

Climate Change, Income Inequality, and Migration in a Spatial Economy

Michael Tueting*

[Latest version](#)

Job Market Paper

This version: October 16, 2025

Abstract

Negative local labor market shocks create strong incentives to migrate, yet low-income households often remain in place. This paper studies how income shapes migration responses to climate change and the resulting welfare effects. Using Brazilian Census data, I first show that higher-income individuals are systematically more mobile than poorer individuals from the same origin. To interpret this pattern and quantify climate impacts, I develop a dynamic spatial general equilibrium model with monetary migration costs and liquidity constraints, embedded in a two-sector trade framework where rising temperatures depress agricultural productivity. The quantification implies sharply regressive climate losses: in already-hot regions, low-income households suffer permanent consumption declines of about 0.6-0.9% per period, with worst-case losses near 3.3%, while richer households are largely insulated. Two mechanisms drive these disparities: climate shocks reduce agricultural wages in hot areas, and monetary migration costs disproportionately burden low-income individuals, limiting their ability to relocate in response. A counterfactual policy that equalizes migration costs across incomes—modeled as a targeted subsidy—raises low-income households' welfare by about 24%, offsets 4% of their baseline climate losses (3% nationally), and increases aggregate output by reallocating labor to more productive regions. Adaptation to climate change depends not only on where productivity shocks occur, but also on who can afford to relocate. Consequently, adverse shocks tend to fall disproportionately on those unable to afford relocation.

*Tueting: Swiss Institute for International Economics and Applied Economic Research, University of St.Gallen, Rosenbergstrasse 22 CH-9000 St. Gallen, michael.tueting@unisg.ch. I am grateful for the advice Bruno Caprettini, Doug Gollin, Roland Hodler, Ferdinand Rauch, and Adam Storeygard provided. I also thank Arnaud Costinot, Juanma Castro-Vincenzi, Federico Esposito, Isabella Manelici, Ferdinando Monte, Christoph Moser, Jose Vasquez, workshop participants at Tufts University, and the University of St.Gallen for their helpful comments on early versions of this paper.

1. Introduction

The distribution of economic activity and opportunity across regions is in constant motion. Trade liberalization, technological change, and increasingly climate change reallocates opportunities across sectors and space. In principle, migration allows households to pursue these opportunities. In practice, many individuals remain in place even when the benefits of relocation are substantial. Understanding this gap between incentives and realized mobility is central to both positive and normative economics: it determines who bears the cost of local downturns, how quickly economies reallocate factors in response to shocks, and which policies are effective when place-based fortunes diverge.

This paper examines how income shapes migration responses to climate change and the resulting welfare consequences. Climate change is a particularly salient spatial shock: rising temperatures unevenly affect agricultural productivity and rural incomes, especially in low- and middle-income countries. Although migration is often viewed as a key adaptation strategy, evidence shows that mobility responses to climate shocks are weak in low-income settings (Cattaneo and Peri, 2016; Bryan, Chowdhury, and Mobarak, 2014). Motivated by this pattern, I develop a dynamic spatial general-equilibrium model that links liquidity constraints to migration decisions, showing how monetary frictions can suppress climate-induced mobility and reshape welfare outcomes. By explicitly modeling liquidity-constraints, my framework departs from canonical models and delivers quantitatively different predictions for mobility, reversing standard policy implications.


Brazil provides a compelling laboratory for this analysis. Spatial inequality is large and persistent: the North and Northeast are less wealthy, hotter, and more rural than the South and Southeast. Agriculture employs up to 40 percent of workers in some northern regions, while non-agriculture dominates in the South and coastal areas. Brazil's climatic gradient is also steep, with Northern regions having an annual mean temperature of around 28°C, 10°C hotter than its' Southern regions.

Two facts motivate the analysis. First, within the same origin, higher-income Brazilians are much more mobile than lower-income Brazilians. Using the 2010 Census, I show that migration rates rise steeply with income: individuals with an annual income of R\$40–50k are almost twice as likely to have moved between 2005 and 2010 as those earning R\$0–10k. This pattern is consistent with migration requiring an upfront *monetary* cost—bus tickets, deposits, foregone earnings—that imposes a disproportionately heavy burden on low-income households. Second, climate change shifts agricultural productivity nonlinearly with temperature: productivity increases up to the low-to-mid 20s °C (peaking around 23–24°C) and declines sharply thereafter. Under standard warming scenarios, productivity in already-hot regions shifts away from the agronomic optimum, while productivity in cooler regions shifts toward that optimum. A rise in temperatures in the North thus depresses incomes and raises the value of mobility while simultaneously constraining the ability to move. The result is a powerful asymmetry: precisely where the incentive to migrate is strongest, the means to do so are weakest.

I build a dynamic spatial general-equilibrium model tailored to these facts. (i) Migration costs are *monetary*, not directly specified in utils: with concave utility, the utility-equivalent

burden of a given cash cost decreases with income. Negative local shocks, therefore, tighten mobility exactly when wage gaps widen. (ii) Regions host *evolving income distributions* rather than representative agents. Agriculture pays a reservation wage; non-agriculture offers a continuum of jobs with heterogeneous productivity levels that pay a wage equal to the marginal product of labor. Workers are matched with a non-agricultural job and accept the best available option; if their non-agricultural wage would be below the agricultural wage, they choose to work in agriculture. If climate change reduces a region's agricultural wage, more workers sort into low-productivity non-agricultural jobs, lowering mean income and raising income dispersion. Panel (a) of Figure I illustrates this mechanism. Initially, everyone with non-agricultural productivity draws above the agricultural wage of workers in agriculture (dark shaded area). If climate change depresses the agricultural wage (left arrow), workers with low non-agricultural productivity who previously worked in agriculture switch to low-productivity non-agricultural jobs. If climate change, however, affects different regions heterogeneously, this creates incentives for workers to relocate across regions, especially for those who have low non-agricultural productivity and face large income losses due to the shift in agricultural wages. If migration costs, however, require a fixed cash outlay, the utility-equivalent burden of moving rises as income falls, muting migration. Panel (b) of Figure I illustrates this mechanism. For a high-income person with income y_h a monetary cost χ implies a strictly lower utility-equivalent burden than for a low-income individual (difference in shaded areas).

FIGURE I
Economic adjustment under productivity shocks.



Notes. Panel (a) illustrates how a negative productivity shock reduces the agricultural reservation wage, real-locating marginal workers into low-productivity non-agricultural jobs. Panel (b) illustrates how a monetary migration cost χ translates into a higher utility-equivalent burden for low-income individuals than for high-income individuals, muting migration even when wage gaps across regions widen.

Applying the model to projected climate change from 2010 to 2100, I find that climate damages are strongly regressive. In already-hot regions, low-income households (those earning below R\$10,000, or about 4,400 USD) face permanent consumption losses of roughly 0.6-0.9%

per period, with worst-hit regions near 3.3%; richer households in the same places lose little. In cooler southern regions, some low-income households benefit as local productivity moves toward the agronomic optimum. A targeted migration subsidy that equalizes the utility-equivalent moving cost across incomes raises the welfare of low-income households by about 24%, offsets roughly 4% of their baseline climate losses (around 3% nationally), and increases output by reallocating labor toward more productive places. At a 3% annual discount rate, the policy's present-value GDP gains amount to 7.9% of initial GDP, against fiscal costs of 8.7%, implying a coverage ratio of about 90%.

Adaptation through mobility depends not only on where productivity changes occur but also on who can act on those incentives. Because migration costs are paid in cash, negative local shocks reduce both earnings and the ability to relocate, making them self-reinforcing. Policies that relax liquidity constraints on mobility, such as targeted moving grants or access to credit, can therefore serve a dual role of redistribution and adaptation instruments. While my application focuses on climate change, the mechanism and the policy logic generalize to spatial downturns more broadly.

I contribute to the literature in multiple ways. First, my framework delivers quantitatively different results from canonical models that treat migration as a fixed utility cost. For example, [O'Connor \(2024\)](#) shows that taxing declining regions can be welfare-improving by pushing workers toward higher-productivity areas—an implication that assumes migration is inhibited by weak incentives. By contrast, my model introduces a market failure: when migration costs are monetary, low-income individuals may face binding liquidity constraints that prevent them from acting on incentives. In this setting, taxation reduces welfare without increasing mobility, and targeted migration subsidies become the optimal policy. This distortion is not merely theoretical—it alters the direction and incidence of policy recommendations. A growing empirical literature supports this constraint-based interpretation of migration responses, which has not yet been widely adopted in structural migration models. In the U.S., [Bastian and Black \(2022\)](#) show that expansions of the Earned Income Tax Credit (EITC) increase out-migration from distressed and rural regions, especially among women, suggesting that even modest income increases can ease constraints and enable efficient relocation. Similarly, [Clark and Cummins \(2024\)](#) show that, even in the 19th century, individuals with higher education were disproportionately likely to leave declining areas in Northern England for the more prosperous South, reinforcing long-run divergence. These findings echo historical patterns from deindustrialization episodes in New England and Northern England, where out-migration was limited despite large and persistent local income losses ([Choi, 2024](#)). Across these contexts, the evidence suggests that individuals do not fail to move for lack of incentive—but for lack of means.

Second, my analysis contributes to the literature on climate, migration, and adaptation. A large empirical literature documents nonlinear temperature effects on crop yields and agricultural income (e.g., [Schlenker and Roberts, 2009](#); [Lobell and Gourdji, 2012](#); [Burke et al., 2015](#); [Feng et al., 2010](#)), and muted migration responses in low-income settings consistent with liquidity constraints ([Cattaneo and Peri, 2016](#); [Bryan et al., 2014](#); [Bazzi, 2017](#)). Brazil-specific work finds similarly weak migration responses where climate damages are largest ([Bastos et al.,](#)

2013; Brunel and Liu, 2025). In my work, I replicate the findings that temperature does affect productivity non-linearly, and provide an explanation why low-income workers respond least.

Third, I build on dynamic trade-migration models (Costinot et al., 2016; Caliendo et al., 2019), and recent work embedding climate shocks in spatial general equilibrium (Conte, 2022; Cruz, 2024; Cruz and Rossi-Hansberg, 2024; Jedwab et al., 2023), but depart by modeling migration costs in monetary terms and incorporating within-region income distributions, allowing inequality to affect—and be affected by—mobility.

Fourth, my results complement evidence that reducing migration frictions raises efficiency (Morten and Oliveira, 2024; Pellegrina and Sotelo, 2021; Bryan and Morten, 2019), but highlight a novel implication: even modest aggregate migration flows may mask large distributional wedges. The incidence of shocks falls most heavily on those who cannot afford to respond.

Hence, my contributions are threefold. (i) Empirically, I document a steep within-origin income gradient in internal migration, both in migration rates and migration distance. (ii) Methodologically, I develop a dynamic spatial general equilibrium model with monetary migration costs and heterogeneous income distributions. (iii) Substantively, I show that climate-induced income losses are highly regressive, and that a targeted migration subsidy delivers large welfare gains for low-income households and raises aggregate output at a fiscal surplus.


The rest of the paper is organized as follows. Section 2 describes data and measurement. Section 3 presents motivating evidence on migration patterns, section 4 estimates the impact of rising temperatures on productivity. Section 5 presents the theoretical framework. Section 6 details the calibration. Section 7 quantifies welfare losses from climate change and documents their spatial and distributional heterogeneity. Section 8 studies an income-targeted migration subsidy and its macroeconomic incidence. Section 9 concludes.

2. Data and measurement

The paper draws data from three main sources: information on employment patterns, household incomes, and migration flows; annual agricultural production figures paired with temperature records; and climate model projections of regional temperatures through 2100. In the following, I first describe the unit of observation and the primary data sources. Then, I provide a brief overview of how key variables used throughout the paper are measured.

Unit of observation The unit of observation is the mesoregion. A mesoregion is a subnational statistical unit in Brazil, which is defined by the Brazilian Institute of Geography and Statistics (IBGE). There are 137 mesoregions, which are grouped into five larger geographic regions: Norte, Nordeste, Centro-Oeste, Sudeste, and Sul. While mesoregions are not administrative units, they offer the advantage of being stable units over time (since 1989) and ease inter-temporal comparisons. They have also been frequently used in related literature (Gollin and Wolfersberger, 2024; Pellegrina and Sotelo, 2021). Figure II (left-panel) shows the mesoregions of Brazil (dark lines) and the respective major geographical regions.

FIGURE II
Regions of Brazil & agricultural income



Notes: Panel (a) shows the five major geographic regions of Brazil—Norte, Nordeste, Centro-Oeste, Sudeste, and Sul—and the 137 mesoregions, as defined by the Brazilian Institute of Geography and Statistics (IBGE). Mesoregion boundaries are shown as thin black lines, and the underlying shapefiles are obtained directly from the IBGE’s official geographic database. Panel (b) shows the average annual income of agricultural workers in each mesoregion in 2010, based on microdata from the 2010 Brazilian Population Census (Censo Demográfico). The sample includes Brazilian-born individuals aged 25–64 with positive reported income and valid demographic, occupational, and geographic information. Income refers to total annualized gross earnings from employment in crop cultivation, expressed in 2010 Brazilian Reais (R\$), where 1 USD = R\$2.3. Crop cultivation is defined using the 5-digit occupation classification. Further details on the sample and mapping of occupation codes are available in appendix A.1.

Socioeconomic data Socioeconomic data on migration flows, employment patterns, and household incomes come from the 2010 Brazilian Census of the Population (*Censo Demográfico*). The sample consists of Brazilian-born individuals aged 25–64 with positive income and non-missing information on occupation, place of residence in 2005, race, education, marital status, and gender, resulting in approximately 6.2 million observations, which correspond to about 60.3 million individuals in the weighted population.¹

Migration flows are constructed from reported municipalities of residence in 2005 and 2010, aggregated into mesoregion-to-mesoregion flows. Table I presents summary statistics on the fraction of respondents who, in 2005 and 2010, were living in the same mesoregion, a different mesoregion within the same state, and a different state.

Sectoral employment is measured as the share of workers in crop cultivation based on the Census 5-digit occupational classification. I use “agriculture” and “crop cultivation” interchangeably and refer to all other activities as “non-agriculture,” consistent with the productivity measure based on crop yields. Income is measured as the annualized total gross monthly income reported in the Census. The average annual income of workers employed in crop cultivation in 2010 is shown in Figure II (right panel) in R\$ (1 US dollar = 2.3 Brazilian Reais).

¹The Census extract includes around 20 million observations. Restricting the sample to the working-age population (25–64) reduces the sample by approximately 10 million observations. Dropping those with missing occupation codes (notably, this includes individuals not in the labor force) reduces the sample to 6.5 million observations. The other restrictions reduce the sample size only minimally to the final dataset of 6.2 million observations.

TABLE I
Patterns in migration

Residence	2005-2010
Same mesoregion	95.6%
Different mesoregion, same state	1.8%
Different state	2.6%

Notes. This table shows the share of individuals who lived in the same mesoregion, a different mesoregion within the same state, or a different state between 2005 and 2010. The measurement is based on microdata from the 2010 Brazilian Population Census (Censo Demográfico). It is defined using Brazilian-born individuals aged 25–64 with positive reported income and valid demographic, occupational, and geographic information. Migration flows are constructed by mapping municipality-of-residence codes in 2005 and 2010 to mesoregion-to-mesoregion transitions using the classification of the Brazilian Institute of Geography and Statistics (IBGE). Further details on the sample and construction of migration flows are provided in Appendix A.1.

Appendix A.1 shows summary statistics on the sample of workers and the mapping of occupation codes to workers employed in crop cultivation.

Agricultural production Data on agricultural production come from the *Produção Agrícola Municipal* (PAM) of the Brazilian Institute of Geography and Statistics (IBGE), which reports municipal-level statistics on more than 60 crops from 1994 to 2014. I aggregate municipal observations to mesoregions, yielding annual measures of crop-specific area planted, physical output, and production value. Total agricultural land is defined as the cumulative area planted with crops, and I use this to construct mesoregion-level land-use measures from 1994 onward.

Measured agricultural productivity is defined as the value of crop output per unit of land at constant prices.² By fixing prices at nationwide reference levels, differences in measured productivity reflect variation in physical yields rather than price heterogeneity across crops, regions, or time. This approach follows standard practice in the literature (e.g., Gollin et al., 2021; Restuccia et al., 2008). In both the empirical analysis and the calibrated model, measured agricultural productivity is adjusted for land use intensity following Costinot et al. (2016). Intuitively, land quality varies within a region, and farmers will cultivate high-productivity fields first, implying that the average (unconditional) productivity of land in a region is weakly smaller than the measured productivity because there is positive selection of land into cultivation.³ Appendix B.2 outlines the procedure and provides further summary statistics.

Temperature data Temperature data come from the Copernicus Climate Change Service (C3S) CMIP6 archive. I use monthly near-surface air temperature from the CMCC-ESM2 Earth system model, available on a 1° x 1° grid, covering the historical period through 2014 and future projections under the SSP2-4.5 and SSP5-8.5 scenarios. Temperatures are converted from Kelvin to Celsius, aggregated to yearly means, and averaged to the mesoregion level

²For each crop c , I compute a nationwide reference price $P_c^* = (\sum_i \sum_t V_{c,it}) / (\sum_i \sum_t Q_{c,it})$, where $V_{c,it}$ is the nominal value of production and $Q_{c,it}$ is the physical quantity in region i and year t , using all years between 1994 and 2014. Constant price output in region i and year t is then $\tilde{V}_{it} = \sum_c P_c^* Q_{c,it}$, and measured productivity is $\hat{A}_{it} = \tilde{V}_{it} / (\sum_c X_{c,it})$, where $X_{c,it}$ is the land area planted with crop c .

