The Curse of Connectivity: Evidence from Indonesia's Village Resettlement Program

Harrison Mitchell

UC San Diego*

August 24, 2026

Abstract

I study how market access affects long-run development by exploiting the Indonesian Transmigration program, which created new villages with variation in initial connectivity. I find that, even today, villages with initially higher levels of market access have lower income per person. These findings stand in contrast to a large literature documenting positive effects of market access. The lower measured incomes are not explained by cheaper local prices, by amenities that compensate residents for lower earnings, or by people and production relocating elsewhere. Rather than limiting structural transformation, higher connectivity accelerates the reallocation of labor out of agriculture. I show that when agricultural productivity rises with the size of the local agricultural workforce, workers do not internalize the productivity loss their exit imposes on those who remain. This scale externality can make reallocation out of agriculture lower long-run incomes, and I provide direct evidence for it using separate variation in the initial settlement size of program villages. Embedding the estimated externality in a spatial general equilibrium model of the Indonesian economy, I show that the model reproduces the negative effect of market access on transmigration village incomes without targeting it, and that removing the externality reverses the effect. In counterfactuals, building roads still raises national income, but a road connecting an isolated village can make it poorer. The policies that raise rural incomes most are the ones that support agricultural employment directly.

*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 determine why some places prosper while others remain poor? A long tradition in economic geography argues that it does (Krugman, 1991; Redding and Sturm, 2008). This view carries real stakes for policy. The World Bank alone maintains an active transport portfolio of roughly \$45 billion, and approved over \$11 billion in new transport lending in the last fiscal year.[1] Whether greater market access raises a location's long-run development, however, is theoretically ambiguous. On one hand, greater market access generates the standard gains from trade and feeds the agglomeration forces that drive growth (Donaldson and Hornbeck, 2016; Desmet et al., 2018). On the other hand, market access reshapes what a location produces, and when it pulls labor away from the activities that carry productivity growth, integration can hinder development (Matsuyama, 1992). When this concern has surfaced in the literature on transport infrastructure, the worry has been that connected locations deindustrialize and specialize further in agriculture (Faber, 2014; Baum-Snow et al., 2020).

In this paper, I bring new evidence to this question from Indonesia's Transmigration program, which between 1952 and 1995 created over a thousand new villages across the outer islands, placed with little regard for the productive potential of their sites. The program lets me observe something rarely visible, the fates of comparable villages that began their lives at very different levels of market access, decades after the fact. I find that villages that began with greater market access are poorer today than otherwise similar villages that began more isolated. The reason is not that connected villages remained trapped in agriculture; more connected villages actually experienced *higher* rates of structural transformation than less connected villages. However, I find that agricultural productivity in these economies increases with the size of the local farming workforce, so each worker who left agriculture lowered the productivity of those who remained, a loss no individual internalizes. The rest of the paper provides evidence for this chain of events. I rule out standard explanations for the income gap, estimate the agricultural scale externality directly using the program's separate variation in settlement sizes, and embed it in a general equilibrium model of the Indonesian economy that reproduces the negative effect of access on income without targeting it. In counterfactuals, roads remain good national investments, but the villages they reach can lose, and support for agricultural employment delivers what connectivity cannot.

The Transmigration program resettled over two million voluntary migrants from the inner islands of Java and Bali into newly constructed villages across the outer islands of Indonesia. Officials, under pressure to meet resettlement targets and provided with limited information about the quality of different potential locations, sited villages wherever land was available, with little strategic regard for connectivity or the potential of a location.[2] Aside from location, the villages started out

[1] World Bank, "Transport Overview," accessed August 2026. The active transport portfolio is nearly \$45 billion, with \$11.2 billion in new transport operations approved in FY2025.

[2] Land-suitability plans were abandoned for nearly the entire duration of the program (Hardjono, 1988; Whitten, 1987), and Appendix A collects contemporary accounts of the siting process.

nearly identical; settlers had no control over which new village they were placed in, received uniform two-hectare plots by lottery upon arrival, and were each given the same initial endowment of agricultural inputs (tools, seeds, and fertilizer). Some villages began near roads, ports, or existing settlements while others began in more isolated locations, independent of settlers' attributes. To measure each village's initial connectivity I hand-digitize the complete pre-program road network of Indonesia, down to individual trails, from 27 detailed US military maps drawn around 1950. I then show that village characteristics are balanced across my measure of market access. The program therefore provides two independent sources of variation that I use throughout the paper: where villages were placed generates variation in initial market access, while how many households were initially settled in each village provides separate variation in the size of the agricultural sector to study the agricultural scale externality.

More than 20 years after their creation, villages that started with *better* market access have significantly *lower* GDP per capita than otherwise similar villages, and the differences are large. I estimate that moving from the 25th to the 75th percentile of initial market access lowers long-run village GDP per capita by roughly 6%. These results are robust to using alternative satellite-derived GDP estimates, to the wages that village residents report in household surveys, and to a novel video-based measure of village material welfare, which I describe below. The finding is surprising because it stands in contrast to a large body of work documenting positive effects of market access on local economic activity (Donaldson and Hornbeck, 2016; Hornbeck and Rotemberg, 2024; Redding and Sturm, 2008; Morten and Oliveira, 2024). In short, greater initial connectivity appears to have been a curse for these communities' long-run development.

