country.

6.1 Counterfactual Design

The estimated structural model provides a framework for evaluating how alternative policies affect structural transformation. I conduct counterfactual simulations by perturbing one policy-relevant structural parameter at a time while holding all other estimated parameters fixed. This allows the simulations to isolate how each intervention affects household labor allocation through labor release and labor reallocation.

To compare policies across countries, I calibrate each intervention using a common fiscal budget of USD 500 million per country. The calibration translates the same policy budget into changes in the structural model

through three steps. First, I define the target population for each policy. Second, I convert the budget into either a transfer per eligible household, the share of households reached, or the share of an implementation gap closed. Third, I map this policy intensity into the corresponding structural parameter. Detailed country-specific calculations are reported in Appendix A.

Policy 1: Relaxing Subsistence Constraints. The first policy provides targeted food assistance to working-age individuals in rural households below the food-subsistence threshold. Because the policy directly relaxes minimum consumption requirements, it is mapped into the subsistence component of the model.

Let B denote the common policy budget. Let POP_c denote the national population in country c, and let $\hat{s}_c^T$ denote the share of the baseline population that is simultaneously working age, rural, and living in a household below the food-subsistence threshold, estimated from the nationally representative LSMS survey. This joint population share differs from the household-level food-subsistence share reported in Table 2 because it is calculated over the baseline population rather than over households.

The number of targeted individuals is

$$N_c^T = \hat{s}_c^T \times POP_c.$$

The transfer per targeted individual is

$$T_c^F = \frac{B}{N_c^T}. \tag{29}$$

I express this transfer relative to the annual food-subsistence benchmark F_c, measured on a per-person basis using the national food poverty line reported by the World Bank. The resulting policy intensity is

$$\lambda_c^F = \min\left\{1, \frac{T_c^F}{F_c}\right\}.$$

The policy is implemented by proportionally reducing the additional subsistence requirement for households below the food-subsistence threshold:

$$\bar{c}_i' = c_0\left[1 + (1-\lambda_c^F)\delta Z_i\right] h_i, \tag{30}$$

where $Z_i = 1$ for households below the food-subsistence threshold and $Z_i = 0$ otherwise.

Policy 2: Improving Agricultural Productivity. The second policy represents agricultural input programs, such as fertilizer subsidies and improved seed distribution. These policies increase agricultural productivity and allow households to satisfy subsistence needs with less agricultural labor.

Let C_A denote the cost of providing an agricultural input package to one household, and let N_c^R denote the

number of rural households in country c.[5] The share of households reached by the policy is

$$s_c^A = \min\left\{1, \frac{B}{C_A N_c^R}\right\}. \tag{31}$$

If treated households experience a productivity gain g_A, the average productivity increase is

$$\lambda_c^A = s_c^A g_A. \tag{32}$$

The policy is implemented as a proportional increase in agricultural productivity:

$$A_i' = (1 + \lambda_c^A) A_i. \tag{33}$$

Policy 3: Reducing Mobility Barriers. The third policy represents investments in rural roads and transport infrastructure. Improved transportation lowers the effective mobility barriers faced by rural households and facilitates access to non-agricultural opportunities.

Let C_R denote the country-specific construction cost per kilometer of rural roads, obtained from the World Bank's Road Costs Knowledge System (ROCKS). Under the common policy budget B, the number of kilometers financed is

$$K_c^R = \frac{B}{C_R}. \tag{34}$$

Let $\bar{K}_c^R$ denote the country-specific rural road gap, measured as the additional kilometers of rural roads required to achieve full rural accessibility using country-specific World Bank transport assessments. The share of the road gap closed by the policy is

$$s_c^R = \min\left\{1, \frac{K_c^R}{\bar{K}_c^R}\right\}. \tag{35}$$

Following the transport literature, I assume that completely closing the road gap removes a fraction $\bar{\rho}^R$ of the estimated mobility barrier.[6] The resulting policy intensity is

$$\lambda_c^R = s_c^R \bar{\rho}^R. \tag{36}$$

The policy is implemented as a proportional reduction in the migration-cost parameter:

$$\kappa_{1c}' = (1 - \lambda_c^R)\kappa_{1c}. \tag{37}$$

Policy 4: Strengthening household land-rights. The fourth policy represents land certification and tenure reforms that strengthen land rights. These reforms reduce institutional barriers to labor reallocation by al-

[5] Similar to the N_c^T, I approximate the number of rural households using national population data and the rural share observed in the household survey.

[6] In the benchmark calibration, we set $\bar{\rho}^R = 0.5$, implying that completely closing the rural road gap removes one-half of the estimated mobility barrier. This assumption reflects that improvements in road infrastructure reduce, but do not eliminate, all barriers to labor mobility, such as information frictions, relocation costs, and social constraints.

lowing households to retain or transfer the value of agricultural land while participating in non-agricultural activities.

Let C_L denote the cost of providing land certification or tenure reform to one household, and let N_c^{NT} denote the number of households with weaker land rights in country c.[7] Under the common policy budget B, the share of weak land-rights households covered by the reform is

$$\lambda_c^L = \min\left\{1, \frac{B}{C_L N_c^{NT}}\right\}. \tag{38}$$

Because the estimated implementation cost of land certification is relatively low, the common policy budget is sufficient to certify all households with weaker land rights in the three study countries.

The policy is implemented by randomly converting a fraction λ_c^L of households from weaker land rights ($T_i = 0$) to stronger land rights ($T_i = 1$):

$$T_i' = \begin{cases} 1, & \text{if household } i \text{ is selected for certification and } T_i = 0, \\ T_i, & \text{otherwise.} \end{cases} \tag{39}$$

The reallocation probability is then recalculated as

$$q_i' = \Lambda\left[\mu_0^{NT}(1 - T_i') + \mu_0^T T_i' + \kappa_1 dist_i\right]. \tag{40}$$

Instead of altering the estimated structural parameters, the policy changes households' institutional environment by expanding the share of households with stronger land rights rights. The estimated reallocation function is then re-evaluated using the updated land status, increasing households' probability of accessing non-agricultural opportunities.

Policy 5: Increasing Labor Absorption Capacity. The fifth policy represents industrial investment that expands the ability of the non-agricultural sector to absorb workers. This policy targets the labor-absorption stage of the model by increasing the probability that workers who successfully access the non-agricultural labor market obtain productive employment.

Because the structural estimates indicate that labor-demand constraints are primarily binding in Ethiopia, I calibrate this policy using evidence from Ethiopia's industrial park investments. The World Bank Competitiveness and Job Creation Project supported the Bole Lemi II and Kilinto industrial parks and generated approximately 19,000 direct jobs with total financing of USD 425 million. Scaling this benchmark to the common policy budget of USD 500 million implies approximately 22,353 new jobs. Because both industrial parks are located in Addis Ababa, I express the job creation effect relative to the unemployed population in

[7] I estimate N_c^{NT} by multiplying the total number of households by the share of households with weaker land rights observed in the nationally representative LSMS survey.

Addis Ababa reported in the 2020 Urban Employment and Unemployment Survey.

Let J denote the number of jobs created by the policy and let U_{AA} denote the unemployed population in Addis Ababa. The implied increase in labor absorption capacity is

$$\lambda^D = \frac{J}{U_{AA}}. \tag{41}$$

The policy is implemented as an increase in the labor absorption probability:

$$p' = \min\left\{1,\ p + \lambda^D\right\}. \tag{42}$$

For Niger and Malawi, I retain the estimated labor absorption probabilities because the structural estimates indicate that labor-demand constraints are already weak in these two countries. Consequently, the labor absorption policy is evaluated only for Ethiopia, where it represents a realistic industrial investment designed to relax the binding demand-side constraint.

Table 14 summarizes the four counterfactual policies and the model objects they modify. The resulting labor allocations are compared with the baseline allocations generated by the estimated structural model. To

Table 14: Calibrated Policy Intensities

Country	Food Subsidy (λ_c^F)	Ag. Productivity (λ_c^A)	Roads (λ_c^R)	Land Reform Coverage (λ_c^L)	Labor Absorption (λ_D)
Niger	0.887	0.150	0.500	1.000	0.000
Malawi	0.397	0.150	0.500	1.000	0.000
Ethiopia	0.463	0.063	0.052	1.000	0.056

Notes: Food subsidy, agricultural productivity, and roads report the calibrated proportional changes in the corresponding model objects. Land reform reports the share of households with weaker land rights that are converted to stronger land rights under the common policy budget. Because the estimated cost of land certification is relatively low, the common budget is sufficient to certify all households with weaker land rights in the three study countries. Labor absorption reports the increase in the labor absorption probability calibrated from the job creation effect of a USD 500 million industrial investment. This policy is calibrated only for Ethiopia because the structural estimates indicate that labor-demand constraints are binding only in Ethiopia.

examine complementarity across stages of structural transformation, I also simulate combinations of policies. For each combination, the total budget is fixed at USD 500 million and divided equally among the selected policies. The corresponding parameter changes are implemented simultaneously, and the model is re-solved to obtain the resulting household labor allocation.

6.2 Niger

Figures 7a and 7b present the counterfactual policy simulations for Niger. Consistent with the structural estimates in Section 5, the results indicate that structural transformation in Niger is primarily constrained during the labor-release stage, while labor absorption is already close to its frontier.

Among the single-policy interventions, improving agricultural productivity generates the largest structural

transformation effect. Agricultural labor declines by 4.6 percentage points, accompanied by a 4.4 percentage point increase in non-agricultural employment. Targeted food subsidies also substantially relax the subsistence constraint, reducing agricultural labor by 1.8 percentage points and increasing non-agricultural employment by a similar magnitude. Neither policy changes the transition probability (π) or the reallocation probability (q), indicating that their effects operate entirely through the first stage of the model by releasing labor from agriculture.

Policies targeting the second stage produce considerably smaller effects. Reducing mobility barriers increases the transition probability by 1.5 percentage points and raises non-agricultural employment by only 0.5 percentage points, suggesting that although mobility frictions exist, they are not the primary obstacle to structural transformation. By contrast, strengthening household land-rights has virtually no effect on labor allocation and slightly reduces the estimated transition probability. This result is consistent with the insignificant land-transferability effect estimated in the structural model, indicating that land institutions do not constitute a major barrier to labor reallocation in Niger.

Finally, increasing labor absorption capacity produces essentially no change in labor allocation. Because the estimated labor absorption probability is already very high ($p \approx 0.96$), expanding labor demand generates little additional structural transformation. This finding confirms that insufficient labor demand is not the binding constraint in Niger.

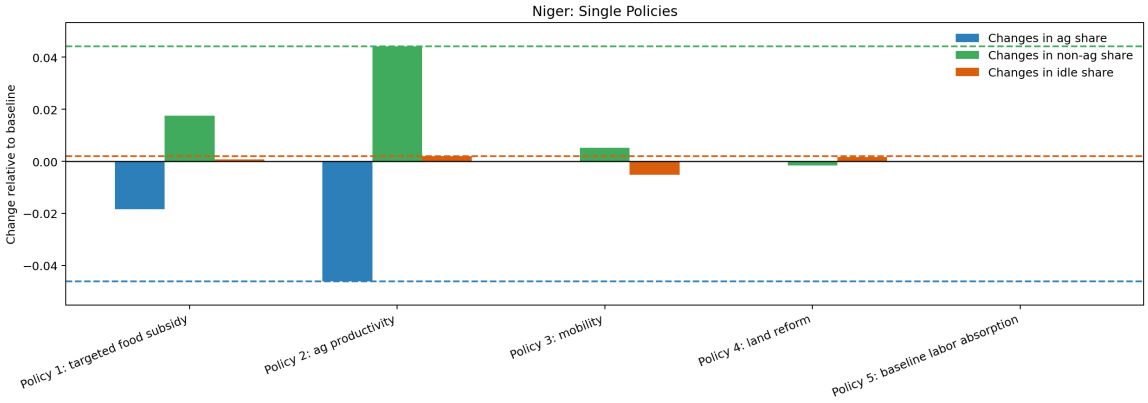
The combined-policy simulations further reinforce these conclusions. Under a fixed total policy budget of USD 500 million, allocating resources to both food subsidies and agricultural productivity improvements generates the largest reduction in agricultural labor. Redistributing part of this budget toward mobility improvements produces only modest additional gains through higher transition probabilities. Similarly, allocating resources to land reform or labor absorption policies yields relatively limited additional improvements beyond those achieved through the agricultural interventions. Even when the fixed budget is distributed across all five policy instruments, most of the increase in non-agricultural employment continues to be attributable to relaxing subsistence constraints rather than to additional investments in labor reallocation or labor absorption.

Overall, the counterfactual results closely mirror the structural estimates. In Niger, structural transformation is primarily constrained by households' ability to release labor from agriculture. Policies that improve agricultural productivity and relax food subsistence constraints therefore generate the largest gains, whereas policies targeting later stages of structural transformation have comparatively limited effects.

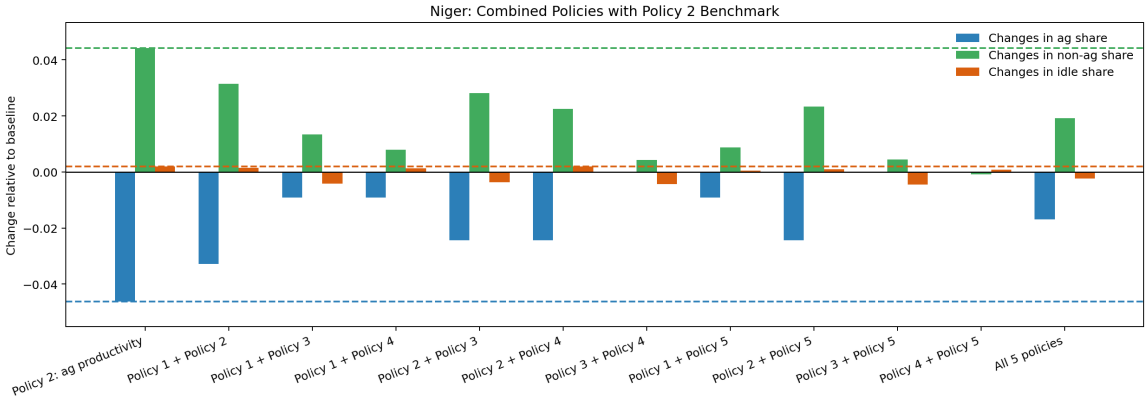
6.3 Malawi

Figures 8a and 8b present the counterfactual policy simulations for Malawi. Consistent with the structural estimates in Section 5, the results indicate that structural transformation in Malawi is overwhelmingly con-

Figure 7: Counterfactual Analysis in Niger



(a) Single Policy



(b) Combined Policy

The counterfactual simulations indicate that policies targeting the labor-release stage generate the largest improvements in structural transformation in Niger. Agricultural productivity improvements and targeted food subsidies substantially reduce agricultural labor and increase non-agricultural employment. Reducing mobility barriers provides modest additional gains by improving access to non-agricultural opportunities, whereas strengthening household land-rights and labor absorption produce relatively limited effects. These results are consistent with the structural estimates, which identify subsistence constraints as the primary barrier to structural transformation and mobility frictions as a secondary constraint in Niger.

strained during the labor-release stage. Once labor is released from agriculture, households face relatively small barriers to reallocation, while labor absorption is already close to complete.

