The Curse of Connectivity:
Evidence from Village Resettlement in Indonesia

Harrison Mitchell

UC San Diego*

September 2026

Click here for the latest version.

Abstract

Does connecting a village to markets make it richer? Using Indonesia's Transmigration program, which created new villages with variation in initial connectivity, I find that villages that began more connected are poorer today. A one standard deviation increase in initial market access lowers GDP per capita decades later by roughly 4 percent. The lower incomes are not explained by lower local prices, compensating amenities, or relocation. Higher connectivity shifts village residents out of agriculture and into local non-farm work. Using separate variation in villages' initial settlement sizes, I estimate a scale externality: agricultural productivity rises with the size of the agricultural workforce. Workers leaving agriculture do not internalize the resulting productivity loss, so reallocation can reduce long-run incomes. A spatial general equilibrium model of Indonesia with the estimated externality reproduces the negative effect of market access on income without targeting it. In counterfactuals, ignoring the externality overstates aggregate gains from road-network improvements by 18 percent and from increases in rural agricultural productivity by 64 percent. An agricultural earnings subsidy at the road program's fiscal outlay delivers roughly twice its rural gains.

*Corresponding author: hamitchell@ucsd.edu

I am grateful to Sam Bazzi, Gabriel Cañedo Riedel, Holt Dwyer, Fabian Eckert, Juan Herreño, Nir Jaimovich, Gaurav Khanna, Jordan Rosenthal-Kay, and Vedant Vohra for their insightful comments and willingness to generously share data.

1 Introduction

Does access to markets make a place richer? A long tradition in economic geography argues that it does (Krugman, 1991; Redding and Venables, 2004). The World Bank alone maintains an active transport portfolio of roughly \$45 billion.[1] Whether greater market access raises long-run development, however, is theoretically ambiguous. Greater market access generates the standard gains from trade and feeds the agglomeration forces that drive growth. But it can also hinder development when it pulls labor away from activities that carry productivity growth (Matsuyama, 1992). In the literature on transport infrastructure, the worry has been that connected peripheral locations lose economic activity to the core (Faber, 2014; Baum-Snow et al., 2020). Which force dominates over the long run is ultimately an empirical question.

I use Indonesia's Transmigration program to show that villages that began with greater market access are poorer today than otherwise similar villages. Between 1952 and 1995, the program resettled over 2 million voluntary migrants into over 1,000 newly constructed villages, placed with little regard for the productive potential of their sites. The villages that started more connected moved more labor out of agriculture, and I trace their lower incomes to a scale externality in agriculture that this reallocation erodes.

Officials, under pressure to meet resettlement targets, sited villages wherever land was available, with little strategic regard for connectivity or the potential of a location. Aside from location, the villages started out nearly identical; settlers had no control over which new village they were placed in, received uniform 2-hectare plots by lottery upon arrival, and were each given the same initial endowment of agricultural inputs (tools, seeds, and fertilizer). Previous work finds no sorting of migrant characteristics across destination villages (Bazzi et al., 2016, 2019), and the migrants themselves were overwhelmingly asset-poor (Kebschull, 1986). The assumption I need is therefore not that placement was random, but that among villages settled in the same year, on the same island, with similar site characteristics, the remaining differences in initial market access are unrelated to the unobserved determinants of long-run income; Section 4 develops this argument and the balance tests behind it.

To measure initial connectivity, I hand-digitize the complete pre-program road network of Indonesia, down to individual trails, from 27 US Army Map Service sheets drawn around 1950. To measure development decades later, I use satellite-derived estimates of village GDP per capita. To validate that measure and trace mechanisms, I draw on population censuses, village and household surveys, and firm listings.

I estimate that moving from the 25th to the 75th percentile of initial market access lowers long-run village GDP per capita by about 5.5 percent, and that one standard deviation of initial access lowers it by about 4 percent. These differences are measured 17 to 67 years after settlement. The results are robust to alternative satellite-derived GDP estimates. Resident wages and a novel

[1] World Bank, "Transport Overview," accessed August 2026.

video-based measure of village material welfare in a small subsample point in the same direction, although the wage estimates are imprecise.

The evidence does not support many conventional explanations. More connected villages experience *more* reallocation of labor out of agriculture and have more firms, weighing against explanations based on agricultural specialization or net firm exit. I find no evidence that out-migration or selective migration explains the results. Food prices do not fall to offset the lower incomes; if anything, they are higher. Connectivity does not improve non-road public goods, so lower output is unlikely to compensate for better measured amenities.

Instead, I uncover evidence for a scale externality in agricultural production. The externality operates across farms rather than within them. Each farm produces under constant returns, but the productivity of all farms rises with the scale of the local agricultural sector, through channels such as knowledge spillovers across farmers, shared capital and equipment rental markets, and the cooperative management of irrigation (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Bassi et al., 2022; Caunedo et al., 2026; Ostrom and Gardner, 1993). A worker who leaves agriculture therefore lowers the productivity of those who remain, a loss no individual internalizes. Consistent with this mechanism, more connected villages use fewer tractors per agricultural worker, and because agriculture is the dominant sector in these village economies, the decline in agricultural productivity lowers village income.

I use a second source of program variation to measure the externality directly. Villages were settled with different initial populations, and I show that this variation is unrelated to the characteristics of the sites. Because initial settlement size predicts the size of a village's agricultural workforce decades later, comparing agricultural output across villages settled at different scales identifies how output responds to the size of the workforce. Output rises with the workforce by more than the private return to labor would predict, and the gap between the social and private returns is the externality I carry into the model.

To formalize the mechanism and quantify its implications, I develop a spatial general equilibrium model of the Indonesian economy. In the model, connectivity pulls workers out of agriculture because rural services must be delivered in person, so villages in more connected locations capture more of that demand and the private return to non-farm work rises with proximity. In a simplified version that can be solved in closed form, I show that proximity harms a village exactly when the externality is strong enough that the income lost from this reallocation outweighs the standard gains from trading with the city. The externality point estimates exceed this threshold.

In the full model, every location in Indonesia carries the same two-sector economy. The externality is based on the settlement-size estimates of Section 6, and trade costs and spending shares are disciplined by survey data. The transmigration villages enter with common fundamentals and differ only in where they sit on the map, so the relationship between market access and income among them is a prediction of the model rather than a target. The model reproduces about 40 percent of the negative effect in the full sample and 80 percent excluding the most isolated 5 percent of villages. Removing the externality reverses the prediction. An externality-free model can fit the observed incomes only by assuming connected villages had systematically worse agricultural fundamentals, which the observed site characteristics do not support.

I then use the model to demonstrate that the externality changes who benefits from connectivity. A road targeted to an isolated village can make that village poorer by drawing workers out of agriculture. Consistent with the simplified model, the loss in income is largest for villages with substantial employment in both sectors; almost entirely agricultural villages can still gain.

For aggregate shocks, I extend the model to incorporate the non-homothetic demand system of Boppart (2014) and Fan et al. (2023), estimating its parameters in household expenditure data. Income growth shifts spending away from food, pulling labor out of agriculture and lowering productivity through the externality. A 10 percent cut in travel times looks 18 percent more valuable in the externality-free economy, and a 5 percent rise in rural agricultural productivity looks 64 percent more valuable.

The bias runs the other way for policies that keep labor in agriculture. An agricultural earnings subsidy at the road program's fiscal outlay raises national income by 0.20 percent and delivers roughly twice the rural gains of roads, while the same subsidy in the externality-free economy is a small distortionary loss. These exercises show that agricultural scale economies change the ranking of policy instruments and the magnitude of aggregate welfare effects.

My paper relates most directly to the literature on the effects of market access on economic outcomes. The central finding of this work, in settings ranging from the American railroad expansion to the postwar division of Germany, is that greater market access raises local economic performance (Donaldson and Hornbeck, 2016; Hornbeck and Rotemberg, 2024; Redding and Sturm, 2008). My contribution is to document a mechanism which may reverse this relationship. In my context, greater market access *lowers* long-run output per capita. The canonical results come from economies where the sectors that gain from connectivity, such as manufacturing, drive growth, while mine is a developing, agriculture-dependent economy where the most common alternative to agricultural work is low-scale services and retail. A large body of evidence finds these sectors to be persistently small, informal and unproductive (La Porta and Shleifer, 2014; Lagakos, 2016). My paper also measures the welfare effects of connectivity more directly. Much of this literature treats population as a sufficient statistic for the welfare effects of market access, but I find that population is a poor predictor of income in my context.

Closer to my setting, road programs raised trade and welfare in Brazil and Indonesia (Morten and Oliveira, 2024; Gertler et al., 2024), Ethiopian rural roads raised agricultural productivity only when paired with extension (Gebresilasse, 2023), and in India a new road moves village workers out of agriculture with no detectable change in consumption or agricultural output (Asher and Novosad, 2020), the release of farm labor into non-agricultural firms that Kaboski et al. (2024) also document for highways. My model supplies what these accounts lack, a mechanism for reallocation without income growth. The workers who leave agriculture lower the productivity of those who remain, and the firms that absorb them in my setting are small informal services rather than a growing manufacturing sector. Closest on the finding is Faber (2014), who finds that connecting peripheral Chinese counties lowered their growth through competition from the core. That channel finds little support here, since connected villages have more firms and higher prices, and losing production to the core still allows gains from cheaper imports, whereas losing agricultural scale is a pure productivity loss.

My contribution to the literature on structural transformation and growth is direct evidence that moving labor out of agriculture is not, by itself, enough for local growth. A long tradition treats labor leaving agriculture as a hallmark of development, and the large measured gaps in output per worker between agriculture and the rest of the economy imply that this reallocation should raise aggregate productivity (Gollin et al., 2014). Recent work qualifies the presumption, showing that developing countries can deindustrialize prematurely (Rodrik, 2016) and that whether reallocation accompanies growth depends on what drives it (Matsuyama, 1992; Bustos et al., 2016). The productivity gap is measured from private returns, but in these villages the social return to agricultural labor exceeds the private one. The villages that move the most labor out of agriculture end up poorer, because the productivity loss among the farmers who remain can outweigh the gains from moving workers toward higher-productivity sectors.

The same estimates complicate the older debate on surplus labor in agriculture. In the dual economy tradition of Lewis (1954) and Ranis and Fei (1961), labor can be pulled out of agriculture at little or no cost to output, while Schultz (1964) argued that this is not the case, and recent experimental evidence finds labor rationing on short-run margins (Breza et al., 2021). My settlement-size instrument provides a long-run version of this test at the village scale, and the estimated elasticity of agricultural output with respect to the agricultural workforce is far from the surplus labor benchmark of zero. Two points reconcile this finding with the modern evidence. First, my elasticity is a long-run object, and the forces behind it would not appear in short-run experimental variation. Second, the existing evidence measures the private return to agricultural labor, while my estimates capture the social return.

To the literature documenting agricultural externalities in developing economies, my contribution is to show that these externalities matter for infrastructure and trade policy. Prior work finds spillovers in the adoption and use of new technologies (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Bandiera and Rasul, 2006; BenYishay and Mobarak, 2019), gains from equipment rental markets and from collective use of capital (Caunedo et al., 2026; Brownstone, 2025; Bassi et al., 2022), and benefits from the cooperative management of irrigation (Ostrom and Gardner, 1993). These otherwise local externalities can reverse the local gains from market access, alter who benefits from infrastructure, and change the aggregate effects of structural transformation.

Sections 2–4 describe the program, data, and empirical strategy. Section 5 presents the effects of market access and Section 6 estimates the agricultural scale externality. Sections 7–9 develop the model, compare it with the data, and evaluate policy. Section 10 concludes.

2 The Transmigration Program and the Outer Islands

This section describes the setting the paper exploits. I describe the Outer Islands at the start of the program, how the transmigration settlements were built and who the settlers were, and how I identify these villages.

The Outer Islands

Indonesia's population has always been concentrated on Java and Bali, the "inner islands," which held roughly two thirds of the country's people on less than a tenth of its land through the period I study. The Outer Islands, principally Sumatra, Kalimantan, Sulawesi, and the eastern provinces, were sparsely settled by comparison. Economic activity there was organized around a small number of provincial capitals and port cities, the plantation belt of eastern Sumatra, and enclaves of oil, gas, and mining. Between these centers lay large areas of forest and swamp with little permanent cultivation. The road network that connected them was thin. The main roads and trails present around 1950, which I digitize in Section 3, changed little until the 1990s (Gertler et al., 2024), so a location's position relative to markets at the time of settlement was fixed for most of the period. The lowest unit of rural administration is the village (*desa*), a settlement of typically 2,000 to 4,000 people with its own elected head, which sits below the subdistrict (*kecamatan*) and the district (*kabupaten*). The village is the unit at which I measure market access and outcomes.

The program

To relieve population pressure on the inner islands and to promote nation-building, the government of Indonesia relocated over 2 million people from Java and Bali to *newly created* villages across the Outer Islands between 1952 and 1995.[2] The program was organized around the five-year development plans of the period (*Repelita*), and the intensity of the program was heavily influenced by the government's ability to pay for it. Transmigration was financed out of central government revenue, and in the 1970s and early 1980s that revenue was dominated by oil and gas; the number of people moved each year tracks the world oil price closely (correlation 0.86), and the collapse of oil prices in 1986 was followed by a sharp contraction of the program (Appendix Figure A2). While activity was concentrated between 1979 and 1989, the average size of villages created each year does not show any trend (correlation -0.02). Villages were created in every year of the period, and I analyze all of them rather than only the 1979 to 1988 cohorts studied in previous work.

As described in numerous sources such as Bazzi et al. (2016), Hardjono (1988), and Kebschull (1986), the transmigration program lacked an organized structure. While in theory the transmigration settlements were supposed to be placed on designated "planned development areas," the

[2] The Transmigration program actually originated in 1905 under the Dutch colonial government. The program greatly expanded after Indonesia became independent, and it is still active today. However, the data I have cover between 1952 and 1995.

plan-as-you-go nature of the assignment process meant that villages were often built wherever there was land available in destination areas. Provincial governments, which had been delegated the final decision on placement, submitted land for settlement, and the unrealistic resettlement targets meant that officials were forced to prioritize resettling individuals as quickly as possible, often without consideration for the location of new villages (Hardjono, 1988). Land suitability maps were abandoned for nearly the entirety of the program, and the maps that existed were unreliable, with updated maps not made available until its last year or two (Hardjono, 1988; Whitten, 1987). What matters for this paper is that a site's position relative to the towns, ports, and roads of the Outer Islands was not what officials were actively considering; they were placing villages where land could be found and settled quickly. Appendix A.1 collects contemporary accounts of the process from a range of sources covering 1960 to 1990.

How a settlement was built

Each settlement was built from the ground up on cleared land. A settlement unit typically housed a few hundred families, laid out around a village center with a school, a health clinic, and a place of worship; each was a village in its own right, administered by the program until handover to the regional administration, on average about 5 years after settlement (World Bank, 1988; Bazzi et al., 2016). Every household received the same allocation: a house and house lot, a 2-hectare field plot often assigned by lottery, and the same package of tools, seeds, and fertilizer, with food support for the first year (World Bank, 1988; Bazzi et al., 2016). Most settlements were food-crop settlements, in which settlers grew upland or wet rice, maize, and cassava on their plots; from the mid-1980s a growing share of settlements were built around a central company plantation, with settlers growing tree crops, principally oil palm and rubber, for the plantation's mill alongside their food plots (World Bank, 1988). Settlers had no control over which new village they were placed in (Kebschull, 1986; Bazzi et al., 2019). Aside from where they were located, the villages therefore started out with the same institutions and infrastructure and, on average, the same endowments of land and inputs per household, which is rarely possible for settlements that formed on their own.

Who the settlers were

Transmigrants were volunteers recruited on Java and Bali. Based on surveys of transmigrant households immediately before migration from Kebschull (1986), they were overwhelmingly asset-poor: most reported no savings, and most held far less land than the plot they received, so few can have had experience operating farms of that size (Appendix Figure A1). Previous work on the program, which studies how the skills settlers brought with them and the mix of ethnic groups within a village affected later outcomes, finds no sorting of migrant characteristics across destination villages (Bazzi et al., 2016, 2019). While that evidence concerns the assignment of people to villages, it does highlight the fact that, on average, villages were populated with similar inhabitants regardless of the location of the settlement.

The settlement register

I identify transmigration villages, their location, the year they were settled, and the number of individuals initially placed in them from the Transmigration Census, a register of settled transmigration sites compiled by the program's directorate. Matching the register to the 2000 administrative village boundaries gives 1,152 villages settled between 1952 and 1995, which is the sample for the rest of the paper. Figure 1 maps them by settlement period. There is substantial variation in where villages were placed across the Outer Islands and that variation in initial position is what the next section measures.

Figure 1: Transmigration Villages by Settlement Period

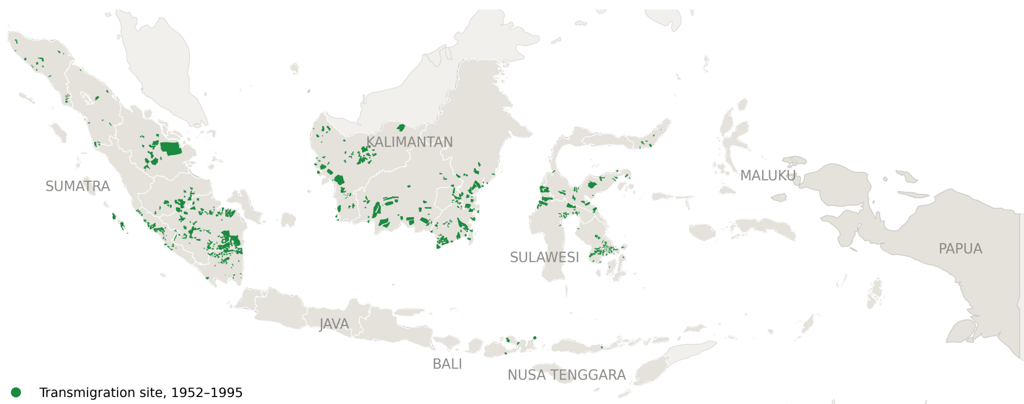

Note: The figure maps the 1,152 transmigration villages in the analysis sample, settled between 1952 and 1995. White lines mark provincial borders as of 2000. The lighter gray areas are neighboring countries, principally Malaysia, Papua New Guinea, East Timor, and the Philippines.

3 Data

The analysis draws on administrative data, individual, firm, and village surveys, satellite-derived outcomes, and a hand-digitized road network. This section first constructs the measure of initial market access, then describes the outcomes, and then the remaining sources and controls. Appendix Table B1 lists every source and Appendix B describes the construction of each variable in detail.

3.1 Initial Market Access

Following Donaldson and Hornbeck (2016), I measure the market access of every location o at the time of the program as

$$MA_o = \sum_d \tau_{od}^{-\theta} N_d, \tag{1}$$

the population N_d of every other location d in Indonesia, discounted by the cost τ_{od} of reaching it. A village that is nearby a large population center or nearby many smaller settlements has high market access. Three inputs are required to construct this measure: the bilateral travel costs τ_{od}, the populations N_d, and the elasticity θ that converts travel cost into a discount.

Travel costs

Travel costs require the road network of the Outer Islands as it stood before the program began in the 1950s. I construct it by hand-digitizing the maps of Indonesia compiled by the US Army Map Service around 1950, which record every road, classified by type from main road down to trail, and every railway in the country across 27 map sheets held in the Perry-Castañeda Library Map Collection. The maps were not drawn in a way that allows the road types to be detected automatically, so I traced each network type by hand and cross-checked roughly half of the sheets for missed segments; Appendix B.11 describes the process and shows an example sheet. A few sheets incorporate limited information from after 1950, but there was little to no road expansion in Indonesia until near the 1990s (Gertler et al., 2024), so the network I digitize is the network the program's villages faced for most of their existence. Because waterways carry goods and people across much of the Outer Islands, I add all navigable waterways from the Humanitarian OpenStreetMap Team. This layer is modern, but treating a river as navigable in 1950 is far less restrictive than treating a modern road as present then.

The digitized network becomes a grid of cells roughly 730 meters on a side, each assigned the speed of the fastest network type that crosses it. Speeds by mode come from the 1993 part of the Indonesia Family Life Survey (IFLS), where each village head reports the distance to nearby towns, the mode of travel, and the time the trip takes. I use the implied speeds for travel by road, by water, and by rail, and a walking speed for cells with no network. Travel times between every pair of locations then follow from a fast marching algorithm over the grid, as in Allen and Arkolakis (2014). As a robustness check I also compute costs under the monetary cost parameters of Donaldson and Hornbeck (2016), which follow Fogel (1964), modified to allow travel by automobile; the two measures of market access have a correlation of 0.89. Appendix B.11 reports the parameters.

Populations

Destination populations N_d come from the 1975 layer of the Global Human Settlement Layer (GHS), a gridded estimate of population built from census counts and satellite-detected built-up area, which I aggregate to village boundaries. I use 1975 because it is the earliest year the layer covers. Because villages were settled between 1952 and 1995, the 1975 weights are strictly pre-settlement only for the later cohorts; for early cohorts they could partly reflect growth induced by earlier program activity. My results are robust to restricting the sample to the 1979 to 1988 villages studied in previous work, for which the weights are unambiguously predetermined (Appendix Table C1). The 2010 GHS layer tracks the 2010 census population closely at the village level, which validates the aggregation of the gridded layer to village boundaries; no village-level population source exists for 1975 against which the earlier layer could be benchmarked directly (Appendix Table B2).

Figure 2: Construction of Market Access

(a) Original Image (b) Digitized Road Network (c) Constructed Market Access

Note: The figure maps log initial market access, Equation (1), for the 1,152 transmigration villages, over the digitized 1950 road and waterway network. Darker shading indicates higher market access.

The elasticity and the resulting measure

I set θ to 8.22, based on Donaldson and Hornbeck (2016). The results are robust to alternatives, including $\theta = 1$, which is essentially a distance-weighted sum of all other locations' population, and to alternative definitions of connectivity.[3] Market access at the time of the program is also strongly correlated with a measure built the same way from Indonesia's modern highway network and a road roughness index from Gertler et al. (2024). Figure 2 demonstrates the process of constructing market access using these data. Panel 2a contains about one eighth of the area of one of the original maps. Panels 2b and 2c show the hand-extracted travel network and final measure of market access, respectively. In this example, connected villages are those connected to the highway system as well as those on Java and in close proximity to Jakarta. In the analysis I use the log of market access, which has a standard deviation of 7.3 log points and an interquartile range of 9.5 across my sample of transmigration villages.

3.2 Outcomes

My primary outcome is village GDP per capita, for which I use the satellite-derived estimates of Rossi-Hansberg and Zhang (2025). These are gridded estimates of GDP that combine nightlights, population, land cover, and other remotely sensed inputs, trained to predict subnational GDP where it is observed, and they perform substantially better at predicting local GDP than nightlights alone. I aggregate them to the 2000 administrative boundaries of each village and observe them annually from 2012 to 2019, so the outcome is measured 17 to 67 years after settlement depending on the cohort. As an alternative, I build GDP per capita myself from harmonized nightlights (Li et al., 2020) and GHS population, converted to GDP using the elasticities estimated by Henderson et al. (2012), Hu and Yao (2022), and Beyer et al. (2022).[4] Section 5.2 adds two other direct outcomes:

[3] These include market access computed only over destinations within 2 hours of the origin, with and without the elasticity, and market access built with an alternative fast marching implementation in QGIS.

[4] I use the Rossi-Hansberg and Zhang (2025) estimates rather than nightlight proxies as the primary measure because the latter rest on an estimated elasticity of light to income, which Gibson et al. (2021) and Bluhm and McCord (2022) have called into question.

the wages that formally employed residents report in household surveys, and a material-welfare index constructed from video footage of a subset of villages.

3.3 Other Sources and Controls

The remaining outcomes come from five sources that I name here and describe when used. The Population Census of 2000 and 2010 is a complete count of individuals with their sector of employment, employment status, and birthplace, from which I build sectoral employment shares, firm ownership, and migration measures at the village level. The village potential census (*Potensi Desa*, PODES) is a census of every village in the country, answered by the village head, that records the village's infrastructure, public goods, land use, and, in the 2003 wave, production and planted area by crop; I use the 1993, 2000, and 2003 waves. The National Socioeconomic Survey (SUSENAS) is the annual household survey, whose consumption modules in 2002 and 2011 record food purchases by item, quantity, and expenditure, and whose employment modules record wages. The 2016 Economic Census lists every medium and large enterprise with its employment. The 2003 Agricultural Census records household landholdings, which the model of Section 7 uses as village land endowments.

The controls are of two kinds. Time-invariant geographic and agro-ecological characteristics of each village location, including potential yields for the main crops, soil, drainage, ruggedness, and climate, follow Bazzi et al. (2016). An additional control based on Bazzi et al. (2016) is a measure of the average agroclimatic similarity between a village's migrant origins and the new destination. The measure captures how much experience migrants have farming in similar conditions, and is a strong predictor of agricultural productivity. District-level characteristics roughly contemporaneous with the program come from the 1980 Population Census and include the sectoral composition of employment, literacy and schooling, and household ownership of durables and housing services. The 1980 measures are predetermined for villages settled after 1980 but not for the earlier cohorts, for which they could reflect the program's own effects; Section 5 reports the estimates with and without them, and the 1979 to 1988 sample is the clean check.

4 Empirical Strategy

This section lays out how I estimate the effect of a transmigration village's initial market access on its long-run outcomes. I state the estimating equation, the assumption under which its coefficient is causal, the control that narrows what the assumption has to cover, and the checks the assumption has to pass.

4.1 Specification

For an outcome y of village v in region r, settled in year c, observed in year t, I estimate

$$y_{vrct} = \beta_1 \ln MA_{vrc} + \beta_2 \ln E(MA)_{vrc} + X'_{vrc}\Gamma + \eta_r + \gamma_c + \alpha_t + \epsilon_{vrct}, \tag{2}$$

where $\ln MA_{vrc}$ is log initial market access from Equation (1), $\ln E(MA)_{vrc}$ is the expected market access control based on Borusyak and Hull (2023) described below, X_{vrc} is the vector of time-invariant village characteristics and 1980 district characteristics of Section 3, and η_r, γ_c, and α_t are region, settlement-cohort, and outcome-year fixed effects, the last included only for outcomes observed in more than one year. My primary specification adopts islands for the relevant region. The coefficient of interest is β_1, the effect of a one log point increase in initial market access.

The island fixed effects absorb the large differences between the islands of the Outer Islands in their fundamentals and their development. The cohort fixed effects do two jobs. The intensity of the program varied greatly over its history (Section 2), villages settled in the same year were built under the same targets, budgets, and administrative practice, and a handful of predetermined characteristics, including some crop-potential measures and the size of settlements, trend across cohorts (Appendix Table A1). And because every outcome is measured in a fixed calendar year, villages settled in the same year have had the same time to develop, so the comparison within a cohort holds the outcome's horizon fixed. All comparisons in the primary specification are therefore among villages placed on the same island and settled in the same year.

Because market access varies smoothly over space, the errors of nearby villages are potentially correlated, and because the outcome is observed in several years while market access is fixed, the errors of the same village are correlated over time. I therefore report standard errors that allow for both, following Conley (1999): the covariance between two observations is weighted by a Bartlett kernel in the distance between their villages, which declines linearly to zero at 100 kilometers, and observations of the same village in different years enter with full weight for every pair of years in the sample.[5] The time kernel is flat over the years in which each outcome is observed, so the within-village component is the same as clustering by village, and the spatial component adds the covariance between villages within the cutoff. Appendix Table C1 reports the estimate at cutoffs of 50, 250, and 500 kilometers, under coarser clustering by district, and under a permutation test that makes no assumption on the error structure.

4.2 Identification

The assumption under which β_1 is causal is that, among villages placed on the same island and settled in the same year, and conditional on expected market access and the village and district characteristics in Equation (2), a village's initial market access is unrelated to the unobserved

[5] The estimator is the spatial heteroskedasticity- and autocorrelation-consistent estimator of Conley (1999) extended to the panel.

determinants of its long-run income. The assumption does not require that villages were created in random locations.

The fact that the characteristics of the villagers and the internal village infrastructure were equalized at the start strengthens the plausibility of this assumption. One of the key determinants of a location's income is the composition of its residents. Because people and institutions often accumulate where access is higher, in most settings a place's market access is jointly determined with the characteristics of those living there. When this is the case, it is not possible to separate the effects of market access and composition when comparing two locations with different levels of market access. Transmigration villages began with the same institutions and infrastructure by design and with settlers who were, on average, alike. Here, market access varied at settlement while the people did not. Any differences in composition that emerge later are therefore consequences of access rather than confounders of it.

The main remaining threat is that the existing network of towns induces a correlation between a village's market access and the fundamentals of the land it is placed on. Suppose two villages are randomly placed on the Outer Islands, one near an existing town and one far from it. The village near town has higher market access, but even though the villages were randomly placed does not guarantee that both villages have the same quality of land. The town was likely formed because the land in that location is productive. If some portion of that productivity also extends to the area where the new village was created, then this introduces a positive correlation between the productivity of the land and initial market access. In the absence of exogenous variation in *changes* in market access, this is the omitted variable that any cross-sectional study of market access confronts, and Appendix C shows in a simulation with randomly placed villages on a non-random network of towns how large it can be and how much of it each of my controls removes.

Three things address the concern. First, the direction of this bias is known. If preexisting towns sit where fundamental productivity is higher, then villages with higher market access are more likely to sit on productive land, which pushes the estimate toward finding *positive* effects of market access on income. I find negative effects, so an omitted variable of this kind cannot explain the finding, and if it is present the true effect is more negative than what I estimate. Second, the village and district controls of Section 3 measure much of what makes land productive directly, potential yields, soil, drainage, ruggedness, and climate, and the balance tests below ask whether market access is related to them. Third, the expected market access control removes the component of a village's access that reflects which part of a province it was assigned to, which is the component most likely to be tied to regional fundamentals. I describe that control next.

4.3 Expected Market Access

Because the location choice of new villages was heavily influenced by provincial and local governments, some of the variation in initial market access reflects which part of the country a village was assigned to be constructed in rather than anything idiosyncratic about its own site. A village assigned to South Sumatra, with its river network and ports, starts with more access than one assigned to the interior of Kalimantan, whatever parcel it lands on. I strip out that component with a simple recentering in the spirit of Borusyak and Hull (2023). I construct a measure of what a village's market access could have been expected to be given where it was assigned, and include it as a control. What remains is the deviation of the village's realized market access from that expectation, which is the variation I exploit.

The measure is built from the set of potential sites. Potential sites include every transmigration village that was actually settled in the same province and every area the program's own planning maps designated as a "Planned Development Area," a candidate settlement location that was never settled. The never-settled sites matter, since they describe where a village could have gone rather than only where villages did go. For each village, I calculate the distance to every other potential site in its province. My baseline measure takes the 20 nearest potential sites, excluding the village itself, and defines expected market access as the average of their log market access. Controlling for this measure isolates the variation in actual initial market access compared to the market access it could plausibly have been assigned. I do not have to claim that a village could have been placed anywhere in its province, only that, among the nearby sites to which it could have been assigned, the particular site it received is uncorrelated with its productive potential. Alternative constructions, an inverse-distance-weighted average and an exponentially distance-weighted average of all other potential sites in the province, give the same results (Appendix Table C1).

A natural alternative would be province fixed effects, which absorb every province-level component of access at once. I use the expected access control as the baseline instead, for two reasons. It absorbs the component of access that the assignment process generated, the position of the province's available land, while keeping the variation in access between provinces that reflects the geography of the Outer Islands, which province fixed effects would discard. And it can be built from the never-settled sites, so it reflects where villages could have gone and not only where they went. Section 5.1 reports the estimates with province and province-by-cohort fixed effects in place of the control; they are smaller, and negative and significant throughout.

4.4 Balance

I then test whether observable characteristics that are likely associated with a location's fundamental productivity differ with a village's initial market access. For each characteristic X I estimate

$$X_{vrc} = \beta_1 \ln MA_{vrc} + \beta_2 \ln E(MA)_{vrc} + \eta_r + \gamma_c + \epsilon_{vrc}, \tag{3}$$

Equation (2) with a characteristic in place of the outcome and without the controls, both with and without the expected access control, so that the table shows what the recentering does.

Table 1 reports the results. The expected market access control does what it is intended to do. Without it, 11 of the 35 characteristics are individually significant at the 10 percent level, and the omnibus test gives $F(35, 284) = 2.152$ with $p < 0.001$. With the control, most coefficients become smaller, only 8 characteristics are significant at the 10 percent level and 1 at the 1 percent level, and the omnibus statistic falls to $F(35, 283) = 1.464$ with $p = 0.050$. The village-level geographic and agro-ecological characteristics are generally balanced. The conventional omnibus p-value should be interpreted with caution though. With many covariates and clustered inference, regression-based joint orthogonality tests can substantially over-reject the null (Kerwin et al., 2024). The significant differences that remain are concentrated among the 1980 district household-amenity measures (piped water, modern fuel, sewage, electricity, and television): villages with higher initial access sit in districts that were somewhat better provided in 1980. Two things should be said about these. They enter the main specification as controls, and Section 5.1 shows that the estimate is unchanged with district fixed effects in their place. And their signs are in the direction that would bias the estimate toward finding positive effects of access, since better-provided districts are richer ones. The other difference worth noting is the number of individuals initially placed, which is a program characteristic rather than a location fundamental and is included as a control throughout. These tests cannot prove that villages are comparable, since they can only speak to the characteristics I observe, but the magnitudes and pattern of the coefficients provide little evidence that initial market access is systematically related to the predetermined site characteristics most likely to confound the estimates.

Table 1: Village Characteristics Across Initial Market Access

	Coefficient on standardized ln(MA)	
Standardized characteristic	**Without** $E[\ln(MA)]$	**With** $E[\ln(MA)]$
Wetland Rice Potential Yield	0.024 (0.033)	0.015 (0.032)
Palm Oil Potential Yield	-0.015 (0.048)	0.011 (0.047)
Coffee Potential Yield	-0.089 (0.104)	-0.051 (0.085)
Cassava Potential Yield	-0.027 (0.039)	-0.019 (0.032)
Vector Ruggedness Measure	-0.050 (0.037)	0.011 (0.034)
Individuals Placed	-0.111*** (0.039)	-0.091** (0.041)
Agrosimilarity Score	0.115** (0.047)	0.076 (0.046)
Share of Land with Poor Drainage	0.027 (0.061)	0.026 (0.057)
Share of Land with Good Drainage	-0.001 (0.057)	0.005 (0.055)
Organic Carbon (%)	0.045 (0.078)	0.020 (0.049)
Topsoil Salinity (Elco) (dS/m)	0.021 (0.016)	0.016 (0.018)
Topsoil Base Saturation (%)	0.024 (0.042)	0.011 (0.041)
Avg. rainfall, 1948–1978	0.083** (0.039)	0.028 (0.038)
Avg. temp (Celsius), 1948–1978	0.059 (0.041)	0.067* (0.038)
ln(District Population, 1980)	-0.038 (0.036)	-0.051 (0.038)
District Own Electricity, 1980	0.069** (0.035)	0.062* (0.036)
District Own Radio, 1980	0.081** (0.039)	0.050 (0.041)
District Own TV, 1980	0.077** (0.035)	0.067* (0.037)

	Coefficient on standardized ln(MA)	
Standardized characteristic	**Without** $E[\ln(MA)]$	**With** $E[\ln(MA)]$
District Literacy, 1980	-0.025 (0.035)	-0.023 (0.039)
District Avg Years Schooling, 1980	0.005 (0.035)	-0.008 (0.037)
District Agricultural Employment Share, 1980	-0.048 (0.041)	-0.057 (0.041)
District Wetland Rice Potential Yield, 1980	-0.020 (0.029)	-0.006 (0.033)
District Palmoil Potential Yield, 1980	0.029 (0.035)	0.030 (0.036)
District Coffee Potential Yield, 1980	-0.059 (0.039)	-0.048 (0.041)
District Cassava Potential Yield, 1980	-0.025 (0.032)	-0.011 (0.035)
District Mining Employment Share, 1980	0.005 (0.039)	0.019 (0.040)
District Manufacturing Employment Share, 1980	-0.018 (0.037)	-0.019 (0.038)
District Trade Employment Share, 1980	0.025 (0.050)	0.057 (0.043)
District Service Employment Share, 1980	0.066* (0.035)	0.067* (0.037)
District Wage Worker Share, 1980	0.002 (0.035)	0.015 (0.037)
District Share HHs with Piped Water, 1980	0.107*** (0.038)	0.105*** (0.039)
District Share HHs with Sewage, 1980	0.072** (0.035)	0.069* (0.037)
District Share HHs with Modern Fuel, 1980	0.088** (0.036)	0.082** (0.038)
District Share HHs with Modern Floor, 1980	-0.064* (0.039)	-0.038 (0.034)
District Share HHs with Modern Roof, 1980	-0.003 (0.035)	-0.020 (0.037)

	Without $E[\ln(MA)]$	**With** $E[\ln(MA)]$
Omnibus test	$F(35, 284) = 2.152$	$F(35, 283) = 1.464$
Omnibus-test p-value	< 0.001	0.050
Observations	1214	1213

Note: Each row reports the coefficient on standardized log initial market access from two regressions, following Equation (3), of the standardized characteristic listed in the row on initial market access. The first column excludes expected log market access, $E[\ln(MA)]$, while the second column controls for it; both specifications include island fixed effects and cohort fixed effects. Standard errors, clustered at the subdistrict (kecamatan) level, are reported in parentheses. The omnibus F-statistics are from corresponding regressions of standardized log initial market access on all of the listed characteristics jointly, without and with $E[\ln(MA)]$, conditional on the fixed effects, and test their joint significance. The reported omnibus p-values are based on the corresponding cluster-robust Wald F-tests. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results of the next section are accompanied by three kinds of robustness checks. I vary the control set and the fixed effects, replacing the island fixed effects with province, province-by-plan-period, province-by-cohort, and district fixed effects. I vary the inference, with coarser clustering and a permutation test that reassigns market access across villages within province-by-cohort cells. And I ask how strong an unobserved confounder would have to be, relative to the observed controls, to eliminate the effect (Diegert et al., 2025).

5 The Long-Run Effects of Market Access

This section presents the effects of initial market access on the long-run outcomes of the transmigration villages. Section 5.1 reports the effect on GDP per capita, its magnitude, and its robustness. Section 5.2 shows that the same pattern appears in three other measures of development, two of which follow residents rather than places. Section 5.3 turns to productivity, and Section 5.4 to the channels that could explain the result.

5.1 GDP per Capita

Table 2 reports estimates of Equation (2) for log GDP per capita, observed annually from 2012 to 2019. Panel A uses the primary measure. Column 1 includes only the island fixed effects and the village controls; column 2 adds the cohort fixed effects; column 3 adds the 1980 district controls; and column 4, my preferred specification, adds the expected market access control. Higher initial market access significantly lowers long-run GDP per capita in every specification, and the estimate barely moves as controls are added. Adding controls that themselves predict income does not shift the estimate, which is the first indication that market access is not capturing a relationship between a location's productivity and income.