Why would better connectivity hinder development in the long run? I investigate a variety of theoretical channels and empirically rule out many conventional explanations. The negative effects are not driven by a lack of structural transformation or industrial activity; on the contrary, more connected villages actually experience *more* reallocation of labor out of agriculture than isolated villages. Nor are the effects explained by the competitive destruction of local industry; more connected villages have more firms, not fewer. The results are also not explained by out-migration or selective migration. Local prices do not fall to offset the lower incomes; if anything, they are higher in more connected villages. Moreover, connectivity does not improve local amenities in this setting, so lower output is unlikely to be an equilibrium compensation for higher amenity values. In short, standard general-equilibrium adjustments (migration, prices, or amenities) neither offset nor explain the connectivity disadvantage. A natural question is why households do not simply move away as opportunities decline. Migration frictions in Indonesia are large and well documented (Bryan and Morten, 2019), and so long as they hold rural populations in place, the forces I document persist and migration can only partially offset them.

Instead, I uncover evidence for a mechanism rooted in the agricultural sector, a scale externality in agricultural production. As documented above, more connected villages reallocate labor out of agriculture. Crucially, agricultural productivity in these village economies depends on the size of the agricultural labor force, so this reallocation erodes the productivity of the farms that remain. The externality I document 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. When the sector shrinks, the channels that sustain this productivity weaken. This mechanism can be tied to well-documented forces, 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). Consistent with this mechanism, I find that more connected villages mechanize less, using fewer tractors per agricultural worker, despite their greater proximity to markets. Because agriculture is the dominant sector in these village economies, the decline in agricultural productivity translates into lower overall factor incomes, lower GDP per capita, and lower welfare. The structural transformation that connectivity promotes is, in this setting, a curse rather than a blessing.

I use a second source of program variation to measure the agricultural scale externality directly. Villages were settled with different sizes of the initial population, 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, comparing agricultural output across villages settled at different scales allows me to estimate how output responds to the size of the agricultural workforce. I find evidence that agricultural output increases with sector size more than the private returns to labor would predict. These results imply that the productivity of the agricultural sector increases with the size of the workforce. This social return is what reallocation out of agriculture erodes, and it 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 I estimate is comfortably above this threshold.

I then take the full model to the data to see how well it can reproduce the negative effect of connectivity on income. Workers in each rural and urban location can work in agriculture or non-agriculture, with agriculture subject to the scale externality of Section 5. The externality comes from the settlement-size estimates, and the trade costs and spending shares are measured in survey data. Transmigration villages are treated as identical except for their location on the map. The relationship between market access and income among them is therefore a prediction of the model rather than a target, and the model reproduces nearly the entirety of the negative effect I find in the data. As a further check, I carry out the model analog to the balance tests that show no relationship between village observables and market access in the data. I let the model choose each transmigration village's fundamental productivity freely so that it fits its observed income exactly, and I show that those recovered fundamentals are unrelated to market access. In the model, as in the data, connected villages do not need worse fundamentals to explain their lower incomes. Removing the externality reverses both answers. Without it, market access makes connected villages richer, as the standard gains from trade would predict, and the model can then fit the observed data only by assuming connected villages were built on systematically worse land, which is counterfactual to the observed proxies of productivity that I observe.

I then use the model to demonstrate that the externality changes who benefits from connectivity. The first counterfactuals speak to the cross-sectional variation I estimate. A road targeted to an isolated village can make that village poorer. Better access increases exposure to non-agricultural demand, which draws workers out of agriculture. The decrease in the size of the agricultural sector erodes agricultural productivity through the scale externality. My model supplies a clear mechanism that can explain other evaluations of rural road programs that have found workers leaving agriculture with little or no gains in income (Asher and Novosad, 2020). Consistent with the closed form solutions I derive in the simplified model, the loss in income is largest for villages that already have some pre-existing non-agricultural employment. The externality therefore changes the incidence of any shock that moves labor into or out of agriculture, since the locations that have workers leaving agriculture lose productivity that a model without the externality never captures.

To then speak to the effects of larger shocks that have aggregate economic impacts, I extend the model to incorporate the standard non-homothetic demand system of Boppart (2014) and Fan et al. (2023), estimating its parameters in the same expenditure survey that measures the budget shares. Income growth now shifts spending away from food, which pulls labor out of agriculture, and the externality prices that reallocation. The consequence is that an analysis without the externality overstates the aggregate gains from policies and shocks that move labor out of agriculture. A ten percent improvement in the entire road network looks eighteen percent more valuable in the externality-free economy than it does with the externality, and a five percent rise in rural agricultural productivity looks sixty-four percent more valuable.[3] The bias runs the other way for policies that keep labor in agriculture. An agricultural earnings subsidy at the same fiscal outlay as a national roads project 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 incidence of transport infrastructure, 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

[3] The road raises national income by 0.86 percent without the externality and 0.73 percent with it, and the productivity gain by 1.21 percent against 0.74 percent.

contribution is to document a mechanism which may reverse this relationship. In my context, greater market access *lowers* long-run output per capita. I argue that the reversal reflects the setting. 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.