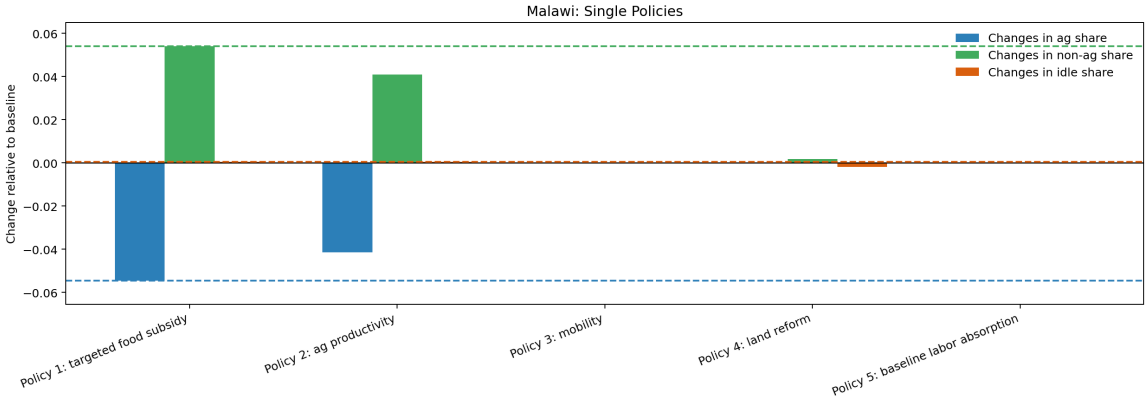
Among the single-policy interventions, targeted food subsidies generate the largest improvement in structural transformation, reducing the agricultural labor share by 5.5 percentage points and increasing non-agricultural employment by 5.4 percentage points. Improving agricultural productivity produces a similarly large effect, reducing agricultural labor by 4.1 percentage points and increasing non-agricultural employment by 4.1 percentage points.

Policies targeting later stages of structural transformation generate substantially smaller effects. Reducing mobility barriers produces virtually no change in labor allocation or transition probabilities, consistent with the insignificant migration-cost parameter estimated in the structural model. Strengthening household land-rights modestly increases the transition probability by 0.7 percentage points, but raises non-agricultural employment by only 0.2 percentage points. Likewise, increasing labor absorption capacity produces no measurable improvement because the estimated labor absorption probability is already close to one.

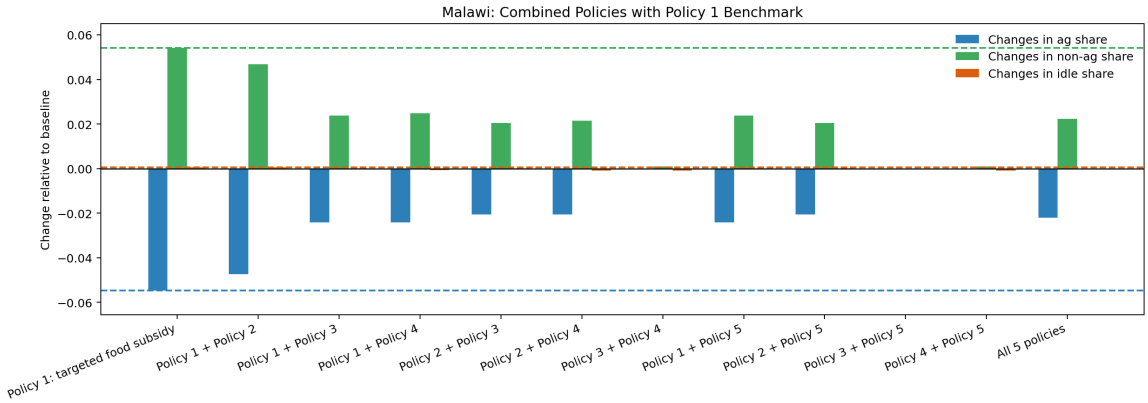
The combined-policy simulations reinforce these findings. Combining targeted food subsidies and agricultural productivity improvements generates the largest reduction in agricultural labor among the two-policy combinations, while adding mobility or labor-absorption policies contributes almost no additional gains. Land reform provides only a small incremental improvement through a higher transition probability. Even when all five policies are implemented simultaneously, nearly all of the increase in non-agricultural employment is attributable to policies that relax the subsistence constraint rather than interventions targeting labor reallocation or labor absorption.

Overall, the Malawi counterfactuals strongly support the structural interpretation of the model. Structural transformation is constrained primarily by households' inability to release labor from agriculture rather than by barriers to labor reallocation or insufficient labor demand. Consequently, policies that relax subsistence constraints through food support or improvements in agricultural productivity generate substantially larger gains than policies targeting later stages of structural transformation.

Figure 8: Counterfactual Analysis in Malawi



(a) Single Policy



(b) Combined Policy

The counterfactual simulations indicate that policies targeting the labor-release stage generate the largest improvements in structural transformation in Malawi. Targeted food subsidies and agricultural productivity improvements substantially reduce agricultural labor and increase non-agricultural employment, whereas policies targeting labor reallocation and labor absorption produce only modest additional gains. These results are consistent with the structural estimates, which identify subsistence constraints as the primary barrier to structural transformation in Malawi.

6.4 Ethiopia

Figures 9a and 9b present the counterfactual policy simulations for Ethiopia. In contrast to Niger and Malawi, the structural estimates indicate that Ethiopia's primary constraints arise after labor is released from agriculture. Although improvements in agricultural productivity can release additional labor, successful structural transformation depends critically on households' ability to access and be absorbed into non-agricultural employment.

Among the single-policy interventions, improving agricultural productivity reduces the agricultural labor share by 2.6 percentage points and increases non-agricultural employment by 1.5 percentage points. However, nearly half of the released labor remains idle, increasing the idle labor share by 1.1 percentage points. Similarly, targeted food subsidies have virtually no effect on aggregate labor allocation. These results indicate that relaxing subsistence constraints alone does not generate substantial structural transformation because released workers continue to face barriers after leaving agriculture.

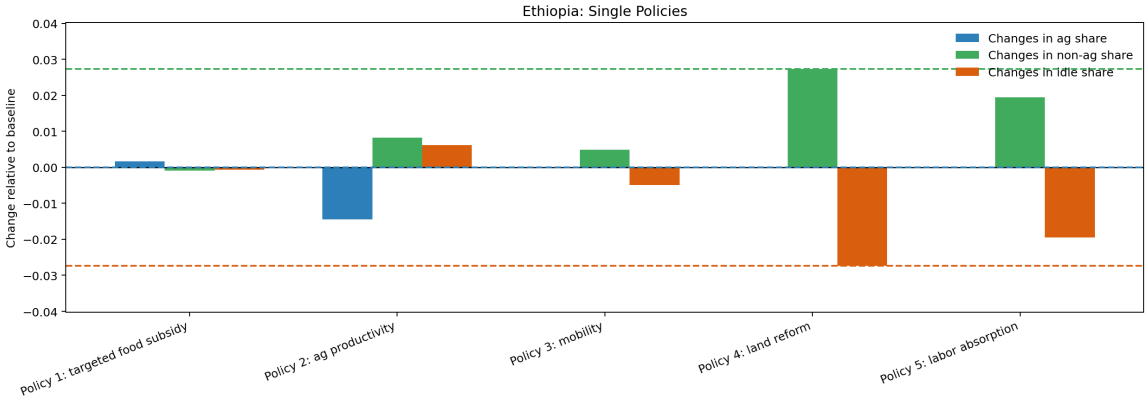
Policies targeting the second and third stages of the model generate substantially larger gains. Reducing mobility barriers increases the transition probability by 1.1 percentage points and raises non-agricultural employment by 0.5 percentage points while reducing idle labor by a similar amount. Strengthening household land-rights produces the largest single-policy improvement, increasing the transition probability by 6.4 percentage points and raising non-agricultural employment by 2.7 percentage points while reducing idle labor by the same magnitude. Increasing labor absorption capacity also generates large gains, increasing non-agricultural employment by 1.9 percentage points and lowering idle labor by 1.9 percentage points without affecting agricultural labor. These findings suggest that improving access to non-agricultural opportunities and expanding labor demand are substantially more effective than further relaxing agricultural constraints.

The combined-policy simulations highlight strong complementarities across stages of structural transformation. Combining agricultural productivity improvements with land reform increases non-agricultural employment by 2.1 percentage points, substantially more than agricultural productivity alone. Likewise, combining land reform with labor absorption produces the largest improvement among the two-policy combinations, increasing non-agricultural employment by 3.3 percentage points while reducing idle labor by an equal amount. When all five policies are implemented simultaneously, non-agricultural employment increases by 3.3 percentage points, idle labor declines by 2.6 percentage points, and the transition probability increases by 6.4 percentage points. These gains substantially exceed those generated by policies targeting labor release alone.

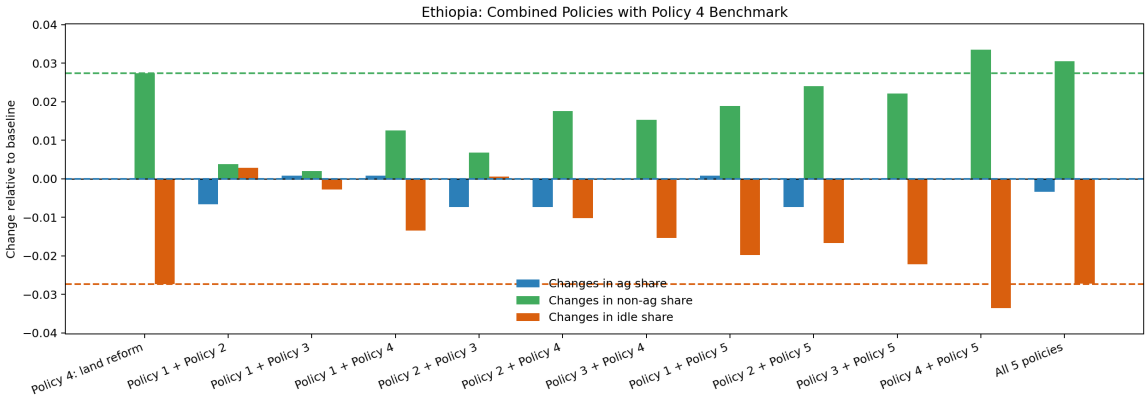
Overall, the Ethiopia counterfactuals provide strong support for the structural estimates. Unlike Niger and Malawi, releasing labor from agriculture is not sufficient to promote structural transformation. Instead, policies that strengthen land institutions and expand labor absorption capacity generate the largest improvements by enabling released workers to successfully transition into productive non-agricultural employment. These findings suggest that demand-side and institutional constraints, rather than agricultural productivity alone,

constitute the primary barriers to structural transformation in Ethiopia.

Figure 9: Counterfactual Analysis in Ethiopia



(a) Single Policy



(b) Combined Policy

The counterfactual simulations indicate that policies targeting the labor-reallocation and labor-absorption stages generate the largest improvements in structural transformation in Ethiopia. Strengthening land transferability and increasing labor absorption capacity produce the largest single-policy improvement and increase non-agricultural employment, whereas policies targeting labor release produce only modest additional gains. These results are consistent with the structural estimates, which identify labor reallocation and labor absorption barriers to structural transformation in Ethiopia.

6.5 Cross-Country Comparison

The counterfactual analysis reveals substantial heterogeneity in the policy drivers of structural transformation across the three countries. These differences closely mirror the structural estimates presented in Section 5 and highlight the importance of identifying the stage at which labor reallocation is constrained.

In both Niger and Malawi, the largest gains arise from policies that relax subsistence constraints and increase the amount of labor released from agriculture. Improvements in agricultural productivity and targeted food subsidies consistently generate the largest reductions in agricultural labor and corresponding increases in non-agricultural employment. By contrast, policies targeting labor reallocation and labor absorption produce relatively small additional gains, reflecting the high estimated transition and absorption probabilities in these countries.

Ethiopia exhibits a fundamentally different pattern. Although agricultural productivity improvements release

additional labor, a substantial fraction of released workers remains idle unless policies also improve access to non-agricultural opportunities or expand labor demand. household land-rights reforms and industrial investments therefore generate considerably larger improvements than policies targeting agricultural productivity alone. These findings are consistent with the structural estimates, which identify labor reallocation and labor absorption as the primary bottlenecks in Ethiopia.

Taken together, the results demonstrate that effective structural transformation policies depend on the location of the binding constraint within the transformation process. Policies that target the relevant stage of the model consistently generate larger gains than policies aimed at non-binding constraints. The structural model, therefore, provides a useful framework for diagnosing country-specific barriers to structural transformation and evaluating policies that directly address those barriers.

7 Conclusion

Structural transformation is a key driver of long-run economic development, but labor released from agriculture does not always enter productive non-agricultural employment. This paper examines this macroeconomic question at the household level, distinguishing three stages: labor release, labor transition, and labor absorption. Combining reduced-form evidence with a structural model, I examine how subsistence needs, mobility costs, land rights, and employment opportunities shape household labor allocation across Niger, Malawi, and Ethiopia.

The reduced-form results show that favorable weather-induced productivity shocks reduce agricultural labor in all three countries. However, the accompanying employment responses differ. In Niger and Malawi, agricultural labor reductions are accompanied by increases in manufacturing employment, while in Ethiopia they are accompanied by a substantial increase in idle labor. The mechanism analysis further shows that subsistence constraints weaken labor release in Niger and Malawi. In Ethiopia, greater distance is associated with stronger agricultural labor reductions and larger increases in idle labor, while land rights are associated with stronger labor release and larger increases in both manufacturing employment and idle labor. Complementary regional evidence suggests that stronger hiring attenuates the idle-labor response. Together, these findings distinguish the ability to leave agriculture from the ability to obtain non-agricultural employment.

I estimate the structural model using simulated method of moments to examine how these constraints jointly determine household labor allocation. The estimates suggest that subsistence constraints primarily limit labor release in Niger and Malawi, while transition barriers and limited employment absorption are more important in Ethiopia. The framework connects the Lewis-type labor-release mechanism with the employment-probability channel emphasized by Harris and Todaro (1970), while explicitly incorporating the intermediate transition stage. This allows the model to capture how mobility costs and land institutions shape the employment gains from labor released by agricultural productivity improvements.

The counterfactual analysis evaluates both individual policies and combined interventions. In Niger and Malawi, policies that relax subsistence constraints, including targeted food subsidies and agricultural productivity improvements, generate the largest increases in non-agricultural employment. Interventions targeting mobility or absorption generate smaller gains under the estimated constraints. In Ethiopia, policies that strengthen household land rights and expand employment absorption perform better than interventions focused solely on agricultural productivity. The combined-policy experiments further examine whether addressing multiple constraints together improves employment outcomes. These comparisons help identify which constraints policies should target to increase productive non-agricultural employment and promote structural transformation, highlighting the importance of country-specific interventions.

Several limitations suggest directions for future research. The model abstracts from general equilibrium effects, including wage adjustments, endogenous firm entry, and feedback between labor supply and labor demand. The counterfactual analysis captures partial-equilibrium responses rather than economy-wide adjustments. Future work could incorporate endogenous labor demand and firm behavior within a general equilibrium framework to study how policy interventions affect both households and firms simultaneously. In addition, extending the model to account for dynamic human capital accumulation, migration decisions, and gender-specific labor constraints would provide a richer understanding of the long-run process of structural transformation.

Overall, by distinguishing labor release, labor reallocation, and labor absorption as separate stages of structural transformation, this paper demonstrates that similar aggregate outcomes can arise from different underlying constraints. Although the empirical analysis exploits temporary weather-induced productivity shocks, the transition evidence shows that labor reallocation generate persistent employment states, providing a behavioral foundation for evaluating how structural constraints shape household labor allocation under alternative policy environments. Identifying these constraints is essential for designing effective development policies and for understanding why structural transformation proceeds at different speeds across countries.