The effect is economically large. Moving from the 25th to the 75th percentile of initial market access lowers GDP per capita by about 5.5 percent, and one standard deviation of access lowers it by about 4 percent. The interquartile effect is nearly 20 percent of the interquartile spread in GDP per capita across the villages. The outcome is measured between 17 and 67 years after settlement depending on the cohort, so these are long-run differences in the level of income, not differences in growth rates, and they are estimated within cohorts, so among villages that have had the same time to develop.

Table 2: Effects of Initial Market Access on GDP per Capita

	(1)	(2)	(3)	(4)
Panel A. ln GDP per capita, Rossi-Hansberg and Zhang (2025)				
ln(Market Access)	-0.0060^{***}	-0.0059^{***}	-0.0057^{***}	-0.0058^{***}
	(0.0022)	(0.0019)	(0.0019)	(0.0020)
DMP breakdown fraction	0.714	0.643	0.643	0.857
Villages	1,152	1,152	1,152	1,152
Panel B. ln GDP per capita from nightlights, Henderson, Storeygard, and Weil (2012) elasticity				
ln(Market Access)	-0.0279^{***}	-0.0219^{***}	-0.0169^{***}	-0.0145^{**}
	(0.0068)	(0.0061)	(0.0060)	(0.0067)
Villages	1,141	1,141	1,141	1,141
Panel C. ln GDP per capita from nightlights, Hu and Yao (2022) elasticity				
ln(Market Access)	-0.0315^{***}	-0.0253^{***}	-0.0203^{***}	-0.0182^{***}
	(0.0068)	(0.0061)	(0.0060)	(0.0067)
Villages	1,141	1,141	1,141	1,141
Panel D. ln GDP per capita from nightlights, Beyer, Hu, and Yao (2022) elasticity				
ln(Market Access)	-0.0358^{***}	-0.0294^{***}	-0.0244^{***}	-0.0227^{***}
	(0.0069)	(0.0061)	(0.0060)	(0.0067)
Villages	1,141	1,141	1,141	1,141
Expected market access control	No	No	No	Yes
Village controls	Yes	Yes	Yes	Yes
District controls (1980)	No	No	Yes	Yes
Cohort fixed effects	No	Yes	Yes	Yes
Island fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes

Note: Each entry is the coefficient on log initial market access from Equation (2), with the controls and fixed effects listed at the bottom of the table. Panel A uses the satellite-derived GDP per capita estimates of Rossi-Hansberg and Zhang (2025), observed annually from 2012 to 2019. Panels B to D use GDP per capita built from harmonized nightlights and GHS population, converted to GDP with the elasticity of the study named in the panel heading. The DMP breakdown fraction is the share of village controls for which the Diegert et al. (2025) breakdown point exceeds the control's own explanatory importance (Appendix C.3). Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Robustness

The estimate is robust on three margins, reported in Appendix Table C1. The first is the fixed effects. Replacing the island fixed effects with province fixed effects gives -0.0034, smaller than the baseline but significant at the 1 percent level, and province-by-plan-period and province-by-cohort fixed effects give -0.0023 and -0.0033. The province-by-cohort specification is demanding, since the median cell contains 2 villages and over a third of cells contain a single village, yet the estimate remains significant at the 1 percent level.

The second is inference. Widening the spatial cutoff of the standard errors to 500 kilometers barely changes them, and clustering by district in place of the spatial correction leaves the estimate significant at the 1 percent level in every specification. A permutation test makes the same point without any assumption on the error structure: none of 1,000 random reassignments of market access across villages within province-by-cohort cells produces a relationship as strong as the one I observe.

The third is sensitivity to omitted variables. Following Diegert et al. (2025), I ask how strongly an unobserved confounder would have to be correlated with market access, relative to the observed controls, to drive the estimate to zero. For over 85 percent of the village controls, the confounder would have to matter more than the control does, a higher fraction than in the application where Diegert et al. (2025) judge a result robust (Appendix C.3).

5.2 Alternative Measures of Development

Panels B to D of Table 2 replace the primary measure with GDP per capita built from nightlights under each of three light-to-GDP elasticities. The estimates are negative under all three and *larger* than the primary estimate, by nearly four times under the elasticity of Beyer et al. (2022), where one standard deviation of market access lowers GDP per capita by 15 percent. The primary measure is therefore the conservative one. The nightlight panels also show why the detailed controls matter. Panel A barely moves across specifications, but in Panels B to D the column 4 estimates are on average 42 percent smaller than those in column 1.

I also find similar negative effects of market access on income in household-level survey data. The SUSENAS records the earnings, hours, and sector of formally employed individuals each year, and although a given transmigration village enters the sample only occasionally, pooling the 2000 to 2007 waves gives wage observations for 502 villages and 5,781 individuals, which I convert to 2007 values. Appendix Table D1 reports the effect of market access on monthly wages, hourly wages, and hours worked. The point estimates for wages are negative in every specification, 2 to 3 percent per standard deviation of access, about half the effect on GDP per capita, with no difference in hours. They are imprecise, since the sample covers formal workers alone, a minority of rural employment.

Finally, I construct a measure of village material well-being from an independent source. Publicly available Indonesian-language documentary and vlog footage covers 47 of the transmigration villages in Kalimantan, and I score frames from each video against a fixed rubric of 7 visible indicators of material welfare (housing quality, road surface, electrification, agricultural assets, commercial activity, public infrastructure, and clothing) using a vision language model, then aggregate the scores to a village-level index (Appendix E). Regressing the index on market access in this small sample, without the controls of the main design, gives a negative and significant slope for the composite index and negative slopes for every component. Applying this estimate to the full sample distribution of initial market access implies that a one standard deviation increase in market access lowers village welfare by 0.35 standard deviations, roughly three times the magnitude of the effect on GDP per capita. Because of the small sample, I read this as corroboration from a source that shares nothing with the satellite or administrative data rather than as an estimate of comparable precision, but the direction is the same.

5.3 Productivity

Lower GDP per capita in more connected villages also appears in measures of village productivity. Table 3 reports the effects of market access on proxies for productivity inside and outside agriculture. In agriculture, PODES records the area planted and the production of each crop in the village, from which I build yields per hectare for the 5 crops with the largest planted area in each village, and separately for rice, the staple that nearly every food-crop settlement grew, and for oil palm, a key cash crop in many settlements (Section 2). PODES also records the number of tractors, which I express per thousand residents and per thousand agricultural workers as a measure of mechanization. Outside agriculture, the Population Census records whether each working individual is an employee or an employer, from which I build the ratio of employees to employers in the village, a measure of the size of the typical non-farm establishment.

Villages with higher market access are less productive on every measure. One standard deviation of access lowers tractors per thousand agricultural workers by 0.40 and tractors per thousand residents by 0.36, about a fifth of the sample mean in each case. The same variation lowers the yield of a village's top 5 crops by roughly 7 percent and the yield of rice by a similar amount, both significant at the 10 percent level, while the palm oil yield is negative but imprecise. Outside agriculture, the ratio of employees to employers falls by about 13 percent of its mean, so the typical firm is smaller. The lower agricultural yields and the lower mechanization are the two facts that Section 6 returns to, since both are what a loss of scale in the agricultural sector would produce.

Table 3: Effects of Initial Market Access on Productivity

	Employees per Employer	Tractors per 1000 Individuals	Tractors per 1000 Ag Workers	ln(Productivity of Top 5 Crops)	ln(Rice Productivity)	ln(Palm Oil Productivity)
ln(Market Access)	-0.6881**	-0.0650*	-0.0715*	-0.0106*	-0.0145**	-0.0015
	(0.3124)	(0.0361)	(0.0418)	(0.0060)	(0.0058)	(0.0177)
Expected MA Control	Yes	Yes	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No	No	No
Unique Villages	1107	1035	1035	980	591	302

Note: Each column reports the coefficient on log initial market access from Equation (2). Column 1 is the ratio of employees to employers among working residents in the 2000 and 2010 Population Census. Columns 2 and 3 are tractors per thousand residents and per thousand agricultural workers, from the 2000 PODES and the 2000 Population Census. Columns 4 to 6 are log yields per hectare for the 5 most planted crops, for rice, and for oil palm, from the 2000 PODES; the samples differ because not every village grows rice or oil palm. Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4 Mechanisms

Why would better connectivity lower long-run income? I examine, in turn, the channels that theory and previous work suggest: that connected villages specialized into agriculture, that they lost activity to competition from the core, that people or production left, that lower income is compensated by better amenities, and that lower measured income is offset by lower prices. None of these explains the result. What the evidence shows instead is that connected villages moved *more* labor out of agriculture, into more numerous and smaller non-farm firms, while their agricultural productivity fell, which is the combination that an agricultural scale externality produces, and which Section 6 then estimates directly.

Sectoral shifts

One explanation for lower output in connected villages is that exposure to competition led productive industry to leave and the village to specialize in agriculture. If better access raised the return to farm output and exposed local industry to competition from more industrial locations, market access would raise the agricultural employment share, consistent with the theoretical arguments of Matsuyama (1992) and the empirical findings of Baum-Snow et al. (2020) and Faber (2014).[6] Table 4 shows the opposite. Column 1 reports the agricultural employment share in the 2000 and 2010 Population Census, and column 2 its change between the two. A one standard deviation increase in market access lowers the agricultural employment share by 2.8 percentage points, and the share falls faster between 2000 and 2010 in more connected villages. Columns 3 to 5 show where the workers go. Manufacturing, services, and retail all gain, but services and retail account for nearly 70 percent of the reallocation, and they are also the largest non-agricultural sectors in these rural economies. Market access therefore spurs structural transformation out of agriculture while lowering output per capita and productivity. Given how often structural transformation is promoted as a route to growth in developing countries, this combination is the central puzzle of the paper, and it is exactly the pattern a model in which agricultural TFP depends on the scale of the agricultural workforce predicts.

Table 4: Effects of Initial Market Access on Structural Transformation

	Agriculture Share	Change in Ag Share	Manufacturing Share	Services Share	Trade Share
ln(Market Access)	-0.0037***	-0.0027***	0.0012***	0.0016***	0.0010***
	(0.0009)	(0.0008)	(0.0003)	(0.0004)	(0.0003)
Expected MA Control	Yes	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	Yes	Yes	Yes
Unique Villages	1151	1151	1151	1151	1151

Note: Each column reports the coefficient on log initial market access from Equation (2), using the 2000 and 2010 Population Census. Column 1 is the agricultural employment share, pooled across the two census years; column 2 is its change between 2000 and 2010. Columns 3-5 are the manufacturing, service, and retail employment share, respectively. Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The shift is toward small-scale services. In these villages, 69.3 percent of rural retail workers

[6] In Faber (2014) increased market access leads to large declines in output in sectors *except* in agriculture, suggesting that perhaps resources were shifted there in response to increased competition.

are self-employed, against 53.0 percent in agriculture and 19.8 percent in manufacturing. The non-farm sector these villages transform into is not industry but small-scale services sold to nearby customers.

Firm entry and size

Connectivity could reduce the number of local firms through competition or raise it through the demand that nearby population brings. Table 5 shows that in this setting it raises it. Across sectors and firm sizes, higher market access increases the number of firms in the village and the share of residents who report operating a firm, formal or informal, in the census. Among villages with medium or large firms, average employment per firm is smaller in more connected villages, the same pattern as the lower ratio of employees to employers in Table 3, and consistent with entry-driven selection in which the marginal entrant is smaller. The evidence also points away from the competition channel of Faber (2014). If market access mainly exposed transmigration villages to more productive external firms, I would expect firm exit and fewer local firms; instead, connected villages have more. Firm counts cannot rule out compositional shifts within the firm population, but the absence of net exit, together with the price evidence below, does not support a competition channel behind the negative effects.

Table 5: Effects of Initial Market Access on Firm Entry and Size

	Employers per 1000	Total Number of Firms	ln(Employment)	ln(Employment)
ln(Market Access)	1.4266**	0.1449	-0.0289**	-0.0244**
	(0.5651)	(0.0941)	(0.0130)	(0.0102)
Expected MA Control	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	No	No
District Controls	Yes	Yes	No	No
Cohort FE	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No
Sector FE	No	No	No	Yes
Unique Villages	1151	947	128	128

Note: Each column reports the coefficient on log initial market access from Equation (2). Column 1 uses the 2000 and 2010 Population Census; column 2 combines the 2003 PODES and the 2016 Economic Census; columns 3 and 4 use the firm-level listing of the 2016 Economic Census. Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Migration and commuting

If initially connected villages offered their residents better opportunities elsewhere, part of the negative effect on GDP per capita could reflect people leaving the village rather than lower resident incomes. I measure out-migration two ways. Because transmigrants came from Java and Bali, the residents of a transmigration village born on the inner islands are, to a close approximation, the original settlers and later arrivals from the same region, so their number and share in the 2000 and 2010 Census fall when settlers leave and proxy for permanent out-migration. The 2003 PODES also records whether residents have left the village to work elsewhere, a direct if coarse indicator. Table 6 shows no relationship between market access and any of these measures. The number and share of inner-island-born residents, their number relative to the original settlement population, and the PODES indicator are all unrelated to initial access.

Commuting is unlikely to explain the result either, though no dataset observes it at a level that can be matched to transmigration villages. The concern is that some of the output of commuters is recorded outside the village, so that high-commuting villages appear poorer than they are. Two facts weigh against it. First, commuting in rural areas is uncommon. In the 2005 Intercensal Population Survey, most districts have rural commuting rates below 2 percent of the working population, so the increase in connected villages would have to be implausibly large to explain the effects I find. Second, commuting can depress measured local production but not outcomes that follow residents, and market access also lowers the wages residents earn wherever they work (Table D1), the material-welfare index built from village footage (Section 5.2), and the nightlight-based measures of GDP per capita (Table 2). These results do not rule out that residents of connected villages travel for work more often; the model of Section 7 suggests they could. But that travel is the structural transformation documented above, and it lowers village income through the agricultural scale externality.

Table 6: Effects of Initial Market Access on Migration

	ln(Inner Islanders)	Share of Inner Islanders	Inner Islanders – Individuals Placed	Commuting in Village (0/1)
ln(Market Access)	-0.0050	-0.0009	-2.8469	0.0006
	(0.0059)	(0.0011)	(5.0745)	(0.0012)
Expected MA Control	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No
Unique Villages	1140	1151	1151	947

Note: Each column reports the coefficient on log initial market access from Equation (2). Columns 1 and 2 are the log number and the share of residents born on Java or Bali in the 2000 and 2010 Population Census. Column 3 is the difference between the number of residents born on Java or Bali in the 2000 or 2010 Population Census and the number of individuals initially settled in the transmigration village. Column 4 is an indicator from the 2003 PODES for whether residents of the village have left to work elsewhere. Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Amenities

Lower measured income need not mean lower welfare if connected villages offer amenities that compensate their residents. I test this with a public goods index built from the 2000 PODES: 18 measures of health care availability and quality, school presence and staffing, safe drinking water, waste collection and sanitation, and road maintenance and quality, each standardized and averaged. Table 7 reports the effect of market access on the full index and on the index split into its road measures and everything else. Market access significantly improves the road measures, which shows that the index can detect differences across villages, but it does not improve the full index, and on every non-road dimension connected villages are no better provided. The index cannot cover every dimension households might value, but the estimates show no compensating improvement in non-road public goods.

Table 7: Effects of Initial Market Access on Amenities

	Public Goods Index Full	Public Goods Index Excl Roads	Public Goods Index Roads Only
ln(Market Access)	0.0032	0.0019	0.0092**
	(0.0025)	(0.0027)	(0.0040)
Expected Market Access Control	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes
District Controls	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes
Island FE	Yes	Yes	Yes
Year FE	No	No	No
Unique Villages	1151	1151	1151

Note: Each column reports the coefficient on log initial market access from Equation (2), using the 2000 PODES. The public goods index is the average of 18 standardized measures of health, schooling, water, sanitation, and roads; column 2 excludes the road measures and column 3 uses them alone. Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Prices and expenditure

Lower GDP per capita would overstate the loss in living standards if connected villages also faced lower prices. The consumption modules of the 2002 and 2011 SUSENAS record, for each food item a household buys, the quantity and the expenditure, so the ratio is a unit value, the price the household paid. The two waves cover 84 transmigration villages and 1,234 households. I build two price indices from the log unit values: one weights items by the household's own expenditure shares, and one by the average expenditure shares of rural households on the same island that do *not* live in transmigration villages, so that the second index compares the cost of a common basket. Unit values reflect quality and outlet as well as price, so I read them as prices under the usual assumption that these do not vary systematically with access within an island (Deaton, 1988); the common-basket index limits the scope for composition to enter.

Table 8 shows that households in more connected villages pay more for food, not less. Both price indices rise with market access, by 0.011 to 0.014 per log point, or 8 to 10 percent per standard deviation, and both are precisely estimated. Purchased quantities and food expenditure rise as well, but that is not evidence of higher consumption. The module records purchases rather than consumption of own-produced food, and as households leave agriculture they buy the food they previously grew, so purchases rise even if total food consumption is unchanged. Housing, the largest non-food category in this setting, tells the same story. Households in more connected villages spend substantially more on it, and the point estimates attribute the increase to the price of housing rather than to the size of dwellings, which does not differ with access, though the two components are individually imprecise (Appendix Table F1). Connected villages do not appear to consume more except through the higher prices they pay. The lower measured production per capita is therefore not offset by a lower cost of living; if anything, the price evidence suggests the GDP per capita estimates understate the real-income consequences of market access in this context.

Table 8: Effects of Initial Market Access on Food Prices, Quantities, and Expenditure

	Price Index^{Ind}	Quantity Index^{Ind}	ln(Total Expenditure)
ln(Market Access)	0.0074**	0.0011	0.0154**
	(0.0034)	(0.0059)	(0.0077)
Expected Market Access Control	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes
District Controls	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes
Island FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Unique Villages	84	84	84
Unique Households	1234	1234	1234

Note: Household-level regressions following Equation (2), using the food consumption modules of the 2002 and 2011 SUSENAS matched to transmigration villages (84 villages, 1,234 households). The price index is the expenditure-share-weighted sums of log unit values, with weights from the household's own shares (Ind); the quantity index is the corresponding weighted sums of log quantities. Column 3 is log food expenditure. Standard errors, in parentheses, allow for spatial correlation following Conley (1999), with a Bartlett kernel and a 100 kilometer cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Land

Land-use records from PODES point to the same squeeze from the supply side. The amount of land under cultivation in more connected villages is lower and decreasing over time, while the land devoted to housing and firm activity does not expand to accommodate the additional non-agricultural activity (Appendix Table F2). Converting land between uses appears difficult in this context, so structural transformation operates through labor rather than land. This is how the model of Section 7 treats land: it works only in agriculture, and non-agricultural activity absorbs labor, not land.

Summary

Taken together, the reduced-form results paint a consistent picture. More connected transmigration villages have lower GDP per capita, lower agricultural productivity, more structural transformation out of agriculture, more firm entry but smaller firms, no sizable migration response, no improvement in non-road public goods, and higher local prices. The evidence does not favor any of the standard channels, agricultural specialization, competition, labor drain, or compensating differentials, as the explanation for the negative effects of connectivity in this setting.

The combination of faster structural transformation and lower agricultural productivity points instead to an agricultural scale externality. Connected villages moved more labor out of agriculture, and their farms became less productive. If farm productivity rises with the size of the sector, every worker who leaves lowers the productivity of those who remain, and the lower yields and fewer tractors of connected villages (Table 3) are what this mechanism predicts. The next section estimates the externality directly, and the rest of the paper quantifies its consequences.

6 Estimating the Agricultural Scale Externality

The results from Section 5 say that connected villages move more labor out of agriculture and farm less productively. This section estimates the link between the two, the elasticity of agricultural productivity with respect to the size of the agricultural workforce, using a second source of variation from the program: villages were settled with different numbers of people. The estimate is the parameter that the model of Section 7 uses for an external scale elasticity.

6.1 Framework and Identification

Let agricultural production of crop c in village v be

$$Q_{vc} = A_{vc}\,(L_v^{ag})^{\gamma_a}\,(H_{vc})^{\gamma_H}\ , \tag{4}$$

where L_v^{ag} is the village's agricultural labor, H_{vc} is the area planted to crop c, γ_a is the private labor share, the elasticity of a farm's output with respect to its own labor, which under competition is the share of output paid to labor, and γ_H is the corresponding elasticity with respect to land. Village-level productivity depends on the scale of agricultural labor through an externality,

$$A_{vc} = \bar{A}_c\,(L_v^{ag})^{\varphi}\ , \tag{5}$$

which no individual farm internalizes. Substituting and taking logs gives the estimating equation

$$\ln Q_{vc} = \delta_c + (\gamma_a + \varphi)\ln L_v^{ag} + \gamma_H \ln H_{vc} + \varepsilon_{vc}, \tag{6}$$

with $\delta_c = \ln \bar{A}_c$ a crop effect. The coefficient $\beta \equiv \gamma_a + \varphi$ on $\ln L_v^{ag}$ is the object of interest: it is how much a crop's output rises with the size of the village's entire farming workforce. With an external estimate of the private share γ_a, the externality is $\varphi = \beta - \gamma_a$.

The data are the village-by-crop records of the 2003 PODES, which report production quantity and area planted for each crop grown in the village in the most recent season, for the transmigration villages of the main sample. Because I observe quantities directly, differences in prices across locations do not enter. Agricultural labor L_v^{ag} is the number of residents working in agriculture in the Population Census, treated as contemporaneous with the 2003 outcomes. PODES does not record labor by crop, so every crop's equation uses the village's total agricultural labor. The coefficient then measures how a crop's output moves with the whole farming workforce, which is the object the model needs, and it reads as a pure scale effect as long as the share of the workforce allocated to a given crop does not itself move with the instrument. Three features address this. The crop effects absorb permanent differences in labor intensity across crops. The planted-area control absorbs differences in the scale of each crop, which is where differences in the allocation of labor across crops would appear. And a direct check finds that initial settlement size does not shift the share of the village's agricultural workers growing cash crops rather than staples in the Population Census; the census distinguishes crop types rather than crops, so the check is coarser than the outcome, but it does not point toward a composition effect.

Generally, agricultural labor is endogenous; sorting, local fundamentals, and equilibrium responses jointly determine agricultural employment and production. I instrument $\ln L_v^{ag}$ with $\ln Z_v$, the number of individuals initially settled in village v by the program, which was determined decades before the 2003 outcomes. The instrument is relevant because larger initial settlements have larger agricultural workforces decades later: the first-stage coefficient is 0.200 and precisely estimated (Appendix Figure G2). Conditional on island and cohort fixed effects, initial settlement size is uncorrelated with the village-level geographic and agro-ecological characteristics, with market access, and with expected market access, so the instrument varies scale within villages that share the connectivity variation of Section 5 (Appendix Table G2). Several 1980 district characteristics, concentrated among measures of infrastructure and urbanization, are correlated with settlement size. I control for them throughout. The qualitative record of Section 2 points the same way, since settlement sizes were set by how many families a cleared site could hold and how many were in the pipeline, not by the quality of the land. Bazzi et al. (2019) also present evidence consistent with the assumption that the number of individuals settled is not related to a location's fundamentals.

I also present evidence that the exclusion restriction holds in this setting. The exclusion restriction is that initial settlement size affects crop output only through the size of the farming workforce it generates. If initially larger settlements received more infrastructure, traders, or services over time, and these changes directly impact crop productivity, then this would violate the exclusion restriction. Two features of the program limit the concern around these channels. First, because villages were created with the same design in terms of the household land allocation and internal infrastructure, differences in land or inputs per farm or access to infrastructure are unlikely to arise. Second, downstream channels that a larger settlement might have attracted, a permanent market, a bank, schools, are for the most part unrelated to settlement size in the 2003 PODES. Conditioning on these potential channels, which is not a fix since they are themselves mediators, also does not substantially change the estimated size of the scale elasticity (Appendix Table G3).

Unlike the village and district controls, which are predetermined, planted area is chosen in the same season as output and varies at the crop level, so it is the one control that could respond to the same shocks as the outcome. The baseline includes it because β is the labor coefficient at given land. The results below therefore report the estimate with and without the area control, and the sensitivity exercise that follows them bounds what an area response could contribute.

The second stage at the village-by-crop level is

$$\ln Q_{vc} = \beta \widehat{\ln L_v^{ag}} + \gamma_H \ln H_{vc} + X_v'\Gamma + \delta_c + \eta_i + \gamma_k + \varepsilon_{vc}, \tag{7}$$

where $\ln L_v^{ag}$ is instrumented by $\ln Z_v$, δ_c are crop effects, η_i and γ_k are island and cohort effects, and X_v holds the village and district controls of Section 3. Standard errors are clustered at the village level, which handles the repeated crop observations within a village. With the Indonesia-specific estimate $\gamma_a = 0.304$ (Fuglie, 2010), I infer $\varphi = \beta - \gamma_a$.

Placing the externality in agriculture rather than in the non-farm sector is a choice the data support. Figure G1 shows that rural wages rise with the size of a village's agricultural sector and are flat in retail. This relationship holds even after including village and demographic controls (Appendix Table G1), and retail services in developing countries show low productivity and no evident scale economies (Lagakos, 2016).

6.2 Results

Table 9 reports the estimates. Column 1 estimates Equation (7) on the main crops with the agricultural controls, and gives $\beta = 1.21$. Column 2 adds the non-agricultural village controls, which raises the estimate to 1.40. Column 3 extends the sample to every crop recorded in PODES, including minor crops grown in few villages, and gives 1.05. Columns 4 and 5 interact the crop effects with island and with cohort effects, so that the comparison is among villages growing the same crop on the same island, or settled in the same year, and give 1.38 and 1.36. Column 6 uses yield per hectare rather than production as the outcome and gives 1.18. Across the six specifications the point estimates of $\beta = \gamma_a + \varphi$ run from 1.05 to 1.40, and every one exceeds the private labor share of 0.304 by a wide margin: the implied externality $\varphi = \beta - \gamma_a$ runs from 0.75 to 1.10.

Table 9: Estimation of the Agricultural Scale Externality

	(1) ln(Production)	(2) ln(Production)	(3) ln(Production)	(4) ln(Production)	(5) ln(Production)	(6) ln(Yield)
$\ln(L_v^{ag})$	1.206***	1.400***	1.052***	1.381***	1.364***	1.176***
	[0.46, 4.59]	[0.53, 6.44]	[0.27, 4.25]	[0.51, 6.41]	[0.52, 5.77]	[0.43, 4.50]
Effective F (Olea-Pflueger)	7.60†	6.58†	8.16†	6.54†	6.88†	7.60†
Implied φ ($\alpha = 0.304$)	0.902	1.096	0.748	1.077	1.060	0.872
Implied φ lower bound (AR, $\alpha = 0.304$)	0.151	0.225	-0.035	0.207	0.218	0.127
Crop Sample	Main	Main	All	Main	Main	Main
Agriculture Controls	Yes	Yes	Yes	Yes	Yes	Yes
Non-Ag Controls	No	Yes	Yes	Yes	Yes	No
FE Type	Crop + Isle + Cohort	Crop + Isle + Cohort	Crop + Isle + Cohort	Crop × Isle + Cohort	Crop × Cohort + Isle	Crop + Isle + Cohort
Observations	1487	1487	2261	1487	1487	1487
Villages	602	602	612	602	602	602

Note: 2SLS estimates with standard errors clustered at the village level. Significance stars reflect the Anderson-Rubin test of $H_0\colon \beta = 0$ (weak-IV robust), not the Wald t-test. $\ln(L_v^{ag})$ is instrumented by ln(individuals initially placed). **Effective** F is the Montiel Olea and Pflueger (2013) first-stage statistic under the same spatial covariance; † flags values below 16.38, the threshold above which conventional Wald inference using $|t| > 1.96$ is approximately valid. **AR 95 percent CI** is the Anderson and Rubin (1949) confidence set, recommended by Andrews et al. (2019) as the primary weak-IV-robust inference object for just-identified models; it is reported in brackets beneath the point estimates. Implied externality $\varphi = \widehat{\beta} - \alpha$ at $\alpha = 0.304$ (Fuglie, 2010). The *lower bound* on φ is computed as (AR lower bound on β) $- \alpha$ and reflects the smallest value of the externality elasticity not rejected by the data at the 5 percent level under weak-IV-robust inference.

Because the instrument is only moderately strong, I rely on the weak-instrument-robust inference recommended by Andrews et al. (2019). The effective F-statistic of Montiel Olea and Pflueger (2013) is between 6 and 8 in every column, below the threshold at which conventional Wald confidence intervals are reliable. The Anderson and Rubin (1949) confidence sets, reported in the table, are bounded in every specification. Their lower bounds on β, between 0.26 and 0.52, sit above the Wald lower bounds, while their upper bounds, between 4.4 and 7.2, sit well above the Wald upper bounds, reflecting the range of large coefficients the data cannot reject once weak identification is taken into account. The data reject $\beta = 0$ in all six specifications and reject $\beta \leq \gamma_a = 0.304$, no externality, in five of six. The table also reports the implied lower bound on φ, which is simply the lower bound of the Anderson-Rubin set for β minus γ_a; it is positive in five of six specifications, between 0.12 and 0.21, and -0.05 in the sixth. The model of Section 7 adopts $\varphi = 0.70$, a value below every point estimate. Appendix M.1 validates the estimator itself on data generated by the calibrated model, where the true coefficient is known: the same 2SLS recovers $\gamma_a + \varphi$ to within 2 to 4 percent, and recovers γ_a when the externality is shut off.

The bracket from the area control is informative about the one channel the exclusion discussion left open. Omitting planted area from the 2SLS gives a coefficient of 1.07, against 1.39 with it, and the two are not statistically different. If the positive but insignificant response of planted area to settlement size is instead read as a genuine response operating through agricultural labor, with elasticity π_H/π_L of area with respect to labor, the specification without the area control identifies $\beta^{\text{no}\,H} = \gamma_a + \varphi + \gamma_H\,\pi_H/\pi_L$. With $\hat{\pi}_H = 0.118$, $\hat{\pi}_L = 0.200$, $\gamma_a = 0.304$, and $\gamma_H = 1 - \gamma_a = 0.696$, the private constant-returns technology the model uses, $\hat{\beta}^{\text{no}\,H} = 1.07$ implies

$\varphi \simeq 1.07 - 0.304 - 0.696\,(0.118/0.200) \simeq 0.36$. Even if the entire point estimate of the land response is attributed to additional agricultural labor, the implied externality remains economically large and positive.

6.3 Magnitude

The externality implies that agricultural output scales roughly one-to-one with the size of the agricultural workforce in a village: even though land is a fixed factor, the externality makes agriculture behave as if it has constant returns to scale in labor at the village level. There is no estimate of scale economies in agriculture at this level to compare against; the nearest benchmark is the sector-level estimates of Bartelme et al. (2025) across manufacturing industries, which average roughly 0.15, four to five times smaller than what I find. A potential reason for the gap is that my variation comes from relatively small movements in the size of villages, while that work uses variation at the country-by-sector level. The average population of a transmigration village is around 2,000, and it seems plausible that scale externalities are non-linear and stronger at lower initial levels. If the externality operates through mechanisms with a minimum viable scale, such as equipment rental markets, the cooperative maintenance of irrigation, or learning from other farmers (Bassi et al., 2022; Foster and Rosenzweig, 1995; Conley and Udry, 2010; Caunedo et al., 2026; Ostrom and Gardner, 1993), then villages of this size are precisely where the marginal worker matters for whether those mechanisms exist rather than how intensively they are used. Appendix J sketches where an externality of this size could come from and measures some potential channels directly in my setting.

7 A Spatial Model with Agricultural Scale Economies

The empirical results establish three facts. Villages with higher initial market access have lower GDP per capita, they reallocate more labor out of agriculture, and agricultural productivity rises with the scale of the local agricultural workforce. In this section I develop a model that ties these facts together, ask whether the estimated externality is large enough to reverse the standard gains from market access, and take the model to the data. Section 7.1 lays out the ingredients. Section 7.2 solves a one-village version in closed form and gives the externality threshold at which the effect of connectivity changes sign. Section 7.3 writes out the full spatial general equilibrium model of the Indonesian economy, and Section 7.4 describes how its inputs are measured and how it is calibrated.

7.1 Ingredients

Two features separate the model from a standard spatial model. First, a village's non-farm service good is traded across space. The provider travels to the buyer or the buyer to the provider, so a sale requires a trip. Importantly, trade costs of services rise more steeply with distance than agricultural goods (Section 7.4). This asymmetry concentrates the demand for non-farm work in high-market-access locations. Second, agricultural productivity rises with the scale of the local agricultural workforce, the externality estimated in Section 6. This effect on agricultural productivity is not internalized when individuals make their sector choice.

The rest of the model is standard, or motivated by empirical evidence. Labor is immobile across locations, since migration does not respond to market access (Section 5.4) and migration frictions in Indonesia are large and well documented (Bryan and Morten, 2019). The non-farm sector is standard in that it produces with constant returns and its productivity does not depend on its size, which is the sense in which the small informal services that absorb labor in these villages do not carry growth. Everything else is standard: workers sort between the two sectors within their village, and trade in food and city goods delivers the usual gains from market access.

Environment

Locations are either rural villages or urban locations, and each one carries the same two-sector economy. There are three goods. Food is homogeneous, traded, and the numéraire. Rural services and the urban good are differentiated by origin location and traded, and both are aggregated by buyers with elasticities σ_n and σ_u, respectively.

Locations are separated by travel times over the network of Section 3. The cost factor between locations ℓ and o is $\tau_{\ell o} = 1 + \text{hours}(\ell, o)$, and for objects that depend on a village's access to urban markets I write τ_v for the travel cost to the village's nearest urban location. Each good faces a sector-specific travel cost, with food shipped at τ^{κ_a}, services at τ^{κ_n}, and the city good at τ^{κ_u}. The distinction that matters throughout is that it is less costly to ship food than services, and the difference in trade costs is what drives the labor reallocation.

Within a location, workers possess sector-specific abilities and sort freely between its two sectors, so the sectoral allocation responds to the relative return to work in each sector. Each location is endowed with L workers and H units of land; land works only in agriculture, and the service and city-good sectors use labor alone, which reflects the land-use evidence of Section 5.4.

Agriculture is the same technology at every location, rural or urban.[7] Production is constant-returns Cobb–Douglas in effective labor $Z_{a,\ell}$ and land, with total factor productivity depending on the scale of the location's agricultural workforce,

$$A_{a,\ell} = \bar{A}_{a,\ell}\, Z_{a,\ell}^{\varphi}, \qquad \varphi > 0, \tag{8}$$

external to the individual farmer and therefore uncompensated in factor markets. Substituting (8) into the production function,

$$Y_{a,\ell} = \bar{A}_{a,\ell}\, Z_{a,\ell}^{\gamma_a+\varphi}\, H_{\ell}^{1-\gamma_a}, \tag{9}$$

[7] The choice to include agriculture as a sector choice in urban locations is motivated by the fact that the agricultural employment share of urban locations is 0.29.

so the location-scale technology is homogeneous of degree $1+\varphi$ in labor and land jointly (in labor alone, land fixed, returns are $\gamma_a+\varphi$, marginally above one at the calibrated values, 1.004; no restriction either way is imposed), while atomistic firms earn zero profits. The agricultural wage per effective unit is $w_{a,\ell}=\gamma_a p_{a,\ell}Y_{a,\ell}/Z_{a,\ell}$, and land rents $(1-\gamma_a)p_{a,\ell}Y_{a,\ell}$ accrue to local households. The externality (8) is the same object estimated in Section 6, whose estimating equation is the log of (9) with $\beta=\gamma_a+\varphi$ the coefficient on agricultural labor. The possibility of multiple equilibria that the unrestricted returns create is handled by the selection protocol described below.

The non-agricultural sector is also the same technology everywhere, but its output differs by location type. Rural locations produce a differentiated service, urban locations a differentiated city good, and both are produced with labor alone,

$$Y_{k,\ell}=A_{k,\ell}\,Z_{k,\ell},\qquad k=n \text{ (rural)},\ k=u \text{ (urban)}, \tag{10}$$

so the gate price of a location's variety is tied to its wage, $p_{k,\ell}=w_{k,\ell}/A_k$.

7.2 A Closed-Form Characterization

Before turning to the full economy, I add enough structure to characterize the contest between the externality and the gains from proximity in closed form. Consider a single village at distance τ from a city with the structure described above. The only addition is that agricultural ability is common across workers and non-agricultural ability is drawn from a Pareto distribution. This allows me to write the response to a connectivity shock only in terms of parameters and the service employment share u, which I can solve for in closed form under these assumptions. Appendix I contains the derivation, and Appendix H solves a many-village version with the full sorting structure numerically.

The village sells food and services at the gate prices p_a and p_n, and because the occupational margin depends on the two producer prices only through their ratio, the relative price of services $\rho=p_n/p_a$ governs sorting. I write $\eta=d\ln u/d\ln\rho$ as the elasticity of the service employment share with respect to the relative price of village services. An increase in the relative service price pulls workers out of agriculture, and the reallocation lowers agricultural TFP due to the externality. The income lost through this channel is given by,

$$\mathcal{E}\equiv s_a\,\varphi\,\frac{u}{1-u}\,\eta, \tag{11}$$

which I call the externality drag. A 1 percent rise in the relative service price reduces effective labor in agriculture by approximately $\frac{u}{1-u}\,\eta$ percent. Through the externality, this lowers agricultural TFP by φ times that reallocation, and then the resulting income loss is weighted by agriculture's share of village income s_a.

Proposition 1 (Sign reversal of the connectivity gradient)**.** *In the simplified model of this subsec-*

tion, derived in Appendix I, the relative price of services falls with distance from the city whenever $\kappa_n \frac{\sigma_n - 1}{\sigma_n} > \kappa_a$, and real village income rises with distance (proximity is a curse) if and only if the externality loss outweighs the two market-access gains,

$$\underbrace{\mathcal{E}\left|\frac{d\ln\rho}{d\ln\tau}\right|}_{\textit{externality loss}} > \underbrace{(s_n - \alpha_n)\left|\frac{d\ln\rho}{d\ln\tau}\right|}_{\textit{terms-of-trade gain}} + \underbrace{\alpha_u(\kappa_a + \kappa_u)}_{\textit{cheaper-city-goods gain}},$$

equivalently if and only if φ exceeds a threshold φ^* available in closed form (Appendix I, equation (I14)). Here s_a and s_n are the village's sectoral income shares, u its service employment share, α_n and α_u its service and city-good consumption shares, and ρ and η are as defined above. For $\varphi < \varphi^*$ proximity raises real income and the standard gains from trade prevail.

The proposition describes the effect of a connectivity shock term by term. Proximity delivers three standard gains. The village sells the services it exports at a better relative price, sells its food at a better farm-gate price, and buys city goods at lower delivered prices. The effect of the agricultural scale externality operates in the opposing direction. At $\varphi = 0$ this term disappears and every channel is a gain from proximity, so any curse in this model must come from the externality. Of note, the mechanism itself requires the sector-specific trade cost asymmetry. The relative service price falls with distance only when $\kappa_n(\sigma_n - 1)/\sigma_n$ is greater than κ_a. This condition guarantees that services are effectively harder to trade than food. Reallocation vanishes only at the knife-edge case where $\kappa_a = \kappa_n(\sigma_n - 1)/\sigma_n$; with equal frictions, $\kappa_a = \kappa_n$, the relative service price would instead rise with distance and proximity would move labor toward agriculture. At the measured values of Section 7.4 this condition holds easily. The sections that follow determine the quantitative strength of each mechanism.