³For example, Rio de Janeiro ranks second based on measured agricultural productivity; however, its share of land under cultivation is only around 2% of its land area. Without adjusting for land-use intensity, this ranking reflects selection into a small set of highly productive plots rather than a high region-wide average.

using land-area weights. The historical data provide annual mesoregion level means from 1994 to 2014, which I match to PAM, and the projections extend through 2100. For the projections, I compute five-year averages and express anomalies relative to 2010.

3. Motivation: Who migrates?

In many low- and middle-income countries, households live close to subsistence, with limited savings and little access to credit. In Brazil, 8.4% of the population lived below the \$3/day poverty line in 2010,⁴ with poverty heavily concentrated in the North and Northeast. For such households, financing large one-time expenses—such as migration—can be prohibitive. Liquidity constraints thus predict that low-income households migrate less, even though they may have stronger incentives to do so to escape structural poverty. Consistent with this view, I document that in Brazil, higher-income individuals are i) substantially more likely to migrate than low-income individuals and ii) tend to move farther relative to less wealthy individuals from the same origin.

To measure this, I construct counterfactual incomes that individuals would have earned in 2010 had they remained in their 2005 mesoregion. These predicted incomes are obtained from region-specific Mincer regressions of 2010 log income on age, age squared, gender, race, marital status, and education. The fitted values are interpreted as earnings capacity in the absence of migration, under two assumptions: (i) conditional on observables, the 2010 wage structure proxies relative opportunities for 2005 residents, and (ii) the small migrant share (about 5%) does not meaningfully affect local wages.

I then estimate how i) migration probabilities, and ii) average distance of actual moves vary across the income distribution. Predicted incomes are grouped into five 10,000 R\$ bins up to 50,000 R\$, with the lowest bin (0-10,000 R\$) omitted.⁵ The estimating equation is:

$$Y_{ijt} = \sum_{b=2}^5 \beta_b \mathbf{1}\{\text{Pred. Inc.}_{ijt} \in b\} + \mu_j + \varepsilon_{ijt},$$

where i indexes individuals, j indexes the 2005 mesoregion of residence, and t indexes the 2010 cross-section (with outcomes measured over 2005–2010). Y_{ijt} is the individual outcome, μ_j are mesoregion fixed effects, and income bin $b = 1$ (R\$0–10,000) is the omitted category.


In panel (a) of Figure III, Y_{ijt} is a binary indicator of migration between 2005 and 2010. β_b represents percentage-point differences in migration probability relative to the omitted bin. In panel (b), Y_{ijt} is the distance moved (in kilometers) between 2005 and 2010, estimated from the sample of movers only. β_b represents differences in kilometers relative to the omitted bin.

Panel (a) reveals a strong upward gradient: relative to the omitted R\$0-10,000 bin, migration probabilities rise by about 0.5 percentage points for R\$10-20,000, 1.86 percentage points for R\$20-30,000, and roughly three percentage points for both R\$30-40,000 and R\$40-50,000.

⁴Measured at 2011 PPP according to the World Bank's Poverty and Inequality Platform. <https://pip.worldbank.org>

⁵The fraction of individuals per bin is: 47% in 0-10K, 38% in 10-20K, 7.8% in 20-30K, 4.2% in 30-40K, and 2.2% in 40-50K.

FIGURE III
Migration & income



(A) Difference in migration rates by income

(B) Difference in distance of moves by income

Notes. The figure plots coefficients from regressions relating migration outcomes between 2005 and 2010 to predicted counterfactual 2010 income bins. Counterfactual incomes are constructed by estimating region-specific Mincer regressions of 2010 log income on age, age squared, gender, race, marital status, and education, on the sample of non-movers, and then predicting what 2005 residents would have earned in 2010 had they stayed. Predicted incomes are grouped into five R\$10,000 bins up to R\$50,000; the lowest bin (0-10,000 R\$) is omitted. The fraction of individuals per bin is: 47% in 0-10K, 38% in 10-20K, 7.8% in 20-30K, 4.2% in 30-40K, and 2.2% in 40-50K. Panel (a) reports percentage-point differences in the probability of migrating (changing mesoregion) relative to the omitted bin. Panel (b) reports differences in the distance moved (kilometers), conditional on migrating, relative to the omitted bin. Vertical lines correspond to the 95% confidence interval of the point estimates. Panel (a) clusters standard errors by origin mesoregion, and panel (b) clusters standard errors by origin and destination. The sample includes Brazilian-born individuals aged 25-64 with positive income and complete demographic and occupational information.

The average emigration rate for those in the lowest income bin is 3.93%, implying that those in the highest income groups are almost twice as likely to have migrated between 2005 and 2010.


Panel (b) shows no statistically significant difference in distance moved for the two lower bins, but among higher predicted incomes, movers go substantially farther—about 50 km more for R\$30-40,000 and roughly 70 km more for R\$40-50,000. The average distance moved by those few movers in the lowest income group is 700 km, implying that the highest income groups move, on average, 10% farther.

Taken together, the results are consistent with liquidity constraints shaping who migrates and how far they go: low-income households face binding upfront cash needs for transport and resettlement, which exclude many high-cost, high-return destinations (extensive margin). Furthermore, even when the costs are payable, the fixed costs represent a larger utility-equivalent sacrifice for them under concave utility, dampening responsiveness to opportunity differences (intensive margin). With distance raising costs, these constraints make long-distance moves feasible primarily for higher-income individuals, reinforcing spatial inequality and limiting migration's role as an adaptation strategy for low-income individuals.

4. The Impact of Climate Change on Agricultural Productivity

To establish the relationship between temperature and agricultural productivity in Brazil, I estimate how yields change with temperature. This relationship is central to the analysis, as it governs how climate change affects rural incomes and, in turn, migration incentives.

FIGURE IV
Estimated relationship between temperature and agricultural productivity.



Notes. The figure shows the estimated relationship between temperature and fundamental agricultural productivity (constant-price output per hectare of land adjusted for land use intensity) in Brazilian mesoregions from 1994 to 2014. The solid blue line shows the binned model estimates; vertical bars are 95% confidence intervals based on standard errors clustered at the mesoregion level to allow arbitrary serial correlation within mesoregions (e.g., in temperature and productivity over time). The dashed orange line shows the quadratic model estimates. Both models include mesoregion fixed effects and region-specific cubic time trends.

Following [Costinot et al. \(2016\)](#), I adjust measured productivity (constant price output per hectare of land) for land use intensity to obtain a measure of fundamental productivity. Fundamental productivity is defined as $A_{it} = \tilde{A}_{it} \cdot \zeta_{it}^{1/\theta}$, where $\tilde{A}_{it}$ is measured productivity, ζ_{it} is the share of land used in agriculture, and θ governs the heterogeneity in land quality, which I set to 2 (the midpoint of available estimates [Sotelo \(2020\)](#); [Costinot et al. \(2016\)](#)). To estimate the impact of temperature on fundamental productivity, I exploit variation in average temperature within mesoregions over time. I estimate a piecewise linear and a quadratic specification of the following form:

$$\ln A_{it} = \sum_{b \in \mathcal{B} \setminus b^*} \beta_b \mathbf{1}\{TG_{it} \in b\} + \mu_i + \sum_{p=1}^3 \tau_{ip} t^p + \varepsilon_{it} \quad (\text{piecewise})$$

$$\ln A_{it} = \beta_1 TG_{it} + \beta_2 TG_{it}^2 + \mu_i + \sum_{p=1}^3 \tau_{ip} t^p + \varepsilon_{it} \quad (\text{quadratic})$$


where TG_{it} denotes the annual mean temperature of mesoregion i in year t , μ_i are mesoregion fixed effects, and $\{\tau_{ip}\}_{p=1}^3$ allows for region-specific linear, quadratic, and cubic time trends. The piecewise specification uses 1°C bins, with $[23, 24)$ as the omitted category.

Figure IV shows the estimated relationship between temperature and productivity. Both the binned and quadratic models trace a concave relationship: productivity increases with temperature up to around 23-24°C and declines sharply thereafter. The quadratic model implies an optimal temperature near 23.5°C, with a 95% confidence interval of [22.8, 24.1]°C.⁶

This nonlinearity implies that the impact of warming is heterogeneous across space. Cooler regions (such as those in the South) may experience productivity gains from moderate warming, while hotter regions (notably in the North and Northeast) are already operating on the downward-sloping segment of the curve and are thus particularly vulnerable to further temperature increases.

The temperature trend in Brazil is strongly upward. Figure V shows the average annual temperature by mesoregion in 2010 (left panel) and the historical and projected annual mean temperatures under the SSP2-4.5 and SSP5-8.5 scenarios (right panel). The solid line shows the median regions' temperature for each scenario, with shaded areas corresponding to the 10th and 90th percentiles, respectively. By 2100, average regional temperatures are projected to rise by 2.9°C under SSP2-4.5 and by 4.1°C under SSP5-8.5, with some regions exceeding 30°C under the high-emission scenario. This implies that many regions will move further into the range where productivity declines sharply with temperature, exacerbating challenges for agricultural workers and rural economies. Linking agricultural productivity to incomes, I estimate that a 1% rise in agricultural productivity is associated with a 0.21% increase in agricultural wages (see Appendix B.2).

FIGURE V
Temperature in Brazil



Notes. Panel (a) shows the average annual temperature in 2010, based on data from the Copernicus Climate Change Service (C3S) CMIP6 archive. Panel (b) shows median regional annual temperatures in Brazil from 1994 to 2100 under the SSP2-4.5 and SSP5-8.5 scenarios, simulated using the CMCC-ESM2 Earth System Model. Temperature data are provided on a $1^\circ \times 1^\circ$ grid and converted from Kelvin to Celsius, aggregated to annual means, and averaged to the mesoregion level using land-area weights. The historical series covers 1994–2014, while future projections extend through 2100. For visualization, projected values are shown as five-year averages, and the shaded areas indicate the 10th and 90th percentiles across mesoregions.

⁶Under the quadratic specification, the optimum is $\hat{T}^* = -\hat{\beta}_1 / (2\hat{\beta}_2)$ (with $\hat{\beta}_2 < 0$); the interval is computed via the delta method using the estimated variance–covariance matrix of the β coefficients.

5. Theoretical framework

I develop a quantitative spatial model to study welfare, migration, and income dynamics in response to climate change. The model consists of two interconnected components: a dynamic discrete-choice problem and a static trade equilibrium.

These components interact as follows: the static trade model determines wages, prices, and consumption levels *conditional on a given allocation of labor and regional fundamentals*. The dynamic component then solves for the optimal reallocation of labor over time, *based on the expected path of real incomes implied by the static equilibrium*.

Dynamic problem: In the dynamic household problem, households take the structure of the economy as given and make forward-looking migration decisions that shape migration flows between regions over time. The framework builds upon the seminal work of [Caliendo et al. \(2019\)](#), which represents the current state of the art in this literature.

I model migration costs as an origin-destination specific monetary cost plus a non-monetary psychic cost of leaving “home”. With concave utility, the monetary cost implies a larger burden for low-income households. By contrast, [Caliendo et al. \(2019\)](#) assume income-invariant utility costs tied to origin $\times$ sector - destination $\times$ sector pairs. This difference has sharp implications. In their model, mobility is unaffected by income shocks (conditional on constant relative utility differences across space). In mine, lower incomes raise the effective burden of moving and reduce mobility. This mechanism captures an important channel of climate change: falling incomes make migration harder, especially for low-income workers.

Static sub-problem: The second component is a static trade model with two sectors and many regions. Unlike standard quantitative spatial models (see [Redding and Rossi-Hansberg, 2017](#)), it embeds region-specific, continuous income distributions that evolve endogenously. Within each region, agriculture offers a single market-clearing wage that acts as a reservation wage. Non-agricultural employment features a continuum of jobs with heterogeneous productivities, with wages equal to the marginal product of labor. Workers draw a productivity realization in this sector; if the implied wage is below the agricultural reservation wage, they work in agriculture instead. This mechanism parallels [Lucas \(1978\)](#) and [Luttmer \(2007\)](#), where individuals draw entrepreneurial productivity and choose between entrepreneurship and wage work. Climate shocks that depress agricultural productivity lower the reservation wage, pushing more workers into low-productivity non-agricultural jobs. The result is higher inequality and lower average income. Climate change enters the model twice: by shifting agricultural productivity and regional income distributions, and by raising the income-dependent utility burden of migration costs.

The rest of this section is structured as follows: First, I describe the dynamic component, highlighting the novel feature I built into the baseline model of [Caliendo et al. \(2019\)](#): monetary migration costs. Second, I describe the static trade model in which climate change interacts with income distribution both across and within regions. [Online Appendix C](#) provides the full derivations, and [Online Appendix F](#) outlines how the full model is solved in practice.

5.1. Dynamic component

As discussed above, many households in Brazil live close to subsistence, with a sizable share below the \$3/day poverty line. For such households, financing large one-time expenditures like migration is prohibitive. I therefore abstract from asset accumulation and model migration decisions as depending solely on income realizations and migration costs.

5.1.1. Setup and timing

The economy consists of a finite set of regions $j, i \in \mathcal{S}$ populated by a mass of households.⁷ Time is discrete, $t \in \{1, \dots, T\}$. Households may migrate at the end of a period, but have to pay a monetary cost $\chi^{j,i}$ as well as a non-monetary cost to do so. The future is discounted at rate $\beta > 0$, and households consume their entire income net of migration costs.

The ability to pay the monetary cost depends on an individual's income. At the start of each period, a household in region j receives a draw from the region's non-agricultural productivity distribution. If the draw implies a wage below the agricultural wage, the household chooses to work in agriculture. Thus the regional income distribution at time t , f_t^i , emerges endogenously from the agricultural wage and the non-agricultural productivity distribution. Productivity draws follow a Calvo-style process: with probability ρ , the household retains its previous draw (and hence its relative position in the non-agricultural distribution); with probability $1 - \rho$, it draws anew from the uniform distribution on $(0, 1)$. The parameter ρ governs **income persistence**.⁸

Households also draw idiosyncratic location preferences. They observe the full wage distributions in all regions and thus know the expected income prospects everywhere. The only uncertainty concerns their own future wage realization within those distributions, as well as future preference shocks.

5.1.2. Instantaneous utility

ASSUMPTION 1: *Agents have logarithmic preferences.*

An individual with income y in region j who chooses to migrate to region i at time t derives utility

$$U(C_t^{j,i}(y)) = \log C_t^{j,i}(y),$$

where consumption equals disposable real income:⁹

$$C_t^{j,i}(y) = \frac{y_t^j - \chi^{j,i}}{P_t^j}.$$

Here, y_t^j denotes income in region j , $\chi^{j,i}$ the monetary migration cost from j to i , and P_t^j the regional price index. Migration costs capture expenditures such as transportation or housing deposits.

⁷ j denotes the origin and i the destination.

⁸This also implies persistence in sectoral choices. A worker with a low initial non-agricultural productivity draw is likely to stay in agriculture for multiple periods for $\rho > 0$.

⁹This follows from the static block presented below.

5.1.3. Utility equivalent costs

Income-invariant monetary costs imply income-dependent utility costs. Let $C_t^j(y) = y_t^j / P_t^j$ denote consumption if a household in j does not migrate. The utility-equivalent cost of migration is

$$\tilde{\tau}^{j,i}(y) \equiv \log \left(\frac{y_t^j}{y_t^j - \chi^{j,i}} \right),$$

which follows from $U(C_t^{j,i}(y)) = U(C_t^j(y)) - \tilde{\tau}^{j,i}(y)$, with $U(C_t^j(y)) = \log(y_t^j / P_t^j)$. This expression is well defined only for $y_t^j > \chi^{j,i}$; whenever $y_t^j \leq \chi^{j,i}$, I set $\tilde{\tau}^{j,i}(y) = +\infty$ (migration is infeasible).

COROLLARY 1: *The utility burden $\tilde{\tau}^{j,i}(y)$ is decreasing in income.*

In other words, monetary costs map into larger utility losses for low-income households. This contrasts with [Caliendo et al. \(2019\)](#), where relocation costs are specified directly in utility terms and remain independent of income. In addition, I allow for non-monetary costs κ^j of leaving “home”. Composite migration costs are

$$\tau^{j,i}(y) = \tilde{\tau}^{j,i}(y) + \mathbf{1}(i \neq j) \cdot \kappa^j.$$

5.1.4. Expected lifetime utility

Let $v_t^j(r)$ denote the lifetime utility of a household residing in region j at time t , conditional on a nonagricultural productivity draw corresponding to rank $r \in (0, 1)$ in the local productivity distribution. The associated income is $y_r^j = (F_t^j)^{-1}(r)$, where $(F_t^j)^{-1}$ denotes the inverse CDF of the regional income distribution that arises endogenously from the agricultural reservation wage and the distribution of nonagricultural productivity draws.

Households consume all income net of migration costs and choose next period’s region to maximize expected lifetime utility. The dynamic problem is

$$v_t^j(r) = U(y_r^j) + \max_{i \in \mathcal{S}} \left\{ -\tau^{j,i}(y_r^j) + \nu v_t^i + \beta \left[\rho \mathbb{E}_t[v_{t+1}^i(r)] + (1 - \rho) \int_0^1 E[v_{t+1}^i(r')] dr' \right] \right\}.$$

where $\tau^{j,i}(y_r^j)$ denotes the effective migration cost from j to i , which depend on income; v_t^i represents an idiosyncratic preference shock for region i ; and ν scales the dispersion of these shocks. The term in brackets captures the expected continuation value, accounting for income persistence governed by ρ and the possibility of a new wage draw. With probability ρ the previous rank is retained so that the lifetime utility of starting in i in $t + 1$ is $E[v_{t+1}^i(r)]$ and with probability $1 - \rho$ the lifetime utility is the average over all ranks $\int_0^1 E[v_{t+1}^i(r')] dr'$.