References

Abate, A., & Kiros, F. G. (2010). Agrarian reform, structural changes and rural development in ethiopia. In *Agrarian reform in contemporary developing countries* (pp. 141–184). Routledge.

Abebe, G., Caria, S., Fafchamps, M., Falco, P., Franklin, S., & Quinn, S. (2021). Anonymity or distance? job search and labour market exclusion in a growing african city. *The Review of Economic Studies*, *88*(3), 1279–1310.

Adamopoulos, T., Brandt, L., Chen, C., Restuccia, D., & Wei, X. (2024). Land security and mobility frictions. *The Quarterly Journal of Economics*, *139*(3), 1941–1987.

Adamopoulos, T., & Restuccia, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, *104*(6), 1667–1697.

Adamopoulos, T., & Restuccia, D. (2020). Land reform and productivity: A quantitative analysis with micro data. *American Economic Journal: Macroeconomics*, *12*(3), 1–39.

Alfano, M., & Görlach, J.-S. (2024). Weather shocks, migration, and in-place adaptation. *University of Lancaster*.

Ali, D. A., Deininger, K., & Ronchi, L. (2019). Costs and benefits of land fragmentation: Evidence from rwanda. *The World Bank Economic Review*, *33*(3), 750–771.

Allen, T., & Atkin, D. (2022). Volatility and the gains from trade. *Econometrica*, *90*(5), 2053–2092.

Alvarez-Cuadrado, F., & Poschke, M. (2011). Structural change out of agriculture: Labor push versus labor pull. *American Economic Journal: Macroeconomics*, *3*(3), 127–158.

Asher, S., & Novosad, P. (2020). Rural roads and local economic development. *American economic review*, *110*(3), 797–823.

Ayele, M. B., & Siba, E. (2017). The land certification program and off-farm employment in ethiopia.

Banerjee, A., & Sequeira, S. (2023). Learning by searching: Spatial mismatches and imperfect information in southern labor markets. *Journal of Development Economics*, *164*, 103111. https://doi.org/10.1016/j.jdeveco.2023.103111

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., & Shapiro, J. S. (2016). Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, *124*(1), 105–159.

Barrett, C. B., Christian, P., & Shiferaw, B. A. (2017). The structural transformation of african agriculture and rural spaces: Introduction to a special section. *Agricultural economics*, *48*(S1), 5–10.

Barrett, C. B., Carter, M. R., Chavas, J.-P., & Carter, M. R. (2019). *The economics of poverty traps*. University of Chicago Press Chicago.

Basu, S. (2025). The effect of drought on household occupation choices in rural india [First View]. *Environment and Development Economics*, 1–20. https://doi.org/10.1017/S1355770X25100314

Benjamin, D. (1992). Household composition, labor markets, and labor demand: Testing for separation in agricultural household models. *Econometrica: Journal of the Econometric Society*, 287–322.

Bentze, T., & Wollburg, P. (2024). *A longitudinal cross-country dataset on agricultural productivity and welfare in sub-saharan africa*. World Bank.

Branco, D., & Féres, J. (2021). Weather shocks and labor allocation: Evidence from rural brazil. *American Journal of Agricultural Economics*, *103*(4), 1359–1377.

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

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

Bustos, P., Caprettini, B., & Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from brazil. *American Economic Review*, *106*(6), 1320–1365.

Byerlee, D., De Janvry, A., & Sadoulet, E. (2009). Agriculture for development: Toward a new paradigm. *Annu. Rev. Resour. Econ.*, *1*(1), 15–31.

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

Christiaensen, L., & Maertens, M. (2022). Rural employment in africa: Trends and challenges. *Annual Review of Resource Economics*, *14*(1), 267–289.

Colmer, J. (2021). Temperature, labor reallocation, and industrial production: Evidence from india. *American Economic Journal: Applied Economics*, *13*(4), 101–124.

De Janvry, A., Emerick, K., Gonzalez-Navarro, M., & Sadoulet, E. (2015). Delinking land rights from land use: Certification and migration in mexico. *American Economic Review*, *105*(10), 3125–3149.

Deininger, K., & Jin, S. (2006). Tenure security and land-related investment: Evidence from ethiopia. *European Economic Review*, *50*(5), 1245–1277.

Deininger, K., & Jin, S. (2008). Land sales and rental markets in transition: Evidence from rural vietnam. *Oxford bulletin of Economics and Statistics*, *70*(1), 67–101.

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

Diao, X., Ellis, M., McMillan, M., & Rodrik, D. (2025). Africa's manufacturing puzzle: Evidence from tanzanian and ethiopian firms. *The World Bank Economic Review*, *39*(2), 308–340.

Diao, X., McMillan, M., & Rodrik, D. (2019). The recent growth boom in developing economies: A structural-change perspective. In *The palgrave handbook of development economics: Critical reflections on globalisation and development* (pp. 281–334). Springer.

Dillon, B., & Barrett, C. B. (2017). Agricultural factor markets in sub-saharan africa: An updated view with formal tests for market failure. *Food policy*, *67*, 64–77.

Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American economic review*, *108*(4-5), 899–934.

Donaldson, D., & Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, *30*(4), 171–198.

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

Fox, M. L., Thomas, M. A. H., & Haines, C. (2017). *Structural transformation in employment and productivity: What can africa hope for?* International Monetary Fund.

Gollin, D. (2023). Agricultural productivity and structural transformation: Evidence and questions for african development. *Oxford Development Studies*, *51*(4), 375–396.

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

Gollin, D., Parente, S., & Rogerson, R. (2002). The role of agriculture in development. *American economic review*, *92*(2), 160–164.

Gollin, D., Parente, S. L., & Rogerson, R. (2007). The food problem and the evolution of international income levels. *Journal of Monetary Economics*, *54*(4), 1230–1255.

Grosh, M. E., & Glewwe, P. (1998). Data watch: The world bank's living standards measurement study household surveys margaret e. grosh and paul glewwe. *Journal of Economic Perspectives*, *12*(1), 187–196.

Harris, J. R., & Todaro, M. P. (1970). Migration, unemployment and development: A two-sector analysis. *The American economic review*, *60*(1), 126–142.

Herrendorf, B., Rogerson, R., & Valentinyi, A. (2022). *New evidence on sectoral labor productivity: Implications for industrialization and development* (tech. rep.). National Bureau of Economic Research.

Hirvonen, K. (2016). Temperature changes, household consumption, and internal migration: Evidence from tanzania. *American Journal of Agricultural Economics*, *98*(4), 1230–1249.

Iizumi, T., Kim, W., & Nishimori, M. (2019). Modeling the global sowing and harvesting windows of major crops around the year 2000. *Journal of Advances in Modeling Earth Systems*, *11*(1), 99–112.

International Labour Organization. (2013). Resolution concerning statistics of work, employment and labour underutilization.

Jessoe, K., Manning, D. T., & 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.

Karayalcin, C., & Pintea, M. (2022). The role of productivity, transportation costs, and barriers to intersectoral mobility in structural transformation. *Economic Modelling*, *108*, 105759.

Krishna, K., & Meenakshi, J. (2022). Agricultural productivity growth and structural transformation in rural india: Some recent evidence. *Journal of Quantitative Economics*, *20*(Suppl 1), 277–302.

Lagakos, D., & Shu, M. (2023). The role of micro data in understanding structural transformation. *Oxford Development Studies*, *51*(4), 436–454.

Lagakos, D., & Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *American Economic Review*, *103*(2), 948–980.

Lewis, W. A. (1954). Economic development with unlimited supplies of labour.

Lobell, D. B., Thau, D., Seifert, C., Engle, E., & Little, B. (2015). A scalable satellite-based crop yield mapper. *Remote Sensing of Environment*, *164*, 324–333.

Manyazewal, M., & Shiferaw, A. (2019). Economic policy and structural transformation in ethiopia. *The oxford handbook of the Ethiopian economy*, 137–158.

McKenzie, D. (2017). How effective are active labor market policies in developing countries? a critical review of recent evidence. *The World Bank Research Observer*, *32*(2), 127–154.

McMillan, M., & Headey, D. (2014). Introduction–understanding structural transformation in africa.

Mensah, E. B., Owusu, S., Foster-McGregor, N., & Szirmai, A. (2018). Structural change, productivity growth and labor market turbulence in africa. *UNU-MERIT Working Paper Series*, *25*, 1–150.

Minoli, S., Jägermeyr, J., Asseng, S., Urfels, A., & Müller, C. (2022). Global crop yields can be lifted by timely adaptation of growing periods to climate change. *Nature Communications*, *13*(1), 7079.

Morsy, H., Shimeles, A., & Nabassaga, T. (2023). Structural change and inequality in africa. *Journal of African Economies*, *32*(Supplement_2), ii228–ii245.

Ngai, L. R., & Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American economic review*, *97*(1), 429–443.

Rodrik, D. (2016). Premature deindustrialization. *Journal of economic growth*, *21*(1), 1–33.

Sacks, W. J., Deryng, D., Foley, J. A., & Ramankutty, N. (2010). Crop planting dates: An analysis of global patterns. *Global ecology and biogeography*, *19*(5), 607–620.

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

Suri, T., & Udry, C. (2022). Agricultural technology in africa. *Journal of Economic Perspectives*, *36*(1), 33–56.

Syrquin, M. (1988). Patterns of structural change. *Handbook of development economics*, *1*, 203–273.

Tombe, T., & Zhu, X. (2019). Trade, migration, and productivity: A quantitative analysis of china. *American Economic Review*, *109*(5), 1843–1872.

Vicente Cateia, J., Bittencourt, M. V. L., Carvalho, T. S., & Savard, L. (2024). The impacts of agricultural productivity on structural transformation, and poverty alleviation in africa: Evidence from guinea-bissau. *Journal of Productivity Analysis*, *61*(3), 305–320.

Wang, H. (2025). Agricultural productivity fluctuations and structural transformation—evidence from rural china. *Applied Economic Perspectives and Policy*, *47*(2), 801–822.

World Bank. (2019). *World development report 2019: The changing nature of work*. https://www.worldbank.org/en/publication/wdr2019

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

A Data and Labor Market Variable Construction

A.1 Survey Timing and Data Collection

A.1.1 Niger

The Niger ECVM survey is conducted in two visits per wave.

- **First visit (planting season):** household characteristics, demographics, labor market participation, and non-farm enterprises
- **Second visit (harvest season):** agricultural production, land, inputs, and output

We combine information from both visits to construct labor allocation and agricultural activity variables.

The 2011 and 2014 survey waves follow this structure, with labor market information collected during the first visit and agricultural information collected during the second visit. As described in Appendix A.2.2, the labor module records individuals' main economic activity during the previous 12 months. Consequently, labor allocation observed during the first (planting-season) visit reflects employment decisions made during and immediately following the preceding agricultural year. This timing is consistent with the empirical specification, which uses the one-year-lagged weather-induced productivity shock as the treatment variable, ensuring that the productivity shock aligns with the observed labor allocation.

A.1.2 Malawi

Malawi data are collected in a single household visit. However, the questionnaire distinguishes agricultural activities across planting and harvest periods through recall questions. Employment and agricultural modules are therefore temporally comparable, although no separate survey visits are conducted.

A.1.3 Ethiopia

The Ethiopia ESS survey uses multiple visits:

- Post-planting visit
- Post-harvest visit
- Household visit

We combine information across visits to construct consistent employment and agricultural variables. For rural households, agricultural modules are administered separately from household employment modules.

A.2 Employment Reference Period

The household surveys record employment using multiple reference periods, including employment during the previous seven days and regular employment during the previous twelve months. This paper uses the

twelve-month reference period because the objective is to study structural transformation, which refers to persistent changes in labor allocation across sectors rather than temporary seasonal fluctuations in employment. Household labor allocation decisions in the theoretical framework represent durable employment choices that evolve over relatively long horizons and correspond more closely to regular employment over the previous year than to employment observed during a single week.

The seven-day measure is well-suited for measuring short-run labor market conditions but is heavily influenced by agricultural seasonality. In predominantly agricultural economies, whether an individual worked during the previous week depends critically on the timing of planting, harvesting, and other seasonal agricultural activities, as well as the survey interview date. Consequently, the seven-day measure captures temporary fluctuations in labor demand rather than persistent occupational choices. Because interviews occur at different times across households and countries, relying on a seven-day reference period would introduce substantial variation unrelated to structural transformation.

Temporary and casual employment nevertheless remains an important source of income in rural Sub-Saharan Africa. Because this paper studies durable labor allocation rather than short-term employment spells, individuals are classified according to their regular employment status over the previous twelve months. Workers who report no regular employment and whose economic activity consists only of temporary or casual work, unpaid household agricultural work, or no work are classified as idle labor. This classification is conceptually related to the ILO notion of labor underutilization, reflecting the idea that intermittent attachment to productive employment represents incomplete labor absorption. From the perspective of structural transformation, these workers have not yet been durably absorbed into productive employment, even though they may participate intermittently in economic activities.

A.3 Definition of Working-age Employment

A.3.1 Definition of Working-Age Population

We define the working-age population as individuals aged 15–85. All labor market variables are constructed at the individual level and then aggregated to the household level.

This broader age range reflects labor market conditions in rural developing economies. Individuals often begin working at younger ages, particularly in family farming and household enterprises, while weak pension systems and limited social safety nets imply that older individuals frequently remain economically active well beyond conventional retirement ages. As a result, labor supply in rural settings includes both younger and older household members.

This approach is consistent with the LSMS-based literature, which often includes all individuals aged 15 and above and does not impose a strict retirement age when measuring labor allocation across sectors. In particular, studies using LSMS and similar microdata emphasize that older individuals continue to participate

in agricultural and informal activities, contributing to household production and labor reallocation decisions (Gollin et al., 2014; Dillon and Barrett, 2017; Young, 2013). Moreover, labor force participation at older ages remains high in low-income countries due to limited pension coverage and continued reliance on family labor (World Bank, 2019).

A.3.2 Definition of Employment

The labor modules of the LSMS-ISA surveys in Niger, Malawi, and Ethiopia all distinguish between agricultural self-employment, household non-agricultural enterprises, wage employment, and temporary or casual work using comparable individual-level questions. Although questionnaire wording and survey implementation differ slightly across countries, we harmonize labor status using a common hierarchy based on reported economic activities rather than survey-specific variable definitions. The same classification procedure is applied consistently across all three countries, ensuring that cross-country differences in underutilized labor reflect differences in reported labor market attachment rather than differences in coding rules.

Employment information is based on individuals' main activity during the last 12 months. We focus on the primary job only; secondary and tertiary jobs are not considered. While the wording of employment questions differs slightly across countries, they are conceptually comparable. The surveys ask about individuals' main occupation during the previous year, such as "Principal employment during the last 12 months" in Niger, "Describe your main employed occupation over the past 12 months" in Malawi, and "Was [NAME] employed in the last 12 months?" in Ethiopia. Individuals reporting a main job are classified as employed workers.

Self-employment is identified using the Non-Agriculture Enterprise modules in the household questionnaires. These modules record up to two owners for each enterprise. We link enterprise owner identifiers to the household roster to identify self-employed individuals. When two owners are reported, enterprise revenue is split equally between them to ensure consistent assignment of labor participation without overstating employment.

An important advantage of the data is that both formal and informal employment are captured. The dataset includes wage employment (formal and informal), household non-farm enterprises, and self-employment. We therefore do not restrict the analysis to formal sector employment.

Labor is first assigned to an employment sector using individual-level information on primary economic activities. The household agricultural information is used only for individuals who ar not assigned a regular employment sector. Consequently, living in an agricultural household does not mechanically classify all household members as agricultural workers. Individuals without regular sectoral employment remain classified as underutilized labor when their reported activities indicate only marginal attachment to employment, such as unpaid household agricultural work, temporary or casual work, or unsuccessful labor market participation.