7.3 The Quantitative Model

The full model embeds the structure of Section 7.1 at every location. The economy consists of 69,225 locations: roughly 190 million people across 48,598 rural villages (1,147 transmigration and 47,451 pre-existing) and 20,627 urban locations, on both the inner and outer islands (the program villages themselves are all located on the outer islands). Every location, rural or urban, carries the same two-sector economy, and all trade flows over the measured travel times. Baseline agricultural productivity $\bar{A}_{a,v}$ is common to all transmigration villages, whose fundamentals were assigned by the program; every other location's fundamentals are recovered from its observed outcomes (Section 7.4).

Households

A representative household lives at every location and has Cobb–Douglas preferences over food, a rural-service composite, and a city-good composite,

$$U_\ell = c_{a,\ell}^{s_{a,\ell}}\, Q_{n,\ell}^{s_{n,\ell}}\, Q_{u,\ell}^{s_{u,\ell}}, \qquad s_{a,\ell}+s_{n,\ell}+s_{u,\ell}=1, \tag{12}$$

where the expenditure shares differ by location type. Rural households spend with shares $(\alpha_a, \alpha_n, \alpha_u)$ and urban households with shares $(\beta_a, \beta_n, \beta_u)$. The service composite aggregates the varieties of all rural origins by CES,

$$Q_{n,\ell} = \Big[\sum_{o\in\text{rural}} \omega_o^{1/\sigma_n}\, q_{n,\ell o}^{\frac{\sigma_n-1}{\sigma_n}} \Big]^{\frac{\sigma_n}{\sigma_n-1}}, \qquad \sigma_n > 1, \tag{13}$$

and the city-good composite aggregates the varieties of all urban origins the same way with elasticity σ_u. Every variety arrives at $\tau_{\ell o}^{\kappa_n} p_{n,o}$, and city varieties at $\tau_{\ell c}^{\kappa_u} p_{u,c}$. The variety weights ω_o convert location populations into village equivalents, with $\omega_o = N_o/\bar{s}$, where N_o is location o's population from the settlement layer and $\bar{s}$ is the average rural village population. A location twice the average size carries twice the variety mass, and for consistency each transmigration village carries $\omega = 1$, but assigning every location one variety instead leaves the estimated model's moments essentially unchanged (Appendix M.4).[8]

Given location income I_ℓ, the sum of factor payments, which by constant returns equals the value of production, demand follows the Cobb–Douglas shares, and the exact cost-of-living index is

$$P_\ell = p_{a,\ell}^{s_{a,\ell}}\; P_{n,\ell}^{s_{n,\ell}}\; P_{u,\ell}^{s_{u,\ell}}, \tag{14}$$

where $P_{n,\ell}$ and $P_{u,\ell}$ are the CES price indices over service and city varieties. Because preferences are homothetic, every weight in (14) is an expenditure share, so income deflated by P_ℓ is the welfare-consistent real income measure used throughout.

Workers draw sector-specific productivities $(z_{i,a}, z_{i,k})$ independently from Fréchet distributions $F_j(z) = \exp(-T_j z^{-\theta})$ with $\theta > 1$ and work in the sector with the higher return, where the non-farm sector is services in a village and the city variety in an urban location. The agricultural employment share and effective labor supplies are

$$\pi_{a,\ell} = \frac{T_a w_{a,\ell}^{\theta}}{T_a w_{a,\ell}^{\theta} + T_k w_{k,\ell}^{\theta}}, \qquad Z_{j,\ell} = L_\ell\, \Gamma\big(1-\tfrac{1}{\theta}\big)\, T_j\, w_{j,\ell}^{\theta-1} \big(T_a w_{a,\ell}^{\theta} + T_k w_{k,\ell}^{\theta}\big)^{\frac{1}{\theta}-1}, \tag{15}$$

and average earnings are equalized across sectors, so employment shares are the same labor-income shares. This property ties the observed level of agricultural employment to the demand structure and makes the employment share the natural target for the sorting scale T_a/T_k.

[8] Scaling variety mass with population is the standard device for carrying locations of very different sizes in a single Armington structure (Ramondo et al., 2016).

Trade

Food is homogeneous and traded, and urban markets anchor the price. The market price of food is the numéraire, equal to one. Every location, rural or urban, grows food and sells at the market price net of transport to its nearest urban market, $p_{a,\ell} = \tau_\ell^{-\kappa_a}$. Cities grow food in the model because they do in the data: the Population Census records an agricultural employment share of 0.26 in the median administratively urban location and 0.36 in the peri-urban locations that the functional classification below counts as urban, so an urban block without agriculture could not reproduce the employment it is asked to match. For urban locations the transport factor is one, so cities sell food at the market price, and a village at travel time to the nearest urban location τ_v receives the farm-gate price $\tau_v^{-\kappa_a}$. Because food is homogeneous, a village's own food demand is satisfied from its own supply at the farm-gate price, and only the excess is then exported. Proximity helps agriculture through this channel.[9] Because cities sit at $\tau_\ell = 1$, the import and export parity prices coincide, so they simply pay the market price for the food they consume. The economy is closed. The national food market clears in the aggregate, and by Walras' law its clearing condition is redundant given the others.

A unit of the service is a unit of labor, so delivering variety o to location ℓ means a provider travels at iceberg cost $\tau_{\ell o}^{\kappa_n}$. The city good is a composite of physical goods, which ship like agricultural goods, and person-embodied town services, which behave like the rural service; κ_u is built from the food and service frictions in proportion to their measured composition rather than assumed, and buyers at v receive city c's variety at $\tau_{vc}^{\kappa_u} p_{u,c}$. Because goods are cheaper to move than people, the city-good friction is much flatter than the service friction. Trade in services does not conflict with the assumption that labor is immobile. Every worker in the model works in their own village and is paid their own village's wage. What the assumption rules out is commuting, where a worker takes a job at another location's wage, which is consistent with the setting, as commuting is rare in rural Indonesia and the cost of regularly traveling to urban areas is high (Section 5.4).

Every buyer, rural or urban, spends the service share of its income (α_n or β_n) on the rural-service composite and allocates it across origins by CES shares. Village v's service revenue is therefore

$$p_{n,v} Y_{n,v} = \omega_v \, p_{n,v}^{1-\sigma_n} \Big[\beta_n \sum_{\ell \in \text{urban}} \tau_{\ell v}^{-\kappa_n(\sigma_n - 1)} \frac{I_\ell}{\Phi_\ell} \; + \; \alpha_n \sum_{\ell \in \text{rural}} \tau_{\ell v}^{-\kappa_n(\sigma_n - 1)} \frac{I_\ell}{\Phi_\ell}\Big], \tag{16}$$

where $\Phi_\ell = \sum_{o \in \text{rural}} \omega_o \big(\tau_{\ell o}^{\kappa_n} p_{n,o}\big)^{1-\sigma_n}$ is buyer ℓ's competition index over all rural locations, with a rural buyer's own variety included at $\tau = 1$, and incomes I_ℓ are equilibrium objects. City-good

[9] For a rural net food importer, this misstates the relevant price by at most the width of the arbitrage band, $2\kappa_a \ln \tau_v$ (the gap between the import-parity price $\tau_v^{+\kappa_a}$ and the export-parity price $\tau_v^{-\kappa_a}$). I use the single farm-gate price rather than switching prices for each village due to the computational difficulty of handling so many nonlinearities. The same farm-gate price enters the cost-of-living index (14), so for an importer the index understates the consumer price of food by at most the same band, which at the food share α_a is under 3 percent of the index where importers cluster. The model is therefore exact for net exporters and an approximation for the importing remainder, which holds 11 percent of the rural population; the accounting audit of Appendix L.4 shows the resulting error is small where the importers cluster.

revenue has the identical gravity structure over city varieties with (κ_u, σ_u), the buyers' city-good budgets α_u and β_u, and the competition index taken over urban locations.

Equilibrium

Given the measured geography, populations, endowments, and parameters, an equilibrium is a wage pair at every location, such that, in every location: (i) workers sort by (15), (ii) the agricultural wage satisfies its first-order condition (9), and (iii) the non-farm labor market clears, service revenue (16) for villages and variety revenue for cities, with every competition index evaluated at the equilibrium price vector. The fixed point has the standard gravity structure. Given prices, I compute the competition indices and demand sums, and given those aggregates each location's two-equation system delivers its wages. The aggregate food market clears by Walras' law.[10]

In the full model the mechanism runs through the same four links as in the one-village case. Demand for a village's service variety decays faster than for agricultural products, so villages near cities and other central locations capture a disproportionate share of spending on rural services. That demand raises the service wage relative to the agricultural wage and workers reallocate out of agriculture. The reallocation lowers effective agricultural labor and, through the externality (8), agricultural TFP. Because food prices do not fall with connectivity, the productivity loss translates into lower factor incomes in agriculture, and the price index (14) is approximately flat because the rising farm-gate food price and the falling city-good index offset at the measured frictions and consumption shares. While this effect is negative, the model still embeds the standard gains from connectivity, since higher market access locations can buy city goods at lower prices and sell their services for higher prices.

7.4 Taking the Model to the Data

Most of the model's inputs are measured directly. The fundamentals of each location are recovered by inverting the model. For every location that is not a program village, I choose the agricultural productivity and the non-farm sector productivity that allow the model to exactly match that location's observed income per capita and agricultural employment share. Inversion like this is the standard way of calibrating a spatial model, and it means that the baseline economy reproduces the observed Indonesian economy by construction. The program villages are the exception in this approach. They were built to be alike and placed without regard to their sites, which is the identification argument of Section 4, so in the model they carry *common* fundamentals rather than fitted ones, and their incomes are solved from the model given only their geography. The relationship between market access and income among them is therefore something the model

[10] Uniqueness cannot be guaranteed because the externality leaves effective returns to scale in agriculture slightly above one. I solve at $\varphi = 0$, where the equilibrium is unique under standard gravity conditions, and follow that branch as φ rises to its calibrated value. I verify that solving the model without this continuation, and after perturbing every price, converges to the identical equilibrium at every location (Appendix L).

predicts rather than something it was fitted to. The rest of this section describes the measured inputs, the inversion, and the one parameter estimated inside the model; Table 10 collects the parameter map, and Appendix K gives the measurement in full.

Geography, endowments, and outcomes

The trade-cost matrix is the travel-time network of Section 3, evaluated between all locations, and the variety weights use the same population layer as the market-access measure. Every program village is assigned the mean recorded labor and land endowment of the program, so that geography alone generates differences among them.[11] All other land comes from the 2003 Agricultural Census of household landholdings.[12] The observed outcomes that the inversion matches are income per capita, from the same village-GDP series as Section 5 normalized to the measured rural income level, and the agricultural employment share, from the 2010 census where locations can be matched and from the 2000 census elsewhere, adjusted for the province-level change between the two waves. Urban locations are matched on the same two outcomes. I classify as urban every administratively urban location and every location whose centroid falls inside a satellite-defined functional urban area, which raises the urban population share from the administrative 44 percent to just under 65 percent, closer to aggregate statistics; Appendix M.5 describes the alternatives and shows that the results do not depend on where the boundary is drawn.

Measured parameters

Technology. The private agricultural labor share is $\gamma_a = 0.304$, the Indonesia-specific estimate of Fuglie (2010); the land share $1 - \gamma_a$ follows. The externality is set to $\varphi = 0.70$, a deliberately conservative reading of the IV evidence of Section 6: the 2SLS point estimates of $\hat{\beta} - \gamma_a$ range from 0.75 to 1.10 across the six specifications, and 0.70 lies below every one of them while remaining inside every confidence interval. Section 8 reports the model across a wide range of externality values and Section 9 reports every experiment at $\varphi = 0$ as well. Because the 2SLS instruments for agricultural *headcount* while the model uses effective labor, Appendix M.1 runs the paper's estimator on data generated by the calibrated model, where the coefficient is known, and recovers $\gamma_a + \varphi$ to within 4 percent of the true value, and γ_a when the externality is shut off. Because φ is imposed from the empirical estimates rather than fitted, everything the model says about income is a prediction of the measured externality, and not a calibration to the effect I find in Section 5.

[11] In the model, any village-level noise in the endowments appears as an endowment cross-section the data do not contain. For instance, a (statistically insignificant) negative relationship between market access and the amount of land a program village has can magnify the negative relationship between market access and income by a noticeable amount. In the data both program moments are *invariant* to controls for recorded placement size and land per worker, so the common scale is the convention the data support.

[12] Because that source misses estate, plantation, and fishery agriculture, a small number of populated locations record near-zero household land. I winsorize non-program land per worker from below at 0.02 hectares per person; without the floor these locations would appear as spuriously cheap service suppliers in the counterfactuals. Appendix K.4 details the gap and reports alternative floors.

Trade frictions. I measure separate trade frictions for services, food, and the city good from the 1993 IFLS community survey, which records travel times and fares to a range of destinations, local wages, doctors' fees, and staple prices for 312 communities (Appendix K.2). Services require a person to travel, so I measure their trade cost as the time and monetary cost of the trip relative to the value it carries. A round trip to the district capital costs a typical community about half the value of a day of work, and a same-day trip to the province capital is too costly for most villages; pricing the trip against a day's wage, or against a doctor's consultation fee, gives service trade-cost exponents of 0.94 and 0.89, and I set $\kappa_n = 0.90$. These same trip costs are why commuting is rare in rural Indonesia. For food, comparing the same trip costs with the value of a 50-kilogram load of rice gives an upper bound on κ_a of 0.10 to 0.12, which overstates the cost of commercial shipments, and the relationship between staple prices and travel time is a precisely estimated null; I set $\kappa_a = 0.05$, and Appendix M.7 varies it between half and twice that value. The city good combines physical goods, which ship at the food friction, with formal health and education services, which move with people at the service friction; in the rural SUSENAS the bundle is 82 percent goods and 18 percent services, so $\kappa_u = 0.82\kappa_a + 0.18\kappa_n = 0.203$. The gap between the service friction and the goods frictions is the intuitive feature of the data that moving people is far costlier than moving goods, and Appendix M.8 shows what happens when either friction is set to the other's value.

Demand. I measure rural expenditure shares on outer-island rural households in the 2002 SUSENAS. Mapping the survey categories into the three goods requires one adjustment: many purchases that appear as local services, such as prepared food and small-scale retail, combine local service labor with purchased food inputs, so I express shares in value-added terms, assigning a fraction $\ell = 0.48$ of these purchases to service labor and the remainder to food, using Economic Census cost data (Appendix K.1). One part of this construction cannot be measured directly. The SUSENAS records what households purchase but not whether the service was sourced within the village or from outside it, and the crosswalk depends on that share, which I denote s^{own} and estimate inside the model below. At the estimated value, $s^{\text{own}} = 0.532$, the rural shares are $(\alpha_a, \alpha_n, \alpha_u) = (0.667, 0.103, 0.230)$. The urban shares are implied by market clearing: given observed incomes and α_n, total spending on rural services must equal total service revenue, which pins β_n, the same accounting for city goods pins β_u, and β_a follows by adding up. This gives a check on the calibration, since the implied urban food share of 0.595 can be compared with the 0.560 obtained by applying the same crosswalk to urban SUSENAS households and adjusting for the measured Engel relationship; the two are close, and Appendix M.6 shows that re-measuring the demand block in other survey years and samples changes little.

Elasticities. Preferences across food, rural services, and the city good are Cobb–Douglas, a restriction that Appendix K.6 tests in the SUSENAS: expenditure shares on city goods do not fall as their delivered price rises with isolation, the implied substitution elasticity is close to one, and a substantially more price-elastic specification is rejected. For the sorting elasticity θ, I impose the choice elasticity of 12.5 (s.e. 1.36) that Bryan and Morten (2019) estimate for Indonesian workers reallocating across alternatives in the same data environment; their margin (migration across locations) is not exactly mine (choice of sector within a village), but it is the closest measured elasticity the mechanism runs on, and Appendix M.2 frees θ and recovers a value just beyond the upper end of their confidence interval with every prediction essentially unchanged. There is no accepted estimate of a rural service substitution elasticity σ_n; applications of the Gervais and Jensen (2019) service-to-manufacturing ratio to standard trade elasticities put substitution across service-industry varieties near 2.8, while substitution across individual service establishments within a city is about 9 (Couture, 2016), and a village's service bundle is likely between an industry composite and a single establishment. I set $\sigma_n = 5$, equal to the city-variety elasticity σ_u. The composition moment below is close to flat in σ_n near this range, so the data would not distinguish values within it, and Appendix M.13 sweeps σ_u and finds similar results throughout.

Table 10: Parameter map

Parameter	Value	How pinned	Source / moment
Externally set or measured			
γ_a	0.304	set	Fuglie (2010)
φ	0.70	from IV	Section 6 2SLS, $\hat{\beta} - \gamma_a$
θ	12.5	imposed	Bryan and Morten (2019) choice elasticity (s.e. 1.36)
σ_n, σ_u	5	set	Gervais and Jensen (2019) direct service range; swept in App. M.13
κ_n	0.90	measured	IFLS 1993 trip costs, two routes (0.94 delivery, 0.89 consumption)
κ_a	0.05	accounting	trip-cost accounting; staple-price slope ≈ 0
κ_u	0.203	composed	consumption-share mix $s_g\kappa_a + s_s\kappa_n$, SUSENAS rural shares (App. K.2)
ℓ	0.48	measured	Economic Census bundle labor content, band [0.42, 0.54]
Implied by identities or inverted from observed outcomes			
$\alpha_a, \alpha_n, \alpha_u$	0.667, 0.103, 0.230	crosswalk	SUSENAS 2002 expenditure block at the estimated (s_{own}, ℓ)
β_n, β_u	0.024, 0.381	adding-up	service and city-good clearing at observed incomes
β_a	0.595	adding-up	residual; measured Engel value 0.560 (overidentifying check)
$(\bar{A}, T_n)$ per rural location	fields	inverted	observed income and agricultural share, all non-program desa
$(\bar{A}, T_u)$ per urban location	fields	inverted	observed income and agricultural share, all urban locations
Internally estimated (Section 7.4)			
s_{own}	0.533	moment	structural-transformation slope b_{ST}
T_a	5.4×10^{-23}	inner	transmigration agricultural employment level (0.648)

Note: The scale of T_a reflects θ in its exponent and has no free interpretation.

Inversion and internal calibration

For each rural location other than a program village, I choose the agricultural productivity and the service-sector fundamental that make the model reproduce its observed income per capita and agricultural employment share; two observed outcomes recover two fundamentals. Urban locations are treated the same way, with an agricultural productivity and a city-good fundamental. The program villages receive common agricultural and service fundamentals, and their outcomes are solved from the model; the one common parameter, the agricultural sorting scale T_a, is chosen so that the model reproduces the program-wide agricultural employment share of 0.648. Appendix L.3 gives the algebra and computational details.

One parameter remains to be estimated inside the model, the local sourcing share s^{own}. Through the value-added crosswalk, s^{own} determines the rural expenditure shares and therefore how strongly access to outside service demand affects the allocation of labor: a lower s^{own} implies a larger service expenditure share α_n, so improvements in market access generate a stronger shift of employment out of agriculture. I estimate it from that relationship. The targeted moment is the coefficient b_{ST} from a regression of the non-agricultural employment share on the same reduced-form measure of market access used in Section 5, among the transmigration villages, which in the data is 0.0053 (s.e. 0.0009).[13] For each candidate s^{own} I update the expenditure shares, re-invert the non-program economy, re-pin T_a, solve the full joint equilibrium, and estimate the same regression in the model. The value that matches the data is $s^{\text{own}} = 0.532$, inside the range of 0.471 to 0.904 that the survey's purchase categories allow. Because the system is exactly identified, both targeted moments are matched exactly (Appendix Table L1). There is also an untargeted check on the estimate: at $s^{\text{own}} = 0.532$ the crosswalk implies an own-village share of service consumption of about 0.35, and the model's equilibrium own-village share, an outcome of the spatial model rather than of the accounting, is about 0.30.

8 The Model and the Cross-Section

This section asks whether the estimated model reproduces the negative relationship between market access and income among the program villages. The relationship between market access and structural transformation is targeted; the relationship with income is not, and there is no guarantee that the model produces a negative one, let alone one of the magnitude in the data. Section 8.1 reports that comparison, Section 8.2 traces the gradient across values of the externality to locate the point at which its sign flips, and Section 8.3 asks whether the model needs worse fundamentals in connected villages to fit the data.

8.1 The Income Gradient

Table 11 compares the slope of log GDP per capita on log market access among the transmigration villages in the data and in the model, estimated the same way in both, with island and cohort fixed effects, on the reduced-form access measure of Section 5. The rows differ only in the sample: every program village, or the sample that excludes the most isolated 5 or 10 percent. The last row reports the model slope with the externality removed.

I find that the model is able to explain between roughly 40 and 130 percent of the effect I find in the data, depending on the exact sample of transmigration villages I examine. A consistent feature of the model is that I am unable to match the high incomes of the most isolated program villages, and my discussion in this section focuses on the results when I ignore the most isolated 5 percent of the sample. In that sample, the model explains just over 80 percent of the relationship between market access and GDP per capita in the transmigration villages. Figure 3 shows the relationship

[13] Both the data and model regressions include island and cohort fixed effects, with standard errors clustered by village in the data. The coefficient is larger than in Section 5 because I am only using the 2010 Population Census in order to have measurements as close to the 2012 measure of GDP per capita that is the other primary input.

between data and model together. The premium earned by the most remote villages is a distinct phenomenon the model does not capture, and Appendix M.3 traces how the slopes move as those villages are excluded; Section 8.3 shows the inversion would load this into fundamentals. Given that none of these slopes are targeted, the success of the model in generating a negative relationship between market access and income highlights the quantitative importance of the scale externality in explaining the effects of market access on income in this setting.

Table 11: Income gradients: model against data

Slope of ln GDP pc	Sample	Data (s.e.)	Model	
reduced-form measure, real	full	−0.0066 (0.0016)	−0.0028	untargeted
reduced-form measure, real	excl. bottom 5%	−0.0047 (0.0017)	−0.0038	untargeted
reduced-form measure, real	excl. bottom decile	−0.0030 (0.0018)	−0.0040	untargeted
reduced-form measure, real, $\varphi = 0$	full	—	+0.0029	untargeted

Note: Regressions include island and cohort fixed effects, district-clustered standard errors. Model slopes are computed in equilibrium on the reduced-form access measure of Section 5, with model income per capita deflated by each village's price index. The $\varphi = 0$ row solves the full economy with the externality removed and the productivity fields re-inverted to reproduce the same observed baseline, with the transmigration agricultural sorting scale re-pinned to match the observed mean program agricultural employment share.

Figure 3: Income and Market Access in the Data and the Model

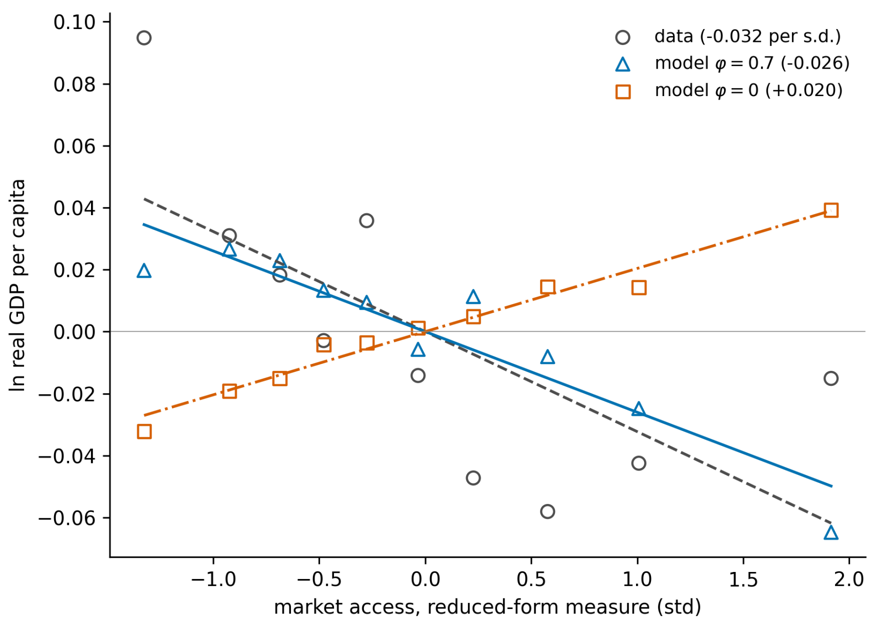

Note: The figure plots the relationship between income and market access for transmigration villages in the data and in the model, excluding the lowest 5 percent of villages in market access. For both data and model, it plots 10 quantile bins of the island-and-cohort fixed-effect residuals for the 1,089 villages, with the fitted slope per standard deviation of log access and its district-clustered standard error in the legend, with and without the externality.

Shutting off the externality reverses the effect, and Figure 3 includes the relationship between market access and income when there is no externality. With $\varphi = 0$ and the world re-inverted to the same observed baseline, the same geography generates a *positive* relationship between market access and income. As the model predicts, connectivity becomes a blessing (better farm-gate prices, cheaper goods, service-export demand) once agricultural productivity no longer depends on agricultural scale. The externality is not one contributor to the curse among several, but the driver of the negative relationship between market access and GDP per capita.

8.2 The Externality Threshold

Proposition 1 says that proximity becomes a curse once the externality exceeds a threshold, and the full model locates that threshold. Figure 4 traces the model's cross-sectional access gradient, the slope of program-village log real income on log market access estimated as in Table 11, across values of the externality. Each point re-solves the whole economy at that φ, with T_a re-pinned and the non-program productivity fields re-inverted so that every φ reproduces the same observed baseline. The gradient reverses between $\varphi = 0.5$ and 0.6, and the calibrated $\varphi = 0.70$ sits inside the curse region. This is the quantitative counterpart of the sign-reversal thresholds of the simplified models, about 0.38 in the closed-form case once the measured frictions are included (Appendix I) and 0.33 to 0.65 across the variants (Appendix H). The point estimates and the calibrated value clear the threshold; the most conservative readings of Section 6, the weak-instrument-robust lower bounds of 0.12 to 0.21 in five of six specifications and the land-response value of 0.36, do not.

Figure 4: The income gradient across externality values

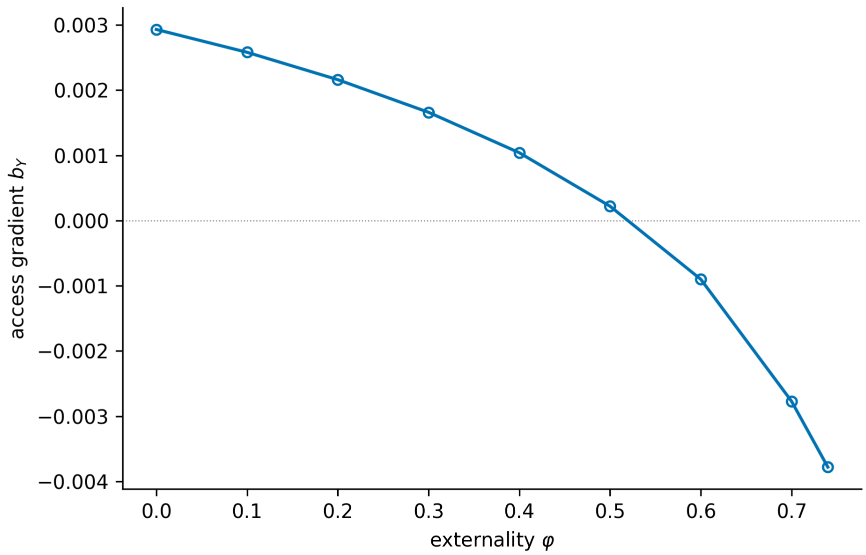

Note: The figure plots the model's income gradient b_Y, the slope of log real income per capita on log initial market access among the 1,147 transmigration villages, with island and cohort fixed effects, on the reduced-form access measure of Section 5. Each point re-solves the whole economy at the stated φ, re-pinning the program villages' agricultural sorting scale so that the baseline reproduces the observed program agricultural employment share. The calibrated value is $\varphi = 0.7$.

8.3 Structural Balance

The balance tests of Section 4 find no evidence of observable factors that are correlated with initial market access and lower income, and the model permits a structural analog of the same test, conditional on the model's structure. Inverting each transmigration village with the same process used for every other location, so that the model reproduces its observed employment share and income exactly, recovers the fundamentals that would replicate the observed relationship between market access and income by construction. I can then ask whether the estimated model needs a relationship between fundamentals and market access in order to fit the data exactly. It does not. Figure 5 shows no relationship between a transmigration village's recovered agricultural productivity and market access, and the same null holds for the service fundamental. This null relationship is a property of the model *with* the externality only. The same exercise carried out in the $\varphi = 0$ world uncovers a strong negative relationship between market access and baseline agricultural productivity: without the scale externality, only worse fundamentals can rationalize connected villages being poorer, and the required imbalance loads on agricultural productivity. Given I observe no relationship in the data between market access and many measures that capture agricultural productivity, excluding an externality in the agricultural sector produces a world that is at odds with the data.

Figure 5: Recovered Agricultural Productivity and Market Access

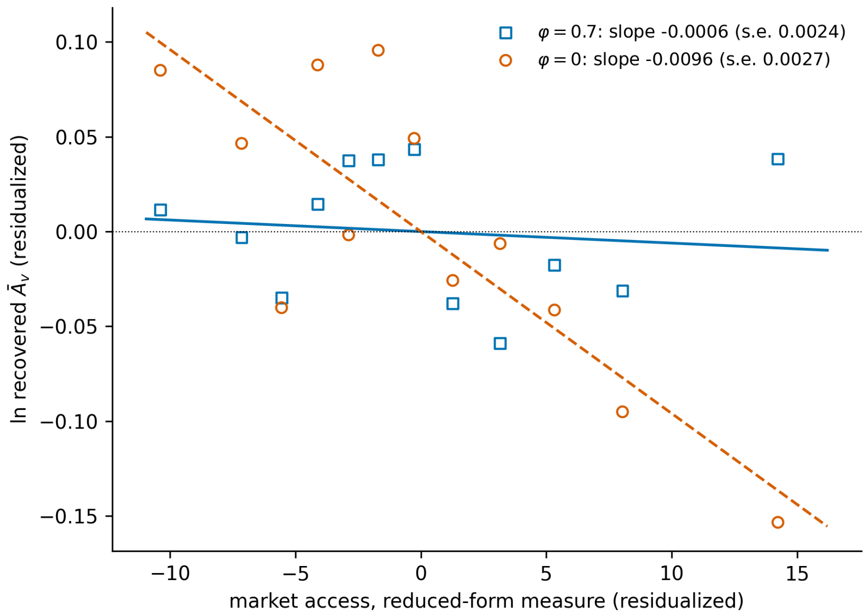

Note: The figure plots the recovered program-village agricultural productivity against market access, residualized on island and cohort fixed effects, in 12 quantile bins with the linear fit, with the slope and district-clustered standard error in the legend, with the externality ($\varphi = 0.7$) and without it ($\varphi = 0$).

9 Counterfactuals

This section uses the estimated model to evaluate policy. Every aggregate experiment re-solves the full joint general equilibrium under the counterfactual shock, so the baseline is the observed economy. The single-village experiments of Section 9.1 instead solve the treated village against the fixed rest of the economy, which is the appropriate frame for a shock to one of 69,225 locations. Fundamentals are held fixed at their baseline inverted values throughout, and every experiment is also solved at $\varphi = 0$, in the externality-free world re-inverted to the same baseline, to quantify what the externality changes.

9.1 A Road to a Program Village

The first experiment improves a single transmigration village's travel times by 10 percent and nothing else, for each of the 1,147 villages in turn. With or without the externality, improved connectivity leads to reallocation out of agriculture and into rural services. With the externality in place, improving the connectivity of transmigration villages *lowers* the improved village's real income by 2.5 percent on average, with 60 percent of villages made worse off, while without the externality every village is made better off (Figure 6a). Table 12 decomposes the effect by channel. In both economies villages benefit from cheaper access to goods, but the increase in external demand for services is the largest channel, and it pulls workers out of agriculture regardless of whether there is an externality. With the externality, the erosion of agricultural productivity that comes with that reallocation outweighs the gains from this and every other channel. Observing structural transformation is therefore not a sufficient condition for improved welfare.

Table 12: The targeted road, by channel and by world

	$\varphi = 0.7$ (measured)		$\varphi = 0$	
Channel	Δ ln real income	Δ agric. share	Δ ln real income	Δ agric. share
service-demand access alone	−4.2%	−4.3 pp	+0.5%	−2.1 pp
import competition, cheaper services	+1.0%	+0.6 pp	+0.4%	+0.4 pp
farm-gate price	+0.3%	+0.2 pp	+0.1%	+0.1 pp
urban-goods price	+0.5%	0.0 pp	+0.5%	0.0 pp
full shock	−2.5%	−3.5 pp	+1.4%	−1.6 pp
share of villages losing	60%		0%	

Note: Population-weighted means across the 1,147 single-village experiments. Each experiment cuts one village's excess travel times by 10 percent while holding all other locations fixed. Each channel row applies the same shock to one trade friction at a time, holding the others at baseline. The $\varphi = 0$ columns run the identical experiments in the externality-free world. Figure 6a plots the distribution of effects across villages.

This experiment is the model's answer to the evidence on village roads. The estimates of Asher and Novosad (2020) for India's rural road construction program show a variant of this pattern, workers exiting agriculture with no rise in incomes. The model reproduces the pattern of workers leaving agriculture with no gains in income (and in this experiment often losses), and supplies a mechanism that can explain the results: the road shifts individuals out of agriculture and toward service production, but the uncompensated externality translates this reallocation into a loss in income. The same experiment also answers a design question about the program. On average, the program placed its villages in isolated areas, with the median transmigration village sitting below the twentieth percentile of the access distribution. Improving every program village's travel times by 10 percent at once lowers the settlers' real income by 1.8 percent with the externality and raises it by 1.3 percent without it, with a national effect near zero either way. Under the externality, the program's haphazard remoteness sheltered the settlements rather than harming them.

Which villages a road hurts is the model's most testable prediction. Figure 6b plots the change in real income from the single-village road against each village's agricultural employment share before the shock. Without the externality every village gains between 1 and 2 percent.[14] With it, villages that were almost entirely agricultural, with shares above roughly 0.8, still gain from a connectivity shock, while nearly every village with a share between roughly 0.1 and 0.75 loses, and the loss is largest, close to 7 percent, near a share of 0.3; a handful of the least agricultural villages, with little agriculture left to erode, also gain. This is the pattern Proposition 1 predicts. A village that has barely begun structural transformation has little non-farm employment for the road to expand, so it keeps the gains from trade and escapes the externality drag, while a village

[14]The slightly larger gains for the least agricultural villages are due to the fact they are able to take advantage of the cheaper imports slightly more than more isolated, primarily agricultural, villages.

with a significant share of workers in both sectors is more responsive to the shock: it experiences a larger decrease in the agricultural employment share, and a larger loss in productivity through the externality. A village that has nearly completed the transition has little agriculture left for the externality to erode.

Figure 6: Effects of a Road Built to a Single Program Village

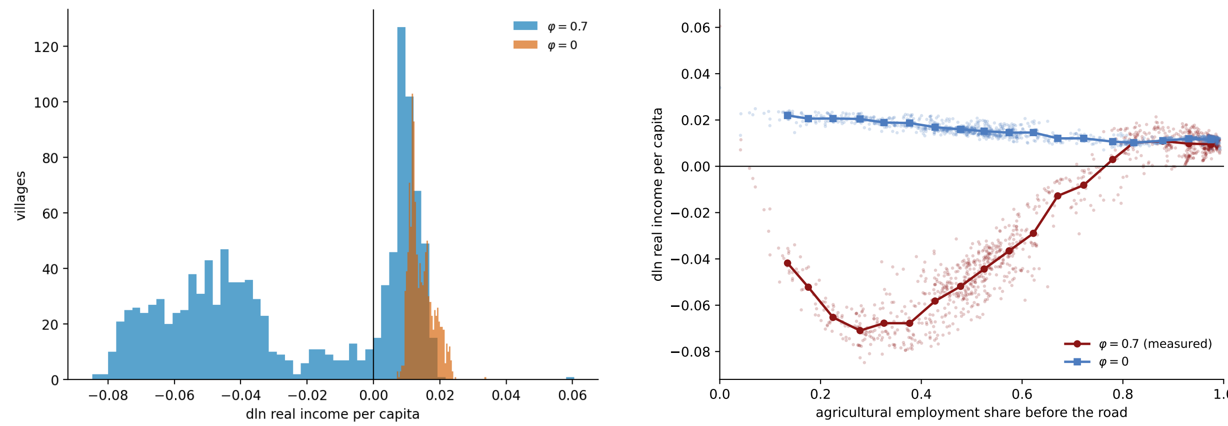

(a) Distribution of income effects

(b) By initial agricultural employment share

Note: Both panels plot the change in log real income per capita across the 1,147 single-village experiments, each of which cuts one village's excess travel times by 10 percent and re-solves that village against the unchanged rest of the economy, with the externality ($\varphi = 0.7$) and without it ($\varphi = 0$). Panel (a) plots the distribution of effects. Panel (b) plots each village's effect against its agricultural employment share before the shock, with lines connecting means within bins of width 0.05 that contain at least 8 villages.

9.2 Aggregate Programs

The remaining experiments improve the entire road network, subsidize agricultural earnings, and raise rural agricultural productivity, shocks of a size that move the aggregate agricultural employment share in the data. The model of Section 7 has homothetic demand, so the composition of national spending does not respond to shocks and even large shocks cannot generate *aggregate* structural transformation. That shuts down a response these policies should produce, and it is the response the externality attaches a cost to. For these experiments I therefore extend the model with the non-homothetic demand system of Boppart (2014), which Fan et al. (2023) use for India, under which households shift spending away from food and toward services and city goods as they grow richer.

Budget shares take the form $s_{g,\ell} = \omega_{g,\ell} + \nu_g \, x_\ell^{-\varepsilon}$, where x_ℓ is real expenditure per capita, $\omega_{g,\ell}$ is the share a household approaches as it grows rich, anchored so that each location reproduces its observed baseline shares, and the loading ν_g is positive for food and negative for the other goods; homothetic demand is the case $\nu_g = 0$. I estimate the Engel curves and the curvature ε with the estimator of Fan et al. (2023) on the same expenditure survey that measures the budget shares, and obtain $\varepsilon = 0.34$, nearly identical to their 0.33 for India (Appendix K.7). Because the system is anchored to each location's observed budget, the calibration and every result of Section 8 are unchanged; the income margin operates only when a counterfactual moves real incomes. On its own the extension does almost nothing: at $\varphi = 0$, switching demand systems changes the aggregate gains from the road and productivity experiments by about 1 percent of their size, since recomposing demand is close to costless without an externality. Whatever the income margin contributes is therefore the interaction of the externality with non-homothetic demand.