ASSUMPTION 2. *Preference shocks v_t^i are i.i.d. across i and t and follow a Type-I Extreme Value (Gumbel) distribution with location 0 and unit scale. The dispersion parameter $\nu > 0$ scales these shocks.*

Under Assumption 2, the dynamic discrete choice problem admits the standard logit closed

form (McFadden, 1974). Let $V_t^j(r) \equiv \mathbb{E}[v_t^j(r)]$ and $\bar{V}_t^i \equiv \int_0^1 V_t^i(r') dr'$. Then

$$V_t^j(r) = U(C_t^j(y_r)) + \nu \log \sum_{i \in \mathcal{S}} \exp\left(\beta \left[\rho V_{t+1}^i(r) + (1 - \rho) \bar{V}_{t+1}^i \right] - \tau^{j,i}(y_r)\right)^{1/\nu}. \quad (1)$$

The first term is the period- t utility from consuming at origin j with current income rank r . The second term, $\nu \log \sum_{i \in \mathcal{S}} \exp(\cdot)^{1/\nu}$, is the *inclusive value*: it summarizes the continuation value of the migration choice set *before* idiosyncratic tastes and next period's productivity draw are realized.¹⁰

5.1.5. Migration flows and labor dynamics

Under Assumption 2, the choice probabilities take the logit form. For a household in j with rank $r \in (0, 1)$ (income $y_r^j = (F_t^j)^{-1}(r)$), the probability of migrating to i at the end of period t is

$$\mu_t^{j,i}(r) = \frac{\exp\left(\left[\beta(\rho V_{t+1}^i(r) + (1 - \rho) \bar{V}_{t+1}^i) - \tau^{j,i}(y_r^j)\right]/\nu\right)}{\sum_{k \in \mathcal{S}} \exp\left(\left[\beta(\rho V_{t+1}^k(r) + (1 - \rho) \bar{V}_{t+1}^k) - \tau^{j,k}(y_r^j)\right]/\nu\right)}. \quad (2)$$

With a large population, $\mu_t^{j,i}(r)$ is the share of type- r households in j that relocate to i .

Intuitively, households are more likely to move to destinations that offer higher expected future payoffs and are accessible at low migration costs. Monetary costs depress mobility at low incomes via (i) an intensive margin, since $\partial \tau^{j,i}(y_r^j) / \partial y_r^j < 0$, and (ii) an extensive margin, since migration is infeasible when $\chi^{j,i} \geq y_r^j$ (so $\tau^{j,i}(y_r^j) = \infty$). This helps explain why poorer households are more likely to stay even when the returns to moving are large. Persistence also matters: with probability ρ households keep their rank and weight destinations favorable to their type; with probability $1 - \rho$ they redraw and weigh average prospects. The logit scale ν governs choice dispersion: as $\nu \rightarrow 0$, choices become nearly deterministic and flows concentrate on the destination that maximizes $\beta(\rho V_{t+1}^k(r) + (1 - \rho) \bar{V}_{t+1}^k) - \tau^{j,k}(y_r^j)$; as $\nu \rightarrow \infty$, probabilities flatten across feasible destinations, while infeasible options remain at zero.

Aggregating over the origin's rank distribution (uniform on $(0, 1)$) yields aggregate migration probabilities from j to i ,

$$\mu_t^{j,i} = \int_0^1 \mu_t^{j,i}(r) dr, \quad (3)$$

which is equivalently $\mu_t^{j,i} = \int \mu_t^{j,i}(F_t^j(y)) f_t^j(y) dy$ by the change of variables $r = F_t^j(y)$, where $f_t^j(y)$ is the income density in j .¹¹

¹⁰Formally, $\nu \log \sum_{i \in \mathcal{S}} \exp(\cdot)^{1/\nu} = \mathbb{E}_{v_t} \left[\max_{i \in \mathcal{S}} \left\{ \beta \left[\rho V_{t+1}^i(r) + (1 - \rho) \bar{V}_{t+1}^i \right] - \tau^{j,i}(y_r) + \nu v_t^i \right\} \right]$.

¹¹For tractability I impose a re-ranking convention: after migration, each destination resorts residents by income, assigns new ranks $r \in (0, 1)$, and then sets incomes by $y_r^i = (F_t^i)^{-1}(r)$. Workers' expectations respect the Calvo process: when evaluating destinations, a household anticipates retaining its current rank r with probability ρ and drawing a fresh rank $r' \sim \text{Unif}(0, 1)$ with probability $1 - \rho$. They do not internalize the subsequent mechanical renormalization of everyone's ranks that occurs when inflows arrive. This implies that the distribution of non-agricultural productivity remains constant across regions. An alternative specification could track individual

Aggregate labor stocks evolve according to

$$L_{t+1}^i = \sum_{j \in S} \mu_t^{j,i} L_t^j, \quad (4)$$

which conserves population ($\sum_i L_{t+1}^i = \sum_j L_t^j$).

5.2. Static trade model

The static block determines wages, prices, land use, and income distributions given the spatial allocation of labor. In the following, I present the minimal setup. The full model and derivations are described in Appendix C.

5.2.1. Setup

Preferences. Households consume an agricultural composite $C_t^{i;a}$ and a non-agricultural good $C_t^{i;n}$ with Cobb-Douglas utility:

$$C_t^i = \frac{(C_t^{i;a})^\gamma (C_t^{i;n})^{1-\gamma}}{\gamma^\gamma (1-\gamma)^{1-\gamma}},$$

where $\gamma \in (0,1)$ is the expenditure share on agriculture. The agricultural composite is a CES aggregate of regional varieties,

$$C_t^{i;a} = \left(\sum_{j \in S} (\psi^j)^{1/\epsilon} (c_t^{j,i;a})^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}},$$

with Armington elasticity ϵ and preference weights ψ^j (normalized to sum to one). Here $c_t^{j,i;a}$ is consumption of the variety from origin j . Transporting variety $c_t^{j,i;a}$ to i incurs iceberg trade costs. The non-agricultural good is freely traded and serves as the numeraire.

Agriculture. Each region i is endowed with a fixed supply of land X^i that consists of a continuum of heterogeneous plots ω . Agriculture produces with Cobb-Douglas technology at the plot level:

$$q_t^i(\omega) = [L_t^i(\omega)]^\alpha [z_t^i(\omega) X_t^i(\omega)]^{1-\alpha},$$

where $L_t^i(\omega)$ and $X_t^i(\omega)$ denote labor and land used on plot ω , and $z_t^i(\omega)$ is plot productivity.

Plots differ in their productivity and operating a plot requires a per-period cost $z^{i0}(\omega)$, paid in units of the non-agricultural numeraire. Plot heterogeneity is governed by a (nested) Fréchet specification. Let $A_t^i > 0$ and $F^{i0} > 0$ be scale parameters for productivity and operating cost,

types across regions without re-ranking; this would alter the non-agricultural productivity distribution through composition effects and yield sorting-type income distributions, at the cost of additional state variables and loss of closed form.

and let $\theta > 1$ be the common shape parameter. The joint CDF of $(z_t^i(\omega), z^{i0}(\omega))$ is

$$\Pr(z_t^i(\omega) \leq z, z^{i0}(\omega) \leq z^0) = \exp \left\{ -o \left[\left(\frac{z}{A_t^i} \right)^{-\theta} + \left(\frac{z^0}{F^{i0}} \right)^{-\theta} \right] \right\},$$

with normalization constant $o \equiv \Gamma(1 + \frac{1}{\theta})^{-\theta}$ so that $A_t^i = \mathbb{E}[z_t^i(\omega)]$ and $F^{i0} = \mathbb{E}[z^{i0}(\omega)]$. Here, A_t^i is region-specific unconditional mean productivity (fundamental productivity). Because high-productivity plots are more likely to be cultivated, the average productivity of cultivated land (i.e., measured productivity $\tilde{A}_t^i$) exceeds A_t^i unless all plots are cultivated.

Non-agriculture. Non-agricultural production consists of a continuum of jobs η , each using labor linearly:

$$q_t^{i;n}(\eta) = A^{i;n}(\eta) L_t^{i;n}(\eta),$$

where $A^{i;n}(\eta)$ is job-specific productivity and $L_t^{i;n}(\eta)$ is labor assigned to job η . Regional productivity distributions are i.i.d. log-normal,

$$\ln A^{i;n}(\eta) \sim \mathcal{N}(\mu^i, (\sigma^i)^2),$$

where μ^i and σ^i are region-specific parameters. μ shifts the mean productivity, while σ governs dispersion.

5.2.2. Temporary equilibrium

A *temporary equilibrium* in period t is a collection of agricultural prices $\{p_t^{j;a}\}_{j \in \mathcal{S}}$, the non-agricultural price p_t^0 , agricultural and non-agricultural wages $\{w_t^{j;a}, w_t^{i;n}\}$, labor allocations $\{L_t^{j;a}, L_t^{i;n}\}$, cultivated land $X_t^{j;a}$, outputs $\{Q_t^{j;a}, Q_t^{i;n}\}$, and regional income distributions $\{f_t^i(y)\}$ such that the following conditions hold:

- (i) **Household demand.** Households maximize Cobb-Douglas utility, choosing $C_t^{j;a}$ and $C_t^{i;n}$ subject to

$$P_t^{i;a} C_t^{i;a} + p_t^0 C_t^{i;n} \leq \tilde{y}_t^i, \tag{5}$$

where $P_t^{i;a}$ is the CES price index for agriculture with

$$P_t^{i;a} = \left(\sum_{j \in \mathcal{S}} \psi^j (p_t^{j;a} \delta^{j,i})^{1-\epsilon} \right)^{1/(1-\epsilon)}, \tag{6}$$

where $\delta^{j,i}$ representing iceberg trade costs. $\tilde{y}_t^i$ denotes household income net of migration costs. CES demand implies import expenditure shares

$$\pi_t^{j,i;a} \equiv \frac{\psi^j (p_t^{j;a} \delta^{j,i})^{1-\epsilon}}{\sum_{k \in \mathcal{S}} \psi^k (p_t^{k;a} \delta^{k,i})^{1-\epsilon}}, \tag{7}$$

where $\pi_t^{j,i;a}$ denotes the share of region i 's agricultural expenditure allocated to variety j .

- (ii) **Labor allocation & wages.** Agricultural wages are equal to the marginal product of labor on the marginal plot:

$$w_t^{i;a} = p_t^{i;a} \alpha \left(\frac{A_t^i (\zeta_t^i)^{-1/\theta} X_t^{i;a}}{L_t^{i;a}} \right)^{1-\alpha}, \quad (8)$$

where ζ_t^i is the share of land cultivated in region i at time t .

Workers accept a non-agricultural job iff $A_t^{i;n}(\eta) \geq w_t^{i;a}$. The agricultural employment share is

$$\lambda_t^i = \Pr(A_t^{i;n}(\eta) < w_t^{i;a}) = \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right), \quad (9)$$

where Φ is the standard normal CDF. Agricultural and non-agricultural labor are then determined as $L_t^{i;a} = \lambda_t^i L_t^i$ and $L_t^{i;n} = (1 - \lambda_t^i) L_t^i$.

Non-agricultural wages are:

$$w_t^{i;n} \equiv E[y \mid y \geq w_t^{i;a}] = \exp\left(\mu^i + \frac{1}{2}(\sigma^i)^2\right) \frac{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i - (\sigma^i)^2}{\sigma^i}\right)}{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right)}.$$

The wage structure implies that the income distribution in region i , denoted $f_t^i(y)$, follows a truncated log-normal distribution with a mass point at $w_t^{i;a}$.

- (iii) **Land use & agricultural output.** Profit maximization implies that plot ω is cultivated if

$$z_t^i(\omega) \Omega_t^i \geq p_t^0 z^{i0}(\omega), \quad (10)$$

where

$$\Omega_t^i \equiv \left[\alpha^\alpha (1 - \alpha)^{1-\alpha} p_t^{i;a} (w_t^{i;a})^{-\alpha} \right]^{\frac{1}{1-\alpha}} \quad (11)$$

captures the effective profitability of agriculture, determining cultivation decisions. The share of cultivated land is

$$\zeta_t^i = \frac{(A_t^i \Omega_t^i)^\theta}{(A_t^i \Omega_t^i)^\theta + (F^{i0} p_t^0)^\theta}, \quad (12)$$

implying that agricultural land area is $X_t^{i;a} = \zeta_t^i X_t^i$. Aggregate agricultural output follows from:

$$Q_t^{i;a} = (L_t^{i;a})^\alpha [A_t^i (\zeta_t^i)^{-1/\theta} X_t^{i;a}]^{1-\alpha}. \quad (13)$$

- (iv) **Non-agricultural output.** Only jobs with productivity above $w_t^{i;a}$ are filled. Output is

$$Q_t^{i;n} = w_t^{i;n} (1 - \lambda_t^i) L_t^i, \quad (14)$$

where $w_t^{i;n}$ is the conditional mean wage of accepted jobs.

(v) **Goods market clearing.** Agricultural and non-agricultural markets clear:

$$p_t^{j;a} Q_t^{j;a} = \gamma \sum_{i \in \mathcal{S}} \pi_t^{j,i;a} Y_t^i, \quad (15)$$

$$\sum_{j \in \mathcal{S}} p_t^0 Q_t^{j;n} = (1 - \gamma) \sum_{i \in \mathcal{S}} Y_t^i, \quad (16)$$

where $\pi_t^{j,i;a}$ are CES import shares, and Y_t^i denotes regional income. I assume that total income in a region equals the wage bill plus agricultural profits, land-clearing costs, and migration costs accrued in region i at time t and that all income is spend locally on the consumption goods.

Given labor allocations across regions $\{L_t^i\}$, the temporary equilibrium determines wages, land use, and income distributions, which feed back into migration decisions in the dynamic block.

6. Estimation and calibration

6.1. Calibration of the baseline economy

The baseline economy is calibrated to 2010 under the assumption that the static trade model is in temporary equilibrium. Some parameters are taken from the literature, while the remaining ones are disciplined by matching model-implied moments to observed data.

Exogenous parameters. The Armington elasticity is set to $\epsilon = 9$ (Allen and Arkolakis, 2014); and the dispersion of land productivity to $\theta = 2$, the midpoint of available estimates (Sotelo, 2020; Costinot et al., 2016). Interregional iceberg trade costs are parameterized as $\delta^{j,i} = (1 + \text{distance}^{j,i})^{0.08}$, implying around 65% higher prices at 500 kilometers from the producing region (Pellegrina, 2022).

Inversion from data. Observed inputs include the land endowments X_0^i , cultivated share of land ζ_0^i , regional labor endowments L_0^i and sectoral allocations $L_0^{i;a}, L_0^{i;n}$, sectoral wages $w_0^{i;a}, w_0^{i;n}$, measured agricultural productivity among cultivated plots, and agricultural value of production $p_0^{i;a} Q_0^{i;a}$.

From these, I back out: the Cobb-Douglas share α from the agricultural wage bill-to-output ratio (0.1579); the distributional parameters (μ^i, σ^i) of non-agricultural wages by matching the agricultural labor share λ_0^i and the observed non-agricultural mean wage; the fundamental land productivity A_0^i by adjusting measured productivity for endogenous selection; producer prices $p_0^{i;a}$ by inverting the agricultural wage equation; and agricultural fixed costs F^{i0} from the optimal land-use condition.

Residual calibration. Given these objects, I recover agricultural and non-agricultural outputs $(Q_0^{i;a}, Q_0^{i;n})$, land-clearing costs B_0^i , agricultural profits R_0^i , and total regional income Y_0^i . The agricultural expenditure share γ is set to match the ratio of agricultural spending to total income (0.1190). Finally, region-specific preference shifters ψ^j are solved as a fixed point to clear the agricultural goods market, and are interpreted as quality/taste parameters.

Table II summarizes all calibrated objects, their interpretation, and their empirical targets. Appendix D.1 presents additional information on the calibration and model fit.

TABLE II
Calibration of the baseline economy, 2010

Object	Interpretation	Target / Source	Value
ϵ, β, θ	Trade elasticities, discounting, heterogeneity in land quality	Literature	9, 0.95, 2
α	Labor share in agriculture	Agricultural wage bill / output	0.1579
γ	Agricultural expenditure share	Agricultural spending / income ratio	0.119
<i>Varying by region:</i>			
μ^i	region parameter (non-agricultural wages)	$w_0^{i;a}, w_0^{i;n}, \lambda_0^i$	9.36 (0.43)
σ^i	Scale parameter (non-agricultural wages)	$w_0^{i;a}, w_0^{i;n}, \lambda_0^i$	0.64 (0.18)
A_0^i	Fundamental land productivity	Corrected observed productivity	660.12 (524.03)
$p_0^{i;a}$	Agricultural producer price	Inverted from wage equation	17.08 (21.34)
F^{i0}	Fixed cost of land clearing	Land share condition	7109.07 (22908.90)
ψ^j	Regional preference shifter	Agricultural goods market clearing	0.01 (0.08)
δ_j^i	Trade costs	Distance data	1.78 (0.12)

Notes: The table summarizes the calibration of the baseline economy to 2010. Exogenous parameters are taken from the literature or distance data. Endogenous parameters and objects are identified using the temporary equilibrium conditions of the static model and matched to Brazilian census and agricultural statistics. For parameters varying by region, the first number reported represents the mean over all regions and the number in paratheses indicates the standard deviation.

6.2. Counterfactual fundamental productivity

To simulate climate change, I map projected temperatures into counterfactual changes in fundamental agricultural productivity. Baseline productivity A_0^i is recovered in the calibration (Section 6.1). Counterfactual values A_t^i are obtained by scaling these baselines according to the estimated quadratic temperature-productivity relationship from Figure IV.