Recent work on transport infrastructure in developing countries provides a useful contrast. Road programs raised trade and welfare in Brazil and in Indonesia itself (Morten and Oliveira, 2024; Gertler et al., 2024), while India's rural roads moved workers out of agriculture with little gain in income or consumption (Asher and Novosad, 2020), and Ethiopian rural roads raised agricultural productivity only when paired with agricultural extension (Gebresilasse, 2023). My model reproduces the roads-without-gains pattern as the equilibrium response to a targeted road; workers leave agriculture while incomes fail to rise, and supplies the mechanism behind it. Kaboski et al. (2024) show how highway programs release labor from farming into non-agricultural firms. I document the same reallocation but show that it carries a cost absent from their framework. The labor that leaves agriculture hurts the productivity of those that remain, and the firms that absorb these workers in my setting are the small informal services discussed above rather than a growing manufacturing sector. Closest on the finding itself is Faber (2014), who finds that connecting peripheral Chinese counties reduced their growth through competition from the better-connected core. That channel finds little support in my setting: connected villages have more firms and higher rather than lower prices. Additionally, losing production to the core can still be consistent with gains from cheaper imports, whereas losing agricultural scale is a productivity loss that can reduce welfare.

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 the movement of labor out of 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 has qualified 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 villages in my setting that move the most labor out of agriculture end up poorer, because with a scale externality each worker who leaves lowers the productivity of those who remain, and this loss can outweigh the gains from moving workers toward higher-productivity sectors.

The same estimates also complicate the debate on the idea of 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. A recent empirical literature finds evidence of 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 for surplus labor. First, my elasticity is a long-run object. The potential forces behind the scale externality I estimate, such as knowledge spillovers, investment in shared capital, and the cooperative management of irrigation, 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. My findings are not inconsistent with a low private marginal product of labor in agriculture, but they imply that the private product alone understates what a village loses when its agricultural workforce contracts.

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). I show that these otherwise local agricultural externalities can have aggregate policy consequences. They can reverse the local gains from market access, alter who benefits from infrastructure, and change the aggregate effects of structural transformation.

This paper also makes several contributions in data and measurement. The hand-digitized road network introduced above is one such contribution, and because the underlying US Army Map Service series covers nearly the entire world at the same level of detail, the approach can be used to construct historical measures of market access in many other settings. My primary outcome, local GDP per capita, is satellite-derived, which connects my work to a recent literature that uses remote sensing to measure economic activity where survey data are scarce (Baragwanath et al., 2024). Satellite measures on their own can show where activity changes but not why, so I pair the state-of-the-art satellite GDP estimates of Rossi-Hansberg and Zhang (2025) with detailed household, firm, and village level data that allow me to validate the measure and to directly test the mechanisms behind the aggregate effects I document.[4] Finally, I develop what is, to the best of my knowledge, a new approach to measuring local economic conditions, in which a vision language model scores publicly available video footage of my villages against a fixed rubric to quantify visible material well-being, and I use it to corroborate the main results from an entirely independent data source.

The rest of the paper proceeds as follows. Section 2 discusses the contextual detail around the

[4] I use the Rossi-Hansberg and Zhang (2025) estimates rather than nightlight proxies because the latter rest on an estimated elasticity of light to income from Henderson et al. (2012), which Gibson et al. (2021) and Bluhm and McCord (2022) have called into question. I show my results are robust to using nightlight proxies.

program and introduces the data I draw on. Section 3 then outlines my empirical strategy and identification argument. Section 4 presents my empirical results. Section 5 estimates the external scale elasticity in the agricultural sector. Section 6 develops the model, first in a simplified version that shows in closed form when the externality reverses the gains from connectivity, and then as a full spatial general equilibrium model of the Indonesian economy. Section 7 describes how I take the model to the data. Section 8 shows that the model replicates the key empirical patterns. Section 9 uses the model to evaluate counterfactual policies, first the cross-sectional experiments the main model is built for and then aggregate programs under an extended demand system. Section 10 concludes.

2 Background and Data

2.1 Transmigration from 1950-1990

In order to find exogenous variation in the initial market access of a location, I leverage a historical government policy in Indonesia. In order to address overpopulation concerns on the inner islands of Java and Bali and to promote nation-building, the government of Indonesia relocated over 2 million individuals from Java and Bali to *newly created* villages between 1952 and 1995. Previous work investigating the role of skill transferability on productivity and inter-group contact on national integration has convincingly shown that there was no selection between migrant characteristics and destination villages (Bazzi et al., 2016, 2019). Households randomly received a two hectare plot of land via lottery at the time of settlement, and, based on surveys of transmigrant households immediately before migration from Kebschull (1986), 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 (Figure 1). As Figure 2 shows, there is significant variation in transmigration village locations across the outer islands, and within a given island.