A uniform road improvement and an agricultural subsidy

A uniform 10 percent cut in all travel times makes every type of location better off, with or without the externality. What the externality changes is the size of the aggregate gain. Table 13 reports the results: national real income rises by 0.73 percent with the externality in place, against 0.86 percent without it, so ignoring the externality overstates the aggregate gain by 18 percent. The gain is smaller because, as real incomes rise, households in both urban and rural locations spend a smaller share of their budget on food, the relative decrease in demand for agricultural output pulls workers out of agriculture, and that lowers agricultural TFP through the externality; without the externality the same reallocation does not lower TFP. Under homothetic demand the same experiment gives 0.83 percent with the externality against 0.85 without, so nearly all of the difference operates through the income response, and almost none through the reallocation across space that the homothetic model already captures.[15] Appendix N.4 works out that maintaining a network improvement of this size would cost at least 0.68 percent of GDP a year, so the improvement is just worth the cost in the favorable case; if raising speeds instead requires other investment, the benefits come to well under half of what the work costs.

Given an externality, a natural question is what policies that target it directly do. Agricultural support programs are commonplace in developing countries, and the model provides a justification for how subsidizing agriculture can raise welfare. I consider two subsidies on agricultural earnings, each financed *within* the subsidized village by a lump-sum tax on its households. The first is the corrective subsidy that maximizes the transmigration villages' income, applied to each of them.[16] It raises their real income by 24.2 percent, and because the externality is then correctly priced, the relationship between market access and income among them turns positive and essentially equal to the $\varphi = 0$ slope of Table 11; but it lowers national income slightly (-0.06 percent), because the subsidized villages expand agriculture, push nearby villages out of it, and reduce their supply of their own service variety. A national planner has no reason to correct the externality for these 1,147 villages and nowhere else, so the second subsidy spreads the annual budget of the road program,

[15] Fixed budget shares do not rule out aggregate reallocation altogether: the national food share can still move with the rural-urban composition of income and with sorting, and the small gap between the two homothetic figures is that channel. What homotheticity removes is the within-location Engel response. Appendix Table N1 decomposes the road by channel and by group.

[16] The corrective rate is $s^* = (\gamma_a + \varphi)/\gamma_a - 1 \approx 2.3$.

0.68 percent of GDP, across every rural village, which comes to a subsidy rate of about 8 percent on rural agricultural earnings. That subsidy raises national income by 0.20 percent, with the gains in rural locations and a small loss in urban ones; the rural gains are several times what the road program generates, and net of the road's cost, which is not recouped, the subsidy delivers roughly four times the national gains. The gains do not survive without the externality: in the $\varphi = 0$ world the same subsidy is a distortion and generates a small loss. These results suggest a reading of agricultural support programs that differs from the usual one. They are usually defended as redistribution and criticized as distortions; in settings where agriculture carries scale economies they can also increase aggregate income.

Table 13: Roads and agricultural subsidies compared

	$\varphi = 0.7$						$\varphi = 0$
Instrument	Gross natl.	Cost	Net natl.	TM	Outer rural	Urban	Net natl.
uniform road network	+0.73	0.68	+0.05	+0.6%	+0.9%	+0.6%	+0.18
agricultural earnings subsidy, same budget	+0.20	≈ 0	+0.20	+2.2%	+1.9%	−0.5%	−0.01
memo: full corrective s^*, program villages only	−0.06	≈ 0	−0.06	+24.2%	−2.0%	−0.1%	—

Note: The gross, cost, and net columns are percent of national income per year; the TM, other rural, and urban columns are within-group percent changes in real income. Other rural is every non-program rural village, on the inner and outer islands. The road row is the uniform 10 percent cut in all travel times, solved under non-homothetic demand. The subsidy rows apply a proportional subsidy to rural agricultural earnings, financed within each subsidized village by a lump-sum tax on its own households. The cost column records real resource costs, which arise only for the road. The first two rows operate at a common fiscal outlay of 0.68 percent of GDP, and the memo row applies the full corrective wedge, which requires an outlay of 1.35 percent. The $\varphi = 0$ column re-solves each experiment in the re-inverted externality-free world.

Agricultural productivity growth

The externality also changes the aggregate effect of the force usually credited with producing structural transformation. I raise the agricultural TFP of every rural location by 5 percent. Table 14 shows that national income rises by 0.74 percent with the externality against 1.21 percent without it, so ignoring the externality overstates the gain by 64 percent. Two things account for the gap. Rural productivity growth shifts agricultural activity from urban locations to rural ones, urban agricultural employment falls by almost 2 percentage points, and the externality erodes the cities' agricultural TFP and income; this channel operates under either demand system, and with homothetic demand the gains are still roughly 25 percent smaller with the externality. The rest is the interaction with non-homothetic demand: as incomes rise, spending shifts away from food, the relative decrease in demand for food translates into more structural transformation, and the externality taxes it. Without the externality the choice of demand system has little effect on the aggregate gains (1.21 against 1.20).

Table 14: Rural agricultural productivity growth

Rural agricultural TFP +5%	TM	Outer rural	Urban	National
Non-homothetic demand				
measured externality ($\varphi = 0.7$)	+4.5	+4.5	−0.8	+0.74
externality off ($\varphi = 0$)	+3.9	+3.9	+0.2	+1.21
Homothetic demand				
measured externality ($\varphi = 0.7$)	+4.6	+4.7	−0.7	+0.89
externality off ($\varphi = 0$)	+4.1	+4.1	+0.1	+1.20

Note: Population-weighted changes in log real income, in percent, and the change in national real income, from a 5 percent rise in the agricultural TFP of every rural location. Other rural is every non-program rural village, on the inner and outer islands. The upper rows solve the model under non-homothetic demand and the lower rows under homothetic demand. The $\varphi = 0$ rows re-solve the identical experiment in the re-inverted externality-free world.

10 Conclusion

This paper uses the placement of Indonesia's Transmigration villages to study the long-run effects of market access. Villages that began more connected are poorer today and farm less productively, and the standard explanations, agricultural specialization, competition, out-migration, compensating amenities, and lower prices, do not account for the gap. What does account for it is a scale externality in agriculture. Connected villages moved more labor out of farming, and each worker who left lowered the productivity of those who stayed. I estimate the externality from program variation in settlement size, and a spatial general equilibrium model of the Indonesian economy that carries it reproduces the negative effect of market access on income among program villages without targeting it.

The model's counterfactuals show that the externality changes both who gains from connectivity and how much there is to gain. A road built to a single village can make that village poorer, and ignoring the externality overstates the aggregate gains from improving the whole travel network and from rural agricultural productivity growth, while a subsidy to agricultural earnings that would be a distortion without the externality raises national income with it.

More generally, this paper makes the point that structural transformation itself is not sufficient for generating growth. The large gaps in output per worker between agriculture and the rest of the economy that motivate much of the development literature are measured from private returns, and they overstate what reallocation delivers when the social return to agricultural labor exceeds the private one, as it does in these villages. Promoting the growth in the size of the non-agricultural sector is a common development goal, but that reallocation can lower incomes when the workers who leave agriculture decrease the productivity of the agricultural sector and also move into sectors with no (or smaller) scale externalities. The policy implication from this paper is not necessarily that policymakers should design programs to stop structural transformation, but that policies that promote externalities in the non-agricultural sector can allow for structural transformation and growth.

References

Alder, Simon, Timo Boppart, and Andreas Müller (2022) "A Theory of Structural Change That Can Fit the Data," *American Economic Journal: Macroeconomics*, 14 (2), 160–206.

Allen, Treb and Costas Arkolakis (2014) "Trade and the Topography of the Spatial Economy," *Quarterly Journal of Economics*, 129(3).

Anderson, T.W. and Herman Rubin (1949) "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 20(1), 46–63.

Andrews, Isaiah, James H. Stock, and Liyang Sun (2019) "Weak Instruments in Instrumental Variables Regression: Theory and Practice," *Annual Review of Economics*, 11, 727–753.

Asher, Sam and Paul Novosad (2020) "Rural Roads and Local Economic Development," *American Economic Review*, 110(3).

Bandiera, Oriana and Imran Rasul (2006) "Social Networks and Technology Adoption in Northern Mozambique," *Economic Journal*, 116 (514), 869–902.

Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare (2025) "The Textbook Case for Industrial Policy: Theory Meets Data," *Journal of Political Economy*, 133 (5), 1527–1573, 10.1086/734129.

Bassi, Vittorio, Raffaela Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume (2022) "Achieving Scale Collectively," *Econometrica*, 90 (6), 2937–2978.

Baum-Snow, Nathaniel, Vernon Henderson, Matthew Turner, Qinghua Zhang, and Loren Brandt (2020) "Does investment in national highways help or hurt hinterland city growth?" *Journal of Urban Economics*, 115(1).

Bazzi, Sam and Christopher Blattman (2014) "Economic Shocks and Conflict: Evidence from Commodity Prices," *American Economic Journal: Macroeconomics*, 6(4).

Bazzi, Sam, Arya Gaduh, Alex Rothenberg, and Maisy Wong (2016) "Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia," *American Economic Review*, 106(9).

——— (2019) "Unity in Diversity? How Intergroup Contact Can Foster Nation Building," *American Economic Review*, 109(11).

BenYishay, Ariel and A. Mushfiq Mobarak (2019) "Social Learning and Incentives for Experimentation and Communication," *Review of Economic Studies*, 86 (3), 976–1009.

Beyer, Robert, Yingyao Hu, and Jiaxiong Yao (2022) "Measuring Quarterly Economic Growth from Outer Space," *IMF Working Paper*, 109.

Bluhm, Richard and Gordon McCord (2022) "What Can We Learn from Nighttime Lights for Small Geographies? Measurement Errors and Heterogenous Elasticities," *Remote Sensing*, 14(5).

Boppart, Timo (2014) "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences," *Econometrica*, 82 (6), 2167–2196.

Borusyak, Kirill and Peter Hull (2023) "Non-Random Exposure to Exogenous Shocks," *Econometrica*, 91(6).

Breza, Emily, Supreet Kaur, and Yogita Shamdasani (2021) "Labor Rationing," *American Economic Review*, 111 (10), 3184–3224.

Brownstone, Steven (2025) "Labor Market Effects of Agricultural Mechanization: Experimental Evidence from India," *Working Paper*.

Bryan, Gharad and Melanie Morten (2019) "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127(5).

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli (2016) "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *American Economic Review*, 106(6).

Caunedo, Julieta, Namrata Kala, and Haimeng Hester Zhang (2026) "Economies of Density and Congestion in Equipment Rental Markets," *Working Paper*.

Chen, Chaoran, Diego Restuccia, and Raül Santaeulàlia-Llopis (2022) "The Effects of Land Markets on Resource Allocation and Agricultural Productivity," *Review of Economic Dynamics*, 45, 41–54.

Collier, Paul, Martina Kirchberger, and Måns Söderbom (2015) "The Cost of Road Infrastructure in Low and Middle Income Countries," *Policy Research Working Paper*.

Conley, Timothy G. (1999) "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 92 (1), 1–45, 10.1016/S0304-4076(98)00084-0.

Conley, Timothy G. and Christopher R. Udry (2010) "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, 100 (1), 35–69.

Couture, Victor (2016) "Valuing the Consumption Benefits of Urban Density," *Working Paper*.

Deaton, Angus (1988) "Quality, Quantity, and Spatial Variation of Price," *American Economic Review*, 78 (3), 418–430.

Diamond, Peter A. (1982) "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90 (5), 881–894.

Diegert, Paul, Matthew Masten, and Alexandre Poirier (2025) "Assessing Omitted Variable Bias when the Controls are Endogenous," *Working Paper*.

Donaldson, Dave and Richard Hornbeck (2016) "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 131(2).

Duranton, Gilles and Diego Puga (2004) "Micro-Foundations of Urban Agglomeration Economies," in Henderson, J. Vernon and Jacques-François Thisse eds. *Handbook of Regional and Urban Economics*, 4, Chap. 48, 2063–2117: Elsevier.

Faber, Ben (2014) "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System," *Review of Economic Studies*, 81(3).

Fan, Tianyu, Michael Peters, and Fabrizio Zilibotti (2023) "Growing Like India–the Unequal Effects of Service-Led Growth," *Econometrica*, 91 (4), 1457–1494.

Fogel, Robert (1964) *Railroads and American Economic Growth: Essays in Econometric History*, Baltimore, MD: Johns Hopkins University Press.

Foster, Andrew D. and Mark R. Rosenzweig (1995) "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 103 (6), 1176–1209.

Fuglie, Keith (2010) "Sources of Growth in Indonesian Agriculture," *Journal of Productivity Analysis*, 33(3).

Gabaix, Xavier and Rustam Ibragimov (2011) "Rank $-$ 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business & Economic Statistics*, 29 (1), 24–39.

Gebresilasse, Mesay (2023) "Rural Roads, Agricultural Extension, and Productivity," *Journal of Development Economics*, 162, 103048.

Gertler, Paul, Tadeja Gracner, Marco Gonzalez-Navarro, and Alex Rothenberg (2024) "Road Maintenance and Local Economic Development: Evidence from Indonesia's Highways," *Journal of Urban Economics*, 143.

Gervais, Antoine and J. Bradford Jensen (2019) "The Tradability of Services: Geographic Concentration and Trade Costs," *Journal of International Economics*, 118.

Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li (2021) "Which Night Lights Data Should We Use in Economics, and Where?" *Journal of Development Economics*, 149.

Gollin, Douglas, David Lagakos, and Michael E. Waugh (2014) "Agricultural Productivity Differences across Countries," *American Economic Review*, 104(5).

Hanson, Arthur (1981) *Transmigration and Marginal Land Development*: Westview Press, Boulder, Colorado, USA.

Hardjono, Joan (1988) "The Indonesian Transmigration Program in Historical Perspective," *International Migration*, 26(4).

Henderson, J. Vernon, Adam Storeygard, and David Weil (2012) "Measuring Economic Growth from Outer Space," *American Economic Review*, 102(2).

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi (2013) "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103 (7), 2752–2789.

Hornbeck, Richard and Martin Rotemberg (2024) "Growth Off the Rails: Aggregate Productivity Growth in Distorted Economies," *Journal of Political Economy*, Forthcoming.

Hu, Yingyao and Jiaxiong Yao (2022) "Illuminating Economic Growth," *Journal of Econometrics*, 228(2), 359–378.

Kaboski, Joseph, Will Lu, Wei Qian, and Lixia Ren (2024) "Melitz Meets Lewis: The Impacts of Roads on Structural Transformation and Businesses," *Working Paper*.

Kebschull, Dietrich (1986) *Transmigration in Indonesia: An Empirical Analysis of Motivation, Expectations and Experiences*, Hamburg, Germany: Transaction Publishers.

Kerwin, Jason, Nada Rostom, and Olivier Sterck (2024) "Striking the Right Balance: Why Standard Balance Tests Over-Reject the Null, and How to Fix It," *Working Paper*.

Krugman, Paul (1991) "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99 (3), 483–499, 10.1086/261763.

La Porta, Rafael and Andrei Shleifer (2014) "Informality and Development," *Journal of Economic Perspectives*, 28 (3), 109–126.

Lagakos, David (2016) "Explaining Cross-Country Productivity Differences in Retail Trade," *Journal of Political Economy*, 124 (2).

Lewis, W. Arthur (1954) "Economic Development with Unlimited Supplies of Labour," *The Manchester School*, 22 (2), 139–191.

Li, Xuecao, Yuyu Zhou, Min Zhao, and Xia Zhao (2020) "A harmonized global nighttime light dataset 1992–2018," *Scientific data 7*, 1.

Matsuyama, Kiminori (1992) "Agricultural productivity, comparative advantage, and economic growth," *Journal of Economic Theory*, 58(2).

Montiel Olea, Josè Luis and Carolin Pflueger (2013) "A Robust Test for Weak Instruments," *Journal of Business and Economic Statistics*, 31(3), 358–369.

Morten, Melanie and Jaqueline Oliveira (2024) "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," *American Economic Journal: Applied Economics*, 16(2).

Oster, Emily (2019) "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, 37.

Ostrom, Elinor and Roy Gardner (1993) "Coping with Asymmetries in the Commons: Self-Governing Irrigation Systems Can Work," *Journal of Economic Perspectives*, 7(4).

Peters, Michael, Youdan Zhang, and Fabrizio Zilibotti (2026) "Skipping the Factory: Service-Led Growth and Structural Transformation in the Developing World," *Working Paper*.

Pingali, Prabhu (2007) "Agricultural Mechanization: Adoption Patterns and Economic Impact," in Evenson, Robert and Prabhu Pingali eds. *Handbook of Agricultural Economics*, 3, Chap. 54, 2779–2805: Elsevier.

Ramondo, Natalia, Andrés Rodríguez-Clare, and Milagro Saborío-Rodríguez (2016) "Trade, Domestic Frictions, and Scale Effects," *American Economic Review*, 106 (10).

Ranis, Gustav and John C. H. Fei (1961) "A Theory of Economic Development," *American Economic Review*, 51 (4), 533–565.

Redding, Stephen and Daniel Sturm (2008) "The Costs of Remoteness: Evidence from German Division and Reunification," *American Economic Review*, 98(5).

Redding, Stephen and Anthony Venables (2004) "Economic Geography and International Inequality," *Journal of International Economics*, 62.

Restuccia, Diego and Raül Santaeulàlia-Llopis (2023) "Land Misallocation and Productivity," *American Economic Journal: Macroeconomics*, 15 (2), 441–465.

Rodrik, Dani (2016) "Premature deindustrialization," *Journal of Economic Growth*, 21.

Rossi-Hansberg, Esteban and Jialing Zhang (2025) "Local GDP Estimates Around the World," *Working Paper*.

Schultz, Theodore W. (1964) *Transforming Traditional Agriculture*, New Haven: Yale University Press.

Whitten, Anthony (1987) "Indonesia's Transmigration Program and Its Role in the Loss of Tropical Rain Forests," *Conservation Biology*, 1(3).

World Bank (1988) *Indonesia: The Transmigration Program in Perspective*, World Bank Country Study, Washington, DC: The World Bank.

Appendix for
"The Curse of Connectivity"

Contents of the Appendix

A Background

A.1 Direct Qualitative Evidence of Program Implementation

This section presents a collection of quotes from a variety of sources that describe the chaotic nature of the transmigration site selection process during my period of study. Since the transmigration program was divided into different five-year plan periods (Repelita) over this period, I document each period separately as well:

Repelita I-II (1969-1979) and pre-periods

There was transmigration activity observed in the transmigration census even before Repelita I, although limited. In particular, there was little government oversight over transmigration site selection during this period:

- "During the first part of the 1960s the government introduced a policy of permitting mass organizations and political parties to participate in transmigration... Although it was hoped that this policy of non-government involvement in the program would enable greater numbers of people to be moved from Java and that the financial and administrative burden upon the government could thus be reduced, the organizations themselves were much more interested in the political advantages that they could gain from transmigration than in the establishment of economically viable settlements. Transmigrants were simply transported to project locations, most of which were badly selected." Hardjono (1988)

Repelita I (1969-1974) was primarily a planning period for the expansion of the transmigration program. Few sites were actually settled during this period. Activity began to ramp up during Repelita II (1974-1979), along with documented instances of the lack of sufficient care in selecting sites. Hanson (1981) reports that transmigration officials during this period had "limited management skills," while followup research on transmigration sites selected during this period revealed that they were frequently chosen in a haphazard manner (Hardjono, 1988). Repelita II was also when reports emerged that unrealistic resettlement targets set by government officials meant that individuals in charge of resettlement lacked sufficient time or incentives to deeply consider potential new village locations:

- "... the emphasis on reaching the [resettlement] targets rode roughshod over moderation and caution, and over the recommendations of planning guidelines and manuals." Whitten (1987)

Repelita III (1979-1984)

By all accounts this was the most chaotic period during the transmigration program. It also was the Repelita with the most individuals resettled. Reports on the program reference multiple instances of the government issuing statements encouraging the acceleration of the resettlement process, particularly between 1981 and 1982. Whitten (1987) additionally reports that the information bureaucrats had available to them at the time to help inform site selection decisions was "often based on poor resource data and maps." Hardjono (1988)'s research on this period provides a perfect summary:

- "Despite occasional references in government statements to regional development and the welfare of transmigrants, the aim was once again to move as many people as possible"
- "As a consequence of the focus on numbers, the land-use plans developed during the 1970s were totally abandoned. Transmigrants were placed on whatever land was submitted by provincial governments for settlement purposes."
- "To sum it up briefly, during Repelita III quality was sacrificed to quantity."

Repelita IV (1984-1989)

While there was limited transmigration activity following Repelita IV, the global decline in oil price and other political concerns led to a drastic reduction in transmigration activity following Repelita IV (Bazzi et al., 2016). Although the transmigration authority attempted to develop new, more accurate land use maps to aid in the selection process of new sites, this only began during the end of this period, and progress was slow (Whitten, 1987). Despite the reported issues that came with setting unrealistic resettlement targets during Repelita III, the government actually increased the relocation target by 50 percent relative to Repelita III[17]. In addition to the haphazard placement process described above in Repelita III, there were additional reports about inter-agency coordination issues that increased the chaotic nature of the program:

- "Due to the rapid expansion [of the transmigration program], a number of program activities were taken from the Directorate General of Transmigration (DGT) and delegated to separate government agencies to speed up the settlement process. Inter-agency coordination problems made it more difficult to carefully match transmigrants... to their outer island settlements." Bazzi et al. (2016)

Two things should be said about what this record establishes. It documents that site selection was rushed, poorly informed, and driven by resettlement targets rather than economic potential, which makes systematic selection of sites on their productive fundamentals implausible. It cannot by itself rule out subtler channels, such as provincial governments choosing which land to submit for settlement; the empirical design of Section 4, with its recentering, balance tests, and controls, attempts to address this. The qualitative record is best read as motivating rather than proving the identification assumptions.

[17]The target was 500,000 households in Repelita III, and 750,000 households in Repelita IV. Neither target was actually reached.

A.2 Settlers, Program Activity, and Village Characteristics over Time

Figure A1 reports the savings and landholdings of transmigrant households before migration from the survey of Kebschull (1986). Figure A2 plots the annual number of transmigrants and the average settlement size against the world oil price. Table A1 asks whether the villages settled in different years differ in their predetermined characteristics. Most of the underlying geographic characteristics of villages show no trend across creation years, though a handful of variables, including some crop-potential measures and the size of settlements, do display statistically significant linear trends, on the order of 0.4 to 0.7 standard deviations over the four decades of the program. These trends are one motivation for the cohort fixed effects used throughout the paper.

Figure A1: Survey Measures of Migrant Characteristics Prior to Relocation

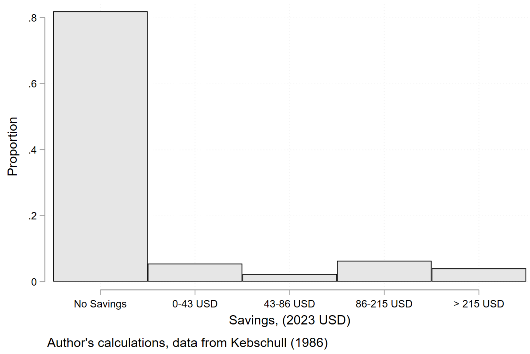

(a) Migrant Savings Immediately Before Migration

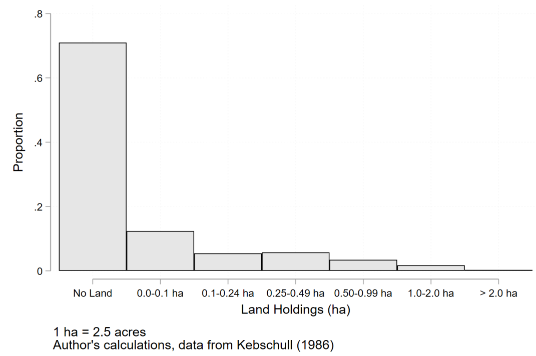

(b) Migrant Landholdings Before Migration

Note: Both panels present results from a pre-transmigration survey of transmigrant households by Kebschull (1986), digitized by the author. Panel (a) plots the distribution of savings immediately before transmigration while panel (b) plots the distribution of household landholdings prior to transmigration.

Figure A2: Transmigration and Oil Prices

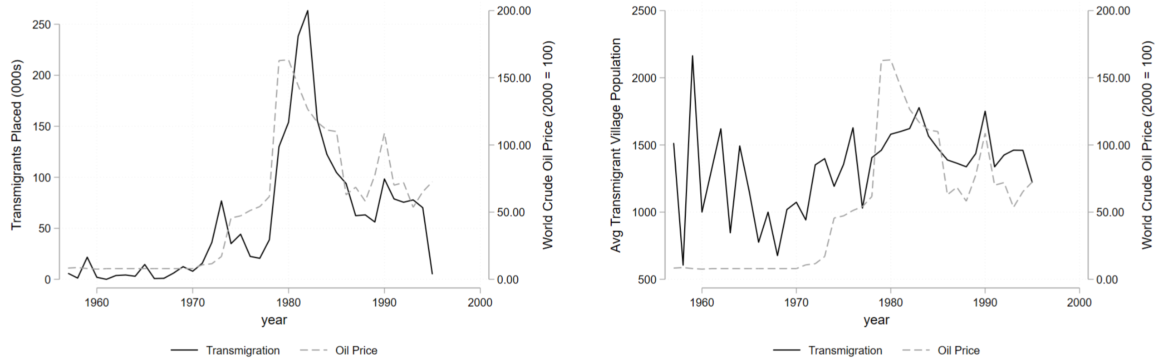

(a) Total Transmigration by Year

(b) Average Transmigration Village Size by Year

Note: Both panels include world crude oil prices from Bazzi and Blattman (2014). Panel (a) also plots the total number of individuals resettled as a result of the transmigration program annually. Panel (b) also plots the average settled transmigration village population after winsorizing village population measures at the 1st and 99th percentiles.

Table A1: Village Characteristics Across Creation Time

Dependent variable (std): Year Created	
Wetland Rice Potential Yield	-0.011** (0.005)
Palm Oil Potential Yield	-0.008 (0.006)
Coffee Potential Yield	-0.010** (0.004)
Cassava Potential Yield	-0.013** (0.005)
Vector Ruggedness Measure	0.010 (0.006)
Individuals Placed	-0.016* (0.009)
Agrosimilarity Score	-0.008 (0.007)
Share of Land with Poor Drainage	-0.001 (0.008)
Share of Land with Good Drainage	-0.002 (0.007)
Organic Carbon (%)	0.007 (0.008)
Topsoil Salinity (Elco) (dS/m)	-0.004 (0.003)
Topsoil Base Saturation (%)	-0.007 (0.006)
Avg. rainfall, 1948–1978	0.008 (0.006)
Avg. temp (Celsius), 1948–1978	0.007 (0.009)
ln(District Population, 1980)	0.001 (0.005)
District Own Electricity, 1980	0.003 (0.005)
District Own Radio, 1980	-0.005 (0.005)
District Own TV, 1980	0.004 (0.005)

Dependent variable (std): Year Created	
District Literacy, 1980	-0.002 (0.005)
District Avg Years Schooling, 1980	-0.005 (0.005)
District Agricultural Employment Share, 1980	-0.005 (0.005)
District Wetland Rice Potential Yield, 1980	-0.006 (0.005)
District Palmoil Potential Yield, 1980	0.005 (0.005)
District Coffee Potential Yield, 1980	0.002 (0.006)
District Cassava Potential Yield, 1980	-0.005 (0.005)
District Mining Employment Share, 1980	-0.000 (0.004)
District Manufacturing Employment Share, 1980	0.005 (0.004)
District Trade Employment Share, 1980	0.003 (0.005)
District Service Employment Share, 1980	0.004 (0.005)
District Wage Worker Share, 1980	0.009* (0.005)
District Share HHs with Piped Water, 1980	-0.000 (0.004)
District Share HHs with Sewage, 1980	0.001 (0.005)
District Share HHs with Modern Fuel, 1980	0.004 (0.005)
District Share HHs with Modern Floor, 1980	0.002 (0.005)
District Share HHs with Modern Roof, 1980	-0.006 (0.005)

F-stat of omnibus test	2.319
Observations	1219

Note: Each row reports the coefficient from a regression of the standardized dependent variable listed in the row on the year a transmigration village was created, controlling for island fixed effects. Standard errors, clustered at the subdistrict (kecamatan) level, are reported in parentheses. The omnibus F-statistic tests joint differences across village creation year for all dependent variables, conditional on island fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Data Appendix: Sources and Construction

This appendix provides detailed descriptions of the data sources summarized in Table B1 and used throughout the analysis, followed by the digitization of the road network and the construction of travel times.

B.1 Local Measures of GDP per Capita

My primary measure of local GDP per capita is the high-resolution estimates of Rossi-Hansberg and Zhang (2025), which combine a diverse set of inputs beyond population and nightlights and perform significantly better at predicting local GDP than previous measures. Figure B4a shows that the measure varies substantially across transmigration villages. The alternative measure combines nightlights from Li et al. (2020), converted to GDP with the elasticities in Beyer et al. (2022), Hu and Yao (2022), and Henderson et al. (2012), with population from the Global Human Settlement Layer. It is subject to some of the issues raised in Rossi-Hansberg and Zhang (2025) but is built from fewer inputs. Throughout my analysis I classify village areas according to their administrative boundaries in 2000.

B.2 Measures of Village Level Productivity

To supplement the satellite-derived measures of GDP per capita above, I construct several proxies for village level productivity across multiple sectors. I measure agricultural productivity directly using village level data from the 2000 and 2003 waves of PODES on agricultural area devoted to different crops along with reported per-crop yields. The top-five-crop productivity measure is the total production of the five most planted crops in the village divided by the total area planted to those crops (tonnes per hectare, in logs), an area-weighted average of the crop-specific yields; because the identity of the top crops can differ across villages, the crop-level analysis of Section 6, which conditions on crop fixed effects, is the specification that isolates within-crop variation. I construct the same measure for rice and palm oil separately given their importance in rural agricultural economies. Figure B4d shows that village agricultural productivity is very heterogeneous. I also measure the number of tractors per capita and tractors per agricultural worker as a proxy for uptake of more productive agricultural practices. Outside of agriculture I use data from two waves (2000 and 2010) of the Indonesia Population Census where working individuals report if they are an employee or employer to construct the ratio of employees per employer. A lower ratio typically implies less productive firms given their smaller average size.

I complement these measures with individual-level wage data from the 2000 to 2007 waves of the SUSENAS, which each year samples individuals within a random set of villages across Indonesia. No transmigration village is guaranteed to be sampled in a given year, but pooling the waves gives wages for formally employed individuals in roughly one third of my villages, which I convert to their 2007 equivalents. Because I observe wages only in this subset of villages and only for

Table B1: Data Description

Sample	Description	Level of Observation	Year(s)	Outcomes
Rossi-Hansberg and Zhang (2025) (Global)	Measures of local GDP per capita from Rossi-Hansberg and Zhang (2025) and population per cell	Village	2012–2019	GDP per capita
Other Satellite Data (Global)	Nightlights harmonized across satellites from Li et al. (2020)	Village	1975–2020	Population, nightlight intensity (combined to construct an alternative GDP per capita measure)
Population Census (Complete Count)	Population census carried out every 10 years that collects basic demographic and socio-economic information	Individual	2000, 2010	(All aggregated to village level) sectoral employment shares, employment status/type, province of birth and province of residence 5 years prior
PODES (∼1148 Village Heads)	Survey of village heads that collects information on a wide variety of economic and social measures of the village	Village	1993, 2000, 2003	Sectoral employment shares (1993); agricultural productivity and tractors (2000); amenities/public-goods index (2000); migration and commuting patterns (2003); small-industry and shop counts (2003); land-use patterns and changes (2000 and 2003); village-by-crop production, area, and yield (2003)
SUSENAS – Prices (∼1234 Households)	Detailed consumption and expenditure module administered at the household level	Household	2002, 2011 (food only)	Quantity and unit price of food items consumed, individual-expenditure-weighted and rural-reference-weighted food price levels, total expenditure, expenditure by non-food category, housing value per square meter
SUSENAS – Wages (∼5781 Individuals)	General socio-demographic survey administered annually	Individual	2000–2007	(For formally employed only) monthly labor earnings, hours worked
Economic Census (Complete Count)	Complete listing of the universe of medium and large firms across Indonesia (plus a sample of micro and small firms in urban areas), with limited firm-level information	Firm	2016	Medium/large firm presence, total employment within the firm
Video-Derived Welfare (111 videos / 47 villages)	Publicly available Indonesian-language documentary and vlog footage of transmigration villages in Kalimantan, with frames scored by a vision language model against a fixed rubric	Village	2021-2022	Composite material-welfare index across seven dimensions of visible welfare (independent validation outcome)
Transmigration Census (Administrative)	Administrative records identifying settled transmigration villages, together with designated Planned Development Areas used to construct expected market access	Village	1952–1995	Village location, year settled, and number of individuals initially settled (instrument for agricultural labor in Section 6)
Geographic Characteristics (per Bazzi et al. (2016))	Time-invariant geographic and agro-ecological characteristics of village locations prior to settlement	Village	Time-invariant	Soil quality, terrain ruggedness, rainfall, temperature, agro-ecological similarity, and location
1980 Population Census (District-level)	District-level characteristics roughly contemporaneous with the transmigration program, used as controls	District	1980	Sectoral employment composition, formal employment rate, housing quality, agricultural potential yield, education and literacy, asset ownership
Historical Road Network (Army Map Service; HOT/OSM)	Hand-digitized maps of Indonesia's pre-program road network (roads by type, railways, trails) compiled by the Army Map Service Corps of Engineers, combined with navigable waterways from the Humanitarian OpenStreetMap Team	Raster / network	∼1950s	Travel-cost layer used to construct bilateral trade costs and pre-transmigration market access
IFLS (1,395 Village Heads)	Survey of village heads that includes a module on travel times to nearby locations along with method of travel and estimated trip time	Village	1993	Travel time by road type (used to calibrate trade-cost parameters)

formal workers, these estimates have less power than the GDP per capita results, but they provide contextual backing to the more abstract satellite measures (Appendix D).

The village-by-crop production data from the 2003 wave of PODES also plays a central role in my estimation of the agricultural externality in Section 6. For each village, PODES reports the production quantity, area planted, and yield for each crop grown in the most recent season, which I combine with census data on agricultural employment to construct the village-by-crop panel used in the instrumental variables estimation.

B.3 Video-Derived Welfare Measure

To complement the satellite and survey-based measures of economic performance, I construct a measure of village material well-being from publicly available video footage of transmigration villages in Kalimantan. I match 111 Indonesian-language documentary and vlog videos from a creator's YouTube playlist to 47 distinct villages in my main analysis sample using location information contained in video titles and captions. Frames sampled from each video are scored by a vision language model against a fixed rubric covering 7 dimensions of visible material welfare, and the frame-level scores are aggregated into a village-level composite welfare index. Because the sample is small and restricted to Kalimantan, this measure serves as a complementary validation of my primary outcomes from an entirely independent data source rather than as a primary outcome. Appendix E provides full detail on the sample construction, scoring rubric, and aggregation procedure.

B.4 Sectoral Employment

Sectoral employment shares come from the 2000 and 2010 waves of the Population Census, which record every individual's demographic information, education, employment status, and sector of employment. Figures B4b and B4c plot the distribution of agricultural and formal employment for the transmigration villages in 2000. While most villages are still primarily agricultural, there is a good deal of variation across transmigration villages, and the most common employment type is informal employment. I supplement these census measures with village-level data on sectoral employment from the 1993 wave of PODES.

B.5 Firm Dynamics

Village-level measures of formal firm presence and employment come from the listing of all medium and large enterprises in the 2016 Economic Census, which also contains a large sample of micro and small enterprises in each urban location and records the owner of the firm, a detailed sector classification, and employment. I collapse this listing to the village level to construct an indicator for whether a village has a medium or large firm and the average employment of those firms. Only 13 percent of transmigration villages have at least one such firm, so I also use the PODES waves, which are administered to village heads, to construct the number and density of small industries and the number of shops in a village. These measures give a more accurate picture of the smaller scale production and trade present in transmigration villages.

B.6 Local Prices and Expenditure

Prices and expenditure come from the 2002 and 2011 waves of the SUSENAS, which contain detailed consumption modules. In both years I construct a price index for food expenditure,[18] and in 2002 I can also construct total expenditure on non-food categories and the price per square meter of housing. Because only these two waves separate prices from quantities, the sample is limited; I observe the food price index for 84 villages and 1,234 households and non-food expenditure for 70 villages and 1,086 households. Other waves of the SUSENAS record only coarser measures of total consumption expenditure and do not separate prices and quantities, so I cannot construct prices for a larger sample.

B.7 Amenities

The public goods index measures each transmigration village's provision of public goods by local village and regional governments, covering health care availability and quality, school presence and staffing, safe drinking water, waste collection and other sanitation, and road maintenance and quality. I standardize each of the 18 measures and take the average.

B.8 Migration and Commuting

Because the census asks each individual their birth location, I construct the two village-level migration variables used in Section 5, the log number and the share of current residents born on the inner islands (Java and Bali), the sending region of the program. Given the nature of transmigration villages, these shares are often large. I measure village-level commuting from the 2003 PODES, which asks whether individuals in the village commute outside it for work. The survey does not record how many individuals commute, but it does indicate the prevalence of commuting at the village level.

B.9 Land Use and Development

Land use comes from the 2000 and 2003 waves of PODES, which record the amount of land in a village devoted to agriculture, housing, and firms. The 2003 PODES also records *changes* in land use over the previous 3 years, so I can compare land allocations across villages in levels and in changes.

[18] In my sample nearly two thirds of total expenditure is devoted to food consumption.

B.10 Time-Invariant Characteristics

The time-invariant characteristics of transmigration villages are the geographic variables of Bazzi et al. (2016) that describe a village's soil quality, terrain, and location, all measured *before* settlement. The Transmigration Census identifies the transmigration villages and records their location, the year they were settled, and the number of individuals relocated to each, which gives the settlement cohorts and, in Section 6, the instrument.

District characteristics roughly contemporaneous with the program come from the 1980 Population Census, from which I build the sectoral employment composition, formal employment rates, housing quality, agricultural potential yield, education and literacy, and asset ownership rates of the district each village was placed in.

B.11 Map Digitization and Travel Times

See main text for the location of the maps, and Figure B1[19] for an example of the un-digitized maps. These maps were created by the US Army Corps of Engineers and feature extremely detailed road network information for most of the world circa 1950. Roads are classified as either (from most navigable to least): Main road, secondary road, other road, or trail. The maps also mark all active and under-construction railways. There is additional information on settlement location, waterways, and topographic detail that I do not digitize. The maps are available via pdf or png format and are split into separate blocks. To cover all of Indonesia requires digitizing 27 separate images. The key used to identify different roadways is not amenable to automatic digitization; for instance, both the main road and secondary road are solid red lines of the same color, and the only way to distinguish them is based on small differences in the thickness of the lines. Trails and other roads are also red, but dashed. I was unable to determine how to separately classify each roadway type and connect the network in an automatic way. Additional complications include the fact that railways are the same color as location names and boundaries, and text often obscures sections of roadways that are easily interpolated by human inspection, but cause issues when trying to automate the process.