Let $\hat{g}(TG) = \hat{\beta}_1 TG + \hat{\beta}_2 TG^2$ denote the fitted quadratic response of $\ln A_{it}$ to temperature. For region i and year t , I compute

$$\ln A_t^i = \ln A_0^i + \left(\hat{g}(TG_{it}) - \hat{g}(TG_{i0}) \right),$$

where TG_{i0} is the region's baseline temperature in 2010.

This procedure preserves the cross-sectional heterogeneity in A_0^i while allowing productivity dynamics to reflect the nonlinear temperature response. Because each region is anchored to its own baseline temperature, cooler regions and hotter regions move along different parts

of the quadratic curve, implying heterogeneous productivity impacts of a common warming shock. Projected temperatures TG_{it} under SSP5-8.5, then generate the counterfactual productivity paths used in the simulations.

In the model, I measure five-year periods from 2010 to 2100, with $t=0$ corresponding to 2010. Temperature projections are averaged over five-year windows to obtain A_t^i for $t = 0, 1, \dots, 18$. My estimates imply that, by 2100, some regions will lose up to 55% of initial productivity, while others (notably those with low starting temperatures) will gain up to 38.5%.

TABLE III
Projected deviation from baseline productivity by year

Year	Median change	Min change	Max change
2030	-1.5%	-8.3%	13.1%
2050	-4.9%	-16.4%	20.8%
2100	-24.2%	-55.3%	38.5%

Notes. This table reports the median, minimum, and maximum change in mesoregions' fundamental agricultural productivity for 2030, 2050, and 2100 (relative to 2010). Counterfactual productivity levels are obtained by applying the estimated quadratic temperature-productivity relationship (from Figure IV) to projected annual mean temperatures under the SSP5-8.5 scenario, simulated with the CMCC-ESM2 Earth System Model from the Copernicus Climate Change Service (C3S) CMIP6 archive. Baseline fundamental productivity for 2010 is recovered in the calibration of the model. Projections use five-year averaged temperatures expressed as anomalies relative to 2010.

The full path of projected deviations from the baseline is shown in Appendix D.2.

6.3. Dynamic part & migration costs

The dynamic model requires parameters $\nu, \rho, \{\chi^{ji}\}_{j,i \in \mathcal{S}}, \{\kappa^j\}_{j \in \mathcal{S}}$, which represent the variance of idiosyncratic location preferences, the persistence of non-agricultural productivity draws, the full bilateral matrix of monetary migration costs, and the origin-specific non-monetary cost of leaving “home”.

Exogenous parameters. Following Cai et al. (2022), I use a discount factor β of 0.86, equivalent to an annual discount factor of 0.97. Rank persistence is fixed at $\rho = 0.9$, consistent with the low intergenerational mobility in Brazil (Britto et al., 2022). The idiosyncratic taste scale is set to $\nu = 1.75$ (Caliendo et al. (2019) find an annual elasticity of 2.02 and discuss that this value must be larger at higher frequencies).¹²

Recovering migration costs. To recover migration costs, I match model-implied aggregate migration probabilities between regions in the first period to the observed migration probabilities in the data. For this, I solve the full dynamic model under alternative cost shifters and

¹²Typically, values of ν are inferred from regressions of migration flows on wage differentials. A weak empirical response of migration to wage gaps is interpreted as evidence of a high ν , implying that idiosyncratic taste dispersion $1/\nu$ is low. However, in my model, low migration responsiveness arises not from strong idiosyncratic preferences but from mobility constraints. Under this interpretation, existing empirical approaches would overestimate ν , attributing limited migration responses to high preference heterogeneity rather than to constraints. This reasoning is consistent with the observation that studies such as Cruz (2024) tend to find higher values of ν for low-income countries. It seems unlikely that individuals in these contexts have stronger idiosyncratic preferences; a more plausible explanation is that they face greater barriers to responding to wage differentials.

select the cost shifter that best aligns the model-implied probabilities with its data counterpart.¹³

I implement the model as a finite horizon problem, whereby the model starts in 2010 and ends in 2100. A period represents five years; hence, there are 18 periods. The state variable is a rank in the non-agricultural productivity distribution r . For tractability, I discretize each region's non-agricultural productivity distribution into 100 equiprobable bins and use each bin's average implied income (y_r) as the representative state when computing continuation values and migration feasibility/costs. To calibrate costs, I assume that workers expect agricultural productivity to evolve according to the pessimistic warming scenario SSP5-8.5.

I assume that migration costs have a variable distance component and an origin-specific fixed cost:

$$\chi^{j,i} = \theta^{\text{dist}} \cdot \text{distance}^{j,i},$$

$$\tau^{j,i}(y) = \begin{cases} -\log(1 - \chi^{j,i}/y) + \kappa^j \cdot \mathbf{1}\{i \neq j\}, & y > \chi^{j,i}, \\ \infty, & y \leq \chi^{j,i}, \end{cases}$$

where θ^{dist} governs how monetary costs rise with distance, and κ^j captures non-monetary “home” preferences at origin j . By construction, migration is disproportionately costly for low incomes, with $\tau^{j,i}(y) \rightarrow \infty$ as y approaches $\chi^{j,i}$ from above, and it becomes infeasible if income does not cover the monetary cost.

Solution algorithm. Conditional on ν and β , I search for the parameters θ^{dist} and κ^j that jointly determine migration costs under the parametric assumption outlined above. Formally, the minimization problem is:

$$\min_{\theta^{\text{dist}}, \{\kappa^j\}_{j \in S}} \sum_{j,i} (\mu_0^{j,i}(\theta^{\text{dist}}, \kappa^j) - \mu_{0,\text{data}}^{j,i})^2,$$

where $\mu_{0,\text{data}}$ are observed bilateral migration probabilities.


The calibration is computationally heavy, as it relies on the model being solved many times for different costs. To keep the implementation feasible, the search over θ^{dist} is implemented through a grid search from R\$10/km to R\$60/km in steps of R\$10/km. Conditional on a given θ^{dist} , I then solve non-linearly for the non-monetary costs of leaving “home”, κ^j , to match each region's overall emigration rate.

Results. I find that a θ^{dist} of R\$50/km best matches bilateral migration probabilities. The fixed costs κ^j have a mean of 6.69, implying a high psychic cost of leaving “home” (equivalent to around 96% of annual income). Figure VI shows the associated model fit. The R-squared of regressing model-implied migration probabilities on the observed probabilities is above 0.99, with a slope coefficient close to one, which is partly driven by the very close fit of regional

¹³I assume that migration probabilities between 2005 and 2010 recovered from the 2010 Census are representative of probabilities from 2010 to 2015.

emigration rates. Excluding the probability of staying in the same region, and focusing on the bilateral migration probabilities to other regions, the R-squared is 0.43 with a slope of 0.77.

FIGURE VI
Migration probabilities: Model vs. data



Notes. This figure compares the model-implied and observed bilateral migration probabilities between Brazilian mesoregions. Migration probabilities are defined as $L^{ji} / \sum_i L^j$, that is, the share of workers from origin j residing in destination i . Observed probabilities are derived from the 2010 Population Census (*Censo Demográfico*) and reflect residential locations in 2005 and 2010. Model-implied probabilities correspond to the first-period migration probabilities generated by the calibrated dynamic spatial model, which infers monetary and non-monetary migration costs by minimizing the distance between simulated and observed probabilities. Monetary costs increase linearly with geographic distance (R\$50 per km), and origin-specific non-monetary “home” preferences are calibrated to match regional emigration rates. The resulting specification closely reproduces observed migration patterns, with an overall R^2 above 0.99; focusing on outflows only, the R^2 is approximately 0.45.

Untargeted moments. In the calibration, I target bilateral migration probabilities (not aggregate flows). In [Table IV](#), I compare the model-implied and observed flows aggregated into three categories: within-mesoregion, between-mesoregion but within-state, and between-states. The model closely matches the observed distribution of flows. Overall, the calibration indicates that the model can replicate key features of observed migration patterns. [Appendix D.3](#) shows further statistics of the model fit and a related discussion of the calibration.

TABLE IV
Fit of migration flows

Category	Model	Data
Same Mesoregion	95.5%	95.6%
Different Mesoregion, Same State	1.8%	1.8%
Different State	2.7%	2.6%

Notes. This table compares the fit of the model-implied and observed migration flows between Brazilian mesoregions. Observed flows represent the mass of workers moving between 2005 and 2010, constructed from the 2010 Population Census (*Censo Demográfico*) based on individuals' reported municipalities of residence in 2005 and 2010, aggregated into mesoregion-to-mesoregion transitions. Model-implied flows correspond to the mass of movers generated in the first period of the calibrated dynamic spatial model, which incorporates the estimated monetary and non-monetary migration costs. Monetary costs increase linearly with geographic distance (R\$50 per km), and origin-specific non-monetary "home" preferences are calibrated to match regional emigration rates. The model closely replicates the observed distribution of worker flows across same-mesoregion, within-state, and cross-state moves.

7. Welfare losses from climate change

Climate change reshapes regional productivity and, through it, the spatial distribution of income and welfare. This section quantifies how these productivity shocks translate into welfare losses across space and along the income distribution. The analysis focuses on three questions: where climate damages are concentrated, which households bear them, and what role migration plays in mitigating losses. I first document the spatial and distributional incidence of climate damages before turning to the role of migration as a potential adaptation channel.

Measurement and counterfactuals. I measure the welfare effects of climate change by contrasting expected lifetime utility in a world where regional productivities remain at their 2010 levels with one where they evolve under projected warming. In the latter, the temperature–productivity mapping from [Section 6](#) implies that already-hot regions become less productive, while cooler regions experience modest gains as they move toward the agronomic optimum.


Welfare differences are expressed in consumption-equivalent terms—the constant proportional, permanent increase in consumption that would make a household under climate change as well off as in the no-climate-change counterfactual. In other words, it measures the uniform percentage change in consumption applied in every period from 2010 to 2100 that equalizes lifetime utility across the two scenarios.

7.1. Quantifying welfare losses

[Figure VII](#) maps welfare changes across the income distribution. Panel (a) displays outcomes for households at the lowest income percentile, and panel (b) for those at the top. Spatial heterogeneity is striking. In the North and Northeast—where rising temperatures sharply depress agricultural productivity—welfare declines by up to 4% in consumption-equivalent terms. In contrast, parts of the South experience welfare *gains* of similar magnitude as moderate warming raises productivity toward the agronomic optimum. These gains reflect favorable local shocks rather than improvements in aggregate welfare.

At the top of the income distribution, welfare effects are negligible across all regions. High-income households are primarily employed in high-productivity non-agricultural jobs and are largely insulated from changes in local agricultural wages and the resulting expansion of low-productivity non-agricultural employment. Their modest welfare losses stem mainly from price effects, as climate change raises the cost of the consumption bundle. Overall, the figure reveals a steep distributional gradient and substantial *within-region* heterogeneity: in many cases, welfare differences within a region are as large as those across regions.

FIGURE VII
Welfare changes by income percentile



Notes: Each panel maps the consumption-equivalent welfare change in 2010 relative to a no-climate-change baseline. Panel (a) reports results for the lowest income percentile in each region and panel (b) for the highest. Welfare changes are expressed as permanent percentage shifts in consumption that equalize lifetime utility across scenarios. Negative values indicate welfare losses; positive values indicate gains. Income percentiles correspond to the 100 discrete income groups in the model and are aggregated into the income bins used in [Table V](#). The figure visualizes the spatial pattern of the welfare differences summarized numerically in the table below.

[Table V](#) summarizes welfare changes by major region and income group in 2010. To facilitate comparison across spatial and income dimensions, each region's 100 income percentiles are grouped into three bins, R\$0–10k, R\$10–20k, and above R\$20k, representing low-, middle-, and high-income households, respectively.

The top panel of [Table V](#) reports population-weighted mean welfare changes by region and income group. Across all regions, losses are concentrated among poorer households: climate damages are markedly regressive. In the North, households earning below R\$10,000 per year lose nearly 0.9% in consumption-equivalent welfare, compared with only 0.3–0.4% among high-income households. Similar gradients appear in the Northeast (–0.6% versus –0.2%) and Center-West (–0.7% versus –0.2%). These gaps reflect poorer households' heavier dependence on local agricultural wages and their limited ability to buffer losses by shifting into non-agricultural work or by migrating: as climate change reduces agricultural productivity, it depresses local reservation wages and draws marginal workers into low-productivity non-agricultural activities, amplifying welfare losses at the bottom of the income distribution.

Regional averages, also shown in the top panel, exhibit clear spatial variation. Mean wel-

TABLE V
Welfare losses by region and income bin

	R\$ 0–10k	R\$ 10–20k	R\$ >20k	All households
<i>Population-weighted means</i>				
Norte	-0.87	-0.36	-0.34	-0.55
Nordeste	-0.57	-0.24	-0.23	-0.40
Centro-oeste	-0.68	-0.21	-0.17	-0.24
Sudeste	+0.09	-0.03	-0.03	-0.02
Sul	+0.77	+0.15	+0.13	+0.24
<i>Maximum losses</i>				
Norte	-3.23	-0.90	-0.90	-3.23
Nordeste	-2.34	-0.67	-0.45	-2.34
Centro-oeste	-2.13	-0.47	-0.43	-2.13
Sudeste	-1.62	-0.35	-0.30	-1.62
Sul	-0.63	-0.11	-0.12	-0.63

Notes: Entries report consumption-equivalent welfare changes (percent) in 2010 relative to a no-climate-change baseline, by major region and income bin. Income bins aggregate 100 model percentiles per region into three observed 2010 income groups. The top panel reports population-weighted mean welfare changes; the bottom panel reports the maximum welfare loss in each region–income cell. Negative values denote welfare losses and positive values denote welfare gains. The final column reports the mean or maximum across all households in each region.

fare losses are largest in the tropical North and Northeast, at -0.55% and -0.40% , respectively. The Center-West experiences smaller average losses (-0.24%). The temperate South, by contrast, records modest welfare *gains* of $+0.24\%$, as moderate warming raises local productivity toward the agronomic optimum and boosts agricultural wages. The Southeast shows near-zero mean effects (-0.02%): in these largely urban regions, incomes depend little on local agricultural productivity, and climate shocks propagate only weakly through local labor markets.

The bottom panel of [Table V](#) highlights the extent of within-region heterogeneity masked by regional averages. It reports the largest welfare loss observed across any income group within each region. In the North, the most adversely affected group loses over 3% in consumption-equivalent welfare, while the corresponding figures for the Northeast and Center-West are around 2–2.5%. Even in the Southeast—where mean effects are close to zero—the most exposed groups experience declines of roughly 1.6%. The South remains largely insulated from major losses, reflecting its cooler baseline temperatures, higher incomes, and more diversified economic base.

Taken together, these results point to a clear pattern: climate change imposes its largest welfare costs on poor households in already-hot regions, while richer and more urbanized areas remain largely insulated. The interaction of spatial exposure, sectoral composition, and income inequality generates a double asymmetry—across both geography and the income distribution—that shapes the aggregate welfare burden of warming.

7.2. Migration as adaptation?

The preceding results show that welfare losses are concentrated among poorer households in hot regions. This subsection examines migration as a potential adaptation mechanism. Ta-

ble VI summarizes how migration opportunities vary across the income distribution under climate change. For each year t , I compute (i) the share of households whose income is insufficient to cover the minimum cost of migrating to any other region and (ii) the mean probability of out-migration. Reported values are population-weighted averages of these measures over the entire simulation horizon (2010–2100). The first column captures the prevalence of liquidity constraints that prevent relocation; the second reports the corresponding realized migration rates in the face of climate change.

TABLE VI
Time-averaged migration constraints and mobility rates under the climate scenario

Income group	Share unable to migrate ($\chi \geq y$)	Mean migration probability
Low income ($\leq 10k$)	52.43%	5.41
Middle income (10–20k)	6.48%	6.39
High income ($> 20k$)	0.38%	11.77

Notes. Each entry reports the population-weighted average across time periods under the climate scenario. Households are classified as *unable to migrate* when their income $y_{i,t}$ falls below the minimum off-diagonal migration cost $\min_{j \neq i} \chi^{ij}$ from their origin. Migration probabilities are defined using off-diagonal elements of the migration matrix μ_t^{ij} .

The table shows that migration opportunities are sharply stratified by income. Roughly half of low-income households are liquidity-constrained, compared with fewer than 7% of middle-income and virtually none of high-income households. Mobility rates mirror these gaps: only about 5% of low-income households migrate in an average year, versus nearly 12% for the high-income group. Hence, migration plays only a minor role as an adaptation mechanism—not because climate shocks fail to generate incentives to move, but because a large share of poor households are unable to finance relocation. The constraint operates primarily on the extensive margin: many low-income households are excluded from migration altogether, rather than marginally adjusting their mobility rates in response to climate shocks.

7.3. Discussion

The pattern of welfare changes reflects the interaction of two central adjustment margins in the model. (i) *Production margin*: Climate change lowers agricultural productivity in hot regions, reducing reservation wages and pushing marginal workers into low-productivity non-agriculture. This process compresses mean income and raises dispersion within affected regions, generating regressive welfare losses. (ii) *Migration margin*: Because migration requires a cash outlay, many low-income households are liquidity-constrained even before the shock. Climate change intensifies incentives to relocate but does little to relax these preexisting constraints, so migration contributes little as an adaptation channel. Together, these forces explain why welfare losses are largest for low-income households in already-hot regions and modest or even positive elsewhere.

When liquidity constraints bind, adaptation through migration is intrinsically limited: the households facing the strongest incentives to move are those least able to do so. The dis-

tributational incidence of climate change is thus shaped as much by financial frictions as by geography. Policies that relax such constraints—through mobility subsidies, credit access, or safety nets—are natural complements to place-based strategies that target climatic exposure (see [Section 8](#)).