We distinguish regular employment from a broader category of underutilized labor (International Labour Organization, 2013), which corresponds to individuals who are not durably absorbed into regular sectoral employment. This category includes working-age individuals who are unemployed or only marginally attached to economic activity through unpaid household agricultural work or temporary and casual employment. Specifically, underutilized labor includes individuals who report only participation in household agricultural activities, casual or temporary work, active job search, willingness to work despite not actively searching, or only temporary employment during the previous twelve months. Individuals with regular agricultural, household enterprise, wage, or salaried employment are classified as employed rather than underutilized. This definition aligns the empirical measure more closely with the concept of underutilized labor represented in the structural model.

A.4 Absent Household Members and Labor Allocation

A potential concern is that rural workers may work outside the village, particularly in nearby towns or cities. All three surveys explicitly track absent household members, allowing us to account for individuals who are temporarily away but still belong to the household. In Niger, the ECVM survey records all household members, including those temporarily absent, and collects information on the duration and reason for absence over the past 12 months. This allows identification of individuals working outside the village while remaining part of the household. In Malawi and Ethiopia, absent household members are also recorded, and information on their employment and activities is collected through proxy responses from the household head or another adult household member. As a result, individuals working outside the household location remain part of the household roster and are included in labor allocation measures.

The surveys also distinguish between temporary and permanent movers. Temporary movers remain members of the household and are included in the labor-allocation measures, whereas permanent movers, for example, individuals who migrate and establish a new household, are no longer part of the original household roster. Accordingly, the analysis captures labor reallocation through local employment, commuting, seasonal migration, informal non-farm work, participation in household enterprises, and temporary migration, but does not observe permanent migration once individuals leave the household. Consequently, the measured non-agricultural labor allocation reflects both local employment opportunities and non-local opportunities accessed while workers remain household members. While permanent migration is not directly observed, temporary migration and commuting remain important channels through which households respond to spatial mobility barriers, allowing the empirical analysis to identify the role of migration frictions in structural transformation. Permanent migration accompanied by household splitting is not observed and therefore represents a limitation of the LSMS household panel data rather than of the empirical framework.

These data features allow us to capture multiple forms of labor allocation, including local employment, migration to nearby towns, seasonal migration, informal non-farm work, and participation in household en-

terprises. Consequently, measured non-agricultural labor participation reflects both local and non-local job opportunities, consistent with the migration mechanism in the model.

B Visual Validation of Household Exposure to Agricultural Activity

The sample construction differs slightly across countries to reflect differences in the meaning of administrative classifications. In Niger and Ethiopia, we restrict the analysis to households classified as rural. In Malawi, however, the administrative definition of "urban" often includes settlements that remain closely integrated with surrounding agricultural land. We therefore include households located in communities with observable agricultural activity.

To verify that sampled households in Malawi are plausibly exposed to agricultural production conditions, we examine high-resolution satellite imagery surrounding the displaced enumeration-area coordinates provided in the LSMS survey, consistent with the 5 km buffer strategy described in Section 2.4.2. Because the World Bank intentionally displaces survey coordinates to protect respondent confidentiality, the imagery is interpreted at the community level rather than as the exact location of individual households. The satellite images show that sampled settlements are embedded within active agricultural landscapes, supporting the use of surrounding environmental conditions to characterize agricultural exposure.

To complement the visual evidence, we also examine labor allocation patterns across administrative classifications. In Malawi, households classified as urban allocate 47 percent of labor to agriculture, substantially higher than the corresponding urban shares in Niger (12 percent) and Ethiopia (7 percent). This difference suggests that the administrative definition of urban areas in Malawi includes settlements that remain closely integrated with surrounding agricultural production, consistent with the community-level evidence from the satellite imagery.

Figure 10 presents representative examples of household locations in Malawi. The imagery reveals that many households, including those located in relatively dense settlements, are situated in close proximity to cultivated plots and agricultural land. In several cases, residential structures are directly adjacent to or interspersed with visibly cultivated fields, indicating a mixed land-use pattern that combines residential and agricultural activity.

These patterns suggest that administrative classifications such as "urban" do not necessarily correspond to a separation from agricultural production in the Malawian context. Instead, households often remain spatially embedded in environments where agricultural activity is prevalent, and thus are plausibly exposed to local agricultural productivity shocks.

Importantly, the key identifying assumption is that household income opportunities, either directly through own production or indirectly through local labor markets, are affected by agricultural conditions. The visual

evidence presented here supports the plausibility of this assumption.



Figure 10: Satellite imagery of representative household locations in different regions in Malawi. The figures illustrate a dense settlement with nearby cultivated plots and a mixed residential–agricultural landscape where households are located adjacent to farmland.

Finally, we verify that the inclusion of urban households in Malawi (Table 15) does not drive the main results. Re-estimating the baseline specifications using only rural households yields effects that are similar in sign and magnitude, although estimated with less precision. This suggests that the main findings are not sensitive to the inclusion of households classified as urban.

Table 15: Sectoral Choice and Weather Productivity Growth

	(1)	(2)	(3)	(4)
	Agriculture	Manufacturing	Service	Idle
$\Delta\theta_{gt-1}$	-0.585	0.568**	-0.104	0.121
	(-1.21)	(2.03)	(-0.39)	(0.81)
Constant	0.835***	0.0440***	0.113***	0.00820**
	(74.89)	(7.34)	(16.65)	(2.60)
Household Fixed Effect	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes
Mean dep. var	0.821	0.049	0.125	0.004
SD dep. var	0.326	0.189	0.276	0.060
Observations	1750	1750	1750	1750
R-squared	0.686	0.615	0.679	0.500

∗ t statistics in parentheses

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

C Weather Productivity Shock Construction

C.1 Exogenous Weather Productivity Growth: Definition

Estimating the effect of agricultural productivity on household labor allocation raises important endogeneity concerns. Observed agricultural productivity is jointly determined with household production decisions, and is affected by unobserved factors such as technology adoption, input use, irrigation, management practices, and local economic conditions (Branco and Féres, 2021; Emerick, 2018; Jessoe et al., 2018). Consequently,

directly using observed agricultural productivity as an explanatory variable would generally produce biased estimates of its effect on labor allocation.

To isolate the exogenous component of agricultural productivity, this paper constructs a weather-induced productivity measure. The idea is to separate productivity changes driven by exogenous weather conditions from those generated by endogenous household decisions. The construction begins with the following agricultural production function at the enumeration-area (EA) level,

$$Y = TFP \cdot \theta \cdot L^{\alpha_1} K^{\alpha_2} T^{\alpha_3} \cdot \varepsilon, \tag{43}$$

where Y denotes agricultural output, TFP represents household-specific productivity, θ captures the productivity effect of weather conditions, L, K, and T denote labor, capital, and land inputs, respectively, and ε is an idiosyncratic error term.

Taking logarithms gives

$$y = \log(TFP) + \log(\theta) + \alpha_1 \log(L) + \alpha_2 \log(K) + \alpha_3 \log(T) + \varepsilon. \tag{44}$$

The empirical analysis focuses on annual changes rather than productivity levels because household labor allocation responds to changes in agricultural conditions over time. Taking first differences yields

$$\Delta y = \Delta \log(\theta) + \Delta H + \Delta \varepsilon, \tag{45}$$

where

$$\Delta H = \Delta \log(TFP) + \alpha_1 \Delta \log(L) + \alpha_2 \Delta \log(K) + \alpha_3 \Delta \log(T). \tag{46}$$

The term ΔH summarizes systematic changes in agricultural productivity arising from non-weather factors, including changes in total factor productivity, labor, capital, and land inputs, while $\Delta\varepsilon$ captures transitory production shocks and measurement error.

Following Colmer (2021), capital and land are assumed to be fixed in the short run. Consequently,

$$\Delta H = \Delta \log(TFP) + \alpha_1 \Delta \log(L), \tag{47}$$

which captures changes arising from household production decisions, including technology adoption, management practices, labor adjustments, crop choices, and other non-weather factors. The objective is therefore to recover the weather productivity component,

$$\Delta\theta \equiv \Delta \log(\theta),$$

which serves as the exogenous treatment variable throughout the paper.

C.2 Exogenous Weather Productivity Growth: Construction

Annual agricultural output is not consistently observed across countries and survey waves. Following Donaldson and Storeygard (2016) and Lobell et al. (2015), annual NDVI is therefore used as a proxy for agricultural productivity because it provides a consistent satellite-based measure of vegetation growth that is highly correlated with crop production.

Using NDVI, Equation 45 can be written as

$$\Delta NDVI_i = \Delta\theta_i + \Delta H_i. \tag{48}$$

The challenge is that only the left-hand side is observed. The weather productivity component $\Delta\theta_i$ must therefore be recovered from exogenous weather variables.

A natural approach would be to estimate

$$\Delta NDVI_i = \beta_1 \Delta Rainfall_i + \beta_2 \Delta Temperature_i + u_i. \tag{49}$$

However, this specification imposes a restrictive linear relationship between rainfall, temperature, and agricultural productivity. Agricultural production depends on nonlinear interactions between rainfall and temperature, including threshold effects and complementarities. Consequently, a flexible nonparametric approach is preferred.

This paper estimates the conditional expectation

$$E(\Delta NDVI_i \mid \Delta Rainfall_i, \Delta Temperature_i) = g(\Delta Rainfall_i, \Delta Temperature_i) \tag{50}$$

using a Random Forest regression model. The estimated function $\hat{g}(\cdot)$ provides the weather-induced component of productivity,

$$\widehat{\Delta\theta_i} = \hat{g}(\Delta Rainfall_i, \Delta Temperature_i), \tag{51}$$

while the residual

$$\widehat{\Delta H_i} = \Delta NDVI_i - \widehat{\Delta\theta_i} \tag{52}$$

captures the component of productivity changes arising from household production decisions, crop choices, measurement error, and other local factors that is not predicted by rainfall and temperature. Accordingly, the weather-induced component should be interpreted as the portion of productivity that is predictable from weather conditions, which may include both direct biophysical effects of weather and systematic household

responses to those conditions.

Random Forest is particularly suitable because it flexibly captures nonlinear relationships and interactions between rainfall and temperature without requiring a parametric functional form. Each regression tree is estimated using a bootstrap sample of the training data, and predictions are averaged across trees to reduce prediction variance. The model is trained using a random 80% sample of the observations and evaluated on the remaining 20% to assess out-of-sample performance.

Finally, the estimated Random Forest model is applied to the full sample to construct the weather productivity measure used throughout the paper. Because the prediction depends exclusively on exogenous weather variables, the constructed treatment is predetermined with respect to household labor allocation decisions. The out-of-sample prediction performs well across all three countries, with R^2 values ranging from 0.29 in Niger, 0.60 in Malawi, and 0.30 in Ethiopia. These values indicate that weather conditions explain a substantial share of annual vegetation growth while leaving considerable variation attributable to household production decisions and other non-weather factors, which is precisely the variation excluded from the treatment measure.

C.3 Why Annual Weather-Induced Productivity?

Crop responses to weather are inherently crop- and location-specific because different crops follow distinct phenological calendars, with substantial variation in planting, flowering, grain-filling, and harvesting dates across regions and management systems (Sacks et al., 2010; Iizumi et al., 2019; Minoli et al., 2022). This heterogeneity is particularly pronounced in the multi-country household surveys used in this paper, where households cultivate different crop mixes under diverse agroecological conditions. Moreover, the productivity effect of a given weather realization depends critically on when it occurs. Insufficient rainfall during crop growth and excessive rainfall during harvest, for example, may both reduce final output even though they represent very different rainfall realizations. Thus, constructing a direct weather-shock measure would require knowing not only the relevant crop calendar for each household, but also the optimal timing and intensity of rainfall and temperature for each crop and location and then measuring deviations from those crop- and location-specific optima. Such information is generally unavailable and would be difficult to harmonize across countries.

This consideration motivates the use of NDVI as an intermediate measure between weather and agricultural productivity. Rather than defining favorable or unfavorable weather ex ante, NDVI measures the realized vegetation outcome after crops have been exposed to the sequence of weather conditions occurring over the production cycle. This is particularly useful in a setting with heterogeneous crop calendars: weather conditions that are beneficial for one crop or at one point in the production cycle may be harmful for another, whereas their consequences are ultimately reflected in realized vegetation growth. The empirical strategy

therefore does not require the researcher to determine whether a particular amount or timing of rainfall constitutes a positive or negative agricultural shock for each crop and location.

The objective of this paper also differs from that of the crop-modeling literature. Rather than estimating the biological response of individual crops to weather at specific phenological stages, this paper seeks to identify the causal effect of agricultural productivity on household labor allocation. The economically relevant treatment is therefore not rainfall, temperature, drought, or another weather realization itself, but the component of realized agricultural productivity that is attributable to exogenous weather. Conceptually, the empirical strategy follows the relationship

$$\text{Weather conditions} \longrightarrow \text{Realized agricultural productivity} \longrightarrow \text{Household labor allocation.}$$

Weather provides the exogenous source of variation, while agricultural productivity is the treatment whose effect on labor allocation is of interest.

The annual frequency follows from this economic object. The analysis focuses on realized agricultural productivity over the production year and on subsequent household labor allocation rather than on the effect of individual weather events. I therefore aggregate rainfall and temperature to the annual level and use annual NDVI to measure realized vegetation growth over the corresponding production period. This provides a common temporal unit that can be applied consistently across households cultivating different crops and across countries with different agricultural calendars. Importantly, this choice does not imply that within-year weather timing is unimportant. Rather, annual aggregation reflects the objective of constructing a comparable measure of annual weather-induced productivity rather than identifying the agronomic effect of weather at each stage of crop development.

To recover the exogenous weather-induced component of annual productivity, I employ a Random Forest machine learning algorithm to estimate the relationship between annual rainfall, annual temperature, and annual NDVI. Existing studies commonly accommodate nonlinear weather effects using temperature bins (Schlenker and Roberts, 2009; Barreca et al., 2016) or polynomial functions of precipitation and temperature (Deschênes and Greenstone, 2007; Burke et al., 2015). These approaches are well suited when the objective is to estimate particular weather-response functions, but they require researchers to choose functional forms, thresholds, and interaction structures ex ante. Random Forest instead allows the mapping from rainfall and temperature into realized vegetation productivity to be nonlinear and interactive without imposing a particular parametric response function.

Importantly, the objective of the Random Forest is not to explain all variation in NDVI or to construct the most detailed possible model of crop production. It is to recover the component of annual vegetation productivity that is systematically predictable from exogenous weather. The baseline specification therefore uses annual

rainfall and annual temperature as predictors. More elaborate measures—such as drought indices, temperature bins, rainfall variability, or researcher-defined extreme-weather indicators—require additional choices about thresholds, reference periods, or functional forms. Restricting the baseline predictor set to consistently measured rainfall and temperature minimizes such researcher discretion, while the Random Forest provides flexibility in how these variables map into realized productivity.

Annual aggregation inevitably sacrifices information about the exact timing and intensity of within-year weather events. Field interviews conducted for this study underscore this limitation: farmers frequently described unfavorable weather as rainfall occurring at the wrong time, such as insufficient rainfall during crop growth or excessive rainfall during harvest, rather than simply unusually low or high annual rainfall. This evidence reinforces the motivation for focusing on weather-induced *productivity* rather than treating a particular rainfall realization as the agricultural shock itself. A rainfall-based shock would require the researcher to determine which rainfall realizations are favorable or unfavorable given local crops and production calendars. By instead using realized NDVI as the productivity outcome and weather as the exogenous predictor, the approach allows the consequences of heterogeneous production environments to be reflected in realized vegetation productivity without requiring a common crop calendar or a common definition of good" and bad" weather across households.