To digitize these maps, I downloaded a copy of each map segment for each network type (27 images X 5 network types = 135 files) so that I could separately identify each network in the least cluttered way possible. I then hand-traced each network type for the entire landmass of Indonesia using a distinct color[20] so that I could later distinguish each type of network clearly. After this I used the coordinate system on the edges of the maps to geo-reference each image so that I could stitch together the 27 images to make up a unified image of Indonesia. Based on the recorded color value of each network type, I was then able to easily extract the paths from each of the 5 unified

[19] Due to limitations on the image size in Overleaf, the included image is more blurry than the actual images I digitized.

[20] I loaded the images in Goodnotes and traced the paths on an iPad, ensuring that roads connected across different image files when needed.

maps of Indonesia (by network type) and record them as separate raster files. I then overlapped all five layers together to generate a single raster layer that included the full, detailed, transportation network of Indonesia. Figure B2 shows the resulting layer for Sumatra.

To estimate the travel times between origin o and all destinations d I implemented a fast marching method in Matlab as described by Allen and Arkolakis (2014) where each grid point is assigned a travel cost based on the network type. Network types are in turn identified by the color of the raster at that particular pixel/grid point. In instances where two lines overlapped (due to intersections of multiple road types), I assigned those pixels to the fastest transportation method. The value of each grid point is meant to represent the cost of traveling over that point. Based on the projection and the map details, each grid point in my network was calculated to be roughly 730m. Therefore, I conducted a simple conversion to convert either a cost per mile (from Donaldson and Hornbeck (2016)) or speed (from my own estimations from the IFLS 1993) to the pixel cost. Figure B3 displays the travel cost parameters. The speed parameters come from my own calculations based on the Indonesia Family Life Survey 1993 Community Survey. In that survey a village head is asked details about the (i) distance to nearby cities (up to 3 distinct cities) (ii) the method of travel and (iii) the time it takes to arrive to that location. I then split the travel types into travel via car (roads), boats (water) and train (railway) and separately estimated their respective speeds. Due to limited observations of travel via train, I instead inputted the speed estimate from the linked source; however, my own estimate of the speed was only slightly slower (roughly 45 km/hr) than the final parameter used. The cost parameters come straight from Donaldson and Hornbeck (2016) except include the inclusion of travel over roads in automobiles, which I assumed to be slightly cheaper than water and train transport for main highways, but slightly more expensive than those two methods for more minor roads.

After calculating the estimated travel times between each transmigration village centroid and all other locations (either a city "block" or village "desa") centroids, I then estimated a transmigration village's initial market access using Equation 1 and the 1975 GHS Layer population estimates. Table B2 compares the GHS layers with Indonesian sources.

Figure B1: Northern Sumatra, Example Map

Figure B2: Digitized Road Network - Sumatra

Note: This figure displays a subset of the travel network as it is input into Matlab to construct travel times and market access. Blue represents ocean and navigable waterways, green is land that has neither water nor any road type. Light blue are railways. Each other color represents different road types from highways to unpaved trails.

Figure B3: Travel Cost Parameters

Type	Code (in QGIS)	Cost (DH) ($/km)	Cost Transform Pixel	Speed Est IFLS (km/hr)	Speed Transform Pixel	Notes
Wilderness (off road)	0	0.3	218.0430107	3.21869	13.54855613	scale down walking speed to 2.5 mph
Main	1	0.004	2.907240142	24.41258	1.786316814	From IFLS 1993
Secondary	2	0.01	7.268100356	24.41258	1.786316814	cont. not separating by road type
Other	3	0.01	7.268100356	24.41258	1.786316814	
Trail	4	0.231	167.8931182	4.82803	9.032380109	assume avg walking speed of 3 mph
Rail	5	0.0063	4.578903224	51.499	0.8467854159	from https://transportgeography.org/conte
Water	6	0.0049	3.561369174	15.8953	2.743490348	From IFLS 1993
Ocean	10	0.01	7.268100356	15.8953	2.743490348	same as water
			Cost*726.8100356		(Speed*(1hr/60min)*(1000m/1km))^-1 * 726.8100356m	
					How many minutes will it take	
					you to travel over the pixel	

Table B2: Validation of GHS Population Measure

	Total Population	Total Population	Total Population	Total Population
Individuals Placed	0.8483	1.2171		
	(0.7184)	(1.0968)		
2010 Census Pop			0.9060***	0.9103***
			(0.1009)	(0.1040)
FE Level	None	Year and Island	None	Year and Island
N	176	176	1015	1015
Unique Villages	176	176	1015	1015

Standard errors in parentheses

Standard errors clustered at the village level

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: This table displays the results from a regression of either the number of individuals placed according to the transmigration census on the GHS predicted village level population in 1975 for transmigration villages created between 1975-1979 (columns 1-2) or the number of individuals within a village from the 2010 population census on the GHS predicted village level population in 2010 (columns 3-4).

Figure B4: Distribution of Village Level Outcomes

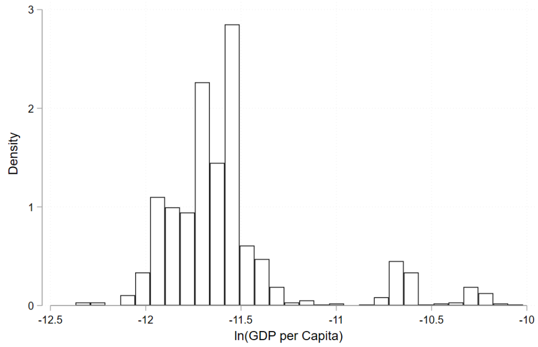

(a) GDP per Capita

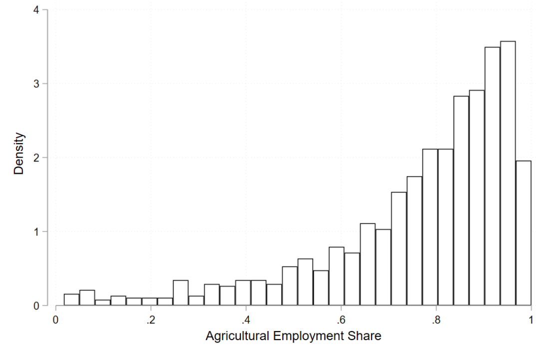

(b) Agricultural Employment Share

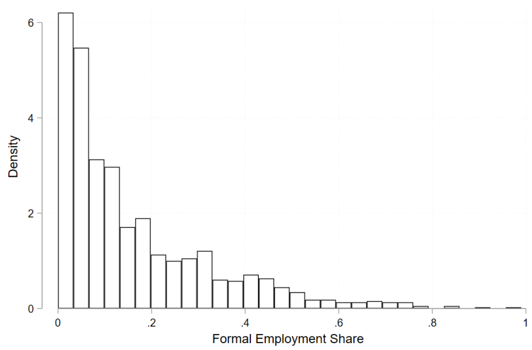

(c) Formal Employment Share

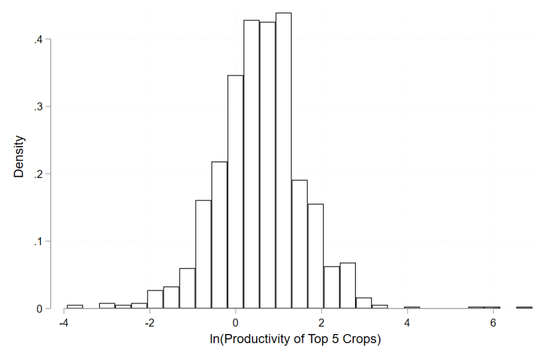

(d) Agricultural Productivity

Note: Each panel plots the distribution of a different village-level variable among the transmigration villages in the analysis sample. Panel (a) plots the log of GDP (expressed in millions of constant 2017 USD) per capita in 2012. Panel (b) and (c) plot the agricultural and formal employment shares from the Population Census in 2000, respectively. Panel (d) plots the log of the productivity (tonnes of yield per hectare planted) of the top 5 crops in each village.

C Empirical Strategy: Simulation, Robustness, and Sensitivity

C.1 Simulation of the Identification Threat

In this section I provide detail around a simulation that demonstrates the potential for omitted variable bias when estimating the relationship for newly created locations between an outcome variable and market access even if the placement of new villages is as-if-randomly assigned. Additionally, the simulation highlights how leveraging the assignment process of new villages and using the selected controls addresses potential bias concerns.

I create a space that is a 500 x 500 grid, then divide it into 9 equally sized provinces, each with 4 districts. Within each district there are then 4 subdistricts. Provinces, districts, subdistricts, and cells each draw a random fundamental productivity value (all the same variance). I specify that provinces, districts, and subdistricts have productivity draws such that there is a correlation between productivities[21]. A cell's total productivity is then the sum of each productivity component. I then create an additional measure I name "observed productivity" that is able to serve as a proxy for the unobserved true productivity[22]. In order to replicate the situation described in the empirical strategy section, I then first place 50 initial "cities" that are more likely to be located in highly productive areas. The potential for bias decreases as you introduce more initial locations, but this choice of initial locations highlights the identification challenge. To replicate the assignment process of the transmigration program, I then randomly place new villages stratified by province. The choice to stratify by province more closely mimics the reality where provincial governments exercise some oversight in the placement process of villages, and the as-if-random assignment occurs even within provinces. Figure C1 displays an example draw of this simulation exercise.

[21] In this simulation example the correlation coefficient is 0.3

[22] The specified R-squared of a regression of observed productivity on unobserved productivity is 0.6.

Figure C1: Map of Simulated Space

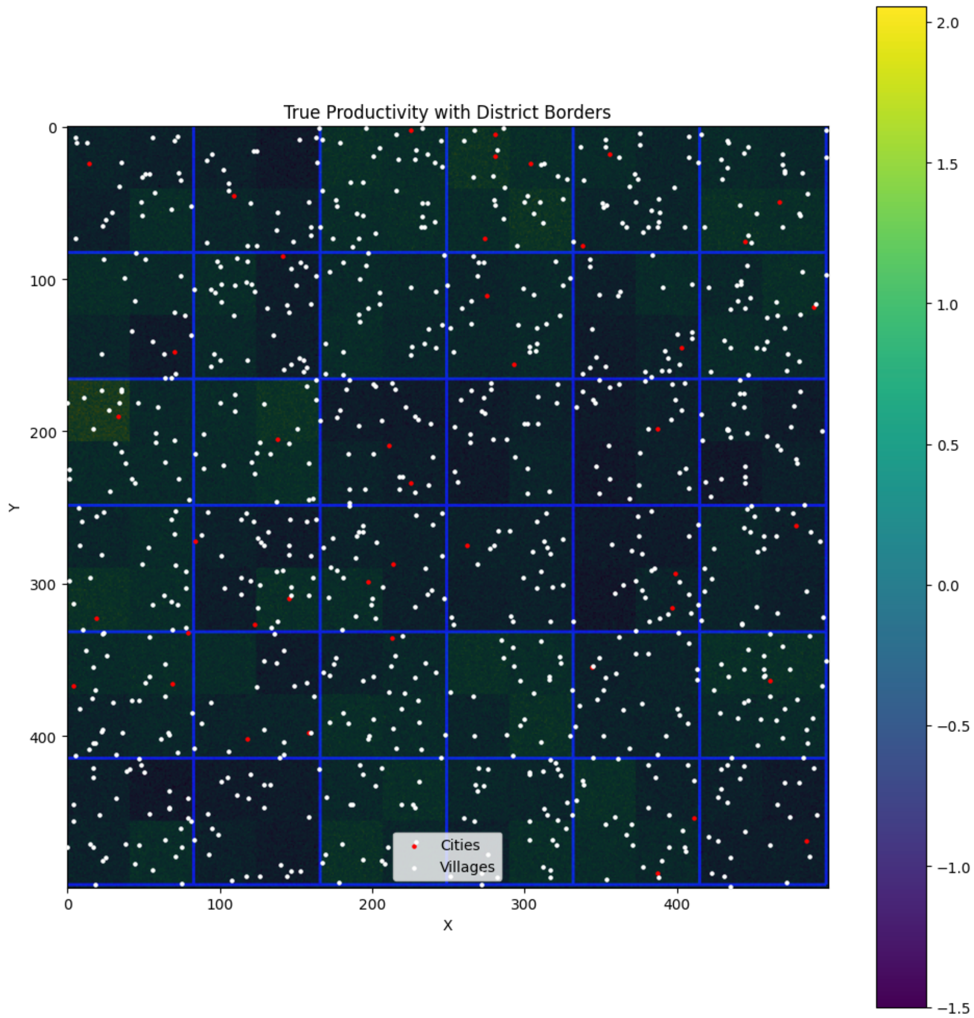

In order to demonstrate the threat to identification that persists even under random assignment of new villages, I then attempt to match the estimation strategy I implement in my actual data. For each new village, I construct its market access to preexisting cities assuming travel costs over each cell are equalized for simplicity. I then estimate the relationship between this measure of market access and the true, unobserved productivity of the new villages. If there was no omitted variable bias, then I would expect a coefficient of 0, meaning that there is no relationship between randomly placed new villages and their fundamental productivity. I also progressively add different controls that are explicitly designed to help control for the potential omitted variable bias present in this setting. I include a control for the average observed productivity of cities within a village's district to proxy for the district level controls I include in the true data. I also construct a measure of expected market access as the province-level average market access of new villages, which corresponds to the province-wide average variant of the empirical control rather than to the preferred closest-20-sites measure. Finally I also include the cell level observed productivity to match the village-level controls

I include. I simulate this process 1000 times and store the estimated coefficients from each iteration. Figure C2 displays the results from my simulation. Notably, in the simple OLS case, there is a positive association between a new village's market access and its unobserved productivity despite the fact that new villages are placed randomly. Importantly, I also find evidence that in this type of assignment process, each of my different controls helps address this threat to identification. In the specification that most closely matches my preferred estimation strategy in the true data (OLS + Both + Expected MA), the interquartile range of my estimates contains 0.

Figure C2: Simulation Results

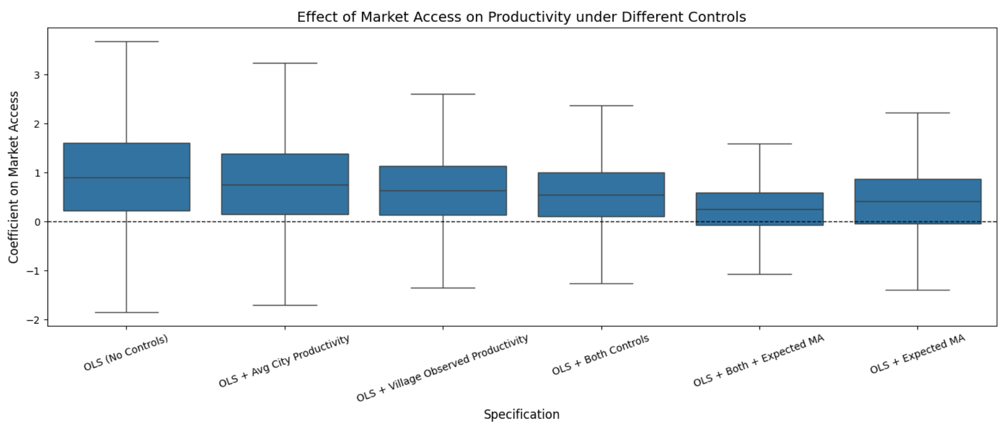

This simulation is an illustration of the potential threats to identification, and its conclusions are conditional on the assumed placement process, productivity structure, and control quality. Given this, it supports two points. First, the identification threat described in Section 4 is real in that even random placement on a network does not guarantee identification in this setting, and the controls of my preferred specification substantially reduce the threat. Second, the direction of the bias in this environment is likely favorable for the findings of this paper. When preexisting locations are more likely to sit in productive places, a randomly placed new village's market access is *positively* related to its fundamental productivity, so omitted variable bias of this form would push my estimates toward finding that higher market access raises GDP per capita. If that is the relevant bias structure, my negative estimates would understate the true negative effects. In a network that is already extremely crowded, which is not the case for the outer islands of Indonesia, this relationship could flip if only the "bad" land is left over in high market access areas. These observations complement, rather than replace, the balance tests and the Diegert et al. (2025) robustness diagnostics reported in the main text.

C.2 Robustness to Fixed Effects and Inference

Table C1: Robustness of the Market Access Estimate to Fixed Effects and Inference

	ln(GDP per Capita)	ln(GDP per Capita)	ln(GDP per Capita)	ln(GDP per Capita)	ln(GDP per Capita)	ln(GDP per Capita)	ln(GDP per Capita)	ln(GDP per Capita)
ln(Market Access)	-0.0057***	-0.0057***	-0.0057***	-0.0057***	-0.0057***	-0.0034***	-0.0023***	-0.0033***
	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0018)	(0.0009)	(0.0007)	(0.0007)
Expected MA Control	Closest 20	Closest 20	Closest 20	Closest 20	Closest 20	Closest 20	Closest 20	Closest 20
Local Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geography FE	Island	Island	Island	Island	Island	Province	Prov x Repelita	Prov x Cohort
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Std. Errors	Conley 100 km	Conley 50 km	Conley 250 km	Conley 500 km	Kabupaten	Conley 100 km	Conley 100 km	Conley 100 km
Permutation p (within-cell)	< 0.001							
Unique Villages	1151	1151	1151	1151	1151	1151	1151	1151

Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

Note: This table presents results from estimating the effect of market access on GDP per capita under alternative approaches to inference and alternative fixed effects. Every column uses the specification of Table 2, with the village and 1980 district controls, the closest-twenty expected market access control, and year fixed effects, on the village-year panel of outcome years 2012 to 2019. Column 1 is the baseline: island and cohort fixed effects with Conley standard errors, computed with a Bartlett kernel that declines with distance and reaches zero at 100 kilometers, admitting all covariances between observations within the cutoff, including those of the same village across outcome years, which enter with full weight. Columns 2-4 repeat the identical specification and sample with the spatial cutoff set to 50, 250, and 500 kilometers; the point estimate is identical across columns 1-4 by construction, and only the standard error changes. Column 5 replaces the spatial standard errors with clustering at the district (kabupaten) level. Columns 6-8 return to the Conley standard errors of column 1 and vary the fixed effects, replacing the island fixed effects with province, province-by-plan-period (Repelita), and province-by-cohort fixed effects. The permutation p-value reports the share of 1,000 random reassignments of market access across villages within province-by-cohort cells that produce an absolute coefficient at least as large as the observed one, and is valid under the sharp null that placement is exchangeable within cells; with 1,000 draws it is a Monte Carlo approximation to the exact permutation distribution. No reassignment does, so the p-value is below 0.001. The sample contains 290 province-by-cohort cells with a median of 2 villages per cell. The 106 cells containing a single village hold 106 of the 1,151 sample villages; they remain in the sample but contribute no identifying variation in column 8 and are unaffected by the permutation reassignments.

C.3 Sensitivity to Omitted Variables

I assess the likelihood that omitted variables determine the estimated relationship by implementing methods proposed by Oster (2019) and Diegert et al. (2025). I elect not to report "Oster's delta" due to concerns raised by Diegert et al. (2025)[23], but instead report the fraction of times that the "breakdown" point from Diegert et al. (2025) is larger than the ρ_k values for my local covariates. The breakdown point is the degree of selection on unobservables, relative to selection on observables, at which the estimated effect would be driven to zero; ρ_k measures the importance of each covariate k relative to all of the other observed covariates. The higher the fraction of covariates whose ρ_k falls below the breakdown point, the less an unobserved confounder that behaved like the observed ones could explain the result. Table 2 reports this fraction for each specification. In the preferred specification with the full control set it is over 85 percent, higher than the same fraction Diegert et al. (2025) report in their empirical application where they conclude the results are robust; it is lower in the specifications without district controls or the expected market access control.

[23] Specifically, the method proposed by Oster (2019) can be incorrect when covariates are correlated with the potential omitted variables.

D Resident Wages

This appendix reports the wage results discussed in Section 5.2. The sample pools the 2000 to 2007 SUSENAS waves and covers formally employed workers in the transmigration villages the survey happens to sample, 5,781 individuals in 502 villages, with wages converted to 2007 rupiah; Appendix B describes the survey. Each column of Table D1 estimates Equation (2) at the individual level, with log monthly wage, log hourly wage, or log hours worked as the outcome, individual controls for age, sex, and education, and survey-year fixed effects; columns 4 and 5 restrict the sample to individuals working in agriculture, manufacturing, and services, and column 5 adds sector fixed effects. The point estimates for wages are negative throughout and hours do not move; the estimates are imprecise, and Section 5.2 discusses them more.

Table D1: Effects of Initial Market Access on Wages

	ln(Wage)	ln(Hourly Wage)	ln(Hours Worked)	ln(Wage)	ln(Wage)
(max) ln_MA_full_noself	-0.0027	-0.0027	-0.0003	-0.0022	-0.0038*
	(0.0023)	(0.0027)	(0.0017)	(0.0023)	(0.0023)
Borusyak and Hull	Yes	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes	Yes
Sample	Full	Full	Full	Limited	Limited
Sector FE	No	No	No	No	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Unique Villages	502	501	503	501	501
Unique Households	5781	5706	5989	5462	5462

Standard errors in parentheses

Standard errors clustered at the village level

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: Individual-level regressions of log monthly wage, log hourly wage, and log hours worked on log initial market access, following Equation (2) with individual controls (age, sex, education) and survey-year fixed effects, using formally employed workers in the 2000 to 2007 SUSENAS waves matched to transmigration villages. Wages are in 2007 rupiah. Columns 4 and 5 restrict the sample to individuals working in agriculture, manufacturing, and services, and column 5 adds sector fixed effects. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Video-Derived Welfare Measure

E.1 Overview

The main analysis identifies the curse-of-connectivity result using satellite-derived GDP per capita estimates from Rossi-Hansberg and Zhang (2025) and a nightlight-based alternative. As a complementary check from a data source independent of both, I construct a welfare measure based on direct visual evidence of village material conditions in a Kalimantan subsample of transmigration villages. The measure is built from publicly-uploaded video footage scored frame by frame against a pre-specified rubric using a vision language model. The estimated relationship between this video-derived welfare measure and initial market access is negative for every welfare outcome considered, and significant at conventional levels for the composite index. The appendix evidence is directionally consistent with the GDP per capita result in Table 2 and was constructed without reference to any of the satellite, census, or PODES data sources used in the main analysis.

E.2 Sample

The source is a curated YouTube playlist of Indonesian-language documentary and vlog content covering Kalimantan transmigration villages, containing 310 videos at the time of writing, created by user Pie'ie Mejink. For each video I extracted candidate location information from the title and the Indonesian-language auto-generated captions. Of the 310 videos in the playlist, 111 of the videos contained location information that enables a high-confidence match to a village in my primary sample. Multiple villages appeared in more than one video, so my final sample of analysis corresponds to 47 distinct villages. Geographic coverage is concentrated in Kalimantan Barat, with smaller numbers of videos in Kalimantan Tengah, Kalimantan Timur, and Kalimantan Selatan.

E.3 The Welfare Rubric

The rubric scores villages based on 7 indicators of visible material welfare in a village environment. These measures are: housing quality, road and path surface, electrification, agricultural assets, commercial activity, public infrastructure, and clothing condition of individuals. Each indicator is scored on a 1-5 scale. Each indicator may be marked "not visible" for any individual frame if, for example, the frame is an interior shot, a talking-head segment or B-roll of unrelated landscape. Irrelevant frames are then not incorporated into the welfare metric.

In order to match the context from the videos so the scoring exhibits real variation, I made two choices to calibrate the model's behavior. First, the scale value of 3 is anchored to "typical rural Kalimantan," so it is not left up to the agent to set its own reference point. This prevents the rubric from compressing scores in the event it selects a poor reference point to compare villages to. Second, the rubric explicitly instructs the model not to assign a score of 1 unless the frame shows real material deprivation. The rubric was developed on a small sample of 4 hand-picked videos before the bulk run and held fixed across all subsequent scoring; the full text is reproduced verbatim in Section E.6.

E.4 Construction of the Village-Level Welfare Measure

Each video is downloaded at 480p resolution and frames are extracted at a fixed 11-second interval then sent in batches of 5 to a vision language model (Claude Sonnet 4.5) with the rubric in the system prompt. The model is required to return a structured JSON object conforming to one record per frame, with typed indicators for: `visible`, `score`, and `evidence` fields for each of the 7 measures. This eliminates free-text parsing and ensures every per-frame record is mechanically validated.

The per-video score for each indicator is the mean of the 1-5 scores across frames in which that indicator was marked visible. The per-video composite welfare index is the equal-weighted mean of the per-indicator means, taken over the indicators observed for that video; when an indicator is never visible in any of a village's videos, the composite for that village averages over the remaining indicators, which is why the component columns of Table E1 have slightly smaller samples than the composite. The equal-weighting choice is a decision: each dimension of welfare contributes once to the composite regardless of how frequently it appears on camera in any particular video, so a village walk-through and a market segment do not receive differentially weighted contributions from their dominant scene types. The per-village welfare measure is the simple mean across all high-confidence-matched videos sharing that village identifier.

For the regression reported in Section E.5, the two composite measures and the 7 individual indicator means are each z-scored across the analysis sample of 47 villages before estimation. Slopes are therefore interpreted as standard deviations of welfare per log point of ln(MA).

I report two composite measures. *Composite* is the equal-weighted mean of the 7 predefined indicators described above. *Composite (incl.)* additionally folds in a free-form "other" channel that captures visually striking welfare-relevant elements not covered by the 7 categorical indicators: luxury vehicles, abandoned structures, large industrial installations under construction, ceremonies, visible child labor, and similar. The per-frame "other" entries are assigned a signed welfare relevance by the same model call, aggregated to a video-level signed score, and combined with the predefined composite as a single pseudo-indicator on the same 1–5 scale. The two composites are highly correlated; the inclusive measure is reported alongside the predefined measure as a robustness check on the decision to fix the indicator set ex ante.

E.5 Results

I regress each z-scored welfare outcome on ln(MA) at the village level. Each observation is a unique village, with welfare averaged across replicate videos when a village is matched by more than one video. Due to the small sample, I exclude controls or fixed effects, and standard errors are clustered at the village level. Table E1 reports the results for the welfare indicators.

All nine slopes are negative, directionally consistent with the main paper's GDP-per-capita finding. The welfare index is consistently lower in high market access transmigration villages. The magnitudes are economically meaningful as well; based on the distribution of market access from the *full* sample of transmigration villages, moving from the 25th to 75th percentile of initial market access corresponds to roughly 0.5 standard deviations lower measured welfare, a figure that combines a slope estimated on the 47-village Kalimantan sample with the full-sample access spread and is therefore indicative rather than exact. This pattern is found in every welfare category except commerce, where I find a slope of essentially zero. This is consistent with the proposed mechanism in the paper itself. The convergence of two welfare measures constructed from entirely disjoint data sources strengthens confidence that the pattern documented in Section 5 is not an artifact of any specific feature of the satellite or administrative data. The video-derived estimates should be read as descriptive corroboration rather than as independent identification: the sample is small (47 villages) and could be selected through filming, playlist inclusion, and location matching, each potentially related to connectivity or visible conditions; the specifications carry none of the main design's controls; and the geographic coverage is restricted to Kalimantan.

Table E1: Video-derived welfare measures regressed on initial market access

	Composite	Composite (incl.)	Housing	Road & Path	Electrification	Agri. Assets	Commerce	Public Infra.	Clothing
ln(MA)	-0.048*	-0.051*	-0.025	-0.043**	-0.030	-0.011	-0.002	-0.063	-0.056
	(0.026)	(0.030)	(0.024)	(0.022)	(0.025)	(0.026)	(0.036)	(0.040)	(0.036)
N (villages)	47	47	47	47	45	46	43	42	47

Note: Each column reports a univariate OLS regression of a z-scored video-derived welfare outcome on ln(MA) at the village level. Welfare outcomes are averaged across replicate videos within villages before z-scoring. *Composite* is the equal-weighted mean of the 7 indicator means; *Composite (incl.)* additionally folds in the free-form "other" channel described in Section E.4. Standard errors are clustered at the village level. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

E.6 Welfare Rubric Reproduced Verbatim

The text below is the system prompt used to score every frame in every video, reproduced word-for-word from the implementation code.

```
You score visible welfare indicators in frames sampled from videos
of Indonesian transmigration villages (mostly Kalimantan).

For each frame, rate each indicator 1-5:
  1 = very poor / severe deprivation
  2 = poor / below regional norm
  3 = average / typical for rural Kalimantan
  4 = good / above regional norm
  5 = excellent / clear prosperity
```

```
If an indicator is NOT VISIBLE in the frame (e.g. interior shot, talking head,
B-roll of only sky/forest/water with no village context), set visible=false
and score=0. Be honest: most B-roll frames will have only 1-2 visible indicators.

Indicators:

- housing_quality: walls/roof materials and condition. Tarp/scrap/bamboo < raw
  wood plank < painted wood < zinc roof + plank < concrete with tiled roof <
  multi-storey painted concrete. Note upkeep and signs of recent investment.
- road_path_surface: dirt < graded dirt < gravel < paved (aspal) < paved with
  drainage. Quality of drainage and ditches matters.
- electrification: power lines, transformers, exterior lights, satellite
  dishes, visible appliances/electronics. Score the density of signs, not
  whether *any* signs exist.
- agricultural_assets: hand tools only < draft animal < small motorized
  (handtractor / pump) < tractor / mechanized rice planter. Visible canals,
  weirs, and pumps count here too.
- commercial_activity: warungs, shops, kiosks, vendors, parked vehicles,
  construction, motorcycle density. Large income-generating builds (swiftlet
  farms, sawmills, brick kilns, logging trucks) push this high.
- public_infrastructure: schools, mosques, churches, clinics (puskesmas),
  government offices, sports facilities. Score on apparent condition.
- clothing_condition: residents' clothing - worn/torn/dirty < clean modest
  village clothing < new/branded/urban-style.

ALSO record an "other" observation per frame: a visually striking element not
covered above that bears on welfare (luxury vehicle, abandoned house,
ceremony, child labour signs, large modern equipment, etc). If nothing
notable, set present=false.

SCENE TYPE: Also tag each frame with the single category that best describes
what the camera is primarily framing. This is purely descriptive | it does
NOT change the 1-5 scoring above. Pick exactly one:
  - residential: homes, yards, residential streets, kampung from the outside
  - commercial: markets, shops, kiosks, vendors, storefronts, commercial buildings
  - public_institutional: schools, mosques/churches, clinics (puskesmas),
    government offices, sports facilities
  - agricultural: paddy fields, plantations, canals, farming activity, livestock
```

```
  - infrastructure: roads in transit, bridges, power infrastructure, when these
    are the subject of the frame (not just incidental background)
  - interior_or_talking_head: inside a building, close-up portrait, on-camera
    interview, title cards / intro graphics
  - b_roll_or_landscape: sky, forest, water, generic natural landscape,
    anything else not informative for welfare

CALIBRATION: Anchor 3 to rural Kalimantan, NOT urban Java. A simple wooden
house with a zinc roof and dirt yard is a 3. Do not score 1 unless the frame
shows real deprivation (collapsing structure, residents in rags, etc).

Return ONE entry per frame in the same order as received.
```

F Additional Empirical Results

Table F1: Effects of Market Access on Housing

	ln(Housing Expenditure)	ln(Housing Expenditure)	ln(House Area)	ln(Price Per Sq Meter)
ln(MA)	0.0130*	0.0118*	0.0030	0.0093
	(0.0074)	(0.0070)	(0.0029)	(0.0074)
Household Size		0.0510***	0.0756***	0.0078
		(0.0082)	(0.0095)	(0.0105)
House Area		0.0094***		
		(0.0018)		
House Area, Sq		-0.0000		
		(0.0000)		
Expected MA Control	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes
Unique Villages	70	70	70	70
Unique Households	1086	1086	1086	1086

Standard errors in parentheses

Standard errors clustered at the village level

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: This table presents results from estimating the effect of market access on housing price and expenditure according to Equation (2). Columns 1-4 use data from the 2002 SUSENAS household expenditure module.

Table F2: Effects of Market Access on Land Use Dynamics

	ln(Ag Land)	ln(Housing + Firm Land)	Change in Ag Land	Change in Housing + Firm Land	New Housing Development
ln(Market Access)	-0.0237***	-0.0155**	-2.7412**	-1.7831*	-0.0000
	(0.0056)	(0.0071)	(1.1285)	(1.0409)	(0.0009)
Expected MA Control	Yes	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	No	No	No
Unique Villages	940	935	944	944	1147

Standard errors in parentheses

Standard errors clustered at the village level

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: This table presents results from estimating the effect of market access on land use dynamics according to Equation (2). Columns 1-4 use data from the 2003 PODES. Column 5 uses data from the 2000 PODES. Cultivated land declines with market access while land devoted to housing and firm activity does not expand; see Section 5.

G The Agricultural Externality: Additional Results

Figure G2 plots the first stage of the instrumental variables design of Section 6. Table G2 reports the balance of initial settlement size with respect to observable village and district characteristics: initial settlement size in this sample is uncorrelated with village-level characteristics, and with both market access and expected market access, so the instrument captures variation in village scale that is orthogonal to the connectivity variation exploited in Section 5. Several of the district-level characteristics from 1980 do show statistically significant correlations with settlement size, concentrated among measures of infrastructure and urbanization; the omnibus F-statistic of 1.32 fails to reject the null of joint insignificance, and the specifications control for these district-level measures directly.

Figure G1 and Table G1 report the wage evidence cited in Section 6.1. The figure is a binscatter of rural workers' log wages against the size of the worker's own sector in the village, separately for agriculture and retail; the table adds village and demographic controls, and the agricultural slope survives while the retail and service slope remains flat.

Figure G1: Sector Size vs. Individual Wages

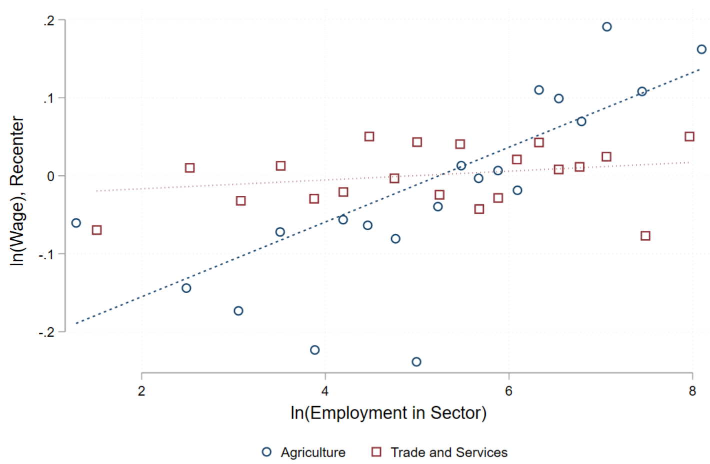

Note: Binscatter of log wages (demeaned) against log sector size. The data here contains all formally employed rural individuals on the Outer Islands that appear in the pooled SUSENAS data.

Table G1: Relationship between Sector Size and Wages

	Agriculture		Trade		Services	
	ln(Hourly Wage)	ln(Hourly Wage)	ln(Hourly Wage)	ln(Hourly Wage)	ln(Hourly Wage)	ln(Hourly Wage)
ln(Employment in Sector)	0.0296***	0.0173***	0.0070	-0.0164	0.0216***	0.0091
	(0.0056)	(0.0052)	(0.0184)	(0.0199)	(0.0060)	(0.0056)
Village Controls	Yes	Yes	Yes	Yes	Yes	Yes
HH Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE Type	Island	Province X Year	Island	Province X Year	Island	Province X Year
Unique Villages	3177	3177	1015	1007	6233	6233
Obs	16646	16645	1687	1678	26864	26864

Standard errors in parentheses
Standard errors clustered at the village level
* $p < .1$, ** $p < .05$, *** $p < .01$

Note: Each column estimates the relationship between log sector size and log hourly wages for all formally employed rural individuals on the Outer Islands that appear in the pooled SUSENAS data. Village controls are the potential yield variables described in the primary specification along with the overall village population, and household controls are gender, age, and years of education.

Table G3 examines specific threats to exclusion. Larger settlements may develop in different ways that could affect agricultural productivity beyond the size of the sector itself. The table reports the association of initial settlement size with candidate downstream channels such as market presence, market access, financial institutions, and schooling infrastructure. Because these are themselves post-settlement outcomes, I read this table as suggestive evidence that exclusion is not violated. Most of these associations are small and insignificant, but I also report specifications that condition on these outcomes. When I control for them, the estimated scale externality increases in magnitude rather than decreases. Larger settlements may also shift crop choices, which the crop fixed effects and the composition test of Section 6 address. While it is impossible to definitively rule out violations of the exclusion restriction, the patterns I can examine do not point toward one.

Figure G2: First Stage: Ag Working Population vs. Initial Settlement Population

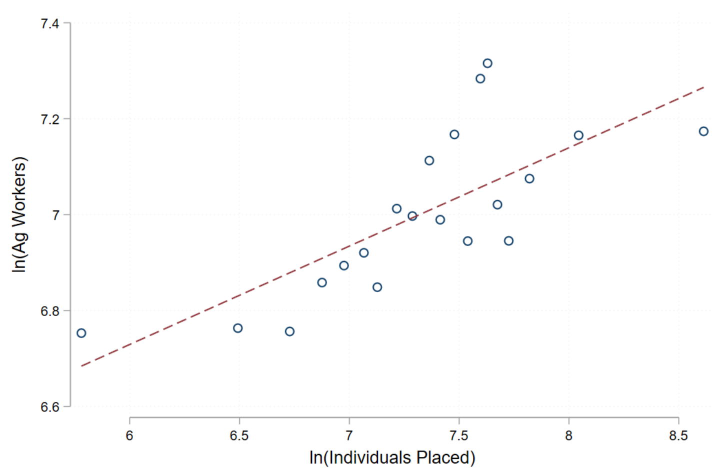

Note: Binscatter of log agricultural working population against log initial settlement population, controlling for island and cohort fixed effects.