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 plan-as-you-go nature of the assignment process meant that villages were often built wherever there was land available in destination areas and land suitability maps were abandoned for nearly the entirety of the duration of the transmigration program (Hardjono, 1988; Whitten, 1987). Appendix A provides a collection of quotes from a wide variety of sources that cover the program between 1960-1990 and provide qualitative support to the idea that transmigration sites were not selected systematically or with consideration of the productive potential of settlement locations. Based on qualitative evidence and quantitative inspection, it does not seem that there was a clear hierarchy of potential village locations that were to be converted into transmigration settlements. This is reflected in the fact that time-invariant transmigration village characteristics show little variation across creation years. The chaotic nature of this program is what opens the door to

Figure 1: Survey Measures of Migrant Characteristics Prior to Relocation

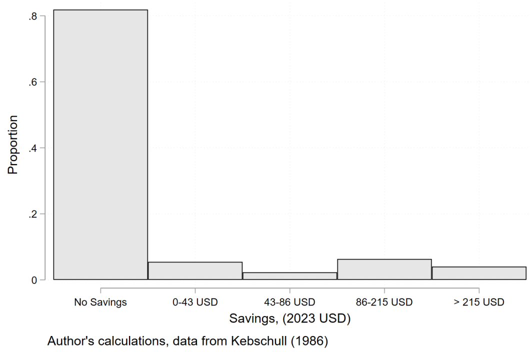

(a) Migrant Savings Immediately Before Migration

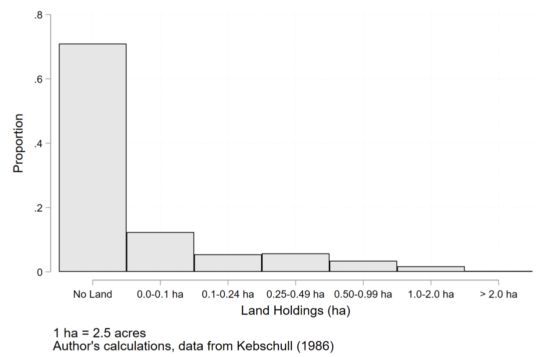

(b) Migrant Landholdings Before Migration

Note: Both panels present results from a pre-transmigration survey from a sample of transmigrant households surveyed by Kebschull (1986), but outcomes were 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 2: Distribution of TM Villages Across Outer Islands

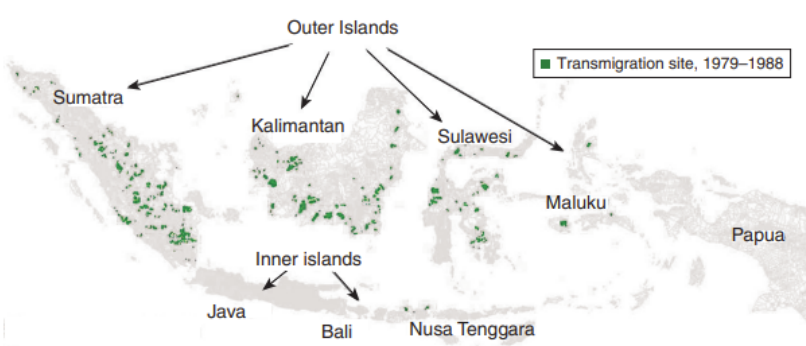

Notes: Each colored location on the map corresponds to a Transmigration village settled between 1979 and 1988. The white areas outlined in gray are other villages.

identifying exogenous variation in the initial market access of these villages, since sites were not chosen for what they were worth.

While previous work examining the effects of this program has focused on transmigration villages created between 1979 and 1988, I analyze outcomes for all transmigration villages between 1952 and 1995. The primary reason for this approach is it allows me to identify the effects of initial market access across a larger time frame; however, results are robust when restricting my sample to the villages studied in Bazzi et al. (2016) and Bazzi et al. (2019). As Table 1 shows, the underlying geographic characteristics of villages are relatively constant over time, though a handful of variables, including some soil and crop-potential measures and the size of settlements, do display statistically significant linear trends across creation years. These trends are one motivation for the cohort fixed effects I use throughout. Previous work has also observed that transmigration program activity, primarily financed by government oil revenues, was influenced by external, global fluctuations in oil price. Figure 3a plots the relationship between world crude oil prices and the number of transmigrants settled each year while Figure 3b instead plots the average population of settled villages each year. I find that the pace of transmigration activity tracks closely with global oil prices (correlation coefficient $= 0.856$) while there is little relationship with the average settlement size (correlation coefficient $= -0.016$). Despite not finding a relationship between settlement size and oil prices, it does appear that the intensity of transmigration activity, and likely outreach, is closely related to oil price changes. Taken together with the slight differences in transmigration villages across time, I primarily limit my analysis in later sections to comparisons among villages of the same "cohort" or year created.

Table 1: 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 district (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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3: Transmigration and Oil Prices