The resulting measure should therefore be interpreted as a conservative measure of annual weather-induced agricultural productivity: the component of realized annual vegetation productivity that can be systematically predicted from exogenous annual rainfall and temperature. It is not intended to capture every agronomic channel through which weather affects individual crops. This distinction is central to the identification strategy, which prioritizes a transparent and comparable source of exogenous productivity variation over a complete model of crop-specific weather responses.

C.4 Validation of the Weather-Productivity Shock Using Household Crop Output

To examine whether the predicted weather-productivity shock captures economically meaningful variation in agricultural productivity, we relate the predicted shock to household-level crop output. Household-level crop output is constructed by aggregating production across all plots cultivating the same crop within a household. We estimate the following specification:

$$ln(Y_{icrt}) = \beta\theta_{it} + \gamma_{cr} + \delta_t + \varepsilon_{icrt}$$

where $ln(Y_{icrt})$ denotes household-level output of crop c produced by household i in community r and year t, θ_{it} is the predicted weather-productivity shock, γ_{cr} are community *times* crop fixed effects, and δ_t are year fixed effects. Standard errors are clustered at the community level.

Table 16 reports the estimated results. Across all three countries, the estimated coefficients are positive,

indicating that favorable weather-induced productivity shocks are associated with higher household crop output. The relationship is statistically significant at the 1 percent level in Niger and Ethiopia and at the 10 percent level in Malawi. These findings provide direct evidence that the Random Forest prediction captures economically meaningful variation in agricultural productivity rather than generic vegetation dynamics. This evidence complements the household agricultural income validation reported in Table 17 and supports the interpretation of the predicted component of NDVI growth as a weather-induced agricultural productivity shock used throughout the empirical analysis.

Table 16: Validation of the Weather-Productivity Shock Using Household Crop Output

	Niger	Malawi	Ethiopia
$\Delta\theta_{gt-1}$	3.133^{***}	12.218^{*}	27.693^{***}
	(1.066)	(6.961)	(7.287)
Observations	9,584	4,887	7,248
Community clusters	181	225	213
R^2	0.523	0.584	0.778
Community $\times$ crop fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes

Notes: The dependent variable is the logarithm of household-level crop output, constructed by summing production across all plots cultivated with the same crop by a household. The weather-productivity shock is the contemporaneous weather-induced component of NDVI growth predicted by the Random Forest model. All specifications include community $\times$ crop fixed effects and year fixed effects. Standard errors, reported in parentheses, are clustered at the community level. Singleton observations are dropped by the fixed-effects estimator.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D Weather Productivity Shock

D.1 Weather Productivity Growth as the Treatment Variable

In this section, we use the weather productivity variable as the exogenous treatment for the household. Before we run into the empirical results, we first need to discuss why the weather productivity is a good treatment variable to estimate an unbiased effect on the labor reallocation decision in the household.

Recall that in Section 2.4, we capture the variation instead of the trends of the weather productivity $\Delta\theta_g$ and total factor productivity change TFP. The positive variation $\Delta\theta_g$ means good weather conditions for crops, while a negative variation represents bad weather conditions for crops. The weather productivity can be decomposed into the long-run productivity changes from climate change, which is a trend, and the weather productivity variation we study, which is the variation around the trend.

In Section 2.4, both weather productivity growth $\Delta\theta_g$ and human productivity growth ΔTFP affect the labor reallocation ΔL. Besides, $\Delta\theta_g$ also affects ΔTFP, indicating that $\Delta\theta_g$ has both direct and indirect (through ΔTFP) effects on labor reallocation. Therefore, we first discuss the correlation (ϕ) between $\Delta\theta_g$ and ΔTFP.

Firstly, $\Delta\theta_g$ and ΔTFP can be complementary ($\phi > 0$). Farmers observe good weather during the growing

season (e.g., timely rains, no drought) after planting. Since land and crop choices are fixed, they cannot expand land or choose other crops, but they can take more effort on the land to take advantage of good weather conditions. They can adjust labor timing, allocate fertilizer more effectively, improve monitoring and weeding, and prevent pests/diseases more effectively. These adjustments are not changes in observable inputs, but improvements in unobserved managerial efficiency, captured by the human TFP in our model. In this case, $\Delta TFP(\Delta\theta_g)$ increases in $\Delta\theta_g$. That is, with both higher weather productivity and human productivity, households achieve higher agricultural income and can decide how to reallocate the labor.

Secondly, $\Delta\theta_g$ and ΔTFP can be substitute ($\phi < 0$). When the weather is good, farmers may exert less unobserved effort. For example, they will reduce field monitoring, reduce the adjustment in fertilizers, and reduce task coordination. These operate through unobserved management quality. Observable inputs remain unchanged, and the decline appears in lower human TFP. In this case, $\Delta TFP(\Delta\theta_g)$ decrease in $\Delta\theta_g$. Although human productivity may decline when favorable weather substitutes for household effort, we generally expected the direct biophysical gains from improved weather conditions to outweight the reduction in effort on average. Consequently, favorable weather is expected to increase agricultural productivity, although the magnitude of this increase depends on the strength of households' behavioral reponses.

Finally, it is possible that $\Delta\theta_g$ and ΔTFP are independence ($\phi = 0$). Weather conditions affect only the biophysical component (θ_g), not human TFP. This serves as the baseline benchmark commonly adopted in production economics and Integrated Assessment Models(IAM). Weather affects labor only through the biophysical contribution to agricultural income. The effect is cleanly interpreted as "weather-driven income change $\rightarrow$ labor reallocation."

The discussion above illustrates that different behavioral responses may alter the magnitude of the relationship between weather productivity and realized agricultural productivity. Nevertheless, under plausible behavioral responses, favorable weather is expected to be positively associated with agricultural productivity on average. Rather than relying on a strict theoretical monotonicity assumption, we empirically validate this relationship in Table 17, which shows that the weather-induced productivity shocks are strongly and positively associated with realized household agricultural income across all three countries.

Table 17: Household Agricultural Income and Weather Productivity Growth

	(1)	(2)	(3)
	Niger	Malawi	Ethiopia
$\Delta\theta_{gt-1}$	4.162**	4.681**	1.874*
	(2.45)	(2.07)	(1.67)
Constant	1.090***	1.632***	0.295***
	(33.25)	(29.54)	(4.81)
Household Fixed Effect	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes
Mean dep. var	1.417	1.729	0.508
SD dep. var	1.396	1.810	1.536
Observations	2546	1708	1442
R-squared	0.708	0.538	0.638

$*$ t statistics in parentheses

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

D.2 Identification of the Total Effect

The reduced-form analysis estimates the total effect of weather-induced agricultural productivity on household labor allocation. Agricultural productivity is determined by both exogenous weather conditions and household-specific productivity factors, such as technology adoption, farming experience, and management practices. Let agricultural productivity be written as

$$\Delta A_{it} = \Delta\theta_{gt} + \Delta P_{it}, \tag{53}$$

where $\Delta\theta_{gt}$ denotes the weather-induced component determined by rainfall and temperature, and ΔP_{it} denotes all other changes in household productivity.

Although household productivity may respond to weather conditions, it can be decomposed into a weather-related component and a remaining component through the linear projection

$$\Delta P_{it} = \phi\Delta\theta_{gt} + X_{it}'\eta + R_{it}, \qquad \mathbb{E}[R_{it} \mid \Delta\theta_{gt}, X_{it}] = 0, \tag{54}$$

where R_{it} captures the component of ΔP_{it} that is orthogonal to weather after conditioning on the household controls and fixed effects included in the empirical specification.

Suppose household labor allocation satisfies

$$s_{ijt} = \beta_1\Delta\theta_{gt-1} + \beta_2\Delta P_{it} + X_{it}'\pi + u_{it}, \qquad \mathbb{E}[u_{it} \mid \Delta\theta_{gt-1}, X_{it}] = 0, \tag{55}$$

where s_{ijt} denotes the household labor share in sector j, X_{it} contains household-level controls, and u_{it} is an idiosyncratic error term.

Substituting Equation (54) into Equation (55) yields

$$s_{ijt} = (\beta_1 + \phi\beta_2)\Delta\theta_{gt} + X_{it}'(\pi + \beta_2\eta) + \varepsilon_{it}, \tag{56}$$

where

$$\varepsilon_{it} = \beta_2 R_{it} + u_{it}.$$

The household controls enter only through the combined coefficient vector $(\pi + \beta_2\eta)$. Consequently, the coefficient on $\Delta\theta_{gt-1}$remains

$$\beta_{\text{tot}} = \beta_1 + \phi\beta_2$$

which captures the total effect of weather-induced agricultural productivity on household labor allocation.

By construction,

$$\mathbb{E}[R_{ht} \mid \Delta\theta_{gt}, X_{ht}] = 0,$$

and the exogeneity assumption implies

$$\mathbb{E}[u_{ht} \mid \Delta\theta_{gt}, X_{ht}] = 0.$$

Therefore,

$$\mathbb{E}[\varepsilon_{ht} \mid \Delta\theta_{gt}, X_{ht}] = 0,$$

which implies that the coefficient on $\Delta\theta_{gt}$ consistently identifies

$$\beta_{\text{tot}} = \beta_1 + \phi\beta_2.$$

Hence, the coefficient on weather shocks consistently identifies the total effect of weather-induced agricultural productivity on household labor allocation, incorporating both the direct effect of weather and the indirect effect through weather-induced adjustments in household productivity.

D.3 Other Endogeneity Concerns

The reduced-form analysis exploits variation in weather-induced agricultural productivity, $\Delta\theta_{gt}$, which is constructed from exogenous rainfall and temperature conditions at the community level. Because weather conditions are determined by nature rather than household decisions, they provide plausibly exogenous variation in agricultural productivity that helps mitigate several common endogeneity concerns.

First, the approach addresses simultaneity between agricultural productivity and labor allocation. Observed agricultural productivity and household labor decisions are jointly determined, since labor input affects agricultural output while expected returns influence labor allocation. By using weather-induced productivity rather than observed productivity, the estimated variation is predetermined with respect to contemporaneous household labor decisions.

Second, the specification reduces omitted-variable bias. Household characteristics such as farming ability, wealth, land quality, infrastructure, credit access, and other persistent local conditions may jointly affect agricultural productivity and labor allocation. Household fixed effects absorb all time-invariant household heterogeneity, while year fixed effects capture common macroeconomic shocks and other aggregate factors affecting all households. The remaining identifying variation therefore comes from within-household changes in weather-predicted agricultural productivity over time.

Third, the use of weather shocks alleviates concerns about reverse causality. Household production decisions, such as irrigation, input use, or labor management, may affect observed agricultural outcomes or vegetation indices, but they cannot influence rainfall or temperature. Moreover, weather-induced productivity is predicted out-of-sample using a Random Forest model trained on satellite-based weather and vegetation data, reducing mechanical feedback between household decisions and the constructed productivity measure.

Finally, the use of satellite-based weather and remote-sensing data minimizes classical measurement error arising from household surveys. Because the weather variables are measured independently of household responses, measurement error in reported agricultural outcomes or labor allocation does not contaminate the construction of the weather-induced productivity measure.

E Empirical Robustness Check

E.1 Alternative definition of "idle labor"

One concern is that the estimated increase in idle labor following weather productivity shocks may reflect changes in educational participation rather than labor market outcomes. Although our baseline definition of idle labor already excludes individuals engaged in employment and focuses on those who are available for work or exhibit weak labor-force attachment, some school-age individuals may still be classified as idle. To address this concern, we construct a more restrictive measure of idle labor that excludes individuals whose years of education increase between survey waves.

The first set of columns in Table 18 reports the results using this "no students" definition. The results are nearly identical to the baseline estimates. Weather-induced productivity shocks continue to increase the idle labor share significantly in Ethiopia, while the estimated effects remain small and statistically insignificant in Niger and Malawi. This suggests that the increase in idle labor observed in Ethiopia is not driven by school-

ing decisions.

As a further robustness check, we restrict the sample to prime-age individuals aged 25–60. This restriction removes nearly all potential schooling transitions while also excluding retirement-age individuals, thereby focusing on individuals with the strongest attachment to the labor market. The second set of columns in Table 18 again produces results that are broadly consistent with the baseline findings. Weather productivity shocks continue to generate a large and statistically significant increase in idle labor in Ethiopia. In Niger, the estimated effect remains small and statistically insignificant. In Malawi, the estimated coefficient becomes positive and marginally significant under the prime-age definition, although it's magnitude remains substantially smaller than that observed in Ethiopia.

Overall, these robustness checks indicate that the baseline findings are not sensitive to alternative definitions of idle labor. Although the prime-age specification yields a small positive idle-labor response in Malawi, the effect is modest compared with the consistently large and statistically significant increase observed in Ethiopia across both alternative definitions. These findings continue to support the interpretation that labor absorption constraints are substantially more severe in Ethiopia, whereas workers released from agriculture in Niger and, to a much greater extent, Malawi are generally more successfully absorbed into productive non-agricultural employment.

Table 18: Alternative Idle Workers and Weather Productivity Growth

	(1)	(2)	(3)	(4)	(5)	(6)
	Niger		Malawi		Ethiopia	
Idle Workers	No Students	Age 25-60	No Students	Age 25-60	No Students	Age 25-60
$\Delta\theta_{gt-1}$	0.0591	0.0869	0.139	0.237*	1.064**	1.107**
	(0.55)	(0.75)	(1.59)	(1.87)	(2.20)	(2.08)
_cons	0.0280***	0.0266***	0.0114***	0.0156***	0.170***	0.164***
	(6.34)	(6.04)	(4.81)	(4.68)	(8.86)	(8.47)
Community Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var	0.028	0.028	0.005	0.006	0.196	0.191
SD dep. var	0.134	0.134	0.053	0.062	0.380	0.376
Observations	3448	3254	1584	1450	1880	1604
R-squared	0.611	0.613	0.507	0.511	0.685	0.692

$*$ t statistics in parentheses

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

E.2 Pre-trend test for weather shocks

A potential concern is that weather productivity shocks may proxy for underlying development trends rather than represent exogenous short-run shocks. For instance, areas with improving climate conditions may also experience long-run development changes—such as infrastructure expansion, policy reforms, or improved market integration—that independently facilitate structural transformation. In this case, the estimated rela-

tionship between weather shocks and labor reallocation could reflect these concurrent trends rather than the causal effects of agricultural productivity changes.

To address this concern, we conduct a placebo test using the weather productivity measure from the current agricultural year, denoted as $\Delta\theta_{gt}$. Although this measure is constructed from weather conditions occurring in the same calendar year as the household survey, it serves as a placebo because agricultural productivity is only fully realized at the end of the harvest season. The LSMS labor module records individuals' labor allocation over the preceding 12 months, implying that the reported labor decisions are largely determined before the current year's harvest outcomes are observed. Consequently, from the perspective of household labor allocation, $\Delta\theta_{gt}$ represents information that has not yet been realized and therefore should not systematically predict current labor allocation. A significant coefficient would instead suggest that the weather productivity measure is capturing underlying trends, such as gradual economic development or policy changes, rather than exogenous weather-induced productivity shocks.

Table 19 reports the results. Across all three countries and all sectoral outcomes, including agricultural, manufacturing, service, and idle labor shares, the coefficients on $\Delta\theta_{gt}$ are uniformly small and statistically insignificant. These findings provide reassurance that our baseline results are not driven by spurious trends or omitted factors correlated with weather productivity over time.

The absence of any systematic relationship between future weather shocks and current labor allocation suggests that differential development trends—such as policy changes or infrastructure improvements—are unlikely to drive the main results. These findings support the interpretation that lagged weather productivity shocks capture exogenous variation in agricultural productivity and provide causal evidence on labor reallocation responses.