Table G2: Village Characteristics Across Initial Settlement Size

Dependent variable (std):	ln(Individuals Placed) (std)
Wetland Rice Potential Yield	-0.012 (0.032)
Palm Oil Potential Yield	0.075* (0.043)
Coffee Potential Yield	-0.013 (0.040)
Cassava Potential Yield	0.025 (0.030)
Vector Ruggedness Measure	0.042 (0.043)
Agrosimilarity Score	0.002 (0.036)
Share of Land with Poor Drainage	0.067 (0.046)
Share of Land with Good Drainage	-0.059 (0.045)
Organic Carbon (%)	0.050 (0.039)
Topsoil Salinity (Elco) (dS/m)	-0.098 (0.075)
Topsoil Base Saturation (%)	-0.027 (0.056)
Avg. rainfall, 1948–1978	-0.061 (0.043)
Avg. temp (Celsius), 1948–1978	-0.027 (0.044)
ln(MA)	-0.023 (0.039)
ln(Expected MA)	-0.032 (0.033)
ln(District Population, 1980)	-0.024 (0.037)
District Own Electricity, 1980	-0.098** (0.046)
District Own Radio, 1980	-0.030 (0.046)
District Own TV, 1980	-0.108** (0.049)

Dependent variable (std):	ln(Individuals Placed) (std)
District Literacy, 1980	-0.062 (0.042)
District Avg Years Schooling, 1980	-0.062 (0.044)
District Agricultural Employment Share, 1980	0.062 (0.046)
District Wetland Rice Potential Yield, 1980	0.016 (0.039)
District Palmoil Potential Yield, 1980	0.012 (0.040)
District Coffee Potential Yield, 1980	0.052 (0.042)
District Cassava Potential Yield, 1980	0.018 (0.039)
District Mining Employment Share, 1980	0.102** (0.040)
District Manufacturing Employment Share, 1980	0.007 (0.038)
District Trade Employment Share, 1980	-0.069 (0.046)
District Service Employment Share, 1980	-0.085* (0.048)
District Wage Worker Share, 1980	-0.052 (0.039)
District Share HHs with Piped Water, 1980	-0.119** (0.049)
District Share HHs with Sewage, 1980	-0.080* (0.046)
District Share HHs with Modern Fuel, 1980	-0.130*** (0.049)
District Share HHs with Modern Floor, 1980	-0.053 (0.045)
District Share HHs with Modern Roof, 1980	0.039 (0.043)

F-stat of omnibus test	1.319
Observations	719

Note: Each row reports the coefficient from a regression of the standardized dependent variable listed in the row on the standardized log number of individuals initially placed in a transmigration village, controlling for island fixed effects and cohort fixed effects. Standard errors, clustered at the subdistrict (kecamatan) level, are reported in parentheses. The omnibus F-statistic tests joint significance across all dependent variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G3: Effects of Initial Settlement Size on Downstream Outcomes

	Presence of Permanent Market	MA, recentered	Total Banks	Banks per Capita	Elementary School Density (per 1000)	Middle School Density (per 1000)	High School Density (per 1000)
ln(Individuals Placed)	0.0619	-0.7123	0.0252	0.0000	-0.0636	-0.0222	-0.0216**
	(0.0412)	(0.5626)	(0.0383)	(0.0000)	(0.0465)	(0.0240)	(0.0105)
MA Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Local Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Unique Villages	601	601	601	601	601	601	601

Standard errors in parentheses

Standard errors clustered at the village level

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: This table presents results from estimating the effect of initial settlement size on a variety of outcomes that could indicate a violation of the exclusion restriction in the 2SLS specification of Section 6. All columns except column 2 use data from the 2003 PODES wave. Column 2 uses the constructed measure of market access, recentered by the expected market access measure described in Section 4.

H The Simple Model

This appendix presents a simplified version of the quantitative model of Section 7. I include the same three goods, technology and externality, and sorting structure, but on a symmetric stylized geography (forty identical villages on a line, one city) instead of the full Indonesian economy. Additionally, parameters are set to illustrative values rather than estimated. The simple model serves three purposes. First, it highlights the mechanism in a setting that is potentially easier to interpret. Second, it is much easier to solve computationally, so I carry out various stress tests with this model. I relax certain assumptions and show how the sign and magnitude of the effect of market access on income changes. Finally, I confirm numerically, under Roy–Fréchet sorting, the sign-reversal threshold that Proposition 1 establishes in closed form under Pareto sorting (Appendix I).

H.1 Environment

Villages $v = 1, \ldots, V$ ($V = 40$) sit at positions x_v evenly spaced on $[1, 4]$ along a line, with the city at the origin; the gross iceberg cost factor between any two locations is one plus the distance between their positions, so the village-to-city factor is $\tau_v = 1 + x_v \in [2, 5]$. Each village has $L_v = 1$ workers and $H_v = 1$ units of land. There are three goods, as in the main text. Food is homogeneous and shipped at iceberg cost τ^{κ_a}; rural services are differentiated by village of origin and traded among villages and to the city at cost τ^{κ_n} under CES aggregation with elasticity σ_n; and an urban good is produced only by the city with a linear technology, taken as the numéraire at source and delivered to villages at cost τ^{κ_u}. City income I_c is exogenous, and the city spends it in fixed shares $(\beta_a, \beta_n, \beta_u)$ on food, rural services, and its own good.

Technology is that of the main text. All village land works in agriculture,

$$Y_{a,v} = \bar{A}_a \, Z_{a,v}^{\gamma_a + \varphi} \, H_v^{1-\gamma_a}, \qquad Y_{n,v} = A_n \, Z_{n,v}^{g_n}, \tag{H1}$$

with the externality $A_{a,v} = \bar{A}_a Z_{a,v}^{\varphi}$ uncompensated, and the rural service produced with labor alone; $g_n = 1$ is the constant-returns baseline, and $g_n < 1$ (a stress-test variant) makes service production decreasing-returns. Workers draw sector abilities from independent Fréchet distributions with shape θ and scales (T_a, T_n) and choose the sector offering the highest private return, exactly as in the main text. Village households have Cobb–Douglas preferences with weights $(\alpha_a, \alpha_n, \alpha_u)$. Food arbitrage sets the village-gate price $p_{a,v} = \tau_v^{-\kappa_a} p_a$, and the economy-wide food price p_a clears the aggregate food market against city demand.

An equilibrium is a pair of wages in each sector per village and the food price ($2V+1$ unknowns) solving each village's agricultural first-order condition, each village's service-variety market clearing, and aggregate food clearing; the urban-good market clears via Walras' Law and is verified numerically. Because $\gamma_a + \varphi$ approaches one, multiple equilibria are possible; the reported equilibrium is selected by continuation in φ from the unique $\varphi = 0$ benchmark, the same device used at scale in the quantitative model.

H.2 Calibration

The parameters I choose here are illustrative and chosen to highlight the workings of the model. They are not estimated in the same way as the full model. The technology block matches the main calibration where the object is common: $\gamma_a = 0.304$ and $\varphi = 0.70$. The sorting dispersion is set at $\theta = 7$; the stress test below reports the model at $\theta = 4$ (strong comparative-advantage dispersion) and $\theta = 30$ (approaching homogeneous workers). The quantitative model's imposed $\theta = 12.5$ sits in the middle of this range.

The trade block sets $\kappa_a = 0.01$, $\kappa_n = \kappa_u = 0.60$, and $\sigma_n = 6$. These values were chosen to be relatively close to the calibrated values and such that the behavior of income and distance to the city is roughly monotone at the equilibrium. The demand weights $(\alpha_a, \alpha_n, \alpha_u) = (0.80, 0.10, 0.10)$ and city shares $(\beta_a, \beta_n, \beta_u) = (0.15,\ 0.10,\ 0.75)$ were selected so that the average agricultural employment share roughly matches the data (0.652 in the model compared to 0.648 in the data).

H.3 Baseline Results

Figure H1 reports the baseline model behavior. Real income per capita falls monotonically toward the city. The closest village is 6.3 percent poorer than the most distant one. The non-agricultural employment share rises toward the city by 16.4 percentage points, and agricultural TFP falls via the effects of the externality. Re-solving the identical economy at $\varphi = 0$ reverses the relationship. With no externality, the closest village becomes 7.1 percent *richer* than the most distant, the standard gains from trade.

Figure H1: Simple Model: Outcomes by Distance from the City

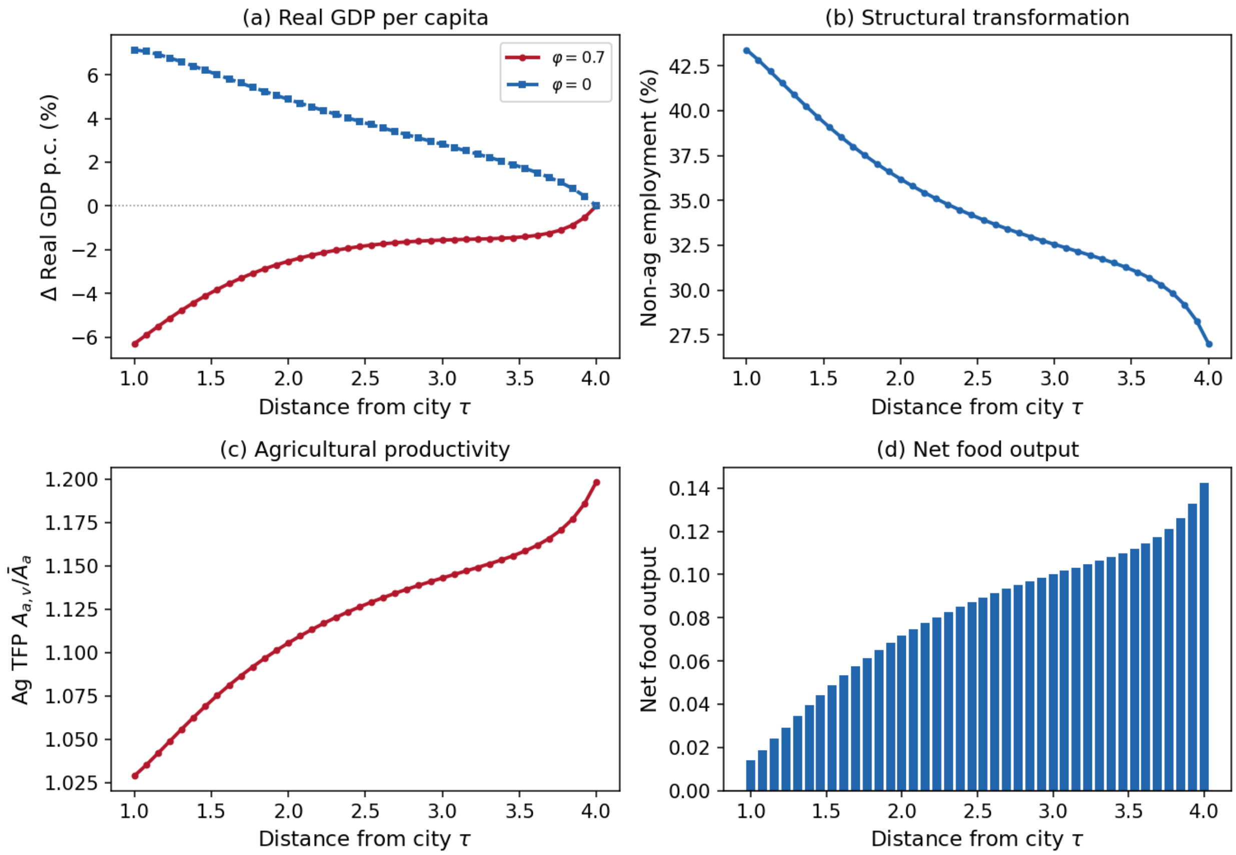

Note: I plot four outcomes for the 40 villages of the simple model against their distance from the city. Panel (a) plots real income per capita relative to the most distant village, at the calibrated $\varphi = 0.70$ and at $\varphi = 0$. Panel (b) plots the non-agricultural employment share, panel (c) the agricultural TFP multiplier $A_{a,v}/\bar{A}_a$, and panel (d) net food output.

H.4 Stress Tests

Table H1 re-solves the simple model under variants that relax, one at a time, assumptions that could influence either the sign or size of the relationship between connectivity and income. Figure H2 plots the corresponding spatial profiles for each test. The table reports the difference in real income with and without the externality, along with the difference in the non-agricultural employment share between the closest and most distant village from the city. Additionally, the table reports the estimated value of φ where the relationship between connectivity and income flips sign; this is the numerical counterpart of the closed-form threshold in the simple model in Proposition 1.

Table H1: Stress test: close-versus-far outcomes by variant

Variant	Δreal (%)	$\Delta\pi_n$ (pp)	$\Delta\text{real}_{\varphi=0}$ (%)	φ^*	avg π_a	Net food
Baseline	-6.30	+16.36	+7.14	0.45	0.652	> 0
Near-homogeneous workers ($\theta = 30$)	-14.78	+24.16	+6.99	0.33	0.652	> 0
High dispersion ($\theta = 4$)	-1.89	+12.54	+7.26	0.59	0.652	> 0
DRS rural services ($g_n = 0.8$)	-1.51	+11.29	+7.27	0.61	0.700	> 0
DRS both sectors ($\rho = 0.8$)	-2.70	+12.09	+7.31	0.56	0.652	> 0
Free food ($\kappa_a = 0$)	-6.88	+17.01	+7.06	0.43	0.652	> 0
Costlier food ($\kappa_a = 0.05$)	-3.96	+13.77	+7.47	0.52	0.652	> 0
Costlier food ($\kappa_a = 0.10$)	-0.97	+10.54	+7.88	0.65	0.652	> 0
Lower substitution ($\sigma_n = 4$)	-2.09	+10.75	+6.90	0.58	0.651	> 0
Higher substitution ($\sigma_n = 8$)	-9.79	+20.99	+7.30	0.38	0.652	> 0
Lower trade costs ($\kappa_n = \kappa_u = 0.40$)	-3.32	+9.91	+4.77	0.49	0.651	> 0
Higher trade costs ($\kappa_n = \kappa_u = 0.80$)	-9.33	+22.76	+9.57	0.43	0.652	> 0
Higher service demand ($\alpha_n = 0.25$, $\beta_n = 0.20$)	-7.05	+14.25	+8.06	0.49	0.407	> 0

Note: Each row re-solves the simple model at $\varphi = 0.70$ under one deviation from the baseline calibration of Section H.2, by continuation from $\varphi = 0$. Δreal is the real-income gap between the closest and most distant village. $\Delta\pi_n$ is the corresponding gap in the non-agricultural employment share. $\Delta\text{real}_{\varphi=0}$ is the income gap with the externality shut off. φ^* is the externality value at which the income gap changes sign, interpolated along the continuation path. Avg π_a is the average agricultural employment share, and the final column reports the sign of aggregate village net food exports.

Figure H2: Simple Model: Real-Income Profiles Across Stress-Test Variants

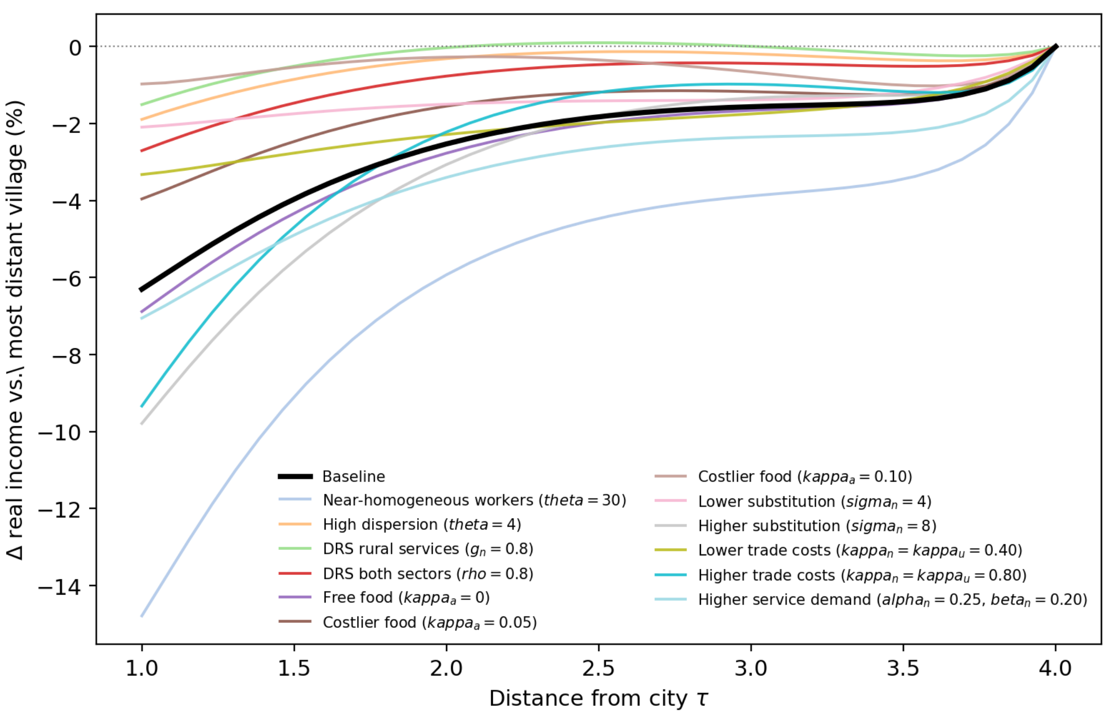

Note: Each curve plots real income relative to the most distant village against distance from the city, solved at $\varphi = 0.70$ under one deviation from the baseline calibration of Section H.2. The baseline is plotted as the heavy black line. The corresponding summary statistics are reported in Table H1.

There are several observations worth commenting on from these results. First, the negative relationship between market access and income is present in every stress-test I included, but the magnitude of the relationship varies significantly. There are parameter combinations where the difference in income is positive, but the point of this exercise is to highlight that the general direction of the relationship is robust to a wide range of parameter combinations. Also, there is no scenario where the difference in income *without* the externality is negative. The externality is the force that makes it possible to generate a negative relationship between connectivity and income.

Second, worker heterogeneity is an important force that determines the magnitude of the relationship between connectivity and income. Making workers more homogeneous ($\theta = 30$) more than doubles the curse, while high dispersion in comparative advantage ($\theta = 4$) reduces the relationship. When comparative advantage is more dispersed, sectoral shifts translate to smaller movements in effective labor than headcount. Since the externality is modeled with respect to effective agricultural labor, this reduces the bite of the productivity loss that arises as villages become more exposed to demand for their non-agricultural output.

Third, the sector-specific trade frictions are key to this mechanism. From a theoretical perspective, the cross-sectional pattern requires services to be effectively harder to trade than food, $\kappa_n(\sigma_n - 1)/\sigma_n > \kappa_a$ (Appendix I); at equal frictions the relative service price would rise with distance and the reallocation would run the other way. The key driving force that generates the cross-sectional relationship between connectivity and non-agricultural employment share is that the demand for non-agricultural goods concentrates nearby the city; without the differences in trade costs, that never occurs. From a quantitative perspective, the size of the curse shrinks as food becomes more costly to trade. At $\kappa_a = 0.10$ the gap nearly closes, while at $\kappa_a = 0$ the curse deepens to -6.9 percent. Given the importance the values of κ have for this model, I also include robustness tests that vary the values of κ_a in the full model (Appendix M.8).

Fourth, the forces that further concentrate city service demand spatially amplify the curse. Raising the substitution elasticity from $\sigma_n = 4$ to 8 moves the gap from -2.1 to -9.8 percent, and raising the trade-cost elasticities from 0.40 to 0.80 moves it from -3.3 to -9.3 percent. Consistent with Proposition 1, the sign-reversal threshold in the last column moves with the same forces (between 0.33 and 0.65 across variants) and never exceeds the calibrated $\varphi = 0.70$, itself below every 2SLS point estimate of Section 6 and inside every weak-instrument-robust confidence set. The final row returns to the demand weights that ignore the employment-level target; the curse is essentially unchanged, so the level calibration of Section H.2 is not what generates it.

The magnitudes in this appendix are illustrative (they depend on a stylized geometry and on levels the toy is not asked to match) and the quantitative statements of the paper rest on the estimated model of Sections 7–9. What the stress test establishes is that the curse's existence needs two key ingredients: the externality itself, and a trade structure in which proximity concentrates service demand without handing food producers a large enough offsetting price premium. Both ingredients are measured in the data.

I Derivation of the Sign-Reversal Threshold

This appendix derives the closed-form connectivity threshold stated in Proposition 1 of Section 7.2. The environment is the simple model of Appendix H with one additional modification that is important for reaching closed-form solutions. The independent Fréchet ability draws are replaced by a system where agricultural ability is common across workers and non-agricultural ability drawn from a Pareto distribution. The Pareto tail makes sectoral labor supplies closed-form in the service employment share, and every response to a connectivity shock then follows by applying the chain rule. The trade structure is still like the full model's. The mechanism is identical to the quantitative model's, and only the sorting algebra differs.[24]

I.1 The Village Block

A village has a unit mass of workers that consumes three goods (food, village non-agricultural services, and a city good) and the technology of the main text. Food is produced with labor and land and features a positive external scale elasticity, and services are produced with a constant returns to scale labor only technology,

$$Y_a = \bar{A}_a \, Z_a^{\gamma_a+\varphi} H^{1-\gamma_a}, \qquad Y_n = A_n \, Z_n^{\gamma_n}, \qquad \gamma_n \in (0,1], \tag{I1}$$

where $\gamma_n = 1$ is the constant-returns baseline and $\gamma_n < 1$ the decreasing-returns variant of Appendix H.[25] The village faces three prices: a farm-gate food price p_a, a village-gate service price p_n, and a delivered city-good price p_u^d; Section I.2 ties all three to distance. Village income is the value of production, $I = p_a Y_a + p_n Y_n$, and the city good enters only through the budget. The occupational margin depends on the two producer prices only through their ratio,

$$\rho \;\equiv\; \frac{p_n}{p_a},$$

the relative price of the two village goods.

Each worker has the same agricultural ability and draws a service ability z from a Pareto distribution with shape $k > 1$ and minimum $\underline{z}$. A worker chooses services when $w_n z > w_a$, so the marginal worker sits at the threshold $z^* = w_a/w_n$. The service employment share, u is then given by $u = (\underline{z}/z^*)^k$. The Pareto partial expectation then gives effective labor in both sectors:

$$Z_a = 1 - u, \qquad Z_n = \frac{k}{k-1}\, \underline{z}\, u^{\frac{k-1}{k}}. \tag{I2}$$

[24] The one-village threshold model is developed in full in a companion derivation, where every elasticity below is confirmed against numerical solutions to 10^{-9} or better. Under independent Fréchet ability sorting the threshold still exists, but I cannot solve for it in closed-form. The numerical φ^* column of Table H1 confirms the crossing in the Fréchet simple model.

[25] Under $\gamma_n < 1$ the residual $(1-\gamma_n)$ share of service revenue is a profit rebated to the village.

Substituting the competitive wage conditions from the FOCs $w_a = \gamma_a p_a \bar{A}_a Z_a^{\gamma_a+\varphi-1} H^{1-\gamma_a}$ and $w_n = \gamma_n p_n A_n Z_n^{\gamma_n-1}$ into the threshold and collecting powers of u reduces the equilibrium to one scalar equation in the relative price,

$$\frac{(1-u)^B}{u^A} = \frac{\tilde{C}}{\rho}, \qquad A = \frac{1+(1-\gamma_n)(k-1)}{k}, \qquad B = 1-\gamma_a-\varphi, \tag{I3}$$

with $\tilde{C}$ a constant collecting the technology levels. Under the stability condition $\gamma_a + \varphi < 1$ (so $B > 0$) the left side falls monotonically from $+\infty$ to 0 on $(0,1)$, so an interior equilibrium exists, is unique, and u rises with ρ. The intuition is that better terms for services pull workers out of agriculture. At the calibrated values $B = -0.004$, marginally below zero, so global uniqueness of this one-village block is not guaranteed; what the argument below uses is only that $\eta > 0$, which requires $A + B\,u/(1-u) > 0$ and holds at $B = -0.004$ for every u below about 0.98, hence at every village. The second, high-u branch that $B < 0$ admits is never selected by the continuation from $\varphi = 0$, and the proposition's characterization is conditional on that branch. Log-differentiating (I3) gives the structural transformation elasticity

$$\eta \equiv \frac{d\ln u}{d\ln \rho} = \frac{1}{A + B\,\dfrac{u}{1-u}} > 0, \tag{I4}$$

and the chain rule delivers the responses of TFP and output: $d\ln A_a/d\ln\rho = -\varphi\frac{u}{1-u}\eta < 0$ and $d\ln Y_n/d\ln\rho = \gamma_n\frac{k-1}{k}\eta > 0$.

The income response collapses by an envelope argument. Competitive sorting maximizes the private value of village production at given A_a, so the labor-reallocation terms cancel by the marginal worker's indifference, leaving only the direct price effects (Hotelling's lemma, $\partial I/\partial p_k = Y_k$) and the externality feedback:

$$d\ln I = s_a\,d\ln p_a + s_n\,d\ln p_n \;-\; \mathcal{E}\;d\ln\rho, \qquad \mathcal{E} \equiv s_a\,\varphi\,\frac{u}{1-u}\,\eta, \tag{I5}$$

where $s_a = p_a Y_a/I$ and $s_n = p_n Y_n/I = 1 - s_a$ are the sectoral income shares. The composite $\mathcal{E}$ is the externality drag. A 1 percent increase in the relative service price reduces effective labor in agriculture by approximately $\frac{u}{1-u}\,\eta$ percent. Because of the externality, this lowers agricultural TFP by φ times that reallocation. Finally, the resulting income loss is weighted by agriculture's share of village income s_a.

The village here consumes only its own service variety; the full model's rural CES over all origins adds cross-village service demand, from which the one-village block abstracts. With Cobb–Douglas consumption the exact price index is $P = p_a^{\alpha_a} p_n^{\alpha_n} (p_u^d)^{\alpha_u}$, welfare equals real income $V = I/P$, and

$$d\ln V = (s_a-\alpha_a)\,d\ln p_a + (s_n-\alpha_n)\,d\ln p_n - \alpha_u\,d\ln p_u^d \;-\; \mathcal{E}\;d\ln\rho. \tag{I6}$$

Each price enters with the village's net exposure to that good.[26] Net food exports, net service exports, and city-good imports sum to zero by the budget identity. At $\varphi = 0$ equation (I6) is the textbook terms-of-trade calculus; the externality contributes the single extra term $-\mathcal{E}\, d\ln\rho$.

I.2 The Three Frictions

Now place a continuum of such villages at distances τ from a single city, the arrangement of Appendix H. Food arbitrage sets the farm-gate price at the market price net of transport, $p_a = \tau^{-\kappa_a}$ in units of the numéraire, so selling food is worth less far from the market. The city good is produced at a price p_u set in the city and delivered at $p_u^d = \tau^{\kappa_u} p_u$, so buying city goods costs more far from the market. And the city allocates a service budget X_n across village varieties by CES with elasticity $\sigma_n > 1$ and iceberg cost τ^{κ_n}, so the free-on-board demand reaching a village at distance τ is

$$D_c(\tau, p_n) = D\, \tau^{\kappa_n(1-\sigma_n)}\, p_n^{-\sigma_n}, \qquad D \equiv X_n/\Phi, \tag{I7}$$

with Φ the city's CES price index over origins. Each village is atomistic and takes D and p_u as given. The village's service market clears,

$$Y_n = \alpha_n \frac{I}{p_n} + D_c(\tau, p_n), \tag{I8}$$

one equation in the local service price $p_n(\tau)$.

Log-differentiating (I8) in $\ln\tau$, writing θ_c and $\theta_l = 1-\theta_c$ for the city and local shares of the village's service sales and using the one-village elasticities of Section I.1 (supply $\varepsilon_S \equiv d\ln Y_n/d\ln\rho = \gamma_n \frac{k-1}{k}\eta$ and the income response (I5), with $d\ln p_a/d\ln\tau = -\kappa_a$ substituted throughout), the service-price and relative-price gradients solve out to

$$\frac{d\ln\rho}{d\ln\tau} = -\frac{\theta_c\left[\kappa_n(\sigma_n-1) - \kappa_a\sigma_n\right]}{\varepsilon_S + \theta_c\,\sigma_n + \theta_l\,(s_a + \mathcal{E})}, \qquad \frac{d\ln p_n}{d\ln\tau} = \frac{d\ln\rho}{d\ln\tau} - \kappa_a. \tag{I9}$$

The denominator captures how much the service price must adjust for the market to clear, and it is always positive. The numerator is what distance does to city demand for the village's services.

Distance increases the cost of delivering services to the city, which shrinks city demand and pushes the village's service price down. But distance also lowers its farm-gate food price, and since ρ measures services against food, a lower food price pushes ρ up. The relative service price falls with distance, and the mechanism operates, if and only if

$$\kappa_n \frac{\sigma_n - 1}{\sigma_n} \;>\; \kappa_a : \tag{I10}$$

The service friction shrinks city demand at rate $\kappa_n(\sigma_n-1)$ rather than $\kappa_n\sigma_n$, because shipping

[26] Here net exposure measures the production share (share of income that comes from a given sector) minus the consumption share for that same good.

τ^{κ_n} units to deliver one raises the village's sales at the same time as the higher delivered price cuts them. The food price carries no such offset, so it enters at the full σ_n. That gap is why the threshold is $\kappa_n(\sigma_n - 1)/\sigma_n$ and not κ_n. Services must be effectively harder to trade than food. This is the friction asymmetry of Section 7.3 appearing as an explicit condition; at the measured frictions ($\kappa_n = 0.90$ against $\kappa_a = 0.05$) it holds by an order of magnitude.

Because distance affects sectoral sorting only through ρ, the employment gradient $d\ln u/d\ln\tau = \eta\,(d\ln\rho/d\ln\tau) < 0$ and the TFP gradient $d\ln A_a/d\ln\tau > 0$ reproduce the reduced-form patterns of Section 5: villages closer to the city do more of their work in services and have lower agricultural productivity.

I.3 The Welfare Gradient and Its Decomposition

Substituting the three price gradients into (I6) decomposes the welfare gradient into four named channels:

$$\frac{d\ln V}{d\ln\tau} = \underbrace{(s_n - \alpha_n)\,\frac{d\ln p_n}{d\ln\tau}}_{\substack{\text{service terms of trade}\\ \text{(<0: proximity gain)}}} \;\underbrace{-\,(s_a - \alpha_a)\,\kappa_a}_{\substack{\text{farm-gate food price}\\ \text{(<0: proximity gain)}}} \;\underbrace{-\,\alpha_u\,\kappa_u}_{\substack{\text{delivered city goods}\\ \text{(<0: proximity gain)}}} \;\underbrace{-\,\mathcal{E}\,\frac{d\ln\rho}{d\ln\tau}}_{\substack{\text{externality drag}\\ \text{(>0: proximity curse)}}}\,. \tag{I11}$$

The first three terms are the standard gains from market access: a connected village sells its services and food for more (for a net food exporter, $s_a > \alpha_a$ and $s_n > \alpha_n$, the empirically relevant case), and buys its city goods cheaper. All three are losses when distance increases. In contrast is the single term the externality adds. At $\varphi = 0$ the relationship between distance and income is unambiguously negative, so the externality is the only potential source of any curse.

Using the trade balance $(s_a - \alpha_a) + (s_n - \alpha_n) = \alpha_u$ to eliminate the food term, the four channels collapse to:

$$\frac{d\ln V}{d\ln\tau} = \big[(s_n - \alpha_n) - \mathcal{E}\big]\,\frac{d\ln\rho}{d\ln\tau} \;-\; \alpha_u\,(\kappa_a + \kappa_u). \tag{I12}$$

Distance hurts the village through exactly two channels. The first works through the relative price ρ. The bracket in front of it is the village's own response, its terms-of-trade gain net of the externality drag, and it multiplies the equilibrium price gradient in (I9). The second is the cost of trading with the city at all. Distance lowers the price the village receives for what it sells, at rate κ_a, and raises the price it pays for what it buys, at rate κ_u, and neither wedge depends on anything the village does. By the trade balance its net exports are exactly what it spends on city goods, α_u, so both rates apply to that same share and the loss is $\alpha_u(\kappa_a + \kappa_u)$ per log point of distance.

The service friction is missing from that second channel because it never sets the village's service price. It weakens the demand the city sends, and the price that follows is the relative-price gradient, so κ_n is already inside the first channel. That second channel involves no externality at all, and it is the ordinary gain from market access that any trade model delivers. It exists only because food and city goods carry frictions of their own. With a single friction, κ_a and κ_u would both be zero, distance would work through the relative price alone, and this level effect would disappear.

I.4 Proof of Proposition 1

Under condition (I10) the relative-price gradient is strictly negative, so (I12) turns into a condition where real income increases with distance if and only if

$$\underbrace{\mathcal{E}\left|\frac{d\ln\rho}{d\ln\tau}\right|}_{\text{externality loss}} > \underbrace{(s_n-\alpha_n)\left|\frac{d\ln\rho}{d\ln\tau}\right|}_{\text{terms-of-trade gain}} + \underbrace{\alpha_u\,(\kappa_a+\kappa_u)}_{\text{market-access gain}}\ . \tag{I13}$$

Condition (I13) is a threshold in φ. Both $\mathcal{E}$ and the denominator of (I9) are linear in η, and $1/\eta = A + (1-\gamma_a)\frac{u}{1-u} - \varphi\frac{u}{1-u}$ is linear in φ, so holding the equilibrium objects (u, s_a, s_n, θ_c) fixed at their values in a given equilibrium, (I13) holds at equality at exactly one externality value, and rearranging (multiplying through by $1/\eta > 0$ and collecting φ) gives the closed form solution for the sign reversal threshold. Because those objects themselves vary with φ in equilibrium, the threshold is a pointwise characterization evaluated at each village's own equilibrium rather than a proof of a unique global crossing; the numerical sweeps below and in Appendix H find a single sign change along the φ path. Define the market-access-to-terms-of-trade ratio and the gain composite

$$b \equiv \frac{\alpha_u(\kappa_a+\kappa_u)}{\theta_c\big[\kappa_n(\sigma_n-1)-\kappa_a\sigma_n\big]}, \qquad S \equiv (s_n-\alpha_n) + b\,\big(\theta_c\sigma_n + \theta_l\, s_a\big);$$

then proximity is a curse if and only if $\varphi > \varphi^*$, with

$$\varphi^* = \frac{S\left[(1-\gamma_a) + A\,\dfrac{1-u}{u}\right] \;+\; b\,\gamma_n\,\dfrac{k-1}{k}\,\dfrac{1-u}{u}}{S + s_a\,(1-b\,\theta_l)}, \tag{I14}$$

provided the denominator is positive, which holds whenever b is small ($b \approx 0.01$ at the measured frictions and shares). With frictionless food and city goods, $\kappa_a = \kappa_u = 0$ so $b = 0$ and $S = s_n - \alpha_n$, and (I14) reduces, using $s_a + s_n = 1$, to

$$\varphi^*\big|_{\kappa_a=\kappa_u=0} = \frac{(s_n-\alpha_n)(1-u)}{s_a\,u\,\eta} = \frac{s_n-\alpha_n}{1-\alpha_n}\left[(1-\gamma_a) + A\,\frac{1-u}{u}\right], \tag{I15}$$

the frictionless expression quoted in Proposition 1. The threshold is a local object, evaluated at each village's own shares, but the sign-reversal logic is exact at every village, which establishes the proposition.

I.5 Numerical Confirmation

The analytical gradients are confirmed against full numerical solutions of the clearing condition (I8) on a distance grid at an illustrative calibration sharing the quantitative model's agricultural labor share ($\gamma_a = 0.304$, $\gamma_n = 0.92$, $k = 4$, $\sigma_n = 4$, $\kappa_n = 0.65$, $\kappa_a = 0.06$, $\kappa_u = 0.10$, $(\alpha_a, \alpha_n, \alpha_u) = (0.85, 0.05, 0.10)$, $\varphi = 0.65$). The analytic gradients (I9) and (I12) match the numerical solutions at every grid point:

	$d \ln \rho / d \ln \tau$ analytic / finite diff.	$d \ln u / d \ln \tau$ analytic / finite diff.	$d \ln V / d \ln \tau$ analytic / finite diff.
$\tau \approx 1.6$	-0.2566 / -0.2566	-0.7166 / -0.7166	$+0.2493$ / $+0.2493$
$\tau \approx 2.8$	-0.2307 / -0.2307	-0.6905 / -0.6905	$+0.1433$ / $+0.1433$

At the model's demand shares (0.667, 0.103, 0.230) and the measured food and city-good frictions ($\kappa_a = 0.05$, $\kappa_u = 0.203$), with the illustrative $\kappa_n = 0.65$, $\sigma_n = 4$, $k = 4$, and $\gamma_n = 0.92$ retained, the welfare gradient crosses zero at $\varphi = 0.376$, and the closed form (I14), evaluated at that equilibrium's own shares, returns the same value; shutting off κ_a and κ_u reproduces the frictionless threshold (I15) (about 0.16). The content of the check is the agreement between the analytic gradients and the finite differences and between the crossing and the closed form; the threshold itself depends on the illustrative sorting and service parameters, and the quantitative model's own reversal is reported in Section 8.2.

J Where the Scale Externality Could Come From

Section 6 estimates the externality as a reduced form: agricultural TFP rises with the size of the agricultural workforce with elasticity φ, and the data say how large φ is but not why. This appendix is an illustration of why. It writes down the simplest village in which a larger workforce raises everyone's productivity, for the three reasons the literature on agglomeration usually separates (Duranton and Puga, 2004): indivisible inputs get shared, markets get thicker, and people learn from one another. The same settlement-size variation that identifies φ also moves things these mechanisms should move, namely whether a village has a tractor or a rice mill, how concentrated its land has become, and how educated its largest farmers are. Where a mechanism leaves a footprint I can see, I use it to size the mechanism's part of φ; where it does not, I say so. The measured parts add up to about 0.166; the rest of the 0.70 the model carries is a residual that the data I have cannot split further. Table J1 collects the facts and the arithmetic.

Throughout, a fact is the coefficient from a regression of an outcome on log individuals initially placed, with the controls of Section 6.1. I divide each by the first stage of 0.200 so that everything is per log point of the agricultural workforce, the units of $\beta = \gamma_a + \varphi$.

J.1 The Village

A village has L farmers and land H. Farmer i works their own farm, so land has diminishing returns at the farm level, and produces

$$y_i = \omega\, \tilde{a}_i\, h_i^{\rho}, \tag{J1}$$

where h_i is the land the farmer operates, $\tilde{a}_i$ is the farmer's productivity, and ω collects the services the farm buys locally. When farmers are alike and hold equal plots, village output is $Y = \omega\tilde{a}\, L^{1-\rho} H^{\rho}$, the technology of the main text with $\gamma_a = 1-\rho$. I therefore set $\rho = 1-\gamma_a = 0.696$, and the division of output between farmers and land is the same as in Section 7.

Three things change when the village is larger. More farmers support more local services, so ω rises. More farmers make it easier to trade land, so land ends up closer to where it is most productive. And more farmers make it more likely that one of them is unusually good, and that the others copy what they do. Each of the three multiplies village TFP, so their elasticities add:

$$\ln Y = \ln A(L) + \gamma_a \ln L + (1-\gamma_a)\ln H, \qquad \varphi = \frac{d\ln A}{d\ln L} = \varphi_S + \varphi_M + \varphi_R, \tag{J2}$$

with φ_S the shared-input part, φ_M the land-market part, and φ_R everything the data do not reach.

J.2 Shared Inputs: Tractors and Rice Mills

A tractor rental business or a rice mill has a fixed cost and needs enough customers to cover it, so a larger village is more likely to have one. Larger settlements are indeed more likely to have a tractor, by 0.35 per log point of the workforce, and a rice mill, by 0.30. If having the service available raises farm TFP by a factor λ_k,

$$\varphi_S = \sum_k \ln\lambda_k\, \frac{d\,\Pr(\text{service } k \text{ present})}{d\ln L}. \tag{J3}$$

The gains are the one input I take from outside. I set $\lambda_t = 1.10$ for tractors, because mechanised land preparation mostly saves labor and affects output through timeliness and cropping intensity rather than yields (Pingali, 2007), and $\lambda_m = 1.08$ for mills, because a mill changes what the paddy fetches and how much is lost after harvest rather than what the field produces. This gives $\varphi_S = 0.056$. Moving the gains across 1.05–1.20 for tractors and 1.00–1.15 for mills moves it between 0.017 and 0.106.