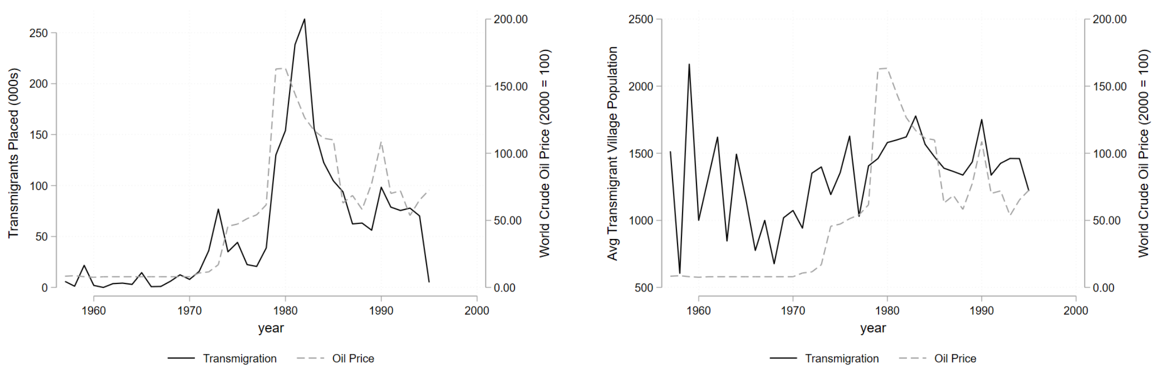

(a) Total Transmigration by Year (b) Average TM Village Size by Year

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.

2.2 Data Used

My analysis draws on a wide variety of data sources that include administrative data, individual, firm, and village level surveys, satellite-derived outcomes, and hand-digitized road network information. Table A1 provides an overview of each source, and Appendix B describes the construction of each in detail. Below I first describe the construction of my key measure of market access, and then summarize the remaining data grouped by their role in the analysis.

2.2.1 Pre-Transmigration Market Access Measure

In order to construct measures of a transmigration village's market access, detailed information on the road network present in the outer islands prior to the transmigration program's start in the 1950s is required. I construct a novel dataset containing extensive maps of Indonesia's road network prior to the resettlement period by hand-digitizing maps of the area compiled by the Army Map Service Corps of Engineers. These maps contain information on the existing roads, classified by type, railways, and trails throughout all of Indonesia, and are split into 27 individual images that have been stored in the Perry-Castañeda Library (PCL) Map Collection. Since the road network data was not drawn in a way that allows for automatic detection of different roads, I digitized the maps by hand. Roughly half of the digitized maps were cross-checked to ensure that no line segments were missed during this process. While in a few cases a given map segment contains some limited information from after 1950 that was used when compiling the image, any concerns over developments in the road network around the 1950s should be assuaged by the fact that there was little-to-no road expansion in Indonesia until near the 1990s (Gertler et al., 2024). Given the potential importance of waterways in transporting goods and people, I also include all navigable waterways in the outer islands from the Humanitarian OpenStreetMap Team. While this layer is a modern source, treating it as available in 1950 is far less restrictive than it would be for roads. Figure 4 shows the resulting cost layer that characterizes the available transportation infrastructure of Sumatra around 1950. See Appendix C for more details about the digitization process.

To estimate a location's market access similar to Donaldson and Hornbeck (2016), I need to construct measures of the trade cost from a given transmigration origin, o, to all other villages in Indonesia, d. Based on the detailed cost layer, I then assign a given speed or cost value for traveling over a given road type. For robustness, I estimate these trade costs under two sets of parameter values for travel; my preferred set of inputs comes from assigning travel times based on reported travel speeds across a variety of methods of transport that I calculate from the 1993 Indonesia Family Life Survey. Additionally, I slightly modify the cost parameters from Donaldson and Hornbeck (2016), which follow Fogel (1964), to allow for travel via automobile. See Appendix C for more details on the map digitization, process of estimating bilateral trade costs, and selection of cost parameters. The last required input needed to construct a measure of market access is estimates of each location's population, which I aggregate to the village level from the 1975 layer of the Global Human Settlement Layer. One caveat deserves note: because villages were settled between 1952 and 1995, the 1975 population weights are only strictly pre-settlement for the later cohorts, and for early cohorts they could partly reflect growth induced by earlier program activity. My results are robust to restricting the sample to the 1979–1988 villages studied in previous work,

Figure 4: 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. See Appendix C for details around the construction of the market access measure.

for which the weights are unambiguously predetermined. Table 2 reports validation of the GHS estimates against Indonesian sources. The 2010 GHS layer tracks the 2010 census population closely and precisely, while the comparison of the 1975 layer with initial settlement counts, though positive, is estimated imprecisely given the small overlapping sample. I then estimate the first-order approximation of market access for location, o as:

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

where N_d is the population of a village and θ is the trade elasticity. Estimates of the trade cost between two locations, τ_{od} come from implementing a fast marching algorithm similar to that described in Allen and Arkolakis (2014). My results are robust to variations in the value of θ (including $\theta = 1$, which is essentially just a distance weighted sum of all other locations' population) as well as the cost structure across the network[5] and alternative definitions of connectivity[6]. Reassuringly, my measure of market access at the time of the transmigration program is strongly correlated with a measure of market access constructed using the same method as above and only Indonesia's highway network combined with a road roughness index from Gertler et al. (2024).

Table 2: 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).

2.2.2 Outcome measures

My primary outcome is local GDP per capita, for which I use the high-resolution satellite-derived estimates of Rossi-Hansberg and Zhang (2025) and a more transparent alternative that combines

[5]The correlation coefficient between initial market access using IFLS speeds and Donaldson and Hornbeck (2016)'s cost inputs is 0.89.