8. Policy Experiment: Income-Targeted Migration Subsidy

Motivation. Climate damages fall most heavily on low-income households in hot regions—the very households facing the steepest migration frictions. To evaluate an adaptation policy that directly addresses this channel, I simulate a counterfactual reform that subsidizes the *monetary* cost of migration for low-income households while leaving non-monetary barriers unchanged. The reform targets liquidity constraints rather than preferences: it equalizes the utility-equivalent migration cost between a low-income household and a higher-income reference type, thereby reducing income-dependent wedges in mobility.

Policy mapping and implementation. Let $\tau^{j,i}(y; \chi^{j,i})$ denote the utility-equivalent cost of moving from origin j to destination i for a household with income y , facing monetary cost $\chi^{j,i}$ and non-monetary cost κ^j embedded in $\tau^{j,i}(\cdot)$. Equalizing utility-equivalent costs between a low-income type y_r and a higher-income reference y_h requires a subsidy $\zeta^{j,i}$ satisfying

$$\tau^{j,i}(y_r; \chi^{j,i} - \zeta^{j,i}) = \tau^{j,i}(y_h; \chi^{j,i}), \quad (17)$$

which, under log utility, implies

$$\zeta^{j,i}(y_r | y_h) = \chi^{j,i} \left(1 - \frac{y_r}{y_h}\right), \quad \text{for } y_r < y_h, \quad (18)$$

and $\zeta^{j,i} = 0$ otherwise. Intuitively, the subsidy scales the out-of-pocket cost in proportion to the income gap relative to the reference type.

In the baseline experiment, the reference income is set to $y_h = \text{R\$ } 20,000$. Thus, only households with $y_r < y_h$ receive a subsidy, while richer households do not. The subsidy is offered every period from 2010 to 2100.

8.1. Welfare gains from subsidy

Table [VII](#) reports consumption-equivalent welfare changes by income bin. Under climate change, relaxing monetary migration costs yields large welfare gains for low-income households and more modest gains for higher-income groups: lifetime welfare rises by 24.09% for households below R\$10k, 5.49% for those with R\$10–20k, and 1.33% for those above R\$20k.

The second row reports the share of baseline climate losses offset by the policy. The subsidy offsets roughly 4.2% of baseline losses for the bottom two income groups and 1.0% for the top, implying that the program cushions damages precisely where they are most concentrated without materially affecting the upper distribution.

TABLE VII

Welfare gains and attenuation of climate losses from a migration subsidy (evaluated at $t = 0$)

	R\$ 0–10k	R\$ 10–20k	R\$ >20k	All households
<i>Climate change in both scenarios</i>				
Subsidized vs. Baseline (CEV, %)	24.09	5.49	1.33	8.22
Share of climate losses removed (%)	4.16	4.10	1.01	2.92

Notes: Entries report population-weighted mean consumption-equivalent welfare changes (percent) under climate change. The first row compares welfare between the subsidized and baseline migration-cost regimes, holding the climate path fixed. The second row reports the share of baseline climate losses offset by the policy, defined as the reduction in welfare losses relative to baseline losses. The last column reports the national average across all households. Negative values (not shown) would denote welfare losses.

Table VIII disaggregates welfare gains by macro-region and income group. Gains are largest in the Northeast, where low-income households experience welfare increases exceeding 36%. This reflects two reinforcing factors: high climate exposure in agriculture and steep baseline migration wedges that strongly limit out-migration absent the subsidy. The North and Center-West also see sizable low-income gains (14% and 9%, respectively), and the South and Southeast show comparable improvements (around 11% in both).

TABLE VIII
Welfare gains by region (evaluated at $t = 0$)

	R\$ 0–10k	R\$ 10–20k	R\$ >20k	All households
<i>Climate change in both scenarios</i>				
Subsidized vs. Baseline (CEV, %)				
Norte	14.11	6.05	2.94	8.30
Nordeste	36.41	13.70	4.94	23.44
Centro-oeste	9.37	4.87	1.29	3.68
Sudeste	11.40	2.66	0.56	2.68
Sul	11.17	4.33	1.09	4.19

Notes: Entries report consumption-equivalent welfare gains (percent) from the migration subsidy under climate change, by macro-region and income bin. Each entry compares welfare under the subsidized and baseline regimes, holding the climate path fixed. The last column reports the regional population-weighted mean. “Income bins” correspond to 2010 income groups used throughout the paper (R\$0–10k, R\$10–20k, and above R\$20k).

To clarify the mechanism behind these welfare gains, Table IX reports how the subsidy alters migration feasibility and actual mobility rates. The first column shows the share of households whose income falls below the minimum monetary migration cost $\min_{j \neq i} \chi^{ij}$, while the second column reports mean migration probabilities averaged over time and population. Under the baseline calibration, roughly half of low-income households face binding liquidity constraints that prevent migration. Introducing the subsidy sharply relaxes these constraints: the share unable to migrate falls from 52% to below 4%, and their average migration probability roughly doubles. Middle-income households experience a smaller but still notable increase in mobility, while high-income households remain virtually unaffected, as they were uncon-

strained to begin with.

TABLE IX
Effect of the migration subsidy on migration constraints and mobility (climate path)

Income group	Share unable to migrate ($\chi \geq y$)	Mean migration probability
Low income ($\leq 10k$)	52.43% $\rightarrow$ 3.90%	5.41 $\rightarrow$ 11.09
Middle income (10–20k)	6.48% $\rightarrow$ 2.56%	6.39 $\rightarrow$ 5.41
High income ($> 20k$)	0.38% $\rightarrow$ 0.38%	11.77 $\rightarrow$ 9.63

Notes: Each entry compares the baseline and subsidized migration-cost regimes under the same climate path. Under the subsidy, all households with baseline income below R\$20,000 are treated as if they had income equal to R\$20,000 when facing monetary migration costs χ^{ij} . “Share unable to migrate” reports the fraction of households whose income falls below the minimum off-diagonal migration cost $\min_{j \neq i} \chi^{ij}$; “Mean migration probability” reports the average off-diagonal migration probability μ^{ij} . Values are weighted by population and averaged over time. Arrows ($\rightarrow$) indicate the change from the baseline to the subsidized regime.

8.2. Fiscal incidence and dynamic payoffs

Table X summarizes the fiscal incidence of the migration subsidy, averaged across all years in which the policy is active (2010–2100). On average, annual transfers amount to about 1.3% of GDP per period. The distribution of payments is sharply progressive: households earning below R\$10,000 per year receive an average per-capita transfer of R\$913, those between R\$10,000 and R\$20,000 receive R\$173, and higher-income households receive none.

This incidence pattern arises endogenously from equilibrium migration decisions under the policy. Because the subsidy is paid only to households that actually migrate, observed transfers reflect both who moves and how large a top-up is required to offset their mobility cost. Low-income households receive higher transfers for two reasons. First, they are more likely to migrate in response to climate shocks once liquidity barriers are relaxed. Second, by construction of the policy rule, the subsidy scales with the *migration cost wedge*—the difference between a household’s income and the benchmark income used to define an unconstrained mover. This design ensures that poorer households, for whom the monetary cost of relocation represents a larger fraction of income, receive proportionally greater support. In equilibrium, fiscal resources are thus concentrated precisely on those facing binding liquidity constraints, achieving strong redistributive targeting at modest fiscal cost.

The preceding analysis shows that the migration subsidy generates large welfare gains for low-income households, mitigating a substantial share of their climate-induced losses by enabling relocation to higher-productivity regions. While these welfare improvements come at a fiscal cost of roughly 1.3% of GDP per period, the policy also delivers significant aggregate output gains over time. Figure VIII traces the dynamic relationship between these fiscal costs and macroeconomic payoffs.

The solid line plots the cumulative present value of additional GDP generated by the subsidy, and the dashed line shows the cumulative present value of subsidy expenditures, both expressed as a share of initial GDP. Over the policy horizon, the program nearly pays for itself: the discounted value of incremental GDP (7.9%) covers about 90% of the present value of total


TABLE X
Average fiscal incidence of the migration subsidy

Income bin	Avg. transfer per capita (R\$)	Share of GDP (%)
R\$ 0–10k	913.3	0.96
R\$ 10–20k	173.1	0.32
R\$ >20k	0.0	0.00
Total	271.5	1.27

Notes: Each entry reports the population-weighted average per-capita transfer (R\$) and the corresponding fiscal outlay as a share of GDP in the same period, averaged across all years in which the subsidy is active (2010–2100). Fiscal shares are calculated as the mean of annual transfer-to-GDP ratios over the policy horizon. The final row reports the population-weighted national average.

transfers (8.7%). Using a five-year discount factor of $\beta = 0.86$ (equivalent to an annual rate of about 3%), the cumulative net present value remains slightly negative within the simulation window, reflecting the front-loaded fiscal cost of assisting early movers. However, the implied internal rate of return on the stream of additional output relative to subsidy spending is roughly 13% per model period, or about 2.5% per year. The shortfall arises because migration-induced output gains accumulate gradually as workers relocate toward more productive regions, while fiscal costs are concentrated at the onset of the program.

FIGURE VIII
Self-financing profile of the migration subsidy



Notes: The solid line shows the cumulative present value of additional GDP generated by the subsidy, and the dashed line shows the cumulative present value of subsidy expenditures. Both are expressed as percentages of initial baseline GDP. Discounting uses a five-year factor of $\beta = 0.86$ (approximately a 3% annual rate). Over 2010–2100, the present value of additional GDP covers about 90% of total transfers, implying an internal rate of return of roughly 13% per model period (2.6% annually).

Mechanism and aggregate effects Aggregate output gains arise from higher output per worker, shown in [Table XI](#). Under baseline migration costs, climate change sharply reduces agricultural productivity: by 2095, agricultural output per worker falls by 12.5%, while non-

agriculture declines only marginally. When the subsidy is introduced, output per worker rises persistently in both sectors. By 2095, output per worker is 6.4% higher in agriculture and over 10% higher in non-agriculture relative to the baseline, reflecting the gradual reallocation of labor from low-productivity, climate-damaged regions to high-productivity destinations. These compositional effects translate large micro-level welfare gains into aggregate efficiency gains, providing the fiscal base that nearly finances the program.

TABLE XI
Output per worker by sector and policy regime (percent differences)

	Year	Agriculture	Non-agriculture	All sectors
<i>Baseline mobility</i>				
Climate – No climate	2030	0.71	-0.05	0.04
Climate – No climate	2050	-0.59	-0.10	-0.16
Climate – No climate	2095	-12.51	-0.40	-1.84
<i>Policy effect under climate change</i>				
(Subsidy – Baseline)	2030	2.74	1.03	1.23
(Subsidy – Baseline)	2050	4.30	1.55	1.88
(Subsidy – Baseline)	2095	6.36	10.48	9.99

Notes: Entries show percent differences in output per worker between scenarios. For each year, “Climate – No climate” compares the climate path to the no-climate baseline within the same migration regime. “Policy effect under climate” reports the difference between the subsidized and baseline regimes under climate change. The final column aggregates sectors using fixed 2010 GDP weights. Output per worker is defined as sectoral output divided by sectoral employment.

8.3. Interpretation and policy perspective

Income-targeted migration assistance operates simultaneously as redistribution and adaptation. The policy does not offset climate-induced productivity shocks directly; rather, it mitigates their welfare incidence by enabling liquidity-constrained households to relocate to higher-productivity regions and sectors. Because these households are precisely those with the largest exposure to climate damages, the subsidy delivers large equity gains at modest fiscal cost and yields gradual macroeconomic dividends over time.

More broadly, the experiment illustrates how the incidence of climate change depends jointly on geography and financial frictions. When migration costs bind, adaptation through mobility is intrinsically regressive: those who would benefit most from moving are least able to afford it. Reducing monetary barriers therefore serves both efficiency and equity, transforming localized adaptation failures into national welfare gains.

9. Conclusion

This paper has shown that income shapes the ability of households to respond to climate change through migration. Using Brazilian Census microdata, I documented that within the same origin, higher-income individuals are far more likely to move. I then developed a dynamic spatial general equilibrium model with income-dependent migration costs to interpret

this fact and quantify the distributional welfare consequences of climate change.

The analysis yields three main findings. First, climate change generates strongly regressive welfare impacts. In already-hot regions, low-income households suffer consumption-equivalent welfare losses of up to 5% per period, while richer households in the same regions remain largely insulated. By contrast, in cooler regions, climate change can raise agricultural productivity, so that some low-income households experience modest gains. Second, the mechanism is twofold: falling agricultural productivity lowers reservation wages and reallocates workers into low-productivity jobs, while lower incomes magnify the utility burden of migration costs, dampening mobility precisely where the incentive to move is greatest. Third, an income-targeted migration subsidy that reduces the monetary costs of moving for low-income households raises welfare, partially offsets climate damages, and pays for itself in the long run through higher output per worker.

Together, these results highlight that adaptation to climate change is shaped not only by geography but also by inequality. Low-income households are most exposed to climate change damages yet are least able to utilize migration as an adaptation strategy. Policies that relax liquidity constraints on mobility can therefore act as both redistribution and climate adaptation, complementing place-based policies that address local productivity shocks. More broadly, the findings suggest that in the presence of migration costs, the distributional incidence of economic shocks depends as much on who is able to respond as on where the shocks occur.

References

- Allen, T. and C. Arkolakis (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics* 129(3), 1085–1140.
- Bastian, J. and S. E. Black (2022). The eitic and internal migration. *American Economic Journal: Applied Economics* 14(4), 1–30.
- Bastos, P., M. Busso, and S. Miller (2013). Adapting to climate change: Long-term effects of drought on local labor markets. Technical report, IDB Working Paper Series.
- Bazzi, S. (2017, April). Wealth heterogeneity and the income elasticity of migration. *American Economic Journal: Applied Economics* 9(2), 219–55.
- Britto, D. G., A. Fonseca, P. Pinotti, B. Sampaio, and L. Warwar (2022). *Intergenerational mobility in the land of inequality*. JSTOR.
- Brunel, C. and M. Y. Liu (2025). Out of the frying pan: Weather shocks and internal migration in brazil. *World Development* 188. Article 106859.
- Bryan, G., S. Chowdhury, and A. M. Mobarak (2014). Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. *Econometrica* 82(5), 1671–1748.
- Bryan, G. and M. Morten (2019). The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy* 127(5), 2229–2268.
- Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature* 527(7577), 235–239.
- Cai, S., L. Caliendo, F. Parro, and W. Xiang (2022). Mechanics of spatial growth. Technical report, National Bureau of Economic Research.
- Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica* 87(3), 741–835.
- Cattaneo, C. and G. Peri (2016). The migration response to increasing temperatures. *Journal of Development Economics* 122, 127–146.
- Choi, J. (2024). The effect of deindustrialization on local economies: Evidence from new england textile towns. Working Paper.
- Clark, G. and N. Cummins (2024). Intergenerational wealth mobility in england: New methods and estimates from linked administrative data. *Working Paper*.
- Conte, B. (2022). Climate change and migration: The case of africa. Technical Report 9948, CESifo Working Paper Series, Munich, Germany.
- Costinot, A., D. Donaldson, and C. Smith (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy* 124(1), 205–248.
- Cruz, J.-L. (2024). Global warming and labor market reallocation. SSRN Scholarly Paper 4946752.
- Cruz, J.-L. and E. Rossi-Hansberg (2024). The economic geography of global warming. *Review of Economic Studies* 91(2), 899–939.
- Feng, S., A. B. Krueger, and M. Oppenheimer (2010). Linkages among climate change, crop yields and mexico-us cross-border migration. *Proceedings of the National Academy of Sci-*

- ences 107(32), 14257–14262.
- Gollin, D., C. W. Hansen, and A. M. Wingender (2021). Two blades of grass: The impact of the green revolution. *Journal of Political Economy* 129(8), 2344–2384.
- Gollin, D. and J. Wolfersberger (2024). Agricultural trade and deforestation: The role of new roads.
- Jedwab, R., F. Haslop, R. Zarate, and C. Rodriguez Castelan (2023). The effects of climate change in the poorest countries: Evidence from the permanent shrinking of lake chad. Technical report, IZA Discussion Papers.
- Lobell, D. B. and S. M. Gourdji (2012). The influence of climate change on global crop productivity. *Plant Physiology* 160(4), 1686–1697.
- Lucas, R. E. J. (1978). On the size distribution of business firms. *Bell Journal of Economics* 9(2), 508–523.
- Luttmer, E. G. J. (2007). New goods and the size distribution of firms. Working paper no. 649, Federal Reserve Bank of Minneapolis.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. New York: Academic Press.
- Morten, M. and J. Oliveira (2024). The effects of roads on trade and migration: Evidence from a planned capital city. *American Economic Journal: Applied Economics* 16(2), 389–421.
- O’Connor, D. (2024). Revitalize or relocate? the welfare effects of place-based policies. *Working Paper*.
- Pellegrina, H. S. (2022). Trade, productivity, and the spatial organization of agriculture: Evidence from brazil. *Journal of Development Economics* 156, 102816.
- Pellegrina, H. S. and S. Sotelo (2021). Migration, specialization, and trade: Evidence from brazil’s march to the west. Technical report, National Bureau of Economic Research.
- Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics* 9(1), 21–58.
- Restuccia, D., D. T. Yang, and X. Zhu (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of Monetary Economics* 55(2), 234–250.
- Schlenker, W. and M. J. Roberts (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.
- Sotelo, S. (2020). Domestic trade frictions and agriculture. *Journal of Political Economy* 128(7), 2690–2738.