Table 19: Household Sectoral Share and Future Weather Productivity Growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Niger				Malawi				Ethiopia			
	Agriculture	Manufacturing	Service	Idle	Agriculture	Manufacturing	Service	Idle	Agriculture	Manufacturing	Service	Idle
$\Delta\theta_{gt}$	0.416	0.134	-0.500	-0.0661	0.530	-0.282	-0.0102	-0.112	-0.0290	-0.0521	-0.0658	0.0354
	(1.00)	(0.49)	(-1.26)	(-0.40)	(1.57)	(-1.22)	(-0.06)	(-1.24)	(-0.06)	(-0.49)	(-0.28)	(0.09)
_cons	0.603***	0.133***	0.228***	0.0299***	0.802***	0.0356***	0.145***	0.00805***	0.701***	0.0224***	0.0922***	0.148***
	(78.04)	(24.91)	(33.24)	(11.08)	(116.61)	(7.18)	(29.57)	(4.60)	(24.87)	(3.36)	(7.87)	(6.62)
Household Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var	0.634	0.118	0.209	0.028	0.776	0.052	0.154	0.005	0.660	0.023	0.094	0.182
SD dep. var	0.406	0.276	0.333	0.135	0.357	0.188	0.301	0.054	0.462	0.134	0.265	0.371
Observations	3780	3780	3780	3780	1910	1910	1910	1910	2240	2240	2240	2240
R-squared	0.681	0.626	0.637	0.610	0.721	0.591	0.698	0.506	0.721	0.632	0.624	0.674

* t statistics in parentheses

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

E.3 Alternative weather measure

To further assess the validity of the weather productivity shock, we examine whether sectoral labor allocation responds to alternative weather measures using growing-season rainfall and temperature. Unlike the

constructed weather productivity shock, raw weather variables do not necessarily map monotonically into agricultural productivity. In particular, a monotonic increase or decrease in rainfall or temperature does not always translate into higher productivity, as both insufficient and excessive weather realizations may harm crop outcomes. To account for this non-monotonic relationship, we allow for nonlinear effects by including quadratic terms for rainfall and temperature. Table 20 reports the regression results, and the corresponding predicted relationships are illustrated in Figures 11-13.

The results show consistent nonlinear patterns across countries. In Niger, rainfall exhibits a statistically significant inverted U-shaped relationship with agricultural labor share: the linear rainfall coefficient is positive while the squared term is negative and significant. This implies that increases in rainfall raise agricultural labor participation at low levels but reduce it beyond a threshold, consistent with diminishing returns or adverse effects from excessive rainfall. Temperature also exhibits a nonlinear relationship, with a negative linear term and a positive quadratic term, implying a U-shaped relationship between temperature and agricultural labor share. This pattern suggests that the marginal effect of temperature depends on the initial temperature level, although the estimated relationship should be interpreted cautiously within the range of temperatures observed in the data.

In Ethiopia, we observe similar nonlinear responses. Rainfall again follows an inverted U-shaped pattern, with both the linear and squared terms statistically significant. Temperature effects are also nonlinear, although the signs differ slightly, reflecting country-specific climatic conditions. Together, these results indicate that agricultural labor allocation responds systematically to weather conditions in Ethiopia.

In Malawi, the estimated coefficients for rainfall and temperature are statistically insignificant. Although the estimated relationships do not contradict the baseline specification, they provide weaker evidence of nonlinear weather effects than those observed in Niger and Ethiopia. The lack of statistical significance may reflect lower climatic variation or reduced statistical power in the Malawi sample, limiting the precision with which nonlinear responses can be estimated.

Overall, the results indicate that nonlinear responses to rainfall and temperature are most clearly evident in Niger and Ethiopia. For Malawi, the alternative specification provides weaker supporting evidence, but the estimates do not suggest a qualitatively different mechanism. Taken together, these findings are consistent with the interpretation that weather conditions affect sectoral labor allocation primarily through agricultural productivity, while highlighting that the strength of the evidence varies across countries.

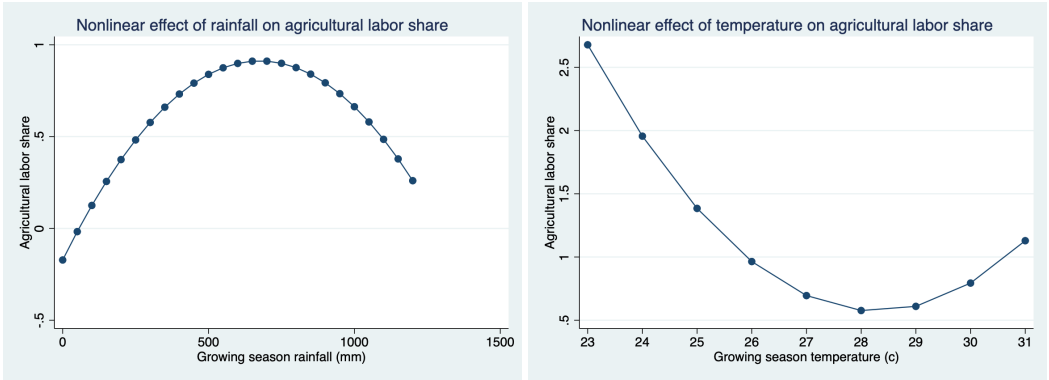
Table 20: Agricultural Household Sectoral Share and Weather Variables

	(1)	(2)	(3)	(4)	(5)	(6)
	Niger		Malawi		Ethiopia	
	Rainfall	Temperature	Rainfall	Temperature	Rainfall	Temperature
rain_mm	0.00321***		0.000295		0.00267*	
	(4.34)		(0.42)		(1.77)	
c.rain_mmc.rain_mm	-0.00000237***		-0.000000105		-0.00000123*	
	(-3.11)		(-0.34)		(-1.73)	
temp_c		-4.273**		-0.0147		0.526**
		(-2.34)		(-0.03)		(1.98)
c.temp_cc.temp_c		0.0755**		-0.000763		-0.00747
		(2.39)		(-0.06)		(-1.64)
_cons	-0.234	60.96**	0.644*	1.496	-0.419	-6.633*
	(-1.40)	(2.32)	(1.73)	(0.25)	(-0.67)	(-1.85)
Household Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var	0.634	0.634	0.793	0.793	0.660	0.660
SD dep. var	0.406	0.406	0.350	0.350	0.462	0.462
Observations	3780	3780	1174	1174	2240	2240
R-squared	0.689	0.682	0.730	0.730	0.726	0.729

$*$ t statistics in parentheses

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

Figure 11: Household Sectoral Share and Weather Variables in Niger

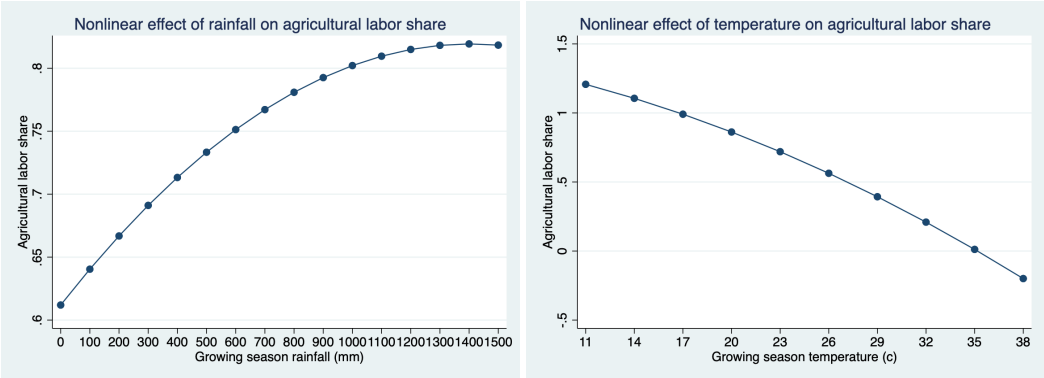


(a) Rainfall **(b)** Temperature

E.4 Alternative subsistence measure

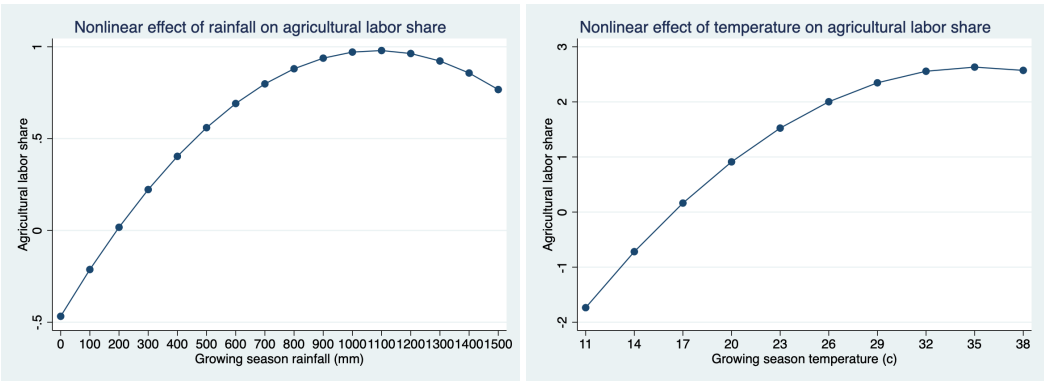
To further assess the role of subsistence constraints, we construct an alternative proxy for subsistence based on household income. While the baseline analysis relies on a survey-based food subsistence indicator, income provides an additional measure of household resource availability. Specifically, we define a relative

Figure 12: Household Sectoral Share and Weather Variables in Malawi



(a) Rainfall **(b)** Temperature

Figure 13: Household Sectoral Share and Weather Variables in Ethiopia



(a) Rainfall **(b)** Temperature

poverty threshold using the bottom 25 percent of the within-country distribution of per capita household income. Households below this threshold are classified as low-income and interpreted as being closer to subsistence, while those above the threshold are considered relatively unconstrained. We then estimate the effect of weather productivity shocks separately for these two groups.

Table 21 reports the results. In Niger and Malawi, productivity shocks lead to larger reductions in agricultural labor share among households above the relative poverty line, while the response among low-income households is weaker and statistically insignificant. These patterns are consistent with subsistence mechanism, whereby households below to subsistence are less able to release labor from agriculture following productivity improvements. In Ethiopia, however, the responses do not differ systematically by income group. Although the estimated effect for higher-income households is statistically significant, the point estimate for lower-income households is of similar magnitude but estimated less precisely. This pattern suggests that income proximity to subsistence is not the primary factor governing labor release in Ethiopia.

Overall, these findings reinforce the heterogeneous mechanisms identified in the main analysis. In Niger and Malawi, lower income, which implies stronger food subsistence constraints, limits households' ability to release labor from agriculture. By contrast, the appendix provides little evidence that income-based subsistence

constraints are the primary barrier in Ethiopia, consistent with the main results showing that Ethiopia's slower structural transformation is instead driven by downstream labor reallocation and labor absorption frictions.

The results are robust to alternative definitions of the relative poverty threshold. When we instead classify households in the bottom 10 percent of the income distribution as low-income, we obtain qualitatively similar patterns across countries. Productivity shocks continue to induce stronger labor reallocation among higher-income households, while the response among low-income households remains limited. Taken together, these results provide additional evidence supporting the subsistence mechanism and reinforce the interpretation that subsistence constraints play an important role in shaping labor reallocation, particularly in more constrained environments such as Ethiopia.

Table 21: Alternative Subsistence Criteria and Weather Productivity Growth

	(1)	(2)	(3)	(4)	(5)	(6)
	Niger		Malawi		Ethiopia	
	$\leq$ 25% HH Income	$>$ 25% HH Income	$\leq$ 25% HH Income	$>$ 25% HH Income	$\leq$ 25% HH Income	$>$ 25% HH Income
$\Delta\theta_{gt-1}$	0.404	-0.718*	-0.547	-1.232**	-1.275	-1.088*
	(0.72)	(-1.91)	(-0.59)	(-2.37)	(-1.10)	(-1.90)
_cons	0.805***	0.561***	0.781***	0.727***	1.001***	0.609***
	(35.40)	(39.49)	(29.26)	(59.92)	(25.56)	(26.83)
Household Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var	0.817	0.567	0.855	0.723	0.954	0.594
SD dep. var	0.368	0.400	0.311	0.380	0.204	0.478
Observations	368	2218	108	1038	98	1622
R-squared	0.700	0.696	0.482	0.747	0.776	0.724

$*$ t statistics in parentheses

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

E.5 Alternative Definition of Working-Age Population

In the baseline analysis, we define the working-age population as individuals aged 15-85, reflecting the fact that in developing countries many individuals remain economically active at older ages because of limited formal retirement systems. To ensure that our findings are not driven by this age definition, we conduct robustness checks using alternative age ranges.

Specifically, we re-estimate the main regressions by restricting the working-age population to ages 15–70 and 15–75. Tables 22 and 23 report the results.

Across both alternative age definitions, the estimated effects of weather productivity shocks remain highly consistent with the baseline results. In Niger, positive weather productivity shocks continue to reduce the agricultural labor share while increasing manufacturing employment, with no statistically significant effect on the service or idle labor shares. In Malawi, productivity improvements likewise reduce agricultural employment and increase manufacturing participation, while the service sector remains largely unaffected. Although the estimated increase in idle labor is positive, it is economically small and only marginally significant,

consistent with the baseline finding that most released workers are successfully absorbed into productive non-agricultural employment.

The Ethiopian results are also highly robust. Weather productivity shocks continue to reduce the agricultural labor share while increasing manufacturing employment and generating little systematic change in service employment. Most importantly, the increase in idle labor remains positive and statistically significant under both alternative age definitions. This confirms that workers released from agriculture are substantially less likely to be absorbed into productive non-agricultural employment in Ethiopia than in Niger or Malawi, consistent with the paper's interpretation that limited labor demand constrains structural transformation.

Overall, both the magnitude and statistical significance of the estimated coefficients remain remarkably stable across alternative definitions of the working-age population. These robustness checks demonstrate that the paper's central conclusions regarding labor release, labor reallocation, and labor absorption are not sensitive to the choice of working-age cutoff, reinforcing the validity of the baseline results.

Table 22: Household Sectoral Share and Weather Productivity Growth (ages 15–70)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Niger				Malawi				Ethiopia			
	Agriculture	Manufacturing	Service	Idle	Agriculture	Manufacturing	Service	Idle	Agriculture	Manufacturing	Service	Idle
$\Delta\theta_{gt-1}$	-0.692**	0.562***	0.0652	0.0983	-0.939*	0.712**	-0.203	0.197*	-1.056**	0.200**	-0.184	0.899*
	(-2.20)	(3.22)	(0.25)	(0.92)	(-1.84)	(2.55)	(-0.67)	(1.83)	(-2.09)	(2.01)	(-0.94)	(1.94)
_cons	0.622***	0.115***	0.220***	0.0264***	0.834***	0.0515***	0.155***	0.0124***	0.707***	0.0200***	0.0936***	0.139***
	(49.72)	(15.34)	(19.76)	(6.47)	(68.90)	(8.47)	(20.46)	(4.60)	(35.11)	(4.60)	(11.46)	(8.05)
Household Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var	0.624	0.118	0.206	0.028	0.815	0.055	0.168	0.005	0.668	0.021	0.101	0.170
SD dep. var	0.402	0.275	0.331	0.135	0.526	0.198	0.323	0.055	0.455	0.128	0.273	0.359
Observations	3722	3722	3722	3722	1844	1844	1844	1844	1860	1860	1860	1860
R-squared	0.675	0.632	0.634	0.611	0.846	0.597	0.703	0.509	0.734	0.615	0.625	0.690

Robust t statistics in parentheses.