[6]These measures include only calculating market access using destinations within 2 hours of the origin (with and without including θ), and constructing market access in a similar way to Equation 1 but using an alternative fast marching package in QGIS.

harmonized nightlights from Li et al. (2020) with population estimates, converted to GDP using the elasticities in Beyer et al. (2022), Hu and Yao (2022), and Henderson et al. (2012). I supplement these with several proxies for village-level productivity: crop yields per hectare and tractor density in agriculture, the ratio of employees to employers outside agriculture, and individual wages for formally employed workers from the National Socioeconomic Survey (SUSENAS). As an independent check, I construct a village-level material-welfare index from publicly available video footage scored by a vision language model (Appendix E). Figure 5 plots the distribution of several of these outcomes; Table A1 summarizes each source and Appendix B details its construction.

2.2.3 Mechanism measures

To examine why connectivity affects output, I draw on data spanning the main candidate channels. I measure sectoral employment and structural transformation from the 2000 and 2010 Population Census and the 1993 PODES; firm entry and size from the 2016 Economic Census and PODES; local prices and expenditure from the consumption modules of the 2002 and 2011 SUSENAS; village amenities from a public-goods index built from PODES; migration and commuting from the census and the 2003 PODES; and land use and its changes from the 2000 and 2003 PODES. Appendix B describes each of these sources in detail.

2.2.4 Controls and time-invariant characteristics

Finally, I assemble a rich set of pre-determined controls. Time-invariant geographic and agro-ecological characteristics of each village location follow Bazzi et al. (2016), and district-level controls roughly contemporaneous with the program come from the 1980 Population Census. I identify transmigration villages and their location, settlement year, and number of individuals initially settled from the Transmigration Census; the number initially settled also serves as the instrument for agricultural labor in Section 5.

3 Empirical Strategy

In this section I outline my empirical strategy for how I will estimate the causal effect of a transmigration village's initial market access on outcomes such as GDP per capita. The assumption I need is that among villages placed in the same province and settled in the same year, and conditional on expected market access and initial local characteristics, a village's initial market access is unrelated to the determinants of its long-run income.

This assumption does not require that officials chose sites randomly, only that whatever determined their choice is captured by the fixed effects and controls I include. Interestingly, random placement itself is not sufficient to identify the effect on its own. As I discuss below, if preexisting towns sit where fundamental productivity is higher, then even villages placed at random would possess a correlation between market access and productivity. What the program does provide is

Figure 5: Distribution of Village Level Outcomes

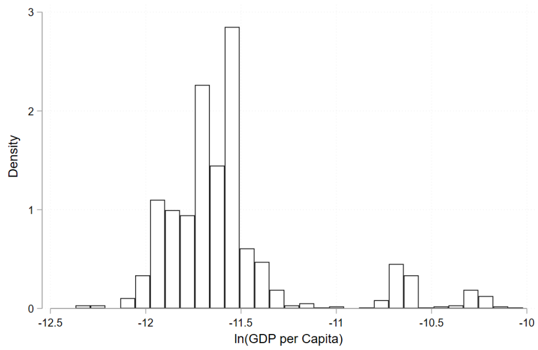

(a) GDP per Capita

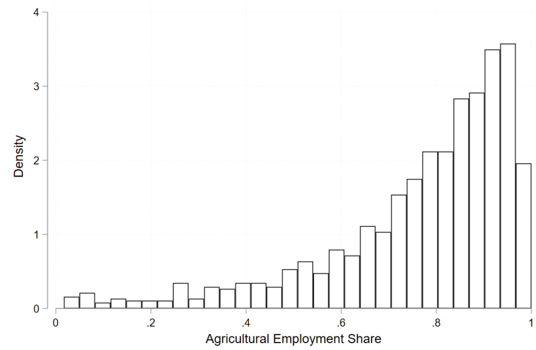

(b) Agricultural Employment Share

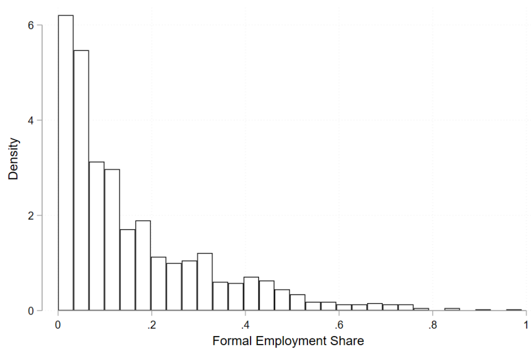

(c) Formal Employment Share

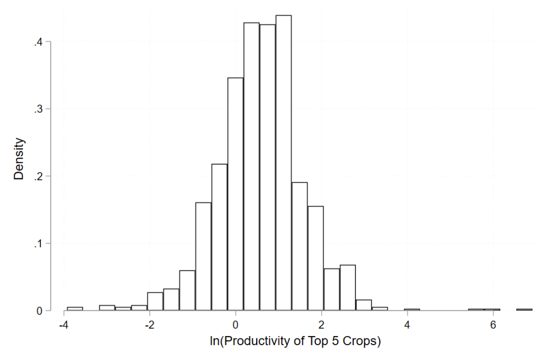

(d) Agricultural Productivity

Note: Each panel plots the distribution of a different village-level variable among the transmigration villages in my final 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.

a large number of villages that were built with similar characteristics at their founding and then placed with little regard for the productive potential of their sites.