Online Appendix

A	Data	ii
A.1	Census data	ii
A.2	Agricultural production data	v
B	Empirical evidence	viii
B.1	Construction of Predicted Counterfactual Incomes	viii
B.2	Climate change and agricultural productivity	ix
C	Derivations	x
C.1	Derivations for the static block	x
C.1	Consumption	x
C.2	Production technology	xi
C.3	Temporary equilibrium	xii
C.2	Exact-hat derivations	xvi
D	Calibration details	xx
D.1	Baseline economic environment and calibration results	xx
D.2	Counterfactual productivity paths	xxi
D.3	Migration costs	xxii
E	Additional quantitative results	xxii
F	Solving the model	xxiii

Appendices

A. Data

A.1. Census data

The primary data source is the 2010 Brazilian Census, which provides detailed information on individuals' demographics, income, occupation, and migration history. The Census is conducted by the Brazilian Institute of Geography and Statistics (IBGE) and is representative at the mesoregion level.

To retrieve the data, I use IPUMS International. While IPUMS provides harmonized variables, I use the original Census occupation codes to map individuals to agricultural and non-agricultural sectors. IPUMS further allows me to directly retrieve an individual's current mesoregion of residence. The place of residence in 2005 is reported as the municipality of residence in 2005, which I map to mesoregions using the official IBGE correspondence. Income is measured as the total monthly income from all sources, reported in Brazilian Reais (R\$) in 2010 (variable INCTOT). I annualize income by multiplying the reported monthly income by 12.

Sample selection. The sample consists of Brazilian-born individuals aged 25–64 with positive income and non-missing information on occupation, place of residence in 2005, race, education, marital status, and gender. This yields approximately 6.2 million microdata observations, which correspond to about 60.3 million individuals in the weighted population.

Mapping sectoral employment. I map Census occupation codes to those working in crop cultivation and non-crop related occupations. Throughout the paper, I refer to the former as "agriculture" and the latter as "non-agriculture". This concept of agriculture is more narrow than the agricultural sector in national accounts, which includes activities such as livestock, forestry, and fishing. However, to align the model with my measure of agricultural productivity (which is based solely on crop production), I focus on crop cultivation. The mapping is based on the 2002 Brazilian Classification of Occupations (CBO) and is detailed in [Table A.1](#).

Summary statistics. [Table A.2](#) shows summary statistics for the sample used throughout the paper. The average age is 40 years with 43% female. The majority of individuals are married or in a union (70%) and identify as White (51%) or Brown (39%). Educational attainment is relatively low, with 30% having less than primary education completed and only 15% having completed university. The average annual income is R\$ 19,334, with a standard deviation of R\$ 59,143. Approximately 8.2% of the sample works in crop cultivation.

TABLE A.1
Mapping of Census occupation codes to agriculture

Code	Occupation
01101	Cultivation of rice
01102	Cultivation of corn
01103	Cultivation of other cereals for grains
01104	Cultivation of herbaceous cotton
01105	Cultivation of sugar cane
01106	Cultivation of tobacco
01107	Cultivation of soybeans
01108	Cultivation of manioc root
01109	Cultivation of other temporary agricultural farm products
01110	Cultivation of salad greens, vegetables and other such plant products
01111	Cultivation of flowers, ornamental plants and greenhouse products
01112	Cultivation of citric fruits
01113	Cultivation of coffee
01114	Cultivation of cocoa beans
01115	Cultivation of grapes
01116	Cultivation of bananas
01117	Cultivation of other permanent farming products
01118	Poorly specified farm crops
01119	Non-specified crop

Notes. The table shows the mapping of Census occupation codes to agriculture. The occupation codes are based on the 2002 Brazilian Classification of Occupations (CBO).

TABLE A.2
Summary statistics

Characteristic	N = 60,310,390
Income	Mean: 19,334 (Std.dev.: 59,143)
Age	Mean: 40 (Std.dev.: 10)
Crop worker	4,962,668 (8.2%)
Gender	
Female	26,113,936 (43%)
Male	34,196,453 (57%)
Marital status	
Married/in union	42,416,733 (70%)
Separated/divorced/spouse absent	8,376,029 (14%)
Single/never married	8,477,258 (14%)
Widowed	1,040,370 (1.7%)
Race	
Asian	654,227 (1.1%)
Black	5,058,423 (8.4%)
Brown (Brazil)	23,700,411 (39%)
Indigenous	160,917 (0.3%)
White	30,736,412 (51%)
Educational attainment	
Less than primary completed	17,892,207 (30%)
Primary completed	15,444,009 (26%)
Secondary completed	17,842,030 (30%)
University completed	9,132,143 (15%)

Notes. The table shows summary statistics for the sample of Brazilian-born individuals aged 25-64 with positive income and non-missing information on occupation, place of residence in 2005, race, education, marital status, and gender. The sample consists of approximately 6.2 million microdata observations, which correspond to about 60.3 million individuals in the weighted population.

TABLE A.3
Ex-post characteristics of migrants vs. non-migrants

Migrated	0 N = 57,634,344 ¹	1 N = 2,676,046 ¹
Age	Mean: 40 (Std.dev.: 10)	Mean: 36 (Std.dev.: 9)
Income	Mean: 19,046 (Std.dev.: 57,856)	Mean: 25,537 (Std.dev.: 81,863)
Crop worker	4,804,920 (8.3%)	157,749 (5.9%)
Gender		
Female	25,090,645 (44%)	1,023,292 (38%)
Male	32,543,699 (56%)	1,652,754 (62%)
Marital status		
Married/in union	40,573,369 (70%)	1,843,363 (69%)
Separated/divorced/spouse absent	7,941,995 (14%)	434,034 (16%)
Single/never married	8,104,096 (14%)	373,162 (14%)
Widowed	1,014,883 (1.8%)	25,487 (1.0%)
Race		
Asian	621,496 (1.1%)	32,731 (1.2%)
Black	4,857,100 (8.4%)	201,322 (7.5%)
Brown (Brazil)	22,664,755 (39%)	1,035,656 (39%)
Indigenous	153,362 (0.3%)	7,555 (0.3%)
White	29,337,631 (51%)	1,398,782 (52%)
Educational attainment		
Less than primary completed	17,248,562 (30%)	643,645 (24%)
Primary completed	14,770,428 (26%)	673,581 (25%)
Secondary completed	17,058,590 (30%)	783,440 (29%)
University completed	8,556,763 (15%)	575,380 (22%)

Notes. The table shows summary statistics for migrants (1) and non-migrants (0) between 2005 and 2010. Importantly, the characteristics are measured ex-post, i.e., in 2010.


Figure A.1 shows the employment share in crop cultivation (left panel) and the average agricultural annual income (right panel) by mesoregion in 2010. The employment share in crop cultivation is highest in the North and Northeast regions, while the average agricultural wage is highest in the South and Southeast regions. I calculate annual agricultural income by multiplying the reported total monthly income of agricultural workers by 12.

FIGURE A.1

Employment share in crop cultivation and agricultural income by mesoregion, 2010

(A) Employment share in crop cultivation

(B) Annual agricultural income by mesoregion



Notes. Panel (a) shows the employment share in crop cultivation by mesoregion in 2010, computed as the number of individuals working in crop cultivation divided by the total number of employed individuals in each mesoregion. Panel (b) shows the average agricultural wage by mesoregion in 2010.


Figure A.2 shows the average annual income (left panel) and the share of respondents who report to have lived in the same mesoregion in 2005 and 2010 (right panel) by mesoregion. The average annual income is highest in the South and Southeast regions, while the emigration rate is highest in the Central regions of Brazil.

FIGURE A.2

Average annual income and emigration rate (2005-2010) by mesoregion

(A) Average annual income by mesoregion

(B) Share of non-movers by mesoregion



Notes. Panel (a) shows the average annual income by mesoregion in 2010. Panel (b) shows the share of respondents who report to have lived in the same mesoregion in 2005 and 2010.

A.2. Agricultural production data

The primary data source for agricultural production is the *Produção Agrícola Municipal* (PAM), a survey conducted annually by the Brazilian Institute of Geography and Statistics (IBGE). The PAM provides detailed information on the area planted or destined for harvest, harvested

area, physical production quantities, and the nominal value of production for a wide range of temporary and permanent crops at the municipality level.

For the empirical analysis, I use PAM data for the years 1994-2010. I restrict attention to crops with strictly positive reported cultivated area in 2010 and exclude crops that display implausible jumps in yield between 2000 and 2001, ensuring a consistent set of crops across all years of the analysis. The total set of crops is: Soybean (grain), Sugarcane, Maize (grain), Coffee (bean) Total, Paddy rice, Cassava, Beans (grain), Herbaceous cotton (seed cotton), Tobacco (leaf), Tomato, Potato, Grape, Pineapple, Cocoa (bean), Papaya, Onion, Coconut, Rubber (coagulated latex), Garlic, Sorghum (grain), Black pepper, Melon, Sweet potato, Peanut (in shell), Yerba mate (green leaf), Oil palm (fruit bunch), Cashew nut, Sisal or agave (fiber), Heart of palm, Oat (grain), Barley (grain), Castor bean, Annatto (seed), Fava bean, Guarana (seed), Walnut (nut), Mallow (fiber), Tea (green leaf), Pea (grain), Flaxseed, Rye (grain), Ramie (fiber), Jute (fiber), Tung nut.

Table A.4 shows the major crops in Brazil by area planted, quantity produced, and value of production for the period 1994-2014. By area planted, the most important crops are Soybean (33.8%), Maize (24.2%), and Sugarcane (12.0%). By value of production, the most important crops are Soybean (28.4%), Sugarcane (18.5%), and Maize (12.9%).

While Soybeans and Maize are primarily used for either human consumption or animal feed, Sugarcane is predominantly used for biofuel production (ethanol) and sugar.

TABLE A.4
Major crops in Brazil, 1994-2014

Crop	Area Planted (%)	Quantity (%)	Value of Production (%)
Soybean (grain)	33.8	7.6	28.4
Sugarcane	12.0	75.4	18.5
Maize (grain)	24.2	7.2	12.9
Coffee (bean) Total	3.9	0.4	8.0
Paddy rice	5.8	1.7	4.9
Cassava	3.2	3.6	4.9
Beans (grain)	7.5	0.5	3.7
Herbaceous cotton (seed cotton)	1.7	0.4	3.6
Tobacco (leaf)	0.7	0.1	3.0
Tomato	0.1	0.5	2.1
Potato	0.3	0.5	1.9
Grape	0.1	0.2	1.3
Cocoa (bean)	1.2	0.0	0.8
Sorghum (grain)	1.1	0.2	0.3
Cashew nut	1.2	0.0	0.1


Notes. The table shows the major crops in Brazil by area planted, quantity produced, and value of production for the period 1994-2014. The percentages are calculated as the share of each crop in the total area planted, total quantity produced, and total value of production over the whole period, respectively.

Since the empirical analysis is conducted at the mesoregion level, I aggregate the municipality-level PAM data using official IBGE correspondence tables. I merge the resulting dataset with land-area information (from mesoregion shapefiles by the IBGE) to compute the share of total land devoted to agriculture in each mesoregion-year.

To obtain a measure of constant-price agricultural output per hectare, I revalue crop-level production using reference prices. Specifically, I construct crop-specific reference prices as the ratio of the total value to the total quantity produced over the full sample. Each crop's production is then revalued at these constant prices and aggregated across crops to obtain total output at the mesoregion level. Dividing by the area planted yields an index of agricultural land productivity at constant prices, which forms the basis of the productivity analysis presented in the main text. [Figure A.3](#) shows the share of land used in agriculture, measured agricultural productivity, and fundamental agricultural productivity by mesoregion in 2010. Fundamental agricultural productivity is defined as constant-price output per hectare of land adjusted for land use intensity with $\theta = 2$, following the model structure. In the "raw" measured agricultural productivity, land use intensity is not accounted for, the some of the highest agricultural productivity levels are observed in urban areas, including Metropolitana do Rio de Janeiro. This is driven by the fact, that these use a small share of their land for agriculture, but the land that is used is highly productive. Once land use intensity is accounted for, fundamental agricultural productivity is highest in the South and Southeast regions, which aligns much more closely with observed agricultural wages in these regions. This validates that the model-based adjustment for land use intensity is not merely a theoretical construct, but has empirical relevance.

FIGURE A.3

Land use, measured productivity, and fundamental productivity by mesoregion, 2010



Notes. Panel (a) shows the share of land used in agriculture by mesoregion in 2010, computed as the area planted for all crops divided by the total land area in each mesoregion. Panel (b) shows measured agricultural productivity, defined as constant-price output per hectare of land. Panel (c) shows fundamental agricultural productivity, defined as constant-price output per hectare of land adjusted for land use intensity with $\theta = 2$.

B. Empirical evidence

B.1. Construction of Predicted Counterfactual Incomes

The motivating evidence in [Section 3](#) is based on the 2010 Brazilian Census, which reports a respondents' current residence (in 2010) and residence in 2005. The sample consists of Brazilian-born individuals aged 25-64 with positive income and non-missing information on occupation, place of residence in 2005, race, education, marital status, and gender, yielding roughly 6.2 million observations. Using the data, I estimate a Mincer regression, for each mesoregion j , of the form

$$\log \text{Income}_{ijt} = \mathbf{X}_{ijt}\beta_j + \varepsilon_{ijt}, \quad (19)$$

where i indexes individuals, $t = 2010$, and $\mathbf{X}_{ijt}$ includes age, age squared, gender, race, marital status, and education.

Using the estimated coefficients $\hat{\beta}_j$ from each mesoregion j , I predict counterfactual incomes for all individuals observed in 2010 as if they had remained in their 2005 mesoregion of residence. For the sample of individuals who did not move, I compare predicted and actual incomes to assess the fit of the Mincer regression. The results are shown in Table B.5, which indicates that the Mincer regression explains a substantial portion of income variation, with an R^2 of 0.43 and a slope coefficient of 0.996 on predicted income.

TABLE B.5
Predicted vs. actual income for non-movers

	Log Actual Income (1)
Constant	6.45×10^{-10} (1×10^{-6})
Predicted Log Income	1.000*** (1×10^{-6})
Observations	5,976,627
R^2	0.43307

Notes. The table shows the results from regressing actual log income on predicted log income for individuals who did not move between 2005 and 2010. The dependent variable is the log of actual income in 2010, and the independent variable is the log of predicted income in 2010 based on the Mincer regression estimated in each mesoregion.

Using the estimated $\hat{\beta}_j$, I then predict 2010 log incomes for individuals observed in 2005 as if they had stayed in their origin mesoregion. I interpret the fitted values from the Mincer regression as counterfactual incomes that individuals would have earned, in expectation, had they remained in their origin mesoregion. Predicted values are exponentiated and grouped into ten 5,000 R\$ bins up to 50,000 R\$.

B.2. Climate change and agricultural productivity

To establish the link between agricultural productivity and wages, I regress log agricultural wages on log fundamental agricultural productivity in Brazilian mesoregions in 2010. The underlying data is shown in Figure II (panel b) and Figure A.3 (panel c). The regression results are shown in Table B.6, which indicates a strong positive relationship between agricultural productivity and wages, with an elasticity of approximately 0.215 (without fixed effects) and 0.128 with state fixed effects.

TABLE B.6
Caption

	Log of mean agricultural income	
	(1)	(2)
Log fundamental productivity	0.2147*** (0.0226)	0.1276*** (0.0192)
Constant	7.314*** (0.1388)	
State fixed effects		✓
Observations	137	136
R ²	0.40001	0.88288
Within R ²		0.28730

Notes. The table shows the results from regressing log agricultural wages on log fundamental agricultural productivity in Brazilian mesoregions in 2010. The dependent variable is the log of the average agricultural wage in each mesoregion. The independent variable is the log of fundamental agricultural productivity, defined as constant-price output per hectare of land adjusted for land use intensity with $\theta = 2$.

C. Derivations

C.1. Derivations for the static block

This part describes how the static problem is set up and how the endogenous objects depend on the allocation of labor in space.

C.1. Consumption

Households get utility from consuming a composite agricultural good and a non-agricultural good:

$$C_t^i = \frac{(C_t^{i;a})^\gamma (C_t^{i;n})^{1-\gamma}}{\gamma^\gamma (1-\gamma)^{1-\gamma}},$$

where $C_t^{i;a}$ denotes consumption of the composite agricultural good and $C_t^{i;n}$ denotes consumption of the non-agricultural good. γ is the constant expenditure share on the agricultural good (Cobb-Douglas).

The composite agricultural good is a constant elasticity of substitution (CES) aggregate of regional varieties:

$$C_t^{i;a} = \left(\sum_{j \in \mathcal{S}} (\psi^j)^{1/\epsilon} (c_t^{j;i;a})^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}},$$

where $c_t^{j;i;a}$ is the consumption of the regional variety of the agricultural good from origin j . ϵ is the elasticity of substitution between different varieties of the agricultural good (Armington elasticity). ψ^j is a preference (or quality) parameter for the agricultural good from origin j that

is constant over time and subject to the normalization $\sum_{j \in \mathcal{S}} \psi^j = 1$.¹ As in [Costinot et al. \(2016\)](#), the non-agricultural good is freely traded, and it is used as the numeraire for the remainder of the paper.

The final consumption bundle satisfies the budget constraint:

$$P_t^{i;a} C_t^{i;a} + p_t^0 C_t^{i;n} \leq y_t^i,$$

where $P_t^{i;a}$ is the CES price index and p_t^0 is the price of the non-agricultural good at time t (normalized to one in the initial period).^{2,3} The regional price index for final consumption is:

$$P_t^i = \left(P_t^{i;a} \right)^\gamma \left(p_t^0 \right)^{1-\gamma}.$$

C.2. Production technology

The economy consists of two sectors: agriculture and non-agriculture. The agricultural sector utilizes land and labor, whereas the non-agricultural sector employs labor only. Each region i is endowed with X^i units of land and L_t^i units of labor. Labor is perfectly mobile across sectors within regions, but not across regions.