Two remarks. The numbers of tractors, mills, and formal agricultural firms also rise with settlement size (Table J1), as free entry with a fixed cost predicts, but the median village has none of them, so the counts mostly restate the presence margin and I do not use them separately. And tractors and mills are simply the two shared inputs I observe. Threshers, pumps, drying floors, storage, seasonal work gangs, and the upkeep of irrigation channels (Ostrom and Gardner, 1993) have the same logic and no data, so φ_S is a floor for this mechanism rather than an estimate of it. The externality is that a farmer who leaves takes demand with them, and the provider who then exits, or never enters, was serving everyone (Bassi et al., 2022; Caunedo et al., 2026).

J.3 The Land Market

Every household in a program village started with the same two hectares, so whatever concentration of holdings exists today came from trades, and a trade needs a counterparty. A household that wants to expand has to find one that wants to sell, and in a larger village that is easier (Diamond, 1982). Let m be the share of households that have reached the land market. Those that have not still hold their original plot; among those that have, I assume land has moved to where it is most productive, that is, in proportion to $\tilde{a}_i^{1/(1-\rho)}$.

Under these assumptions village TFP is linear in participation,

$$\frac{Y(m)}{Y(0)} = 1 + m\,(g-1), \qquad \text{so} \qquad \varphi_M = \frac{g-1}{1+\bar{m}\,(g-1)}\,\frac{dm}{d\ln L}, \tag{J4}$$

where g is the factor by which TFP would rise if all of the village's land were reallocated efficiently. Three numbers are needed. For g I take a doubling, in the middle of the range the misallocation literature finds, from about 1.4 in Ethiopia (Chen et al., 2022) to 3.6 in Malawi (Restuccia and Santaeulàlia-Llopis, 2023). The level of concentration then says how far the village has got. With land among the households that trade as dispersed as a doubling implies, a top-decile share of 0.30 means $\bar{m} = 0.25$, so a quarter of households have traded; the same numbers imply a Gini of 0.36 against 0.38 observed, a check the setup was not fitted to. Finally, the top-decile share rises by 0.115 per log point of the workforce, which the same setup converts into $dm/d\ln L = 0.137$. Together, $\varphi_M = 0.109$; with g at the ends of the literature's range it is 0.071 or 0.198.

As with shared inputs, the land market is the one market I can see. Hired labor at planting and harvest, credit, and the traders who buy the crop all work better with more participants, and none of them is measured here.

J.4 What Is Left

The two measured parts sum to $\varphi_S + \varphi_M = 0.166$, and to between 0.09 and 0.30 across the ranges above. Against the $\varphi = 0.70$ the model carries that leaves $\varphi_R = 0.53$, and against the point estimates of Section 6 more. I cannot split the residual with the data I have, so I say what it plausibly contains. Part of it is the unmeasured shared inputs and thicker markets just listed, which belong to the two channels above and would raise their floors. The rest is knowledge. A larger village is more likely to contain an exceptional farmer, and farmers copy the practices of their most successful neighbors (Foster and Rosenzweig, 1995; Conley and Udry, 2010). The census offers a hint that scale does operate at the top: the agricultural employers of larger settlements, the farmers who run the larger operations, have 0.23 more years of schooling per log point of settlers, while the schooling of self-employed farmers and of farm employees does not move (Table J1). To give a sense of magnitude, suppose farmers' practices converge a fraction λ of the way to the best farmer's and productivity has a Pareto upper tail, so that the best of L draws improves with L at a rate set by the tail. At the tail I measure in land holdings (1.8, the rank-minus-one-half slope of Gabaix and Ibragimov, 2011), this channel would fill the residual on its own at $\lambda \approx 0.76$, that is, if most of the gap to the best farmer closes.[27] Over the decades these villages have existed, on soils the settlers did not know, that is not implausible, but it is not something these data establish, and I read the residual as the sum of all of the above rather than as evidence for any one of them.

One property of the measured parts is worth carrying back to the main text. A tractor or a mill matters most where a village is on the margin of having one at all, and participation in the land market cannot exceed one, so both measured channels are strongest among small villages and fade in large ones. That is consistent with the noisy evidence in Section 6 that the externality is larger among smaller settlements, and it is a reason to treat the constant φ that the aggregate experiments of Section 9.2 apply to the crowded inner islands as an upper bound.

Table J1: Settlement size, the mechanism outcomes, and the parts of φ they size

Outcome	Mean	Coef. on $\ln Z_v$	Per log point of L	Part of φ
Any tractor (0/1)	–	0.070	0.35	0.033
Any rice mill (0/1)	–	0.060	0.30	0.023
Top-decile share of owned land	0.30	0.023	0.115	0.109
Gini of owned land	0.38	0.013	0.065	(check)
Number of tractors	2.24	1.020	2.28	(entry)
Number of rice mills	1.20	0.220	0.92	(entry)
Number of formal agricultural firms	0.33	0.180†	2.73	(entry)
Years of schooling, agricultural employers	–	0.230	1.15	(hint)
Years of schooling, self-employed and employees	–	≈ 0	≈ 0	(hint)
Shared inputs, φ_S				0.056
Land market, φ_M				0.109
Measured, $\varphi_S + \varphi_M$				0.166
Residual to $\varphi = 0.70$, φ_R				0.534

Note: Coefficients are from regressions of each outcome on log individuals initially placed with the controls and fixed effects of Section 6.1; counts are in levels, presence indicators are linear-probability coefficients, and schooling is years among agricultural workers by employment status. The response per log point of the workforce divides by the first stage 0.200 and, for counts, by the sample mean. The parts of φ use (J3) at $\lambda_t = 1.10$ and $\lambda_m = 1.08$, and (J4) at a reallocation gain of a factor of two; the Gini is the check described in the text, the counts are consistent with free entry but not used, and the schooling rows are the hint about the residual. † not significant at the 10 percent level.

K Measurement of Model Inputs

This appendix describes how I construct each measured input to the model of Section 7. I describe the value-added demand shares, the three trade frictions, the income anchors, the land and urban endowments, the test of the top-level substitution elasticity, and the Engel slopes used in the

[27] With $\tilde{a}_i = a_i^{1-\lambda} a_{(1)}^{\lambda}$ and land in proportion to $\tilde{a}^{1/(1-\rho)}$, the channel's part of φ is $\gamma_a \lambda / ((1-\lambda)\zeta_h)$.

extension of Section 9.2.

K.1 Value-Added Expenditure Shares and the Consumption Crosswalk

The model's demand shares are measured in value-added terms, while the survey measures spending, so I develop a crosswalk that translates one into the other. Because the rural households the shares should describe are those of the transmigration villages, I use the 2002 SUSENAS consumption module and restrict the sample to rural households on the Outer Islands, dividing their spending between food, services, and goods.

The difficulty in converting this to a value-added measure is that the local services bundle service labor together with purchased intermediates. I measure the split from the Economic Census cost schedules of "household-scale food-service producers." Raw and auxiliary materials, which I assume are predominantly locally supplied food, account for 85 percent of intermediate spending, and fuel, transport, and utilities make up the rest. The labor value-added share of output, ℓ, runs from 0.42 to 0.54 across regions and establishment scales, and I take 0.48 as the baseline.

The crosswalk assigns the labor content ℓ of own-village service spending to service demand and moves the remainder into the food share. That remainder also contains the capital value added of the service producer and non-food intermediates such as fuel, transport, and utilities, so assigning all of it to food overstates α_a and understates α_n relative to a complete value-added decomposition; the crosswalk preserves total spending but reclassifies those components toward food. Pricing those components at the village-gate food price rather than at delivered city-good prices strengthens the negative effect of market access on income slightly, since the farm-gate food price falls with distance while delivered city-good prices rise. Appendix M.9 bounds how much this matters: repricing the measured non-food portion at delivered city-good prices attenuates the income gradient by under 10 percent, and moving ℓ changes the results little.

Classifying every service not explicitly consumed in a city as locally produced gives an own-village share of 0.904, but the survey cannot tell where a service was produced, so this treats every locally purchased service as locally produced and is an upper bound. At $\ell = 0.48$ the crosswalk evaluated at that bound puts the outer-island rural shares at $(\alpha_a, \alpha_n, \alpha_u) = (0.694, 0.076, 0.230)$. At the estimated $s^{\text{own}} = 0.532$ they are (0.667, 0.103, 0.230), and these are the values the model carries. A higher sourcing share shifts spending from services toward food, so it lowers α_n and raises α_a, while α_u is untouched because city goods do not pass through the crosswalk.

In the baseline model I measure the value-added shares on outer-island rural households and apply them to every rural location, inner islands included. Appendix M re-measures the entire demand block described above in the 2012 SUSENAS on both a national sample and the outer-island rural sample. Re-estimating the sourcing share on each of these leaves the untargeted income gradient close to its baseline value, so the negative relationship between connectivity and income is not driven by a particular snapshot of household preferences.

The urban budget shares are not measured from the survey; the aggregate accounting identities imply them (Section 7.4), so the survey provides an overidentifying check. I measure the urban food share two ways. Running the same crosswalk on outer-urban households gives 0.522, which is what urban households actually spend. Taking the rural shares instead and moving them to urban income levels, using the relationship between budget shares and expenditure estimated in Appendix K.7, gives 0.560. The second is the object the model's assumption predicts, since each location's shares are fixed at its own income level, and it sits within roughly 3 percentage points of the accounting-implied $\beta_a = 0.595$.

K.2 Trade Frictions

I measure the three sector-specific trade frictions separately. The food and rural service frictions come from an accounting exercise and from the relationship between prices and travel times, and the city-good friction is built as a combination of the other two.

I measure κ_a and κ_n from the 1993 wave of the Indonesia Family Life Survey community and facility module. For each community surveyed, I observe one-way travel times and fares to 9 destinations.[28] I also observe local wages, doctors' fees, and food prices. The median community daily wage is Rp 2,800 and the median rice price Rp 625 per kilogram. Travel cost enters throughout as $\tau = 1 + t$, with t the one-way time in hours, which is the model's own geometry convention and makes the iceberg factor equal one at zero distance. Every exponent below is specific to this metric and to hours as the unit.

Person-embodied services, $\kappa_n = 0.90$. A unit of the rural service is a unit of labor, so delivering it means a person travels, and the friction is that trip's cost as a share of the value the trip carries. The survey lets me price this two ways, and Table K1 reports both.

Route 1 prices a day of service labor delivered to the city. For a round trip of one-way time t and one-way fare f against a daily wage w, the share of a day's gross value consumed in transit is

$$\text{cost share} = \underbrace{2t/8}_{\text{time}} + \underbrace{2f/w}_{\text{fare}}, \tag{K1}$$

with the workday set at 8 hours, so an hour in transit costs $w/8$ and the wage cancels from the time term. Where the cost share reaches one the day trip is uneconomic, and a weekly variant amortizes one round trip over 6 workdays instead. Route 2 prices an urban service consumed by a villager. For the doctor and hospital rows of the facility roster, the cost share is the round-trip travel cost, fare plus transit time valued at the village wage, divided by the transaction value, with the private doctor's posted consultation charge of Rp 3,750 as the denominator.

The two give the same picture. A day trip to the district capital consumes about 40 percent of a day's value and is uneconomic for a quarter of villages. The province capital is infeasible as a day trip for 72 percent of villages, and even a 6-day stay leaves it uneconomic for 38 percent. A doctor

[28] The destinations are the nearest bus terminal, market, public telephone, post office, health facility, and bank, and the subdistrict, district, and province capitals.

visit costs 15 percent of a consultation's value at the median and a hospital visit 40 percent.

I estimate the exponent by pooling every destination pair within a community and fitting $\ln(\text{iceberg}) = \kappa_n \ln(1+t)$ with community fixed effects and standard errors clustered by community. Route 1 gives $\kappa_n = 0.94$ (s.e. 0.02) across all 1,812 delivery pairs, and Route 2 gives 0.89 (s.e. 0.05) across the 995 doctor and hospital pairs. These are conceptually distinct measurements, one denominated in a day's wage and borne by the provider who travels out, the other in a service fee and borne by the buyer who travels in, and they agree. Because there is no spatial price gradient to discipline κ_n, that agreement is the discipline instead. I take $\kappa_n = 0.90$.

Two things bound how much weight the exact figure carries. First, the mapping from a trip's cost share c to the model's iceberg factor is not unique. The headline 0.94 uses $\ln(1 + c)$, which is defined for every pair including the infeasible ones; the mapping implied by the delivered-share normalization, $-\ln(1 - c)$, is defined only on feasible trips, and because the trip cost is a fixed fare plus a time cost linear in t rather than a power law, on those trips the fitted exponent runs between 1.6 and 2.1 without an intercept and returns toward 0.9 once an intercept absorbs the fare, while the semi-log form that matches the fare-plus-time structure most closely gives 0.88 per hour. Second, because the mechanism runs on village-to-city service trade, matching the friction's level at the trips that carry that trade gives a range of roughly 0.4 to 1.2. Both regression estimates and the baseline sit inside it.

Table K1: Person-embodied service frictions from trips (IFLS 1993)

Trip (community medians)	t (h)	fare (Rp)	cost share	pairs	κ_n (s.e.)
Route 1: delivering a day of service labor (cost share $= 2t/8 + 2f/w$)					
Kabupaten capital, day trip	0.50	600	0.40		
Province capital, day trip	2.50	2,375	infeasible		
Province capital, weekly stay	2.50	2,375	0.38		
Pooled fit				1,812	0.94 (0.02)
Route 2: consuming an urban service (cost share = round-trip cost / fee)					
Doctor visit	0.17	200	0.15		
Hospital visit	0.42	500	0.40		
Pooled fit				995	0.89 (0.05)
Baseline					0.90

Note: Trip rows are community medians across the 312 communities. Cost shares are the trip's cost as a share of the value it carries, computed village by village and reported as medians over the villages for which the trip is feasible. Route 1 values travel time at the median daily wage of Rp 2,800; the province-capital day trip is infeasible for 72 percent of villages, and the weekly row amortizes one round trip over 6 workdays. Route 2 divides the round-trip cost by the private doctor's consultation fee; at the seventy-fifth percentile the doctor and hospital shares are 0.28 and 0.81. The pooled fits regress ln(iceberg) on $\ln(1 + t)$ across every destination pair, feasible or not, with community fixed effects and standard errors clustered by community. Across alternative fee denominators the Route 2 estimate ranges from 0.69 to 1.27.

Food, $\kappa_a = 0.05$. Food is homogeneous and the villages are net exporters, so spatial arbitrage sets the farm-gate price at the market price net of transport, $p_{a,v} = \tau_v^{-\kappa_a} p_a$. The survey brackets the friction from two sides (Table K2).

The upper side comes from freight. The same round trip that delivers a day of service labor can instead carry a 50-kilogram rice load, and the round-trip passenger fare over the load's value at the measured rice price is 4 percent to the district capital and 16 percent to the province capital. Those are an order of magnitude below the person-embodied frictions from the identical trips. Converting them to the model's exponent, $\kappa_a = \ln(1 + \text{cost})/\ln(1 + t)$ per village, gives medians of 0.10 to the district capital and 0.12 to the province capital. Both are upper bounds, for two reasons. The 50-kilogram headload is the smallest commercially meaningful lot, since a motorbike carries several times that and a truck far more, spreading the same trip over more value. And cargo tariffs run below the passenger fares I use here.

The lower side comes from prices. Regressing village rice prices on $\ln(1 + t)$ with province fixed effects and standard errors clustered by district gives gradients of -0.006 (s.e. 0.051) to the nearest market, -0.035 (s.e. 0.022) to the district capital, and -0.004 (s.e. 0.015) to the province capital. All are statistically zero and, if anything, faintly negative. That is the farm-gate pattern for a cheap-to-ship staple that villages produce themselves, and it is inconsistent with a friction anywhere near the freight bound.

The baseline $\kappa_a = 0.05$ sits inside this interval, in the lower half toward which the null price gradients point and below the 0.10 to 0.12 freight bound. The moment system disciplines it from both sides within [0.024, 0.12] (Appendix M.7).

City goods, $\kappa_u = 0.203$. The city good is a composite, neither a single shipped box nor a single carried service, so rather than give it a friction of its own I build κ_u from the two frictions I do measure. It is a bundle of physical goods such as fuel, vehicles, clothing, durables, toiletries, and medicines, which ship as freight at the food friction κ_a, together with a smaller set of person-embodied town services, formal health and education, which move with people at the service friction κ_n. Under Cobb–Douglas within the bundle the delivered-price index is $\tau^{s_g \kappa_a + s_s \kappa_n}$, so the representative iceberg exponent is the consumption-share-weighted average of the two. In SUSENAS rural households the bundle is 82 percent physical goods and 18 percent person-embodied services, which gives $\kappa_u = 0.82 \times 0.05 + 0.18 \times 0.90 = 0.203$ (Table K2). It ranges from roughly 0.18 to 0.22 across the service-definition and network-utility scenarios.

Building it this way has three consequences. It uses only the two frictions I measure directly. It makes the goods side a mild upper bound, since manufactured imports have higher value-to-weight than the rice on which κ_a is measured. And because it comes from these primitives, κ_u does not depend on σ_u. In the model κ_u prices a real trade channel rather than only the cost of living, because city goods are supplied by the endogenous urban block. The gap between the person friction of 0.90 and this composite goods friction of 0.203 is the asymmetry the mechanism runs on, and Appendix M.8 shows what breaks when either friction is set to the other's value.

Table K2: Goods frictions from prices and freight

Measurement	Statistic	Value
Food, κ_a: homogeneous, village-exported; farm-gate $p_{a,v} = \tau_v^{-\kappa_a} p_a$		
Freight accounting, kabupaten capital	round-trip fare / 50 kg rice value	0.04
Freight accounting, province capital	round-trip fare / 50 kg rice value	0.16
Rice-price gradient, nearest market	$d \ln p_a / d \ln(1+t)$, province FE	-0.006 (0.051)
Rice-price gradient, kabupaten capital	"	-0.035 (0.022)
Rice-price gradient, province capital	"	-0.004 (0.015)
Baseline		$\kappa_a = 0.05$
City goods, κ_u: a bundle of physical goods and person-embodied services		
Physical goods (fuel, vehicles, clothing, durables)	ship as freight, $\kappa_a = 0.06$	share 0.82
Person-embodied services (health, education)	move with people, $\kappa_n = 0.90$	share 0.18
Representative city good	$s_g \kappa_a + s_s \kappa_n$	$\kappa_u = 0.21$

Note: The freight accounting prices the round-trip passenger fare against a 50-kilogram rice load valued at the community's measured rice price. Rice-price gradients are estimated at the village level, in the same $\tau = 1 + t$ units as the service frictions. The city-good friction is a consumption-share-weighted mix of the goods and service frictions, $\kappa_u = s_g \kappa_a + s_s \kappa_n$, with the goods and service shares (s_g, s_s) measured on SUSENAS rural households. Across the service-definition and network-utility scenarios κ_u ranges from 0.19 to 0.23.

K.3 Income Anchors

I measure the rural and urban per-capita income anchors on the same outer-island samples as the demand shares, and the urban-to-rural ratio is 1.75. Location-level incomes are observed inputs to the inversion, so the anchors do only one job. The rural anchor converts the village-GDP series into the model's wage units, which is the normalization of Section 7.4.

K.4 Land Endowments and the Household-Holdings Gap

Transmigration villages carry their program land allocations. Outer villages carry agricultural land from the 2003 agricultural census collapsed to the village level, matched to the geometry for 74 percent of locations (85 percent population-weighted), with district-median land-per-worker imputed to the remainder and an upper cap at 1.5 times polygon area. The census concept is *household* landholdings, which misses estate, plantation, and fishery agriculture. The gap appears as populated locations, with median population around 2,400, that have near-zero recorded household land, and they are concentrated in the Sumatran and Sulawesi plantation belts and the fishery provinces.

These locations would become spuriously cheap service suppliers in the model. With no measured land, agriculture cannot absorb their labor, so service wages fall toward zero, and because a CES competition aggregate is set by its cheap tail they would dominate the service market far beyond their size. Lower-winsorizing land-per-worker at 0.02 hectares per person, which touches 7,178 locations or 12 percent of the outer sample, removes the artifact and moves the outer-island aggregate agricultural share onto its census value. I report floors of 0.05 and 0.10 as robustness checks. The transmigration villages themselves are land-rich under the household concept, at a median of 0.50 hectares per person against 0.30 for outer locations, which is consistent with the program allocation.

K.5 Urban Endowments

The endogenous urban block requires the same endowment pair as the villages. Urban workforces are location populations from the settlement layer, floored at 10 workers for a small number of tiny administrative units. Urban agricultural land comes from the same 2003 agricultural census. About 7,200 of the 8,759 administratively urban locations match directly, with the district median imputed to the remainder, and the functionally urban locations of Section 7.4 are ordinary location polygons that carry their own census land. The same land-per-worker floor applies. Urban agricultural shares are far from zero, at 0.26 for the median administrative city and 0.36 for the functionally urban locations, which is why the census land matters. It is what lets the model's cities produce the food their measured employment says they produce, and lets the national food account close without phantom imports (Appendix L.4).

K.6 The Top-Level Substitution Elasticity

The Cobb–Douglas assumption across the three consumption groups is a testable restriction, and the same survey that measures the shares tests it. The test conditions on total expenditure throughout, so it disciplines the price margin of demand and not the Engel slopes, which are measured from the expenditure variation the test holds fixed (Appendix K.7). Under CES(σ_c) over food, services, and city goods, the log ratio of the city-good to the food expenditure share is taste $-\ (\sigma_c - 1)\ln(P_u/P_a)$. City goods carry a higher trade friction than food, so their price rises relative to food in more isolated villages, and that makes the restriction's sign sharp. Under Cobb–Douglas the share ratio does not respond to isolation at all. Under price-elastic demand it tilts away from city goods as they grow dearer, sharply so at $\sigma_c = 2$, the value at which the road counterfactuals' rural gains would largely vanish.

In the 2002 SUSENAS I match household shares over the (a, n, u) groups to travel times at the ten-digit village code, and regress the log share ratio on travel time to the nearest city with log total-expenditure controls, district fixed effects, and survey weights. Table K3 reports the results. The coefficient is statistically indistinguishable from zero and, if anything, of the opposite sign to the price-elastic prediction, and the implied σ_c sits at or just below one on every specification. Cobb–Douglas is therefore not rejected, while the price-elastic value $\sigma_c = 2$, which predicts $b_\tau = -0.05$ per hour, is rejected at roughly four standard errors. The result survives restricting the city-good group to core tradables, which removes administered-price health, education, and telecommunication services, and it passes a placebo on those administered services. The design's contribution is to test the top-level restriction using within-country variation in isolation in a developing economy.

Table K3: The top-level substitution elasticity σ_c

Specification	b_τ (s.e.)	implied σ_c (s.e.)	N
$\ln(s_u/s_a)$, all rural	+0.0087 (0.0144)	0.83 (0.29)	33,462
$\ln(s_u/s_a)$, outer islands	+0.0107 (0.0150)	0.79 (0.30)	18,982
core goods/food, excl. durables	+0.0021 (0.0183)	0.96 (0.37)	33,449
core goods/food, incl. durables	+0.0062 (0.0190)	0.88 (0.38)	33,449
placebo: administered services/food	−0.0268 (0.0349)	n/a	26,092

Note: 2002 SUSENAS households, 97 percent of them rural, matched to travel times at the ten-digit village code. Each row regresses the stated expenditure-share object on travel time to the nearest city, controlling for log total expenditure, with district fixed effects and survey weights. b_τ is the travel-time coefficient and σ_c is its conversion through the measured delivered-goods price gradient. A value of $\sigma_c = 2$ implies $b_\tau = -0.05$ per hour. Regressing each share level separately rather than the ratio gives implied values of σ_c of 0.95, 1.00, and 1.05 for city goods, food, and services. The placebo row uses administered-price services, whose delivered price does not vary with isolation.

K.7 The Engel Slopes and the Extended Demand System

This section describes the demand system used in the extension of Section 9.2. I take in turn the measured Engel slopes, the estimated curvature, the preferences behind them, the anchoring convention, and the evidence on functional form.

Slopes. The slopes come from the same 2002 SUSENAS sample as the share levels. For each of the three consumption groups I regress the household's budget share on log per-capita expenditure with desa fixed effects, survey weights, and district-clustered standard errors, which is the Working-Leser form. The slope is the level response of the share at the rural mean, and I measure the curvature of the Engel curve separately below. Table K4 reports the estimates. The raw slopes are −0.166 for food, +0.037 for local services, and +0.129 for the city bundle per log point of expenditure, and they sum to zero by construction.

Table K4: The measured Engel slopes

	Food	Services	City bundle
Budget share	0.586	0.179	0.235
Engel slope b_g^{raw}	−0.1655	+0.0367	+0.1288
	(0.0041)	(0.0029)	(0.0042)
Value-added slope b_g	−0.1553	+0.0265	+0.1288
Observations		34,423	

Note: 2002 SUSENAS rural households; desa fixed effects, survey weights, district-clustered standard errors. The value-added row applies the crosswalk transform of equation (K2).

The extension's goods are value-added concepts while the survey measures spending on final goods, so the raw slopes pass through the same crosswalk as the levels (Appendix K.1). This is the empirical counterpart of the mapping described below, under which PIGL demand over final goods produced with Cobb–Douglas intermediates implies PIGL demand over value added with the same curvature, so the crosswalk moves the slopes between the two systems and leaves ε alone. A fraction f of marginal service spending is service labor and the remainder is food ingredients, with

$f = \alpha_n/e_n = 0.72$ the value-added retained fraction at the estimated sourcing share and $e_n = 0.143$ the raw service expenditure share, so

$$b_n = f\, b_n^{\text{raw}}, \qquad b_a = b_a^{\text{raw}} + (1-f)\, b_n^{\text{raw}}, \qquad b_u = b_u^{\text{raw}}, \tag{K2}$$

which gives the slopes the extension carries, $(b_a, b_n, b_u) = (-0.155,\, 0.027,\, 0.129)$. The slopes are measured across households spanning several log points of expenditure, so the income movements the counterfactuals generate stay far inside the measured support, and the shares stay far from their bounds at every reported solution. The urban comparison of Appendix K.1 provides an overidentifying check, since extrapolating the rural anchors along the estimated curve across the measured urban-to-rural income gap reproduces the accounting-implied urban food share within roughly 3 percentage points; the type intercepts nearly lie on one common curve.

The curvature ε. I estimate the curvature with the estimator of Fan et al. (2023), regressing the log budget share on log per-capita expenditure with desa fixed effects and district-clustered standard errors. Under the PIGL form with the asymptotic food share set to zero, the convention of Fan et al. (2023), the log food share is linear in log expenditure and this recovers ε as the negative of its slope. With a positive asymptotic food share the same regression identifies $\varepsilon(\bar{s}_a - \omega_a)/\bar{s}_a$ instead, which at the anchoring described below equals the Working-Leser slope divided by the share and carries no separate information about ε; the curvature is then identified only from how the slope changes with expenditure, which is the functional-form evidence at the end of this section. On the same sample it gives -0.340 (s.e. 0.011) for food, so $\varepsilon = 0.34$, with luxury elasticities of $+0.19$ for services and $+0.51$ for the city bundle (Table K5). Because the estimator is identical, the comparison with the literature is direct. Fan et al. (2023) obtain 0.33 by least squares and 0.395 instrumenting expenditure on Indian household data, Boppart (2014) obtains 0.18 to 0.25 on United States household data, and Alder et al. (2022) obtain 0.34 to 0.90 on long-run sectoral shares for four rich countries. Indonesian rural households in 2002 sit almost exactly on the Indian value.

Table K5: The PIGL curvature and the functional-form evidence

	Food	Services	City bundle
Panel A. Share elasticities and the curvature			
$d \ln s_g / d \ln e$	-0.340	$+0.188$	$+0.511$
	(0.011)	(0.020)	(0.014)
Implied ε		0.340 (0.011)	
Panel B. Functional form			
Quadratic term	-0.022	-0.003	$+0.025$
	(0.004)	(0.004)	(0.004)
Per-equation exponent κ_g	$+0.28$	-0.14	$+0.43$
Common κ, unrestricted		$+0.31$ (0.042)	
Common κ, restricted to $(-1, 0]$		0.00 (the PIGLOG corner)	

Note: 2002 SUSENAS rural households, desa fixed effects, survey weights, district-clustered standard errors. The elasticity panel regresses log budget shares on log per-capita expenditure, the estimator of Fan et al. (2023); ε is the negative of the food coefficient. The functional-form panel reports the quadratic Working-Leser coefficients and the profiled NLS exponent described in the text.

Preferences. The extension's budget shares follow from the PIGL indirect utility

$$V_\ell(p, e) = \frac{1}{\varepsilon}\Big(\frac{e}{A_\ell(p)}\Big)^{\varepsilon} - \sum_g \nu_g \ln p_{g,\ell}, \qquad A_\ell(p) = \prod_g p_{g,\ell}^{\omega_{g,\ell}}, \tag{K3}$$

by Roy's identity, with the composite CES price indices of Section 7.3 standing in for the service and city-good prices by two-stage budgeting. Two properties of this class are the reason I use it rather than another non-homothetic system. It aggregates, so a representative household per location is still a legitimate object even though demand is non-homothetic at the household level. And combined with Cobb–Douglas intermediates it maps demand over final goods into demand over sectoral value added with the same curvature ε. The second property is what lets me estimate ε on household spending over final goods and then use it in a model whose sectors are value added, and it is not a generic feature of non-homothetic demand (Herrendorf et al., 2013). The aggregation property is established in Boppart (2014). The value-added mapping is not: Boppart (2014) is explicit that moving from final goods to sectoral value added requires a model of the production process, and Herrendorf et al. (2013) show that the two perspectives differ in general. What delivers it here is the Cobb–Douglas intermediate structure of the crosswalk (Appendix K.1), under which the value-added shares are affine in the same $x^{-\varepsilon}$ as the final-good shares, the same route taken by Fan et al. (2023) and Peters et al. (2026). The share equations are

$$s_{g,\ell} = \omega_{g,\ell} + \nu_g\, x_\ell^{-\varepsilon}, \qquad \textstyle\sum_g \omega_{g,\ell} = 1, \quad \sum_g \nu_g = 0, \tag{K4}$$

with x_ℓ real expenditure per capita. The logarithmic price term in (K3) is the $\sigma_c \to 1$ case of Boppart (2014)'s price aggregator, the case the price test of Appendix K.6 supports, so the system is Cobb–Douglas in prices and PIGL in income, and the main model's homothetic demand is the case with the loadings ν_g set to zero. The loadings are pinned by the measured slopes at the rural mean, so the Engel slope at income x is $b_g\,(x/\bar{x}_r)^{-\varepsilon}$ and flattens as households grow richer; at the urban mean the implied slopes are about 14 percent flatter than the rural ones. The intercepts $\omega_{g,\ell}$ anchor each location at its observed baseline budget and are the asymptotic budget shares; at the rural mean they are (0.21, 0.18, 0.61), summing to one automatically because the slopes sum to zero. The PIGLOG/Working-Leser system is the $\varepsilon \to 0$ limit.

The anchoring convention and what identifies it. The invariance of the baseline to the demand parameters rests on this convention, so I state it precisely. The slopes are identified within desa, where the fixed effects compare richer and poorer households in the same village facing the same prices. The estimate is therefore a pure income response, uncontaminated by the spatial variation in delivered prices that the model prices through its own trade structure. The fixed effects absorb the cross-village variation in share levels, and the location-specific intercepts $\omega_{g,\ell}$ are the model counterpart of those fixed effects, since every location sits at its observed budget at its own baseline income and the Engel curve governs movements away from that point. Applying the within-village slope to counterfactual changes in a location's income is the experiment the estimate matches. Applying it to baseline level differences across locations would extrapolate onto the margin the design deliberately excludes.

Two things follow. First, the baseline equilibrium, the inversion, and every targeted moment are invariant to (b_g, ε), which operate only when counterfactuals move real incomes. Second, the extension says nothing about whether baseline budget composition varies with income across locations. What that omits is bounded by the curse's income spread times the food slope, roughly a percentage point of budget shares across program villages, and the 3-percentage-point urban check above bounds it across types.

This construction involves one approximation. Because shares are non-linear in expenditure, the aggregate share of a population with dispersed incomes is not the share of the household at mean income, and the exact aggregate carries a term in the dispersion of income within the location (Boppart, 2014; Peters et al., 2026). The model has one household per location while the slopes applied to it are estimated across households, so the income margin here is the response of a household at the location's mean income rather than the exact aggregate response. The two differ by a factor below one that falls with within-location income dispersion, and at the estimated ε that factor is roughly 3 to 5 percent for log income standard deviations of 0.5 to 0.7. Because the anchoring already fixes each location's baseline budget at its observed value, the approximation affects the slope and not the level, and it scales the income margin uniformly rather than distorting comparisons across experiments or across values of φ. At the magnitudes the counterfactuals produce it moves the reported aggregates by well under a hundredth of a percentage point.

Functional form. Write κ for the exponent on expenditure, so that shares are affine in x^{κ}. The Boppart (2014) class is $\kappa = -\varepsilon \in (-1, 0]$ and PIGLOG is $\kappa = 0$. Three pieces of evidence discipline the choice.

A quadratic term added to the Working-Leser regressions is statistically significant for food

(-0.022, $t = -5.7$) and the city bundle ($+0.025$, $t = 6.5$) and zero for services. That says the magnitude of the food-share elasticity rises with income, from 0.20 at the rural tenth percentile to 0.37 at the ninetieth. The Boppart (2014) form with a positive asymptotic food share implies the opposite direction, a magnitude that falls as the share approaches ω_a, so the disagreement is with the form's curvature rather than with a constant-elasticity restriction. Profiled nonlinear least squares over the full share system puts the unrestricted exponent at $\kappa = +0.31$ (s.e. 0.04), and constrained to the admissible region of the Boppart (2014) class the estimate is the PIGLOG corner $\kappa = 0$. The three candidate fits differ by two tenths of a percent of residual variance end to end, with per-equation exponents of $+0.28$ for food, -0.14 for services, and $+0.43$ for the city bundle, so no one-parameter member fits the curvature exactly.

The choice is second order. Predicted demand responses coincide at every income move the counterfactuals in the paper generate; the three fits separate only at income changes an order of magnitude larger, where for a change of roughly 60 percent in a location's income the predicted food-share response is $+0.17$ under the operative form, $+0.15$ under PIGLOG, and $+0.12$ under the unrestricted exponent. I carry the Boppart (2014) form because it is the standard in this literature and because its estimator delivers the cross-country comparison above. Its shares are bounded only locally, as for any system affine in x^{κ}, but they stay far from their bounds over the measured support and at every reported solution. Because the system is anchored at each location's observed budget and the loadings are pinned by the measured level slopes, the income response the counterfactuals run on is the measured slope at baseline whatever the curvature; ε governs only how that slope flattens as income rises, which is why the three fits coincide at counterfactual-sized income moves.

Welfare. With (K3) in hand welfare is exact, but the tables need a deflator and there are two candidates that sit far apart. The price aggregator $A_\ell(p)$ in (K3) weights prices by the asymptotic shares $\omega_{g,\ell}$, while the reported real incomes deflate by the observed baseline shares $\bar{s}_{g,\ell}$. I therefore set out which one the tables use. Write $\Delta x_\ell = \Delta \ln e_\ell - \sum_g \omega_{g,\ell} \Delta \ln p_{g,\ell}$ for the change in the model's own real expenditure. Equating utility across price vectors, the equivalent variation at baseline prices is

$$\Delta_\ell^{EV} = \frac{1}{\varepsilon} \ln\Big[e^{\varepsilon\,\Delta x_\ell} - \varepsilon \sum_g \big(\bar{s}_{g,\ell} - \omega_{g,\ell}\big)\, \Delta \ln p_{g,\ell}\Big], \tag{K5}$$

using $\nu_g x_\ell^{-\varepsilon} = \bar{s}_{g,\ell} - \omega_{g,\ell}$ at the baseline. Expanding to first order, the ω terms cancel exactly and (K5) collapses to $\Delta \ln e_\ell - \sum_g \bar{s}_{g,\ell} \Delta \ln p_{g,\ell}$, which is the object the tables report. The asymptotic shares therefore never enter the welfare measure, and the familiar result that the first-order welfare effect of a price change is its baseline expenditure share survives the non-homotheticity intact. What the curvature adds is the second-order term $\varepsilon\big(\Delta x_\ell B_\ell - \frac{1}{2} B_\ell^2\big)$, with $B_\ell \equiv \sum_g (\bar{s}_{g,\ell} - \omega_{g,\ell}) \Delta \ln p_{g,\ell}$, which at the estimated ε and the price movements these counterfactuals generate is on the order of thousandths of a percentage point, against reported effects of tenths of a percent.

L Computation

This appendix documents how the 69,225-location joint equilibrium is represented, solved, selected, and audited, and how the estimation is organized around the solver.

L.1 Exact Dense Representation

The service demand system requires, at every price vector, each buyer's competition index over all rural origins and each origin's demand sum over all buyers, bilateral objects on a 69,225 × 60,466 travel-cost matrix: all 69,225 locations as buyers against the 60,466 locations as potential service origins. The stored matrix keeps every location column regardless of classification; under the baseline functional-urban classification the 11,868 reclassified location polygons are removed from the rural supplier pool by a mask (their variety weight is set to zero and their price left inert), so the active rural-origin set is the 48,598 rural villages while the matrix dimensions are classification-invariant.

The natural way to economize would be sparsity: truncate each buyer's sourcing at a travel-time cutoff and bound the discarded tail. That shortcut fails here, and measurably so: computing the exact tail share of every buyer's competition index at baseline prices shows that at every cutoff up to 8 hours one-way, population mass at distance grows faster than the $\tau^{-\kappa_n(\sigma_n-1)}$ kernel decays, leaving median tail shares above 1 percent and, for isolated buyers, near 100 percent. The implementation therefore uses the exact dense system: the two required orientations of the travel-cost matrix are stored as single-precision arrays of about 17 gigabytes each, the kernel powers are built in memory, and a full evaluation of all demand aggregates is two matrix-vector products taking 1 to 2 seconds on a workstation. There is no truncation error to defend.

Two conventions matter for exactness. The fast-marching travel-time matrix is numerically asymmetric (median 0.2 percent, ninety-ninth percentile 3.3 percent), so each location's statistics are computed from its own matrix row (buyers' competition indices from buyer rows, origins' demand sums from origin rows), which is why both orientations are stored; this makes the general-equilibrium solver reproduce the per-village estimation battery exactly when the rest of the world is frozen. And transmigration-village terms are carried in dedicated double-precision blocks outside the single-precision products, so the objects the moments are built on never see single-precision rounding.