I first present the evidence describing the site selection process, drawing on the program design described in Section 2 and the qualitative accounts collected in Appendix A, and then test it directly against a wide range of village characteristics. I then discuss the remaining assumptions required to estimate a causal effect of initial market access on village outcomes and the remaining threats to identification. Finally I outline the robustness checks I implement to address concerns around any remaining bias in my empirical strategy.

3.1 How Villages Were Sited

The comparability I need rests on two features of the program. The first is that the villages were built to be alike. Settlers had no control over which new village they were placed in, received uniform two-hectare plots by lottery upon arrival, and were each given the same initial endowment of tools, seeds, and fertilizer. This program design allows me to compare villages that genuinely started with similar individuals and starting endowments, which is rarely possible for settlements that formed on their own. Aside from where they were located, the villages started out nearly identical.

The second is that where they sat was chosen with little attention to what made a site valuable. From reading qualitative reports, it becomes clear that the widespread delegation of transmigration management duties often allowed provincial governments the final decision on the placement of transmigration villages that were created (Bazzi et al., 2016). Additionally, 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). Furthermore, despite land-use maps being often ignored (Hardjono, 1988), during the end of the transmigration period I study it was also revealed that the maps were unreliable and inaccurate, with updated maps not being made available until the last year or two of the program (Whitten, 1987). Officials were therefore placing villages wherever land was available, without the information they would have needed to place them strategically even had they wanted to. This qualitative evidence cannot establish the assumption on its own, so I test it against the data below.

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. I introduce a simple recentering approach motivated by the work of Borusyak and Hull (2023) to strip out that component. I construct a measure of expected market access, and include it as a control. What remains is the deviation in that village's realized market access relative to what could be expected, which is the variation I exploit.

This narrows what I have to assume. I do not have to claim that a village could have ended up anywhere in its province. I need only that among the nearby sites it could plausibly have been assigned to, which one it received is uninformative about the productive potential of the site. My preferred method of constructing this measure is to control for the average market access of the closest 20 "potential sites" of a given Transmigration village. Potential sites include actually settled Transmigration villages in the same province and also areas that were designated as "Planned Development Areas," which were maps developed by the government agency in charge of the transmigration program to identify potential transmigration village locations, but were never settled. The never-settled sites matter here, since they describe where a village could have gone rather than only where villages did go. I can construct alternative measures with different weighting schemes of nearby sites, including an inverse-distance-weighted average of all sites within a province and a simple average of all sites within a province, and my results are robust to all of these measures.

I then empirically test to see if I am able to observe differences in a wide range of observable variables that are likely associated with a location's fundamental productivity and a village's initial market access according to the following specification:

$$X_{vdic} = \beta_1 ln(MA)_{vdic} + \beta_2 ln(E(MA))_{vdic} + \eta_i + \gamma_c + \epsilon_{vdic}. \tag{2}$$

For each observable characteristic X in village v located in province d on island i part of transmigration cohort c, I test for differences associated with a transmigration village's initial market access while controlling for expected market access, constructed from nearby potential sites as described above, and island and cohort fixed effects. The island fixed effects are added to account for the large differences across heterogeneous islands, while the cohort fixed effects are motivated by the large variation in transmigration activity across the history of the program. Additionally I perform an omnibus test of difference and record the F-statistic associated with that regression. Table 3 reports the results from the estimations of Equation 2. I find economically small, precisely estimated null effects for the vast majority of included variables, and the omnibus test does not reject. Only one variable has a p value less than 0.01, out of the thirty-five variables tested. These tests cannot prove that villages are comparable, since they can only speak to the characteristics I observe, but they are the check the assumption has to pass and it passes it. One marginally significant imbalance worth noting is the number of individuals initially placed, which is slightly negatively associated with initial market access. Settlement size is an observed program characteristic rather than a location fundamental, and Section 5, which uses it as an instrument, examines its balance directly in a slightly different sample. I include individuals placed as a control in the regressions.

My final specification to estimate the causal effects of a location's initial market access on an outcome is given by:

$$y_{vdict} = \beta_1 ln(MA)_{vdic} + \beta_2 ln(E(MA))_{vdic} + \gamma_c + \eta_i + X_{vdic} + X_d + \alpha_t + \epsilon_{vdict} \quad (3)$$

where the only differences between Equation 3 and Equation 2 are the inclusion of local village controls and initial, 1980 district controls, and also the inclusion of the time dimension t for outcome variables that I observe for multiple years along with their associated time fixed effect α_t.