For the agricultural production technology, I closely follow [Costinot et al. \(2016\)](#); [Gollin and Wolfersberger \(2024\)](#); [Sotelo \(2020\)](#) and assume that land consists of a continuum of heterogeneous plots ω in each region i that differ in their productivity and fixed costs to operate.⁴ Agricultural firms make profits that are spent locally on consumption. The costs incurred for operating a plot (e.g., slashing vegetation or clearing forest) are likewise spent locally on consumption.

The non-agricultural sector consists of a continuum of jobs η that differ in productivity and wages. Each job in region i produces a homogeneous good using labor, with productivity varying across jobs and regions.⁵ The sector is perfectly competitive.

Agriculture Production technology in agriculture is Cobb-Douglas, using land and labor:

$$q_t^i(\omega) = \left[L_t^i(\omega) \right]^\alpha \left[z_t^i(\omega) X_t^i(\omega) \right]^{1-\alpha},$$

¹Because I lack intra-national trade flow data, I assume ψ^j to be origin-specific, rather than origin-destination specific (as in [Costinot et al., 2016](#)). I set this parameter to clear the agricultural goods market in the calibration. In a planned extension, where $\mathcal{S}$ includes foreign countries, I allow this parameter to vary between country pairs to match country-to-country trade flows.

²The CES price index is given by

$$P_t^{i;a} = \left(\sum_{j \in \mathcal{S}} \left(\psi^j \right)^{1/\epsilon} \left(p_t^{j;i;a} \right)^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}},$$

where $p_t^{j;i;a}$ is the price of the agricultural variety produced in j in destination market i .

³In the counterfactual scenarios, p_t^0 adjusts to guarantee market clearing and must not be one.

⁴This fixed cost is needed to reconcile the fact that not all land is used for the cultivation of crops. If land were costless, agricultural producers would utilize all land for the cultivation of crops. Incomplete specialization of labor is guaranteed by the outside option for labor: working in non-agriculture.

⁵This structure allows me to capture the stark income inequality within the non-agricultural sector across Brazilian regions, unlike a representative-agent framework where a single wage would prevail.

where $q_t^i(\omega)$ is the output of plot ω in region i at time t , $L_t^i(\omega)$ is the amount of labor used on plot ω in region i at time t , and $X_t^i(\omega)$ is the amount of land used on plot ω in region i at time t . The parameter α represents the share of labor in production. $z_t^i(\omega)$ is the plot's productivity.

As in [Gollin and Wolfersberger \(2024\)](#), a representative farmer must pay a monetary cost $z^{i0}(\omega)$ expressed in units of the non-agricultural good to cultivate a plot of land. This cost needs to be paid every period a plot is used.

Productivity $z_t^i(\omega)$ and opening costs $z^{i0}(\omega)$ are randomly distributed across plots ω in region i and drawn from a nested Fréchet distribution with parameters $\{A_t^i, F^{i0}, \theta\}$. Formally, the distribution of land quality $z_t^i(\omega)$ and investment in land $z^{i0}(\omega)$ is given by:

$$\Pr \left(z_t^i(\omega) \leq z_t^i, z^{i0}(\omega) \leq z^{i0} \right) = \exp \left(-\zeta \left(\left(z_t^i / A_t^i \right)^{-\theta} + \left(z^{i0} / F^{i0} \right)^{-\theta} \right) \right),$$

where $\theta > 1$ is the shape parameter of the distribution. ζ is a normalization constant that is set such that $A_t^i = E[z_t^i(\omega)]$ and $F^{i0} = E[z^{i0}(\omega)]$. A_t^i and F^{i0} are the *unconditional* average land quality and initial investment in land across all plots in region i at time t .⁶

Non-agriculture The non-agricultural production function for a job η in region i is given by:

$$q_t^{i;n}(\eta) = A_t^{i;n}(\eta) L_t^{i;n}(\eta),$$

where $q_t^{i;n}(\eta)$ is the output of job η in region i at time t , $A_t^{i;n}(\eta)$ is the productivity of job η in region i , and $L_t^{i;n}(\eta)$ is the amount of labor used in job η in region i at time t .

Productivity $A_t^{i;n}(\eta)$ is drawn from a log-normal distribution with parameters μ^i and σ^i , which captures the heterogeneity in job productivity across potential jobs in non-agriculture in region i :

$$\ln A_t^{i;n}(\eta) \sim \mathcal{N}(\mu^i, (\sigma^i)^2).$$

C.3. Temporary equilibrium

Following [Caliendo et al. \(2019\)](#), I refer to the equilibrium of the static model as a *temporary equilibrium*. In this temporary equilibrium, the allocation of labor across regions is fixed, and the model determines the equilibrium prices and wages in each region such that all markets clear.

Labor income & allocation Households are randomly matched to jobs in the non-agricultural sector. Each household receives a job offer associated with income $A_t^{i;n}(\eta)$, where η is a random draw from the distribution of job offers in region i . The household then compares this offer to the agricultural wage $w_t^{i;a}$, which serves as a reservation wage. If the job offer is below the agricultural wage, the household chooses to work in agriculture; otherwise, they accept the job offer in the non-agricultural sector.

⁶Formally, ζ is given by: $\zeta = \Gamma \left(1 + \frac{1}{\theta} \right)^{-\theta}$, where $\Gamma(\cdot)$ is the gamma function.

The fraction working in agriculture is

$$\lambda_t^i = \Pr(A^{i,n}(\eta) < w_t^{i;a}) = \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right),$$

where $\Phi(\cdot)$ is the standard normal cumulative density function (CDF). Agricultural employment is $L_t^{i;a} = \lambda_t^i L_t^i$.

The resulting income distribution, $f_t^i(y)$, is a truncated log-normal distribution with a point mass at the agricultural wage $w_t^{i;a}$:

$$f_t^i(y) = \lambda_t^i \delta(y - w_t^{i;a}) + (1 - \lambda_t^i) \cdot \frac{1}{y \sigma^i \sqrt{2\pi}} \exp\left(-\frac{(\ln y - \mu^i)^2}{2(\sigma^i)^2}\right) \cdot \frac{\mathbb{I}(y \geq w_t^{i;a})}{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right)}, \quad (20)$$

where $\delta(\cdot)$ is the Dirac delta function and $\mathbb{I}(\cdot)$ is the indicator function. The income distribution is time-varying and depends on the agricultural wage $w_t^{i;a}$.

The average non-agricultural income among those who accept a non-agricultural productivity draw is⁷

$$w_t^{i;n} \equiv E[y \mid y \geq w_t^{i;a}] = \exp\left(\mu^i + \frac{1}{2}(\sigma^i)^2\right) \frac{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i - (\sigma^i)^2}{\sigma^i}\right)}{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right)}.$$

Agriculture In each region, agricultural firms choose the amount of land to use in agriculture $X_t^i(\omega)$ and the amount of labor to use in agriculture $L_t^i(\omega)$ to maximize profits.

The firms must pay a monetary cost $z^{i0}(\omega)$ to use land in agriculture, which is expressed in units of the non-agricultural good. Agricultural firms are maximizing:

$$\max_{L_t^i(\omega), X_t^i(\omega)} \pi_t^i(\omega) = p_t^{i;a} q_t^i(\omega) - w_t^{i;a} L_t^i(\omega) - p_t^0 z^{i0}(\omega) X_t^i(\omega).$$

The first-order condition with respect to labor yields:

$$w_t^{i;a} = \alpha p_t^{i;a} \left[\frac{z_t^i(\omega) X_t^i(\omega)}{L_t^i(\omega)} \right]^{1-\alpha}.$$

Rearranging for labor and substituting back into the profit function allows to write:

$$\pi_t^i(\omega) = X_t^i(\omega) \left[z_t^i(\omega) \Omega_t^i - p_t^0 z^{i0}(\omega) \right], \quad (21)$$

with $\Omega_t^i \equiv \left(\alpha^\alpha (1 - \alpha)^{1-\alpha} p_t^{i;a} (w_t^{i;a})^{-\alpha} \right)^{\frac{1}{1-\alpha}}$.

This expression indicates that profit is linear with respect to the amount of land used.

⁷This setup has an appealing implication. When regional agricultural wages decline, non-agricultural wages tend to fall on average, but their variance increases. Hence, climate change, or a decline in agricultural incomes in general, leads to an increase in income inequality within a region. Both are driven by agricultural versus non-agricultural incomes, as well as by increasing inequality within non-agricultural incomes.

Hence, the optimal land use decision for each plot is a corner solution: either the firm sets $X_t^i(\omega) = 1$ (and operates the plot) or $X_t^i(\omega) = 0$ (and leaves the plot unused).

From eq. 21 (as in [Gollin and Wolfersberger, 2024](#)), the implicit net return to land is

$$r_t^i(\omega) = z_t^i(\omega)\Omega_t^i - p_t^0 z^{i0}(\omega).$$

This implicit return reflects the net revenue generated by the land after paying labor and monetary costs. It is: i) strictly positive for inframarginal plots (which are operated and yield profits), ii) zero for the marginal plot (which breaks even), iii) negative for unused plots (which are not worth operating at current prices).

Together with the distributional assumption on productivity, the fraction of land that is profitable to be cultivated, and hence, used in agriculture in region i at time t is

$$\begin{aligned}\zeta_t^i &\equiv \Pr \left(z_t^i(\omega)\Omega_t^i \geq \max \left\{ p_t^0 z^{i0}(\omega), z_t^i(\omega)\Omega_t^i \right\} \right) \\ &= \frac{(A_t^i \Omega_t^i)^\theta}{(A_t^i \Omega_t^i)^\theta + (F^{i0} p_t^0)^\theta}.\end{aligned}\tag{22}$$

This equation shows that the share of land used in agriculture in region i increases in agricultural profitability, $A_t^i \Omega_t^i$, and decreases in the monetary cost of making land available for agriculture, $F^{i0} p_t^0$.

Average productivity is endogenous and depends on the fraction of land used. In particular, farmers will first exploit plots with high productivity and low operational costs.⁸

$$E[z_t^i(\omega)|\omega \text{ optimal}] = A_t^i (\zeta_t^i)^{-1/\theta}, \quad E[z^{i0}(\omega)|\omega \text{ optimal}] = F^{i0} (1 - \zeta_t^i)^{1/\theta}, \tag{23}$$

where $E[z_t^i(\omega)|\omega \text{ optimal}]$ is the expected productivity of the plots that are cultivated, and $E[z^{i0}(\omega)|\omega \text{ optimal}]$ is the expected cost of cultivating a plot.

The output is thus:

$$Q_t^{i;a} = \left(L_t^{i;a} \right)^\alpha \left(A_t^i (\zeta_t^i)^{-1/\theta} X_t^{i;a} \right)^{1-\alpha}, \tag{24}$$

where $L_t^{i;a} = \lambda_t^i L_t^i$ and $X_t^{i;a} = \zeta_t^i X_t^i$.

The sum of fixed costs, B_t^i , and total profits of agricultural firms, R_t^i , are:

$$B_t^i = E[z^{i0}(\omega)|\omega \text{ optimal}] \cdot \zeta_t^i X_t^i = F^{i0} (1 - \zeta_t^i)^{1/\theta} \zeta_t^i X_t^i \tag{25}$$

$$R_t^i = p_t^i Q_t^{i;a} - \left(w_t^{i;a} L_t^{i;a} \right) - \left(F^{i0} (1 - \zeta_t^i)^{1/\theta} \zeta_t^i X_t^i \right). \tag{26}$$

⁸In general, the average productivity of agricultural land is larger than the unconditional average productivity of all plots in a region unless all land is cultivated ($\zeta_t^i = 1$). This also implies that the measured productivity described in ?? will not reflect A_t^i , but rather $A_t^i (\zeta_t^i)^{-1/\theta}$.

Wages are equal to the marginal product of labor on the marginal plot:

$$p_t^{i;a} \frac{\partial Q_t^{i;a}}{\partial L_t^{i;a}} = w_t^{i;a}$$

$$w_t^{i;a} = p_t^{i;a} \alpha \left(\frac{A_t^i (\zeta_t^i)^{-1/\theta} X_t^{i;a}}{L_t^{i;a}} \right)^{1-\alpha}. \quad (27)$$

Non-agriculture The non-agricultural production is implicitly determined by the agricultural wage. Only jobs that have a productivity $A^{i;n}(\eta)$ above the agricultural wage $w_t^{i;a}$ are filled.

Total output in the non-agricultural sector is then given by:

$$Q_t^{i;n} = \int_{A^{i;n}(\eta) \geq w_t^{i;a}} A^{i;n}(\eta) L_t^{i;n}(\eta) d\eta, \quad (28)$$

where $L_t^{i;n}(\eta)$ is the amount of labor used in job η in region i at time t . Without loss of generality, the amount of labor used in a job η can be normalized to one. The total amount of labor in the non-agricultural sector is given by: $L_t^{i;n} = (1 - \lambda_t^i) L_t^i$. Hence, eq. 28 reduces to:

$$Q_t^{i;n} = w_t^{i;n} (1 - \lambda_t^i) L_t^i, \quad (29)$$

which shows that the total output in the non-agricultural sector is equal to the average non-agricultural income $w_t^{i;n}$ times the amount of labor in the non-agricultural sector $(1 - \lambda_t^i) L_t^i$.

Market clearing The sum of the wage bill, profits, and land-clearing costs gives total income in a region i at time t :

$$Y_t^i = \underbrace{(y_t^i L_t^i)}_{\text{Wage bill}} + \underbrace{(R_t^i)}_{\text{Profits}} + \underbrace{(B_t^i)}_{\text{Land-clearing}}, \quad (30)$$

where y_t^i is the average income in region i at time t , which is given by:

$$y_t^i = \lambda_t^i w_t^{i;a} + (1 - \lambda_t^i) w_t^{i;n}.$$

What remains to be done is to derive the trade between regions. The Cobb-Douglas preferences over the composite agricultural and the non-agricultural good imply expenditure shares of γ and $1 - \gamma$, respectively. Within agricultural expenditures, the expenditure share on variety j is:

$$\pi_t^{j,i;a} = \frac{\psi^j \left(p_t^{j;a} \kappa^{j,i} \right)^{1-\epsilon}}{\sum_{k \in \mathcal{S}} \psi^k \left(p_t^{k;a} \kappa^{k,i} \right)^{1-\epsilon}}, \quad (31)$$

where $\kappa^{j,i}$ are the iceberg trade costs between j and i .

Let $S_t^{j;a} \equiv p_t^{j;a} c_t^{j;a}$ be the value of agricultural production of region j at time t .⁹ Then, market clearing requires:

$$S_t^{j;a} = \gamma \sum_{i \in \mathcal{S}} \pi_t^{j,i;a} Y_t^i. \quad (32)$$

Non-agricultural market clearing is achieved when the economy-wide value of non-agricultural production is equal to the demand for non-agriculture:

$$\sum_{j \in \mathcal{S}} S_t^{j;n} = (1 - \gamma) \sum_{i \in \mathcal{S}} Y_t^i. \quad (33)$$

C.2. Exact-hat derivations

This section presents the hat algebra system used to solve for counterfactual equilibria in the spatial model. By expressing equilibrium conditions in proportional changes relative to the baseline, the system captures how wages, land use, sectoral labor allocation, trade shares, and market clearing change in response to changes in the agricultural productivity $\hat{A}_t^i$ and changes in the allocation of labor across regions $\hat{L}_t^i$. Since the dynamic model requires real wage levels, these are recovered ex post from baseline values and proportional changes.

- Agricultural wages:

$$\begin{aligned} w_t^{i;a} &= p_t^{i;a} \alpha \left(\frac{A_t^i (\zeta_t^i)^{-1/\theta} X_t^{i;a}}{L_t^{i;a}} \right)^{1-\alpha} \\ &= \alpha \left(X_t^i \right)^{1-\alpha} p_t^{i;a} \left(A_t^i \right)^{1-\alpha} \left(\zeta_t^i \right)^{\frac{\theta-1}{\theta}(1-\alpha)} \left(\lambda_t^i \right)^{\alpha-1} \left(L_t^i \right)^{\alpha-1} \end{aligned}$$

where α and X_t^i are constant.

In changes:

$$\hat{w}_t^{i;a} = \left(\hat{p}_t^{i;a} \right) \left(\hat{A}_t^i \right)^{1-\alpha} \left(\hat{\zeta}_t^i \right)^{\frac{(\theta-1)(1-\alpha)}{\theta}} \left(\hat{\lambda}_t^i \right)^{\alpha-1} \left(\hat{L}_t^i \right)^{\alpha-1}$$

- Land use:

$$\zeta_t^i = \frac{(A_t^i \Omega_t^i)^\theta}{(A_t^i \Omega_t^i)^\theta + (F_t^{i0} p_t^{i0})^\theta}$$

where F_t^{i0} is a constant.

In changes:

$$\hat{\zeta}_t^i = \frac{(\hat{A}_t^i \hat{\Omega}_t^i)^\theta}{\zeta_t^i (\hat{A}_t^i \hat{\Omega}_t^i)^\theta + (1 - \zeta_t^i) (\hat{p}_t^{i0})^\theta}$$

⁹I use S for “sales” to avoid confusion with land X .

- Net profitability:

$$\Omega_t^i = \left(\tilde{\alpha} \frac{p_t^{i;a}}{(w_t^{i;a})^\alpha} \right)^{\frac{1}{1-\alpha}}$$

In changes:

$$\hat{\Omega}_t^i = \left(\hat{p}_t^{i;a} (\hat{w}_t^{i;a})^{-\alpha} \right)^{\frac{1}{1-\alpha}}$$

- Sectoral labor allocation:

$$\lambda_t^i = \Phi \left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i} \right)$$

in changes:

$$\hat{\lambda}_t^i = \frac{\Phi \left(\frac{\ln(w_t^{i;a} \hat{w}_t^{i;a}) - \mu^i}{\sigma^i} \right)}{\Phi \left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i} \right)}$$

- Trade shares:

$$\pi_t^{j,i;a} = \frac{\psi^j \left(p_t^{j;a} \kappa^{j,i} \right)^{1-\epsilon}}{\sum_k \psi^k \left(p_t^{k;a} \kappa^{k,i} \right)^{1-\epsilon}}$$

where κ is constant.