L.2 Fixed Point, Continuation, and Audits

Given the kernel operators, the equilibrium is computed by a damped two-operator iteration: service prices imply competition indices and demand sums; each location's two-equation system (sorting, agricultural first-order condition, service clearing) then solves in closed form up to a scalar per location, vectorized across all locations; updated wages imply updated prices. Adaptive damping starts at one half. The convergence tolerance of 5×10^{-8} in maximum log-wage change is set just above the single-precision noise floor of the matrix products (about 1.5×10^{-6} relative, which propagates to roughly 3×10^{-8} in the fixed-point residual); tightening beyond it wastes iterations wandering the noise floor. A full solve from a warm start converges in 50 to 90 iterations, about 11 minutes.

Selection follows the protocol of Section 7.3: continuation in φ from the externality-free equilibrium, in 13 steps refined near the top of the path, warm-starting each step. At the endpoint the audit battery runs: a perturbation restart (log-wage noise injected and re-solved) and a full cold start with no continuation. At the calibrated point both return the identical equilibrium, with maximum log-wage discrepancies of 4×10^{-8} for the perturbation and 3×10^{-7} for the cold start and no entry of the rural wage vectors or the urban price and wage vectors differing beyond 10^{-6}; the audit covers both blocks of the joint system. Uniqueness at system scale cannot be certified in general; what is certified is that the reported equilibrium is not an artifact of its construction path, and the audit battery re-runs at every reported point, including every counterfactual.

L.3 Estimation and the Inversion

The estimation never solves for the fundamentals inside its loops; it exploits the architecture instead. The inversion does not involve θ, but it does involve the demand shares, so each trial of the sourcing share s updates $(\alpha_a, \alpha_n, \beta_n)$ through the crosswalk and the adding-up identity and re-inverts the non-program world once (one fixed point over competition indices at observed incomes), cached for the trial. The composition moment is then evaluated in the joint equilibrium of Appendix L.4: the agricultural scale T_a concentrates out by bisection on the program agricultural-employment level, every joint solve warm-starting from its predecessor so that an evaluation costs a handful of damped iterations rather than a cold start, and the outer root closes the structural transformation moment in s. The two-parameter variant of Appendix M.2 adds an inner root that sets θ on the composition moment before the outer root closes the income moment.

A small-open alternative evaluates the program side as a *battery*: each of the 1,147 program villages solves as its own two-equation system (sorting, agricultural first-order condition, service clearing) against the fixed world objects it faces (its demand access, its competition index, its farm-gate price, its city-good price index), which takes seconds and remains the frame for the single-village experiment of Section 9.1. The program's general-equilibrium footprint is visible in the estimate itself: the battery moment is exact at a sourcing share of 0.67 rather than the joint frame's 0.53, which is why the baseline moment is evaluated in the joint equilibrium, the frame every counterfactual solves in. Table L1 reports the targeted moments and the estimated parameters.

Table L1: Targeted moments and estimated parameters

Moment	Data (s.e.)	Model
Agricultural employment share, TM	0.648	0.648
b_{ST}, reduced-form measure, full sample	+0.0053 (0.0009)	+0.0053
Imposed or set:	$\theta = 12.5$ (Bryan and Morten 2019); $\sigma_n = \sigma_u = 5$	
Estimated:	$s_{\text{own}} = 0.533$ (T_a concentrated out)	

Note: Moments computed identically in data and model: non-agricultural employment share (b_{ST}) on the reduced-form measure of market access, island and cohort fixed effects, district-clustered standard errors, 1,147 transmigration villages; the model moment is computed in the joint equilibrium (Appendix L.4). Estimation on the functional-urban inverted world with the city-good price index from the inverted urban block.

L.4 The Joint Equilibrium, Certification, and the Food Account

Counterfactuals solve the full two-block system: rural service prices and city variety prices update jointly, with the city-good gravity system over 20,627 suppliers stored in half precision (about 6 gigabytes) and applied through single-precision chunked products; the whole solve holds in about 45 gigabytes and converges in 30 to 70 damped iterations. Certification re-solves the entire economy forward from the recovered fundamentals and compares to observed incomes: mean absolute error 0.6 percent across rural locations and 0.6 percent across urban ones. The food market, closed by Walras' law in the solver, is audited by direct accounting at the certified baseline: total food spending matches total farm-gate revenue to 0.2 percent (within the audit's own reconstruction tolerance), iceberg losses on shipped food are 0.45 percent of production, and the sell-side farm-gate convention, exact for net exporters, covers locations holding 89 percent of the rural population. The net-importing remainder, concentrated among the most transformed, best-connected locations, including a third of the program villages, is mispriced by at most the arbitrage band $2\kappa_a \ln \tau$, under 4 percent where the importers cluster.

M Quantitative Robustness

This appendix collects the robustness of the quantitative results to every input that could reasonably be measured differently: the mapping of the externality estimate from the data into the model, the elasticity pair (θ, σ_n), the variety weighting, the demand block's vintage and sample, the food friction, the friction asymmetry, the bundle labor share, the outer land floor, the outer-externality extrapolation, and the food-market closure. The organizing finding is that the model's untargeted income gradient is an *envelope*: most inputs move offsetting margins, since a change in a measured share moves both village income and the cost-of-living weights and the recalibration of the internal parameters absorbs the remainder, so the gradient moves by far less than the inputs do.

M.1 The Estimator Run on Model-Generated Data

The model imposes $\varphi = 0.70$ through the mapping $\varphi = \hat{\beta} - \gamma_a$, which reads the 2SLS coefficient on agricultural headcount as the model's coefficient on effective agricultural labor. Because workers sort on comparative advantage, instrumented headcount and effective labor need not move one for one, so I validate the mapping by simulation. Each of the 1,147 program villages is solved in the per-village battery frame with its *recorded* settlement size L_0 in place of the homogenized mean, under two land conventions (the common mean $\bar{H}$ held fixed, and the recorded two-hectare-per-household allocation), in the baseline world and in the re-pinned $\varphi = 0$ world; the sorting scale T_a is held at its baseline value throughout, since a common constant cannot absorb the cross-village variation that identifies the regressions. On the model-generated data I then run the estimator of Section 6 exactly: log agricultural output on log agricultural headcount $\pi_{a,v} L_v$, instrumented by $\ln L_0$, island and cohort fixed effects, standard errors clustered by village, with and without the land control.

Table M1 reports the results. The benchmark column, output regressed on effective labor, returns $\gamma_a + \varphi$ exactly, a mechanical identity that certifies the construction. The paper's estimator returns 0.979 and 0.968 against the true 1.004, with Anderson–Rubin sets [0.977, 0.983] and [0.964, 0.974], and in the externality-free world returns 0.316 and 0.318 against a true 0.304: with the land control in place the estimator recovers the structural coefficient to within 2 to 4 percent in both worlds, and does not manufacture an externality where none exists. The one exception is instructive: with recorded land and no land control, the externality-free world returns 0.482, with an Anderson–Rubin set that excludes the true 0.304, because recorded land co-moves with settlement size (0.17 log points per log point of persons placed) and the omitted land loads onto the labor coefficient. That is the model's version of why the empirical specification controls for planted area.

The deviation from the truth is the Roy-selection wedge itself, and the table names it. Because $\ln Q_a = (\gamma_a + \varphi) \ln Z_a + (1 - \gamma_a) \ln H_v$ holds identically, the estimates satisfy $\hat{\beta} = (\gamma_a + \varphi)$ times the 2SLS of $\ln Z_a$ on headcount once the land control is partialled out, and the passthrough column shows effective labor moving 0.96 to 0.98 for one with instrumented headcount at the calibrated externality, and 1.04 to 1.05 for one at $\varphi = 0$. The sign tracks the sorting response: at $\varphi = 0.70$ agricultural returns are nearly constant in labor, larger villages sort *into* agriculture (first stages above one), and the marginal entrants dilute average farming ability, while at $\varphi = 0$ diminishing returns push the marginal workers out and the sector concentrates on those best suited to it. The wedge is small because composition scales with $1/\theta$; at $\theta = 12.5$ even sizable sorting responses move average ability by a few percent. Its direction means the empirical $\hat{\beta}$, if anything, understates $\gamma_a + \varphi$, so the calibrated $\varphi = 0.70$ is conservative on this margin as well.

Two features of the table deserve interpretation. First, the model's first stages (1.44 to 1.80 at the calibrated externality, 0.63 to 0.67 at $\varphi = 0$) are far from the empirical 0.20 (s.e. 0.07): in the model the settlement size is the village's permanent labor endowment, while in the data five decades of demographic churn separate placement from the observed workforce, so only part of the initial size differences survives. Because 2SLS divides the reduced form by the first stage, the recovered coefficient is invariant to this attenuation; the exercise validates the estimand mapping, not the persistence process, and whether the churn is itself selective is an exclusion question outside the model, treated by the balance and mediation evidence of Section 6. Second, dropping the land control moves the records-row estimate from 0.968 to 1.006, the model's version of the with-and-without-area bracket of Section 6, though the two pairs are not the same experiment: the model's land object is the endowment allocation, which in the settlement records moves only 0.17 log points (s.e. 0.03) per log point of persons placed within fixed-effect cells, while the empirical control is chosen planted area.

Table M1: The paper's estimator run on model-generated data

	Truth	Benchmark	IV, land control	IV, no control	First stage	Z_a passthrough
$\varphi = 0.70$, H fixed	1.004	1.004	0.979	0.979	1.439	0.976
			(0.002)	(0.002)	(0.039)	(0.001)
			[0.977, 0.983]	[0.977, 0.983]		
$\varphi = 0.70$, H records	1.004	1.004	0.968	1.006	1.801	0.964
			(0.003)	(0.006)	(0.103)	(0.003)
			[0.964, 0.974]	[0.995, 1.017]		
$\varphi = 0$, H fixed	0.304	0.304	0.316	0.316	0.672	1.039
			(0.001)	(0.001)	(0.011)	(0.002)
			[0.315, 0.317]	[0.315, 0.317]		
$\varphi = 0$, H records	0.304	0.304	0.318	0.482	0.627	1.048
			(0.001)	(0.026)	(0.016)	(0.003)
			[0.317, 0.320]	[0.429, 0.531]		

Note: Each row solves the 1,147 program villages with their recorded settlement sizes L_0 in place of the homogenized mean, under the stated land convention and externality, and then runs the 2SLS specification of Section 6 on the resulting model-generated data. That specification regresses log agricultural output on log agricultural headcount, instrumented by $\ln L_0$, with island and cohort fixed effects and standard errors clustered by village reported in parentheses. Brackets report Anderson–Rubin 95 percent confidence sets. Truth is the model's labor coefficient $\gamma_a + \varphi$. The benchmark column regresses output on effective labor Z_a. First stage is the coefficient of log headcount on $\ln L_0$ within the fixed effects. The Z_a passthrough column runs the same 2SLS with $\ln Z_a$ as the outcome. The sorting scale T_a is held at its baseline value in every row.

M.2 Estimating Both Elasticities

The main text imposes $\theta = 12.5$ from Bryan and Morten (2019) and estimates the sourcing share s_{own} alone on the structural transformation slope. Freeing θ as a second parameter in the per-village battery frame of Table M2, with σ_n held at its set value and the income-access slope on the reduced-form measure as the second moment, moves the estimate to $\theta = 16.3$ and $s_{\text{own}} = 0.84$; the composition slope stays matched throughout, with the sourcing share climbing from 0.67 as θ rises. The data's full-sample income slope, -0.0066, is steeper than the model reaches at any θ up to 30, where it attains only -0.0047: the sorting margin does not reproduce the extreme isolation tail, so the second moment is the interior slope that drops the 5 percent most isolated villages, -0.0047, which the model matches at $\theta = 16.3$. Two things follow. First, the freed elasticity lies above the imposed one, just past the upper end of the Bryan and Morten (2019) confidence interval (9.8 to 15.2), so imposing the migration estimate is conservative for the sorting margin and the curse it drives, not generous. Second, the untargeted predictions barely move: the model's own value-added home share is 0.30 at both the imposed and the freed point. The income moment carries a standard error in this appendix and nowhere else; in the main text it is a prediction, never a target.

Table M2: Calibration perturbations

Perturbation	s^*_{own}	b_{ST}	b_Y (full)	b_Y (trim)	χ
Panel A. Own-share re-estimated at each perturbation					
baseline	0.673	+0.0053	−0.0070	−0.0084	0.30
unit weights $\omega_o = 1$	$0.904^\dagger$	+0.0060	−0.0076	−0.0083	0.41
$\kappa_a \times 0.5$	$0.904^\dagger$	+0.0054	−0.0089	−0.0104	0.31
$\kappa_a \times 2$	$0.471^\dagger$	+0.0043	−0.0020	−0.0031	0.30
$\ell = 0.42$	0.603	+0.0053	−0.0070	−0.0084	0.30
$\ell = 0.54$	0.761	+0.0053	−0.0070	−0.0084	0.30
Panel B. Own-share held at 0.673*; friction symmetry broken*					
κ_n set to goods friction	0.673	−0.0009	+0.0045	+0.0045	0.001
κ_u set to person friction	0.673	+0.0053	+0.0040	+0.0021	0.30

Note: Each row re-solves the model under the stated change to a calibration input. All rows use the per-village frame, in which the own-share estimate is $s^*_{\text{own}} = 0.673$, while the joint-equilibrium estimator of the main text gives 0.533 on its own moment. b_{ST} is the structural transformation slope on the reduced-form access measure, which is the estimation target. b_Y is the income-access gradient on the frozen-kernel measure, reported for the full sample and with the bottom decile trimmed. χ is the model's value-added home share. In Panel A the own-share is re-estimated to hold the target at each perturbation, and rows marked † are those where the estimate reaches a bound of its admissible range $[0.471, 0.904]$ and the target is left slightly off. Panel B holds the own-share at its baseline value and sets the person-embodied service friction κ_n and the goods friction κ_u to each other's values.

M.3 The Isolation Tail

Figure M1 traces how the income slopes in the data and in the model move as the most isolated villages are excluded, from 0 to 10 percent. The data slope flattens as the isolated tail is dropped, because the most remote program villages earn a premium the model does not capture, while the model slope steepens slightly; the two cross near the bottom decile. Section 8.3 shows what the inversion makes of the premium: it loads into the recovered fundamentals of the most isolated villages, which is where the model's null balance result carries its largest residuals.

Figure M1: Income-Access Relationship as Isolated Villages Are Excluded

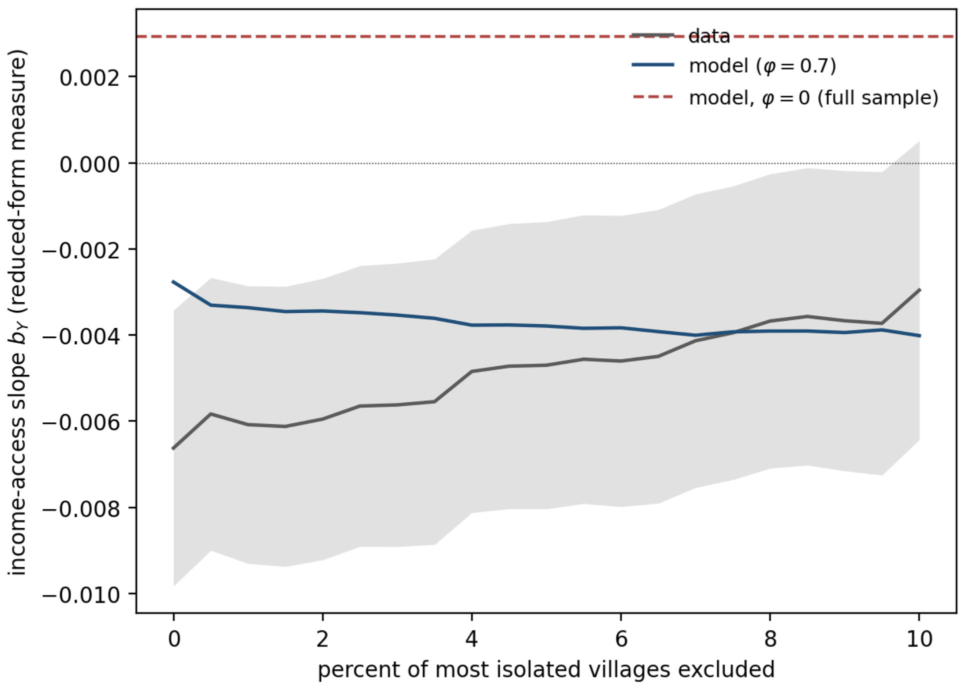

Note: The figure plots the relationship between income and market access for transmigration villages in the data and in the model as the villages lowest in market access are excluded, from 0 to 10 percent, with 95 percent confidence intervals for the data estimates. The red horizontal line marks the full-sample model slope at $\varphi = 0$.

M.4 Variety Weights

The variety weights $\omega_o = N_o/\bar{s}$ make a large location a proportionally larger mass of service varieties rather than one oversized variety, which keeps the inversion invariant to administrative aggregation: splitting a location into two polygons changes no equilibrium object. Dropping the weights and giving every location one variety is the unit-weights row of Table M2: re-estimating under unit weights pushes the own-share to the top of its range and leaves the income gradient essentially where it was ($b_Y = -0.0076$ against the baseline -0.0070). In the two-parameter variant of Appendix M.2 the same change moves only θ (to 17.5) along the weakly identified dimension, with σ_n and every untargeted income slope unchanged. The weighting is a reporting convention, disciplined by replication invariance, not a quantitative driver.

M.5 The Urban Classification

Section 7.4 classifies as urban the administratively urban locations plus every location whose centroid falls inside a satellite-defined functional urban area (the GHS-FUA layer). The administrative classification alone produces an implausibly low urban share for Indonesia, 44 percent against roughly 60 percent in aggregate statistics, because it assigns many peri-urban and urban areas on the outskirts of cities to the rural side; the functional definition brings the share to just under 65 percent. Table M3 reports the estimated point under alternative classifications.

Table M3: The estimated point under alternative urban classifications

Classification	Urban pop. share	b_{ST}	b_Y (estimation)	b_Y (reduced form)
GHS-FUA (baseline)	0.65	+0.0136	−0.0128	−0.0050
outcome-based	0.62	+0.0141	−0.0129	−0.0050
FUA + 10 minutes	0.73	+0.0136	−0.0127	−0.0049
administrative +	0.48	+0.0076	−0.0084	−0.0037

Note: Model slopes are computed at $(\theta, \sigma_n) = (12.5, 5)$. For each classification the service world and the urban block are re-inverted at that point, with the city-good friction held at its consumption-weighted value of $\kappa_u = 0.203$ across classifications. Slopes are computed in the per-village estimation frame of Appendix L.3: b_{ST} and b_Y (estimation) are on the model-consistent access measure, and b_Y (reduced form) is on the reduced-form measure of Section 5. The targeted moment of Table L1 is the reduced-form b_{ST} in the joint equilibrium, which is why the b_{ST} here differs from 0.0053; the baseline row's reduced-form joint-equilibrium counterpart appears in Table 11. "Outcome-based" reclassifies near-city locations whose agricultural share falls below the urban 75th percentile. "Administrative +" applies a minimal correction to the administrative classification, reclassifying locations below the urban median agricultural share and within 30 minutes of a city. Data slopes are estimated as in Table 11, with island and cohort fixed effects and district-clustered standard errors.

Every functional definition of urban, satellite-based, outcome-based, or widened by a commuting buffer, delivers slopes of the same sign and similar size, with the structural transformation slope between 0.0136 and 0.0159 and the reduced-form income gradient between −0.0050 and −0.0062: nothing qualitative hinges on where exactly the boundary is drawn, provided functionally urban locations are not counted as rural service suppliers. Under the administrative classification the same qualitative patterns appear at smaller magnitudes: with functionally urban locations left in the rural supplier pool, they absorb part of the city service demand before it reaches the countryside, and the structural transformation slope roughly halves. One caveat applies to all functional layers: the satellite footprint reflects the contemporary built environment rather than a pre-program one, so the classification conditions the model on the observed urban system; the administrative row, whose definition is closest to predetermined, is the relevant bound on what this choice contributes.

M.6 The Demand-Block Envelope

Every survey-based demand parameter can be re-measured at a different vintage or sample. Re-running the inversion and the own-share estimation on the value-added crosswalk of the 2012 SUSENAS, for outer-island rural households and for the national sample, matches the composition slope at every block and leaves the untargeted income gradient and own-village share close to their baseline values: a decade of preference drift and a change of sample move nothing the paper uses.

The crosswalk itself is not optional, and the accounting shows why. Feeding the model raw expenditure shares instead of their value-added conversion roughly doubles rural own-spending on services, and the economy-wide adding-up identity, which sets the urban service budget from total measured service revenue net of rural own-demand, pays the price. At the raw 2002 shares the implied urban service share collapses from 0.037 to 0.005 and the urban composition gradient inverts (-0.026 against the measured $+0.0147$); at the raw 2012 shares the implied share is negative and the inversion's fixed point does not exist. Raw expenditure weights assign the countryside more service spending than the economy's measured service income can absorb; the crosswalk is an existence condition, not a preference.

M.7 The Food Friction

The trip-cost accounting puts κ_a at 0.05, and the moment system disciplines it from both sides (the κ_a rows of Table M2). At half the measured value the own-share estimate rises to the top of its admissible range and the untargeted gradient steepens slightly ($b_Y = -0.0089$); at double the value the estimate falls to the bottom of the range, the target can no longer be held exactly, and the gradient flattens toward zero (-0.0020). A food friction far from the measured value forces the sourcing share against a bound rather than reproducing the data. The simple model shows what is at stake without this discipline: at stylized parameters, pushing the food friction to 0.10 can extinguish that model's curse entirely (Appendix H), which is why κ_a is measured and its measurement cross-checked against the null staple-price slope.

M.8 The Friction Asymmetry

Section 7.3 states that the gap between the person-embodied service friction ($\kappa_n = 0.90$) and the goods friction ($\kappa_u = 0.203$) is central to the results; Panel B of Table M2 quantifies how, holding the own-share at its baseline value. Setting κ_n down to the goods friction, so that services ship as cheaply as goods, destroys the mechanism's spatial structure: the own-village share collapses toward zero, the structural transformation slope changes sign, and the income gradient turns positive, because geography no longer concentrates city demand on nearby villages. Setting κ_u up to the person friction instead leaves the composition slope untouched but turns the income gradient positive on its own ($b_Y = +0.0040$ against the baseline -0.0070): expensive goods make proximity valuable on the consumption side, and the gains-from-trade channel overwhelms the externality. The measured flatness of spatial goods prices rules the second world out, and the measured cost of moving people makes the first counterfactual; the size of the curse rests on both measurements.

M.9 The Bundle Labor Share and the Own-Consumption Crosswalk

The Economic Census cost structures put the bundle labor share in a band from 0.42 to 0.54, and the value-added crosswalk that turns measured expenditure shares into the model's demand block depends on it. Re-estimating the own-share at the band endpoints (the ℓ rows of Table M2) moves the sourcing estimate as the crosswalked shares shift, from 0.60 to 0.76 in the per-village frame of that table, whose baseline estimate is 0.673 rather than the joint-equilibrium 0.532 of the main text. Every reported moment, targeted and untargeted, is unchanged to the stated precision: the income gradient stays at -0.0070 throughout. The moment that identifies the own-share pins the service value-added share α_n, so the band feeds the estimate and not the predictions. The crosswalk's composition assumption (pricing the ingredients content at the village-gate food price) is the curse-favorable case; repricing the measured non-food portion of the intermediates basket at delivered city-good prices attenuates the real gradient by under 10 percent, the relevant cost-of-living weight being about 0.06.

M.10 Outer Land and the Extrapolated Externality

The outer land floor of Appendix K.4 and the externality's extrapolation to non-program agriculture are null directions at the baseline in the inverted-world architecture. Both enter only through the recovered fundamentals, and the inversion absorbs any choice of either exactly: whatever floor or outer φ is imposed, the recovered $(\bar{A}, T_n)$ adjust so that the model reproduces observed outcomes, and no baseline moment, targeted or untargeted, can move. The choices matter only for counterfactual responses, where they scale how strongly non-program agriculture reacts to shocks. That margin is bracketed far more aggressively by the externality sweep of Section 8.2 and Appendix M.12, which re-solves every road experiment for the entire economy from $\varphi = 0$ to 0.77 and finds the aggregate flat and the incidence conclusions monotone throughout.

M.11 The Food-Market Closure

The joint model is closed: food is the numéraire, the national food market clears by Walras' law, and there is no open-economy bound left to report. The closure question that the earlier exogenous-city architecture had to bracket is settled by construction and audited by direct accounting (Appendix L.4): the value balance holds to 0.3 percent and iceberg losses are 0.45 percent of production. What remains of the closure as a modeling choice is the sell-side farm-gate convention, whose scope the same audit quantifies: it is exact for the locations holding 89 percent of the rural population, and for the net-importing remainder it errs by under 4 percent, within the arbitrage band.

M.12 The Road Across Externality Values

Table **??** in Section 8.2 solves the whole economy across the externality under homothetic demand, which isolates the relative-access margin, and the cross-sectional reversal near $\varphi = 0.55$ is discussed there. Two further patterns hold across the range. First, the relative-access margin's aggregate is flat in φ: under homothetic demand the uniform 10 percent road raises national income by about 0.8 percent at every φ, so the externality never taxes this component of the aggregate anywhere in the range. Under the extended demand system of Section 9.2 the aggregate does fall with the externality, from 0.86 percent at $\varphi = 0$ to 0.73 at the calibrated value, and the decline is entirely the income margin. Second, what the externality moves is incidence: the program villages' share of the uniform gains falls with φ, from +0.95 percent at $\varphi = 0$ to +0.64 at the calibrated value (+0.65 at $\varphi = 0.74$), the redistribution that Section 9 documents experiment by experiment.

The representable range is bounded above near $\varphi \approx 0.78$: the continuation solves every step up to $\varphi = 0.77$, but the transmigration villages' system fails to converge at $\varphi = 0.80$, where the labor exponent with land fixed reaches $\gamma_a + \varphi = 1.10$ and the increasing-returns region defeats the solver. This is a computational boundary at the estimated elasticities, not a claim about non-existence in general, and every counterfactual in the paper sits far below it.

Table M4: Counterfactual robustness: demand elasticities

	uniform road (national)	program villages
baseline ($\sigma_u = 5$, Cobb–Douglas)	+0.83	+0.64
Panel A. City-good elasticity σ_u		
$\sigma_u = 8$	+0.79	+0.61
$\sigma_u = 12$	+0.75	+0.58
Panel B. Demand margins ($\sigma_u = 5$)		
measured Engel curves	+0.75	+0.53
$\sigma_c = 2$ (rejected)	+0.37	−0.45

Note: Each row re-solves the 10 percent uniform road of Section 9.2 under homothetic demand. The baseline row reproduces the homothetic figure of Section 9.2 and the $\varphi = 0.70$ row of Table **??**. Columns report the change in national real income and in the program villages' real income, in percent. Panel A varies the city-good substitution elasticity σ_u, holding the city-good friction at its consumption-weighted value. Panel B sets the top-level substitution elasticity to $\sigma_c = 2$, the value the test of Appendix K.6 rejects. Results under non-homothetic demand are reported in Appendix Table N1.

M.13 The Urban-Differentiation Sweep

The city-good elasticity σ_u is set, not estimated, so the results are reported across $\sigma_u \in \{5, 8, 12\}$ with the friction held fixed at its consumption-weighted value ($\kappa_u = 0.203$, Appendix K.2), a construction that does not depend on σ_u. The sweep runs under homothetic demand, where the city-good channel it varies operates. Panel A of Table M4 reports the 10 percent uniform road across the sweep: the national gain falls modestly from +0.83 to +0.75 percent of national income and the program villages' share from +0.64 to +0.58, the entire movement carried by the city-good channel, which alone runs +0.59, +0.53, and +0.48 percent as σ_u rises. No sign changes anywhere in the sweep.

M.14 The Rejected Price Elasticity

Price substitution is the demand margin that could move the road results from the other side, and it is disciplined by measurement (Panel B of Table M4). At the rejected $\sigma_c = 2$ the 10 percent uniform road's national gain roughly halves (to +0.37 percent) and the program villages turn from a gain to a loss (−0.45 percent), because price-elastic demand triggers the same reallocation the externality penalizes; the test of Appendix K.6 rejects the price-elastic case in direction, so at

$\sigma_c \approx 1$ this margin is shut. The exercise is evaluated in the homothetic frame, the case its null hypothesis describes. The income margin, the demand margin the data do support, is part of the extended model of Section 9.2; shutting it there raises the road's national gain from 0.73 to 0.83 percent and the program villages' share from 0.55 to 0.64, and at $\varphi = 0$ shutting it moves the aggregate by almost nothing, which confirms that the tax comes from the externality rather than from the demand system itself.

M.15 The Service-Price Slope

The spatial slope of village service prices is +0.006, estimated from a specification that strips out the delivered-goods and composition content that inflates the broad non-food index to +0.01, and it enters the general-equilibrium calibration (Section 7.4) as a held-out check, not a targeted moment. The comparison between the model's held-out slope and this estimate is informative about structure: the model's village-gate service price is a pure labor cost, so any component of the measured slope that reflects premises costs is a channel the model deliberately omits, and the model's slope should undershoot the data by roughly that component. The housing evidence of Section 5 suggests that such a component exists: housing expenditure rises with market access, with the point estimates attributing most of the rise to prices, though residential prices are an imperfect proxy for commercial premises costs and the price component is imprecisely estimated.

N Additional Counterfactuals

This appendix holds the exhibits that support Section 9 without belonging in it: the distribution of the single-village road effects, the channel decomposition of the uniform road, a connectivity experiment that lowers the cost of distance itself rather than travel times, and the costing of the road network.

N.1 The Single-Village Road: Distribution of Effects

Figure N1 plots the distribution of the effects of the single-village road of Section 9.1 on income and on the agricultural employment share across the 1,147 experiments, with and without the externality.

Figure N1: Effects of a Road Built to a Single Program Village

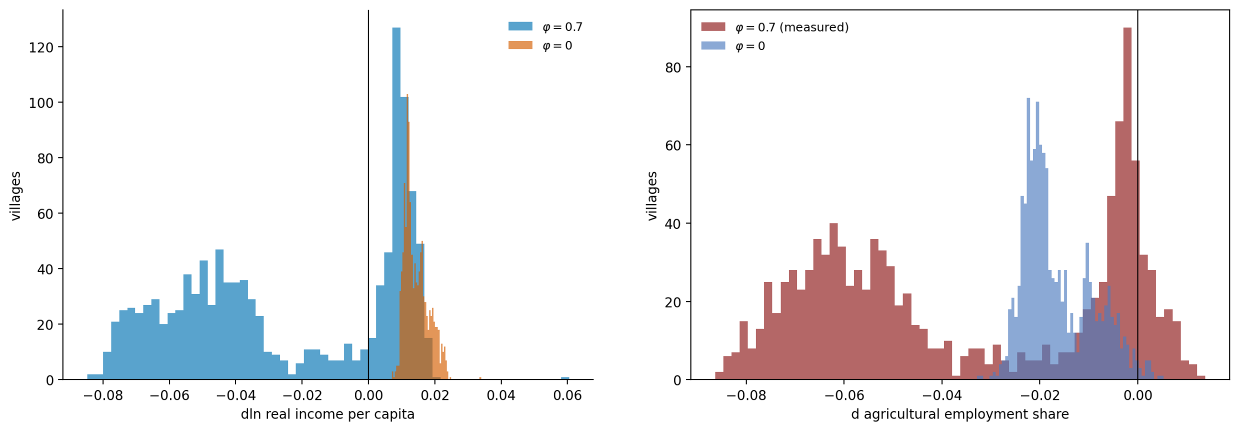

Note: The figure plots the distribution of effects across the 1,147 separate experiments, each of which cuts one village's travel times by 10 percent and re-solves that village against the unchanged rest of the economy. The left panel plots the change in log real income per capita and the right panel the change in the agricultural employment share, with the externality ($\varphi = 0.7$) and without it ($\varphi = 0$).

N.2 The Uniform Road by Channel and by Group

Table N1 decomposes the uniform 10 percent travel-time reduction of Section 9.2 by trade friction and by group, under both demand systems and with and without the externality. Without the externality the policy is close to distributionally neutral, with the program villages, other rural villages, and the cities all gaining about 0.9 percent, close to the national figure. With the externality the other rural villages keep their full gain at 0.9 percent, while the program villages keep only 0.6 percent and the cities 0.6 percent. A policy that improves every link in the country by the same proportion still has unequal incidence once the externality is in place, and the locations that give up the most are the isolated settlements the program created. That redistribution holds at every value of φ the model can solve (Appendix M.12) and across different values of the elasticity of substitution across city varieties (Appendix M.13).

Table N1: A uniform 10 percent travel-time reduction, by channel and by group

	$\varphi = 0.7$ (measured)				$\varphi = 0$			
	TM	Outer rural	Urban	National	TM	Outer rural	Urban	National
Non-homothetic demand								
farm-gate channel (κ_a)	+0.1%	+0.1%	0.0%	0.00%	+0.1%	0.0%	0.0%	+0.01%
service channel (κ_n)	0.0%	+0.5%	+0.1%	+0.22%	+0.4%	+0.5%	+0.2%	+0.25%
urban-variety channel (κ_u)	+0.4%	+0.4%	+0.5%	+0.52%	+0.4%	+0.4%	+0.7%	+0.59%
all channels	**+0.6%**	**+0.9%**	**+0.6%**	**+0.73%**	**+0.9%**	**+0.9%**	**+0.9%**	**+0.86%**
Homothetic demand								
all channels	+0.6%	+1.0%	+0.7%	+0.83%	+1.0%	+1.0%	+0.8%	+0.85%

Note: Population-weighted changes in log real income, in percent, from a uniform 10 percent cut in all travel times. Other rural is every non-program rural village, on the inner and outer islands. Urban real income is deflated by the city block's own consumption index. The upper block solves the model under non-homothetic demand and the lower row under homothetic demand. Each channel row re-solves the equilibrium with one friction's kernel system shocked at a time. National is the income-weighted total across every location, gross of construction costs, which are applied in Appendix N.4. The $\varphi = 0$ columns re-solve the identical experiments in the re-inverted externality-free world.

N.3 Cutting the Friction Parameters

The experiments of Section 9 all cut travel times; a natural closing question asks what a different kind of progress does. Roads scale travel times inside a fixed cost-of-distance technology, whereas cutting each friction parameter κ by 10 percent rotates the technology itself: it lowers delivered costs in proportion to $\ln \tau$, so the reduction grows with distance and the experiment is isolation-targeted by construction, what containerization, logistics, and telecommunication do rather than what asphalt does. Table N2 reports the results.

Table N2: Cutting the friction parameters ($\kappa \times 0.9$)

	$\varphi = 0.7$ (measured)				$\varphi = 0$	
Friction cut	TM	Outer rural	Urban	National	TM	National
Non-homothetic demand						
κ_a (farm-gate)	+0.2%	+0.1%	0.0%	0.00%	+0.1%	+0.02%
κ_n (services)	−0.4%	+1.0%	+0.2%	+0.51%	+0.7%	+0.61%
κ_u (goods)	+1.3%	+1.1%	+1.3%	+1.49%	+1.3%	+1.66%
all three	+1.1%	+2.2%	+1.4%	+2.01%	+2.1%	+2.30%

Note: Population-weighted changes in log real income, in percent, from cutting each friction parameter by 10 percent. Each row cuts one parameter, and the final row cuts all three. All rows are solved under non-homothetic demand. The $\varphi = 0$ columns re-solve the identical experiments in the re-inverted externality-free world.

Both friction cuts deliver large gains. The goods-friction cut is worth +1.5 percent of national income, because the rotation concentrates its cost reductions on long-haul inter-island trade, exactly where the value lies; this is the margin that logistics, shipping, and port technology operate on. The service-friction cut is worth +0.5 percent, delivered by any technology that lets services reach buyers with less person-travel, from communications to cheaper passenger transport. These experiments carry no resource cost and rotate the cost-of-distance schedule rather than cutting travel times, so their gross gains are not directly comparable with the costed road program of Appendix N.4.

Cutting all three frictions together raises national income by 2.0 percent against 2.3 percent in the externality-free economy, the same overstatement the other aggregate experiments display.

The service-friction cut also has the most distinctive incidence of any experiment in the paper: it is the one that harms the program villages on average while the other groups gain on average. Other rural villages and the cities gain 0.2 to 1.0 percent while the transmigration villages lose 0.4 percent, their isolation protection stripped as remote service delivery spreads exposure to city demand everywhere. The externality is what makes them lose, since without it the same cut turns their loss into a gain of 0.7 percent. It is also the one experiment that visibly compresses the cross-sectional curse, with the trimmed gradient flattening by 6 percent; the group averages and the gradient change are consistent with the curse shrinking by pulling the protected down rather than lifting the exposed up, though they do not by themselves establish that no exposed village is lifted.

N.4 The Cost of the Road Network

The gains reported in Section 9.2 are annual flows, while roads are stocks that depreciate. Pricing the 10 percent experiment against Indonesian cost data therefore requires three inputs: a per-kilometer cost for the works that raise speeds, a mapping from those works to the 10 percent travel-time reduction, and an annualization. The natural anchor for the first two is the Gertler et al. (2024) study of Indonesian road maintenance. For costs, they price periodic maintenance at the ROCKS overlay benchmark of roughly \$93,000 per kilometer, and they measure pavement deterioration of 6.3 percent per year, which implies an overlay is needed about every 8 years. For the mapping, their estimates of how speed falls with roughness imply that resurfacing a rough road raises achievable speeds on the order of 10 to 20 percent; I read the 10 percent experiment as the conservative end of what clearing the roughness backlog delivers, applied across the network. The annualization then prices resurfacing the whole network on its replacement cycle. Indonesia's road network is roughly 540,000 kilometers across national, provincial, and district classes, and resurfacing it at \$93,000 per kilometer, spread over the 8-year cycle, comes to about \$6.3 billion per year, or 0.68 percent of Indonesia's roughly \$920 billion GDP. Against the model's +0.73 percent gross flow at the measured externality this is a benefit-cost ratio just above one. That gross flow is itself generous to the road program, because the uniform experiment improves every link in the country, including the city-to-city connections that a maintenance budget of this size would not purchase, so the benefits it delivers are an upper bound on what the costed program buys. The same arithmetic at the realistic works mix documented by Collier et al. (2015), whose unit costs for developing-country road works run well above the overlay benchmark, gives a ratio below one, and at the betterment costs actually incurred on eastern Indonesian national roads, roughly \$850,000 per kilometer under EINRIP for the widening and realignment that raises design speeds once pavements are already smooth, the ratio falls well below one half.

The published claim that road maintenance in Indonesia pays two to three times its cost sits well above any of these ratios. Part of that is the cost side, since that claim prices the cheapest category of work. The rest is that the two exercises measure different quantities. A local estimate asks what happens to a district when its own roads improve and its neighbors' do not, which includes any activity it draws away from them; the experiment here improves every road at once, so activity that simply moves between districts nets out of the national total. Establishing how much of the gap that accounts for would require running the single-district experiment across all districts and comparing the sum of those local effects with the aggregate reported here. I do not do that, and I therefore do not attribute the gap to either source.