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

Table 23: Household Sectoral Share and Weather Productivity Growth (ages 15–75)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Niger				Malawi				Ethiopia			
	Agriculture	Manufacturing	Service	Idle	Agriculture	Manufacturing	Service	Idle	Agriculture	Manufacturing	Service	Idle
$\Delta\theta_{gt-1}$	-0.669**	0.578***	0.0489	0.0980	-0.917*	0.707**	-0.191	0.197*	-0.996***	0.168**	-0.253*	0.955***
	(-2.14)	(3.32)	(0.19)	(0.92)	(-1.82)	(2.54)	(-0.63)	(1.84)	(-2.76)	(2.45)	(-1.83)	(2.90)
_cons	0.624***	0.114***	0.223***	0.0261***	0.838***	0.0507***	0.155***	0.0122***	0.689***	0.0219***	0.0880***	0.156***
	(51.21)	(15.46)	(20.44)	(6.48)	(69.08)	(8.28)	(19.68)	(4.58)	(48.77)	(7.28)	(16.40)	(12.97)
Household Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var	0.630	0.118	0.208	0.028	0.822	0.055	0.167	0.005	0.659	0.023	0.094	0.181
SD dep. var	0.405	0.276	0.332	0.134	0.522	0.197	0.322	0.054	0.461	0.135	0.265	0.369
Observations	3758	3758	3758	3758	1880	1880	1880	1880	2214	2214	2214	2214
R-squared	0.680	0.630	0.634	0.612	0.846	0.596	0.703	0.509	0.729	0.634	0.626	0.685

Robust t statistics in parentheses.

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

E.6 Multiple Distance Measures

To examine whether our results depend on how spatial frictions are measured, we conduct robustness checks using alternative distance indicators. In the baseline specification, mobility barriers are captured using dis-

tance to a specific economic center. However, different measures of remoteness may capture distinct aspects of market access and employment opportunities. We therefore re-estimate the model using three alternative distance measures: distance to the nearest road, distance to the nearest market, and distance to the nearest population center. Each measure is interacted with weather productivity shocks to capture how spatial frictions influence labor reallocation.

Tables 24–26 report the results for Niger, Malawi, and Ethiopia, respectively. Across countries, the qualitative patterns are broadly consistent with the baseline findings, although the precision of the estimates varies across alternative distance measures

In Niger, the interaction terms generally suggest that greater distance attenuates the response of non-agricultural employment to productivity shocks, although the estimates are imprecisely estimated across the alternative distance measures.

In Malawi, productivity shocks continue to increase non-agricultural participation, and the interaction terms remain broadly consistent in sign across the alternative measures of remoteness. However, statistical significance varies across specifications, suggesting that the estimated effects are sensitive to the particular distance measure and the relatively limited spatial variation in the Malawi sample.

In Ethiopia, the interaction terms generally indicate that remoteness weakens households' responses to productivity shocks, consistent with the presence of spatial mobility barriers. Although the estimated effects differ in precision across alternative measures of distance to roads, markets, and population centers, the overall pattern remains qualitatively consistent with the baseline specification.

Overall, these robustness checks suggest that the main conclusions are not driven by the particular measure of remoteness adopted in the baseline analysis. While the statistical significance of individual interaction terms varies across alternative distance measures, the estimated coefficients generally point in the same direction and continue to support the interpretation that spatial frictions can impede labor reallocation following productivity shocks.

Table 24: Multiple Distance Measures and Weather Productivity Growth in Niger

	(1)	(2)	(3)	(4)	(5)	(6)
	Road		Market		Population Center	
	Non-agriculture	Idle	Non-agriculture	Idle	Non-agriculture	Idle
$\Delta\theta_{gt-1}$	0.337	-0.0172	0.746	-0.0368	1.193**	-0.165
	(0.91)	(-0.17)	(1.41)	(-0.25)	(2.41)	(-1.28)
$\Delta\theta_{gt-1}$ × distance to road	0.0121	0.00434				
	(1.27)	(0.99)				
$\Delta\theta_{gt-1}$ × distance to market			-0.00165	0.00149		
			(-0.30)	(0.70)		
$\Delta\theta_{gt-1}$ × distance to population center					-0.00626*	0.00268
					(-1.77)	(1.58)
Constant	0.341***	0.0272***	0.340***	0.0270***	0.340***	0.0271***
	(28.23)	(6.34)	(27.82)	(6.38)	(27.69)	(6.31)
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.327	0.028	0.327	0.028	0.327	0.028
SD dep. var.	0.392	0.135	0.392	0.135	0.392	0.135
Observations	3,780	3,780	3,780	3,780	3,780	3,780
R^2	0.679	0.610	0.678	0.610	0.679	0.611

Notes: The dependent variables are the household shares of labor allocated to non-agricultural employment and idle labor. Each pair of columns uses a different measure of spatial accessibility: distance to the nearest road, market, or population center. All regressions include household and year fixed effects. Standard errors are clustered at the community level. t-statistics are reported in parentheses.

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

Table 25: Multiple Distance Measures and Weather Productivity Growth in Malawi

	(1)	(2)	(3)	(4)	(5)	(6)
	Road		Market		Population Center	
	Non-agriculture	Idle	Non-agriculture	Idle	Non-agriculture	Idle
$\Delta\theta_{gt-1}$	0.720	0.270*	0.686	0.0710	0.858	0.247
	(1.47)	(1.86)	(1.06)	(0.55)	(1.12)	(1.19)
$\Delta\theta_{gt-1}$ × distance to road	-0.0569*	-0.0124				
	(-1.77)	(-1.47)				
$\Delta\theta_{gt-1}$ × distance to market			-0.0317	0.0157		
			(-0.62)	(1.27)		
$\Delta\theta_{gt-1}$ × distance to population center					-0.0173	-0.00167
					(-0.88)	(-0.37)
Constant	0.189***	0.0123***	0.189***	0.0125***	0.186***	0.0121***
	(21.66)	(4.82)	(20.33)	(4.62)	(20.16)	(4.55)
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.207	0.005	0.207	0.005	0.207	0.005
SD dep. var.	0.340	0.054	0.340	0.054	0.340	0.054
Observations	1,878	1,878	1,878	1,878	1,878	1,878
R^2	0.697	0.510	0.697	0.510	0.697	0.509

Notes: The dependent variables are the household shares of labor allocated to non-agricultural employment and underutilized labor. Each pair of columns uses a different measure of spatial accessibility: distance to the nearest road, market, or population center. All regressions include household and year fixed effects. Standard errors are clustered at the community level. t-statistics are reported in parentheses.

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

Table 26: Multiple Distance Measures and Weather Productivity Growth in Ethiopia

	(1)	(2)	(3)	(4)	(5)	(6)
	Road		Market		Population Center	
	Non-agriculture	Idle	Non-agriculture	Idle	Non-agriculture	Idle
$\Delta\theta_{gt-1}$	0.142	0.116	0.176	1.244**	0.0129	0.331
	(0.47)	(0.21)	(0.54)	(2.12)	(0.04)	(0.49)
$\Delta\theta_{gt-1}$ × distance to road	-0.0119	0.0432*				
	(-1.34)	(1.84)				
$\Delta\theta_{gt-1}$ × distance to market			-0.00309	-0.00368		
			(-1.17)	(-0.65)		
$\Delta\theta_{gt-1}$ × distance to population center					-0.00318	0.0201
					(-0.66)	(0.85)
Constant	0.113***	0.144***	0.112***	0.159***	0.110***	0.152***
	(11.99)	(8.07)	(12.19)	(9.45)	(12.44)	(8.97)
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.116	0.182	0.116	0.182	0.116	0.182
SD dep. var.	0.291	0.371	0.291	0.371	0.291	0.371
Observations	2,240	2,240	2,240	2,240	2,240	2,240
R^2	0.662	0.689	0.662	0.683	0.661	0.686

Notes: The dependent variables are the household shares of labor allocated to non-agricultural employment and underutilized labor. Each pair of columns uses a different measure of spatial accessibility: distance to the nearest road, market, or population center. All regressions include household and year fixed effects. Standard errors are clustered at the community level. *t*-statistics are reported in parentheses.

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

F Structural Estimation: Model Fit

This section evaluates the fit of the structural model by comparing empirical moments from the data with their simulated counterparts. The model is estimated using a Simulated Method of Moments (SMM) approach, targeting a set of moments that capture both aggregate labor allocation and heterogeneity across households.

The estimation targets three groups of moments. First, we match the average sectoral labor shares, including agricultural, manufacturing, and idle labor shares. These moments capture the overall distribution of labor across sectors. Second, we match moments by distance bins, which proxy for spatial frictions and access to non-agricultural opportunities. These moments capture how labor allocation varies with geographic access and help identify mobility barriers in the model. Finally, we match moments by climate bins, which capture heterogeneity in weather-driven productivity shocks.

Figures 14– 16 compare the empirical and simulated moments for Niger, Malawi, and Ethiopia. Overall, the model reproduces the main empirical patterns across the three countries. The simulated moments closely match the empirical moments used in estimation, including agricultural labor shares, idle labor shares, and

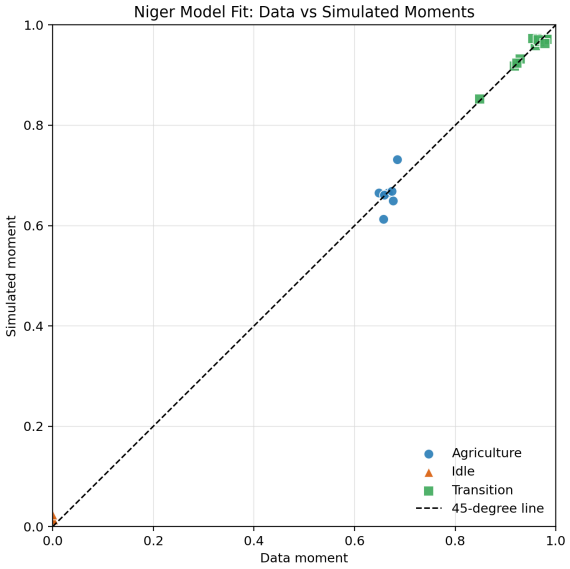


Figure 14: Model Fit in Niger

transition probabilities into productive non-agricultural employment.

In Niger, the model closely matches the aggregate labor-allocation moments and reproduces the observed variation across both distance and climate bins. In particular, it captures the relatively low idle labor share and the strong responsiveness of transitions into productive non-agricultural employment following favorable productivity shocks.

In Malawi, the model also fits the data well, particularly in terms of aggregate labor shares and their variation across climate conditions. The simulated moments closely match the empirical distribution across both distance and climate bins. While there are small deviations in some bins, especially for manufacturing shares, the model successfully captures the overall responsiveness of labor allocation to productivity and spatial conditions.

For Ethiopia, the model reproduces the main empirical patterns of labor allocation, including relatively high idle labor and weaker transitions into productive non-agricultural employment. Although the simulated moments deviate from the empirical moments for some agriculture and idle-labor bins, the model captures the qualitative relationships targeted by the estimation and provides a reasonable approximation for the counterfactual analysis.

Taken together, the model fit indicates that the structural framework is able to jointly match aggregate labor allocation and key dimensions of heterogeneity across households. In particular, the model successfully captures cross-country differences in labor reallocation patterns, including the stronger absorption of labor in Niger and Malawi and the accumulation of idle labor in Ethiopia. These results provide support for the model's ability to capture the underlying mechanisms driving structural transformation and justify its use for counterfactual policy analysis.

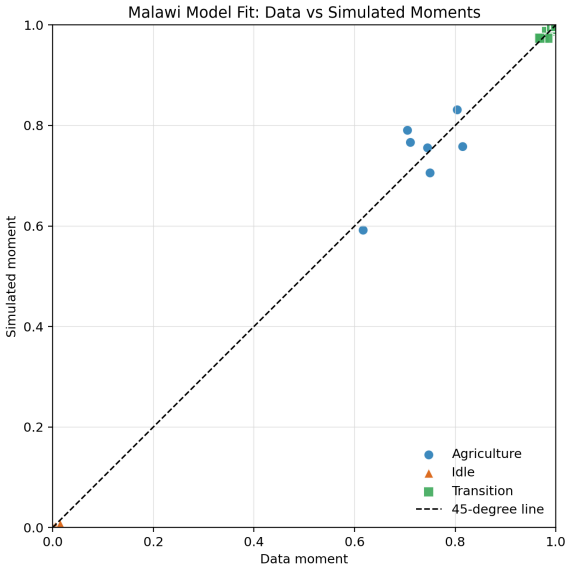


Figure 15: Model Fit in Malawi

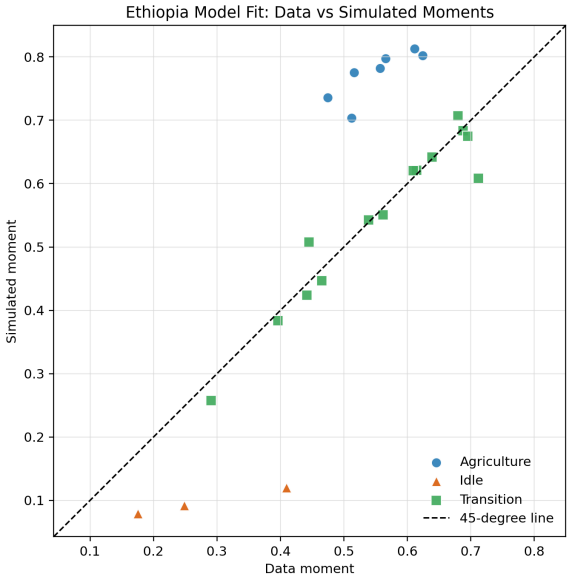


Figure 16: Model Fit in Ethiopia

G Policy Calibration

This appendix documents the calibration of the five counterfactual policy interventions. Each policy is assigned the same benchmark budget of USD 500 million to facilitate comparisons across interventions. The calibration follows a common two-step procedure.

First, the policy budget is converted into a country-specific implementation intensity using external information on policy costs together with country-specific demographic statistics. Depending on the intervention, the implementation intensity measures the share of households covered, kilometers of roads constructed, or jobs created.

Second, the implementation intensity is mapped into the corresponding structural parameter of the model. Food subsidies reduce household subsistence requirements, agricultural input programs increase agricultural productivity, rural road investments reduce migration costs, land reforms convert households from weaker to stronger land rights, and industrial investment increases the labor absorption probability. All demographic statistics are obtained from the World Bank World Development Indicators and the LSMS household surveys, while policy costs are collected from published policy evaluations and World Bank project documents.

G.1 Policy 1: Targeted Food Subsidy

The first policy represents targeted food transfers to working-age individuals in rural households whose food consumption falls below the estimated subsistence threshold. The objective is to relax food subsistence constraints and thereby reduce the amount of household labor that must remain in agriculture.

Let B denote the common policy budget. Let s_c^T denote the share of the baseline survey population that is simultaneously working age, rural, and living in a household below the food-subsistence threshold. This is a population-weighted target share and therefore differs from the household-level food-subsistence share reported in Table 1. The number of targeted individuals is calculated as

$$N_c^T = POP_c \times s_c^T, \tag{57}$$

where POP_c denotes the national population in the baseline year.