In changes:

$$\hat{\pi}_t^{j,i;a} = \frac{\left(\hat{p}_t^{j;a} \right)^{1-\epsilon}}{\sum_k \pi_t^{k,i;a} \left(\hat{p}_t^{k;a} \right)^{1-\epsilon}}$$

- Agricultural market clearing:

$$X_t^{j;a} = \gamma \sum_i \pi_t^{j,i;a} Y_t^i$$

In changes:

$$\hat{X}_t^{j;a} = \frac{\sum_i \hat{\pi}_t^{j,i;a} \hat{Y}_t^i \pi_t^{j,i;a} Y_t^i}{\sum_i \pi_t^{j,i;a} Y_t^i}$$

- Labor market clearing (non-agricultural wage):

$$w_t^{i;n} = \exp\left(\mu^i + \frac{1}{2}(\sigma^i)^2\right) \frac{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i - (\sigma^i)^2}{\sigma^i}\right)}{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right)}$$

In changes:

$$\hat{w}_t^{i;n} = \frac{1 - \Phi\left(\frac{\ln w_t^{i;a} + \ln \hat{w}_t^{i;a} - \mu^i - (\sigma^i)^2}{\sigma^i}\right)}{1 - \Phi\left(\frac{\ln w_t^{i;a} + \ln \hat{w}_t^{i;a} - \mu^i}{\sigma^i}\right)} \frac{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i}{\sigma^i}\right)}{1 - \Phi\left(\frac{\ln w_t^{i;a} - \mu^i - (\sigma^i)^2}{\sigma^i}\right)}$$

- Total income:

$$Y_t^i = w_t^{i;a} L_t^{i;a} + w_t^{i;n} L_t^{i;n} + R_t^i + B_t^i$$

In changes:

$$\hat{Y}_t^i = \frac{\hat{w}_t^{i;a} w_t^{i;a} \hat{\lambda}_t^i \lambda_t^i \hat{L}_t^i L_t^i + \hat{w}_t^{i;n} w_t^{i;n} (1 - \hat{\lambda}_t^i \lambda_t^i) \hat{L}_t^i L_t^i + \hat{R}_t^i R_t^i + \hat{B}_t^i B_t^i}{Y_t^i}$$

- Quantity produced in agriculture:

$$Q_t^{i;a} = \left(\lambda_t^i L_t^i\right)^\alpha \left(A_t^i \left(\zeta_t^i\right)^{\frac{\theta-1}{\theta}} X_t^i\right)^{1-\alpha}$$

where X_t^i is constant.

In changes:

$$\hat{Q}_t^{i;a} = \left(\hat{\lambda}_t^i \hat{L}_t^i\right)^\alpha \left(\hat{A}_t^i \left(\hat{\zeta}_t^i\right)^{\frac{\theta-1}{\theta}}\right)^{1-\alpha}$$

- Agricultural production value (in changes):

$$\hat{p}_t^{i;a} \hat{Q}_t^{i;a}$$

which must, in the counterfactual equilibrium, be equal to $\hat{X}_t^{i;a}$ for markets to clear. I.e., we have the counterfactual agricultural goods market-clearing condition:

$$\begin{aligned} \hat{p}_t^{j;a} &= \frac{\hat{X}_t^{j;a}}{\hat{Q}_t^{j;a}} \\ \hat{p}_t^{j;a} &= \frac{\frac{\sum_i \hat{\pi}_t^{i,j;a} \hat{Y}_t^i \pi_t^{i,j;a} Y_t^i}{\sum_i \pi_t^{i,j;a} Y_t^i}}{\hat{Q}_t^{j;a}} \end{aligned}$$

- Non-agricultural market clearing:

$$\hat{p}_t^{j0} = \frac{\hat{X}_t^{j;n}}{\hat{Q}_t^{j;n}}$$

- Profits:

$$R_t^i = p_t^{i;a} Q_t^{i;a} - w_t^{i;a} L_t^{i;a} - B_t^i$$

In changes:

$$\hat{R}_t^i = \frac{\hat{p}_t^{i;a} \hat{Q}_t^{i;a} p_t^{i;a} Q_t^{i;a} - \hat{w}_t^{i;a} \hat{\lambda}_t^i \hat{L}_t^i w_t^{i;a} \lambda_t^i L_t^i - \hat{B}_t^i B_t^i}{R_t^i}$$

- Income from fixed costs:

$$B_t^i = F_t^{i0} \left(1 - \zeta_t^i\right)^{1/\theta} \zeta_t^i X_t^i$$

where F_t^{i0} and X_t^i are constants.

In changes:

$$\hat{B}_t^i = \left(\frac{1 - \hat{\zeta}_t^i \zeta_t^i}{1 - \zeta_t^i} \right)^{1/\theta} \zeta_t^i$$

The algorithm to solve the general equilibrium of the static model nests two main loops. In the *outer loop*, I update the vector of agricultural producer-price changes $\{\hat{p}^{j;a}\}$.

Wrapped inside is an *inner loop* over regions. For each region i , I treat the current $\hat{p}^{i;a}$ as given and solve for the corresponding change in the agricultural wage $\hat{w}^{i;a}$ by finding the root of

$$\hat{w} - M^i(\hat{w}) = 0,$$

where

$$M^i(\hat{w}) = \hat{p}^{i;a} (\hat{A}^i)^{1-\alpha} [\hat{\zeta}^i(\hat{w})]^{\frac{(\theta-1)(1-\alpha)}{\theta}} [\hat{\lambda}^i(\hat{w})]^{\alpha-1} (\hat{L}^i)^{\alpha-1}.$$

Because land-use shares $\hat{\zeta}^i$ and labor-allocation shares $\hat{\lambda}^i$ both depend non-linearly on $\hat{w}^{i;a}$, I solve this one-dimensional fixed-point problem for each i .

Once each $\hat{w}^{i;a}$ is determined, I back out the non-agricultural wage change $\hat{w}^{i;n}$, recompute aggregate income changes $\hat{Y}^i$, and update trade shares $\hat{\pi}^{j,i;a}$. Feeding these into the market-clearing condition gives a new price guess

$$\hat{p}_{\text{new}}^{j;a} = \frac{\sum_i \hat{\pi}^{j,i;a} \hat{Y}^i \pi^{j,i;a} Y^i}{\sum_i \pi^{j,i;a} Y^i} \Big/ \hat{Q}^{j;a},$$

which I under-relax as

$$\hat{p}^{j;a} \leftarrow \rho \hat{p}_{\text{new}}^{j;a} + (1 - \rho) \hat{p}^{j;a},$$

to ensure stable convergence.

I repeat both loops until the maximum changes in $\hat{p}^{j;a}$ and $\hat{w}^{i;a}$ fall below the chosen tolerance, yielding the full counterfactual equilibrium in prices, wages, quantities, and trade flows.

D. Calibration details

D.1. Baseline economic environment and calibration results

This section presents additional results from the static calibration. By construction, the model exactly matches key moments in the baseline year 2010, as shown in [Table D.7](#). In particular, by model inversion, the model exactly matches in each region: i) total labor, ii) average wages, iii) the employment share in agriculture, iv) the average agricultural wage, v) the non-agricultural average wages, vi) total land endowments, and vii) the fraction of land used in agriculture.


TABLE D.7
Exactly matched moments in 2010

Parameter	Mean	Min	Max
<i>Exactly matched</i>			
<i>Labor:</i>			
L^j	440222	11612	7160612
w^j	15576.17	7444.09	39283.89
λ^j	0.1243	0.0030	0.3193
$w^{j;a}$	5861.50	2813.27	10649.65
$w^{j;n}$	16770.86	8800.74	39554.08
<i>Land:</i>			
X^j	6558425	283872	50469934
ζ^j	0.1321	0.0001	0.6959

Notes: The table shows key moments that the static model exactly matches in the baseline year 2010. I report the unweighted mean across regions, as well as the minimum and maximum values.

In the model calibration, the Cobb-Douglas parameter α is set to match the agricultural wage bill-to-output-ratio. I.e., $\alpha = \frac{\sum_i w^{i;a} L^{i;a}}{\sum_i V_t^{i;a}}$, where $V_t^{i;a}$ is the value of agricultural production in region i reported by the PAM dataset. I calibrate a value of $\alpha = 0.1579$ to match the data. However, because of this it is not strictly necessary that the model also matches the agricultural wage bill-to-output ratio in each region. To match the agricultural wage bill-to-output ratio in each region, I would need to introduce additional heterogeneity in the model, for example by allowing for region-specific α_i . [Figure D.4](#) shows that, under the assumption of one α at the national level, the model fits the regional agricultural value of production reasonably well, with an R^2 of 0.5.

FIGURE D.4
Model fit: Value of production




Notes: The figure shows data vs. model fit for the log value of agricultural production in 2010. The R^2 of the fit is 0.5.

D.2. Counterfactual productivity paths

Figure D.5 shows the counterfactual agricultural productivity paths for each region. The paths are constructed by combining the estimated climate impacts on agricultural productivity from Section 4 with the RCP 8.5 temperature projections from the Copernicus Climate Change Service (C3S) CMIP6 archive. I use monthly near-surface air temperature projections from the CMCC-ESM2 Earth System Model.

FIGURE D.5
Counterfactual agricultural productivity paths




Notes: The figure shows the counterfactual agricultural productivity paths for each region. The paths are constructed by combining the estimated climate impacts on agricultural productivity from ?? with the RCP 8.5 temperature projections from the Copernicus Climate Change Service (C3S) CMIP6 archive. I use monthly near-surface air temperature projections from the CMCC-ESM2 Earth System Model.

D.3. Migration costs

Figure D.6 reports model-implied migration probabilities by income, averaged across regions using national population weights. The model successfully reproduces the qualitative pattern that mobility rises with income, but it overstates the gradient relative to the empirical evidence. Low-income households migrate slightly less than observed, while high-income households migrate too frequently. This discrepancy likely reflects a mechanical feature of the extreme-value preference structure adopted from the spatial-equilibrium literature. In the logit formulation, workers draw idiosyncratic location preferences whose variance is governed by the dispersion parameter ν . At high incomes—where expected utility differences across locations are small and pecuniary migration costs are negligible—these random preference draws dominate, generating excess migration in the upper tail. The structure preserves tractability and closed-form choice probabilities, but it likely inflates high-income mobility relative to reality.

FIGURE D.6
Migration probabilities by income (model)



Notes: The figure shows the model-implied average emigration rate by income.

E. Additional quantitative results

Destination patterns. Beyond increasing overall mobility, the subsidy also changes *where* households migrate. Table E.8 decomposes average migration probabilities into moves toward rural and urban destinations, weighted by population and averaged over time. Under the baseline climate scenario, migration flows of low-income households are primarily directed toward rural regions, reflecting both affordability constraints and proximity to agricultural employment opportunities. Once the subsidy is introduced, these patterns shift markedly: low-income households' migration to rural destinations more than doubles (from 3.5% to 9.2%), indicating that the removal of liquidity constraints facilitates movement even within the rural economy. At the same time, their urban migration rate remains nearly

unchanged, suggesting that the policy does not disproportionately encourage rural–urban exodus. Middle-income households also exhibit modest increases in rural mobility, whereas high-income households—who were largely unconstrained—show little change or slightly lower overall migration. Together, these findings confirm that the subsidy primarily enhances *adaptive migration within and between rural areas*, rather than triggering large-scale rural–urban shifts.

TABLE E.8

Effect of the migration subsidy on rural and urban destination choices (climate path)

Income group	Migration to rural regions (%)	Migration to urban regions (%)
Low income ($\leq 10k$)	3.54 $\rightarrow$ 9.20	1.87 $\rightarrow$ 1.89
Middle income (10–20k)	2.07 $\rightarrow$ 2.53	4.32 $\rightarrow$ 2.88
High income ($> 20k$)	2.87 $\rightarrow$ 2.54	8.90 $\rightarrow$ 7.09

Notes: Each entry compares the baseline and subsidized regimes under the same climate path. “Migration to rural (urban) regions” reports the population- and time-weighted average probability of migrating to destinations classified as rural (urban) based on whether their agricultural employment share λ_a lies above or below the national median. Arrows ($\rightarrow$) indicate the change from the baseline to the subsidized regime.

F. Solving the model

To determine the evolution of labor allocations and spatial equilibrium in response to forward-looking migration decisions, I solve a dynamic problem where households choose their migration paths based on current fundamentals and expectations about the future. In the final period, agents make one last consumption decision and do not live beyond it. This simplifies the backward recursion and removes any artificial incentive to migrate based on infinite future gains. I further assume that productivity remains constant after the year 2100, while migration remains possible until 2200. This avoids migration patterns driven purely by end-of-horizon effects, an artifact common in finite-horizon models.

1. **Static equilibrium:** For each period $t \leq T$, and given a guess of labor allocations L_t^j , I solve for the static spatial equilibrium. This yields region-specific agricultural wages $w_t^{j,a}$, price indices P_t^j , non-agricultural wages, sectoral labor shares, income, and output. These prices and wages, together with the sectoral composition of employment, determine the income distribution in each location. The static equilibrium is solved using “hat-algebra”. The full procedure is described in [Online Appendix C](#).
2. **Income distributions and consumption utility:** In each region j , income is lognormally distributed with parameters μ^j, σ^j , but is truncated below by the agricultural wage $w_t^{j,a}$, which represents a fallback income level for marginal workers. Additionally, a fraction λ_t^j of households earn exactly the agricultural wage, giving rise to a point mass at $w_t^{j,a}$. The resulting income distribution follows from [eq. 20](#).

The expected utility from consumption is then given by:

$$\mathbb{E}_y \left[\log \left(\frac{y}{P_t^j} \right) \right] = \lambda_t^j \cdot \log \left(\frac{w_t^{j,a}}{P_t^j} \right) + (1 - \lambda_t^j) \cdot \int_{w_t^{j,a}}^{\infty} \log \left(\frac{y}{P_t^j} \right) \cdot \tilde{f}_t^j(y) dy$$

3. **Terminal condition:** In the final period T , agents make one final consumption decision and do not live beyond it. As a result, the value function in T is simply the expected utility from consumption in that period:

$$V_T^j = \mathbb{E}_y \left[\log \left(\frac{y}{P_T^j} \right) \right]$$

This removes the need to account for continuation values beyond the horizon, ensuring that migration incentives in earlier periods are not distorted by artificially inflated future utilities.¹⁰

4. **Backward induction:** Starting from the terminal condition, I solve the Bellman equation recursively for $t = T - 1, \dots, 0$. The value function consists of two terms:

$$V_t^j = \mathbb{E}_y \left[\log \left(\frac{y}{P_t^j} \right) + \nu \log \left(\sum_{i \in \mathcal{S}} \exp \left(\frac{\beta V_{t+1}^i - \bar{\tau}^{j,i|y}}{\nu} \right) \right) \right]$$

where $\bar{\tau}^{j,i|y} = \tau^{j,i|y} + \kappa^{j,i}$ captures the total disutility from migration.

5. **Migration probabilities:** Conditional on income y , the probability that a household in region j chooses destination i is determined using [eq. 3](#). The unconditional migration flow from region j to i aggregates over the full income distribution:

$$\mu_t^{j,i} = \lambda_t^j \cdot \mu_t^{j,i|w_t^{j,a}} + (1 - \lambda_t^j) \cdot \int_{w_t^{j,a}}^{\infty} \mu_t^{j,i|y} \cdot \tilde{f}_t^j(y) dy$$

These flows summarize the dynamic migration decisions made by forward-looking agents.

6. **Labor propagation:** The population evolves forward based on the endogenous migration flows:

$$L_{t+1}^j = \sum_{i \in \mathcal{S}} \mu_t^{i,j} \cdot L_t^i$$

This law of motion maps past labor allocations and migration decisions into future population distributions.

7. **Convergence check:** At each iteration, I compare the newly computed labor allocations $\{L_t^j\}$ to the previous guess. If the difference is small (e.g., below a pre-defined toler-

¹⁰Alternatively, if agents were assumed to remain in their final location and consume forever in a stationary environment, the terminal value would reflect the present discounted value of perpetual utility: $V_T^j = \frac{1}{1-\beta} \cdot \mathbb{E}_y \left[\log \left(\frac{y}{P_T^j} \right) \right]$. While this approach can capture long-run incentives, it tends to generate sharp spikes in value functions near the horizon and can distort migration flows in the penultimate period.

ance), the algorithm converges. Otherwise, the updated labor path is used to re-solve the sequence of spatial equilibria, and the process repeats.

This iterative algorithm ensures that both household migration decisions and local wages and prices are jointly consistent over time, resulting in a fully dynamic spatial equilibrium. Expectations are computed precisely over region-specific income distributions that include both a continuous upper tail and a discrete mass at the agricultural reservation wage. Each household's dynamic optimization problem is solved using backward induction, and the resulting migration flows guide the spatial redistribution of labor over time.

Additional references

- Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica* 87(3), 741–835.
- Costinot, A., D. Donaldson, and C. Smith (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy* 124(1), 205–248.
- Gollin, D. and J. Wolfersberger (2024). Agricultural trade and deforestation: The role of new roads.
- Sotelo, S. (2020). Domestic trade frictions and agriculture. *Journal of Political Economy* 128(7), 2690–2738.