The transfer per targeted individual is

$$T_c^F = \frac{B}{N_c^T}. \tag{58}$$

Let F_c denote the annual food-subsistence expenditure benchmark per person. The policy intensity is defined as the transfer relative to this benchmark:

$$\lambda_c^F = \min\left\{1, \frac{T_c^F}{F_c}\right\}. \tag{59}$$

The policy is implemented for households classified as food-subsistence constrained by proportionally reducing the additional subsistence wedge:

$$\bar{c}_i' = c_0 \left[1 + \left(1 - \lambda_c^F\right) \delta Z_i\right] h_i, \tag{60}$$

where $Z_i = 1$ denotes a food-subsistence-constrained household and $Z_i = 0$ otherwise.

Table 27 reports the calibrated values of λ_c^F. The calibrated policy intensity differs across countries because both the annual food subsistence expenditure and the number of targeted households vary substantially. Countries with higher subsistence costs or a larger subsistence-constrained population receive a smaller transfer per household under the common policy budget, resulting in a smaller proportional reduction in the subsistence requirement.

Table 27: Calibration of Policy 1: Targeted Food Subsidy

Country	Target Share	Target Population N_c^T (million)	F_c (USD)	T_c^F (USD)	λ_c^F
Niger	0.120	2.058	273.75	242.90	0.887
Malawi	0.273	4.789	262.80	104.40	0.397
Ethiopia	0.022	2.527	427.05	197.88	0.463

Target share is the share of the baseline survey population that is simultaneously working age, rural, and living in a household below the food-subsistence thre
It is therefore a person-level joint share rather than the household-level food-subsistence share reported in Table 1. The targeted population is calcula
$N_c^T = POP_c s_c^T$, where POP_c is national population in the baseline year. F_c denotes the annual food-subsistence benchmark per person. The transfer per ta
individual is $T_c^F = B/N_c^T$, where $B = \$500$ million, and the calibrated policy intensity is $\lambda_c^F = \min\{1, T_c^F/F_c\}$.

G.2 Policy 2: Agricultural Productivity Improvement

The second policy represents an agricultural input program that increases farm productivity through fertilizer, improved seeds, or related input packages. The policy is calibrated as a budget-constrained productivity intervention targeted to rural households.

Let B denote the common policy budget, C_A denote the cost of providing an input package to one rural household, and N_c^R denote the number of rural households in country c. The share of rural households reached by the policy is

$$s_c^A = \min\left\{1, \frac{B}{C_A N_c^R}\right\}. \tag{61}$$

If each treated household experiences a proportional productivity gain g_A, the average productivity increase is

$$\lambda_c^A = s_c^A g_A. \tag{62}$$

In the benchmark calibration, I set the cost per household to $C_A = 80$ USD and the productivity gain among treated households to $g_A = 0.15$. These values are consistent with evidence from agricultural input programs in Sub-Saharan Africa, including fertilizer and improved seed subsidy programs.

The number of rural households is constructed from national population statistics and LSMS survey data. Let POP_c denote the national population, $\bar{h}_c$ denote average household size in the baseline survey, and s_c^R denote the share of households classified as rural in the baseline LSMS survey. Total households are

$$HH_c = \frac{POP_c}{\bar{h}_c}, \tag{63}$$

and rural households are

$$N_c^R = s_c^R HH_c. \tag{64}$$

The policy is implemented as a proportional increase in agricultural productivity,

$$A_i' = (1 + \lambda_c^A) A_i. \tag{65}$$

Table 28 reports the country-specific calibration. Because the common budget covers all rural households in Niger and Malawi under the benchmark cost, the calibrated productivity increase equals the treated-household gain of 15 percent in both countries. In Ethiopia, the rural target population is much larger, so the same budget covers only part of the rural household population and implies a smaller average productivity increase.

Table 28: Calibration of Policy 2: Agricultural Productivity Improvement

Country	Total Households HH_c	Rural Share s_c^R	Rural Households N_c^R	Coverage	λ_c^A
Niger	3.953	0.637	2.517	1.000	0.150
Malawi	4.302	0.708	3.046	1.000	0.150
Ethiopia	35.420	0.419	14.844	0.421	0.063

Notes: Household counts are reported in millions. Total households are computed as national population divided by average household size in the baseline s
Rural share is calculated from the LSMS baseline sample. Rural households are $N_c^R = s_c^R HH_c$. Coverage is $B/(C_A N_c^R)$, capped at one. The calibrated produc
increase is $\lambda_c^A = \min\{1, B/(C_A N_c^R)\} g_A$, with $B =$ \$500 million, $C_A = 80$ USD per household, and $g_A = 0.15$.

G.3 Policy 3: Rural Road Infrastructure

The third policy represents investments in rural roads and transport infrastructure. In the structural model, these investments reduce spatial mobility barriers by lowering the effective cost of reaching non-agricultural opportunities.

Let B denote the common policy budget and let C_c^R denote the country-specific cost per kilometer of rural road construction. The number of kilometers financed by the policy is

$$K_c^R = \frac{B}{C_c^R}. \tag{66}$$

This calibration is intended to reduce the magnitude of the estimated distance friction toward zero. For countries where the estimated distance coefficient is positive but statistically insignificant, such as Malawi, the proportional adjustment has a negligible mechanical effect, reflecting that spatial mobility barriers are not

empirically important in the estimated model.

Let $\bar{K}_c^R$ denote the country-specific rural road gap measured in kilometers. For Niger and Malawi, the World Bank transport assessments report rural accessibility gaps as percentages rather than missing road length. To obtain a road-gap measure in kilometers, I first approximate each country's rural road stock and then multiply it by the reported accessibility gap. Specifically, for Niger, the rural road stock is proxied by the sum of regional roads (2,075 km) and unclassified roads (6,952 km), giving 9,027 km. Since approximately 40 percent of agricultural and agro-pastoral zones are reported to be isolated, the rural road gap is calculated as

$$\bar{K}_{NGR}^R = 0.40 \times 9,027 = 3,611 \text{ km}.$$

For Malawi, the classified road network is 15,551 km, of which secondary and tertiary roads account for approximately 75 percent, yielding an estimated rural road stock of 11,663 km. Because only 23.1 percent of the rural population has access to an all-season road, the rural road gap is calculated as

$$\bar{K}_{MAL}^R = (1 - 0.231) \times 11,663 = 8,965 \text{ km}.$$

For Ethiopia, the road gap is measured directly from the Roads Development Program as the difference between the planned and existing rural road network:

$$\bar{K}_{ETH}^R = (63,542 + 108,797) - (33,618 + 79,806.9) = 58,914.1 \text{ km}.$$

The share of the rural road gap closed by the policy is

$$s_c^R = \min\left\{1, \frac{K_c^R}{\bar{K}_c^R}\right\}. \tag{67}$$

Because road investments reduce only the transportation component of mobility barriers, I assume that complete closure of the rural road gap removes only a fraction $\bar{\rho}^R$ of the estimated distance friction. In the benchmark calibration, I set $\bar{\rho}^R = 0.5$. The resulting policy intensity is therefore

$$\lambda_c^R = s_c^R \bar{\rho}^R. \tag{68}$$

The policy is implemented by proportionally shrinking the estimated distance-friction parameter toward zero:

$$\kappa_{1c}' = (1 - \lambda_c^R)\kappa_{1c}. \tag{69}$$

Road construction costs are obtained from the World Bank Road Costs Knowledge System (ROCKS), using country-specific average costs for rural gravel reconstruction and closely related road works.

Table 29: Calibration of Policy 3: Rural Road Infrastructure

Country	Road Cost C_c^R (USD/km)	Road Gap $\bar{K}_c^R$ (km)	Financed Roads K_c^R (km)	Gap Closed s_c^R	$\bar{\rho}^R$	λ_c^R
Niger	15,960	3,611	31,335	1.000	0.500	0.500
Malawi	37,617	8,965	13,291	1.000	0.500	0.500
Ethiopia	81,500	58,914	6,135	0.104	0.500	0.052

Notes: $B = \$500$ million. Financed roads are calculated as $K_c^R = B/C_c^R$. The share of the road gap closed is $s_c^R = \min\{1, K_c^R/\bar{K}_c^R\}$. The benchmark calibratio $\bar{\rho}^R = 0.5$, so that full closure of the rural road gap removes one-half of the estimated distance friction. The final policy intensity is $\lambda_c^R = s_c^R \bar{\rho}^R$, and the counterf migration-cost parameter is $\kappa'_{1c} = (1 - \lambda_c^R)\kappa_{1c}$.

G.4 Policy 4: Strengthening household land-rights

The fourth policy represents land certification and tenure reforms that strengthen the land rights. These reforms reduce institutional barriers to structural transformation by allowing households to retain or transfer the value of agricultural land while participating in non-agricultural activities. Unlike the previous policies, land reform changes households' institutional environment rather than directly modifying an estimated structural parameter.

Let B denote the common policy budget and let C_L denote the cost of providing land certification or tenure reform to one household. Let N_c^{NT} denote the number of households with weak land rights in country c. The policy coverage is

$$\lambda_c^L = \min\left\{1, \frac{B}{C_L N_c^{NT}}\right\}, \tag{70}$$

where λ_c^L is interpreted as the share of weak land-right households that receive the reform.

The number of households with weak land is calculated as

$$N_c^{NT} = HH_c \times s_c^{NT}, \tag{71}$$

where HH_c denotes the total number of rural households and s_c^{NT} is the share of households classified as having weak land rights in the baseline LSMS survey.

The benchmark calibration follows the World Bank's estimates of land certification costs. The average implementation cost is approximately USD 50 per household, including cadastral mapping, registration, certificate issuance, and administrative costs. Under the common policy budget of USD 500 million, the available resources are sufficient to provide land certification to all households with weak land rights in the three study countries. Consequently,

$$\lambda_c^L = 1$$

for Niger, Malawi, and Ethiopia.

The policy is implemented by converting households with weak land rights into stronger land rights:

$$T_i' = \begin{cases} 1, & \text{if } T_i = 0, \\ T_i, & \text{otherwise.} \end{cases} \tag{72}$$

For each household, the probability of accessing non-agricultural opportunities is then recalculated as

$$q_i' = \Lambda\left(\mu_0^{NT}(1 - T_i') + \mu_0^T T_i' + \kappa_1 dist_i\right). \tag{73}$$

The policy therefore changes household outcomes by expanding the share of households with stronger land rights while leaving the estimated structural parameters unchanged.

Table 30: Calibration of Policy 4: household land-rights

Country	Weak land-right Share (s_c^{NT})	Weak land-right Households (N_c^{NT}, million)	Coverage (λ_c^L)	Converted Share
Niger	0.941	2.37	1.000	100%
Malawi	0.126	0.38	1.000	100%
Ethiopia	0.281	4.17	1.000	100%

Notes: Weak land-right households are identified as households with land type 0 in the baseline LSMS survey. Land types 1 and 2 are classified as stronge rights. The number of weak land-rights households is calculated as

$$N_c^{NT} = s_c^{NT} N_c^R,$$

where N_c^R is the total number of rural households reported in Table 28. The benchmark implementation cost is USD 50 per household. Under the common budget of USD 500 million, the available resources are sufficient to provide land certification to all households with weak land rights in each country. Consequ the policy coverage is $\lambda_c^L = 1$, and all Weak land-right households are converted to stronger land rights before the household-specific probability of accessing agricultural opportunities is recalculated.

G.5 Policy 5: Increasing Labor Absorption Capacity

Policy 5 represents industrial investment that expands the ability of the non-agricultural sector to absorb workers. Unlike the previous policies, which relax supply-side constraints faced by households, this intervention targets the demand side of structural transformation by increasing the probability that workers who successfully reallocate into the non-agricultural labor market obtain productive employment.

The calibration proceeds in two steps. First, the common policy budget is translated into the number of jobs created using evidence from recent industrial investment projects. Second, the implied job creation is converted into an increase in the labor absorption probability in the structural model. The structural model does not explicitly model vacancy creation or job matching. Instead, the effect of industrial investment is summarized by an additive increase in the probability that a worker seeking employment is successfully absorbed into productive non-agricultural employment.

Following the World Bank Competitiveness and Job Creation Project, I use the construction of the Bole Lemi II and Kilinto Industrial Parks in Ethiopia as the benchmark investment. The project provided approximately USD 425 million in financing and created approximately 19,000 direct jobs. Scaling this benchmark to the

common policy budget of USD 500 million implies

$$J_c = \frac{500}{425} \times 19{,}000 = 22{,}353 \tag{74}$$

new jobs.

To account for differences in labor market size across countries, the same investment is expressed relative to the unemployed population in each country. Let U_c denote the number of unemployed workers in the benchmark labor market. The calibrated increase in labor absorption capacity is

$$\lambda_c^D = \frac{J_c}{U_c}. \tag{75}$$

Because the model-implied pool of reallocated workers seeking employment is not directly observed, I use the observed unemployed population in the benchmark labor market as its empirical proxy. Unemployed workers constitute the closest observable group of individuals who have entered the labor market but have not yet been absorbed into productive employment. Under the stylized assumption that newly created jobs are allocated proportionally among workers seeking employment, the ratio J_c/U_c approximates the increase in the probability that a reallocated worker successfully obtains a productive non-agricultural job.

$$p_c' = \min\left\{1, p_c + \lambda_c^D\right\}, \tag{76}$$

where p_c denotes the estimated labor absorption probability from the structural model.

The objective of the counterfactual analysis is to isolate the contribution of each structural friction to the process of structural transformation. Accordingly, each policy perturbs a single model parameter while holding all other mechanisms fixed. The labor-absorption policy should be interpreted as a stylized, non-spatial simulation that abstracts from the geographic placement of industrial parks and their interaction with household-specific migration costs. Instead, the policy is modeled as a uniform increase in the labor-absorption probability faced by workers who have successfully entered the non-agricultural labor market.

For Ethiopia, the benchmark labor market is Addis Ababa because both industrial parks are located in the city. The unemployed population is obtained from the 2020 Urban Employment and Unemployment Survey (UEUS). For Niger and Malawi, the unemployed population is obtained from the most recent ILOSTAT labor force statistics. Using country-specific labor market sizes allows the same investment budget to generate different increases in labor absorption capacity across countries.

Table 31: Calibration of Policy 5: Labor Absorption Capacity

Country	Investment	Jobs Created	Reference Labor Market	λ_c^D
Niger	\$500 million	22,353	—	—
Malawi	\$500 million	22,353	—	—
Ethiopia	\$500 million	22,353	Addis Ababa UEUS	0.056

G.6 Interpretation of Policy Calibration

The purpose of the policy calibration is to compare the effectiveness of alternative interventions under a common fiscal budget rather than to evaluate the exact impacts of specific government programs. To ensure comparability across policies and countries, each intervention is assigned the same benchmark budget of USD 500 million. Country-specific policy intensities therefore arise from differences in demographic structure, infrastructure needs, land institutions, and labor market conditions rather than differences in public expenditure.

The calibration follows a common two-step procedure. First, the benchmark budget is translated into a policy implementation intensity using external information on program costs together with country-specific demographic or geographic statistics. Second, the implementation intensity is mapped into the corresponding structural parameter of the model. Food subsidies reduce the household subsistence requirement, agricultural input programs increase agricultural productivity, road investments reduce migration costs, land reforms convert households from weak to stronger land rights, and industrial investment increases the probability that workers are absorbed into productive non-agricultural employment.

The resulting policy parameters are summarized in Table 9 and are held fixed throughout the counterfactual analysis. The common policy budget is translated into country-specific changes in the relevant structural parameters using the calibration procedures described above. Consequently, differences in simulated policy outcomes reflect both the estimated structural parameters and the calibrated policy intensities implied by country-specific baseline conditions. Because the same calibration framework is applied consistently across countries under a common fiscal budget, the counterfactual experiments provide a comparable basis for assessing how country-specific structural constraints influence the effectiveness of alternative development policies.