Table 3: Village Characteristics Across Initial Market Access

Dependent variable (std): ln(Market Access) (std)	
Wetland Rice Potential Yield	0.015 (0.032)
Palm Oil Potential Yield	0.011 (0.047)
Coffee Potential Yield	-0.051 (0.085)
Cassava Potential Yield	-0.019 (0.032)
Vector Ruggedness Measure	0.011 (0.034)
Individuals Placed	-0.091** (0.041)
Agrosimilarity Score	0.076 (0.046)
Share of Land with Poor Drainage	0.026 (0.057)
Share of Land with Good Drainage	0.005 (0.055)
Organic Carbon (%)	0.020 (0.049)
Topsoil Salinity (Elco) (dS/m)	0.016 (0.018)
Topsoil Base Saturation (%)	0.011 (0.041)
Avg. rainfall, 1948–1978	0.028 (0.038)
Avg. temp (Celsius), 1948–1978	0.067* (0.038)
ln(District Population, 1980)	-0.051 (0.038)
District Own Electricity, 1980	0.062* (0.036)
District Own Radio, 1980	0.050 (0.041)
District Own TV, 1980	0.067* (0.037)

Dependent variable (std): ln(Market Access) (std)	
District Literacy, 1980	-0.023 (0.039)
District Avg Years Schooling, 1980	-0.008 (0.037)
District Agricultural Employment Share, 1980	-0.057 (0.041)
District Wetland Rice Potential Yield, 1980	-0.006 (0.033)
District Palmoil Potential Yield, 1980	0.030 (0.036)
District Coffee Potential Yield, 1980	-0.048 (0.041)
District Cassava Potential Yield, 1980	-0.011 (0.035)
District Mining Employment Share, 1980	0.019 (0.040)
District Manufacturing Employment Share, 1980	-0.019 (0.038)
District Trade Employment Share, 1980	0.057 (0.043)
District Service Employment Share, 1980	0.067* (0.037)
District Wage Worker Share, 1980	0.015 (0.037)
District Share HHs with Piped Water, 1980	0.105*** (0.039)
District Share HHs with Sewage, 1980	0.069* (0.037)
District Share HHs with Modern Fuel, 1980	0.082** (0.038)
District Share HHs with Modern Floor, 1980	-0.038 (0.034)
District Share HHs with Modern Roof, 1980	-0.020 (0.037)

F-stat of omnibus test	1.398
Observations	1213

Note: Each row reports the coefficient on standardized log initial market access from a regression, following Equation 2, of the standardized characteristic listed in the row on initial market access, controlling for expected market access, island fixed effects, and cohort fixed effects. Standard errors, clustered at the district (kecamatan) level, are reported in parentheses. The omnibus F-statistic is from a regression of initial market access on all of the listed characteristics jointly, conditional on expected market access and the fixed effects, and tests their joint significance. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.2 Identification

The balance tests speak to the characteristics I observe. This section addresses the concern they cannot reach, which is the one that would survive even if placement had been random.

The identification challenge I face is similar to those confronted by others that have studied the effects of market access on a variety of outcomes. Put simply, in the absence of exogenous variation in *changes* in market access, there is still the potential for omitted variable bias to influence my results. Consider the following: if there exists some (partially) unobserved fundamental productivity that varies across space, *and* preexisting cities and towns are more likely to be located where that value is higher, then randomly placed new villages that are closer to more locations (i.e., high market access) are also more likely to be placed in a location with a higher fundamental value. Crucially, if the preexisting population distribution exists as described, then even perfectly random placement of new locations is subject to this potential omitted variable bias. Reassuringly, I carry out multiple checks to determine that this is likely not an issue in my setting.

Three things address this potential concern. The results in Table 3 suggest the bias is likely not present, though this relies on a "no selection on unobservables" argument that one is never able to explicitly prove. The construction of expected market access and the inclusion of local and regional controls both reduce the scope for it directly, and in Appendix D I outline the identification challenge associated with this scenario and document how much of the bias each control removes.

The third is the most useful, because it does not require the assumption to hold. The bias I have described would mean that high initial market access locations are more likely to be productive locations, which would push my estimates toward finding *positive* effects of market access on GDP per capita. I find negative effects. Omitted variables of this kind therefore cannot explain what I find, and if they are present at all the true effects of market access in this setting are even more negative than what I estimate.

3.3 Robustness Checks

Despite finding no statistical evidence that my results could be subject to omitted variable bias, I perform a variety of robustness checks when estimating the effect of initial market access on my primary outcome of interest, GDP per capita. Most simply, I vary my regression specification by varying the inclusion of my controls and assess how the estimated effect changes. While not a formal test of robustness, the fact my results remain virtually unchanged across specifications is reassuring. More formally, I assess the likelihood that omitted variables determine my 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)[7], 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. This value measures the importance of each covariate k relative to all of the other observed covariates. The higher this fraction, the smaller the potential for omitted variables to explain the results. To further highlight that the inclusion of expected market access and also initial district level controls address concerns associated with omitted variable bias, I report these same values for specifications that do not include various controls.

Finally, I assess the robustness of my inference and of the assignment assumptions directly. Because market access varies smoothly over space, spatially correlated errors could understate my standard errors, and because the placement argument operates within provinces, a reader may prefer specifications that absorb the assignment strata themselves. Appendix Table C1 re-estimates the effect of market access on GDP per capita under progressively coarser clustering, under fixed effects that absorb provinces and their interaction with settlement periods (Repelita), and under a randomization inference procedure that randomly reassigns market access across villages within

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

province-by-cohort cells. That last exercise is the one place random assignment is imposed, and it is imposed as a null rather than relied on. I do not need it to identify the effect, but asking how often reshuffled access reproduces my estimate is a direct check on whether the result could be an artifact of how the villages happen to be arranged in space. The program also administers two cross-cutting treatments, where a village was placed and how many households were settled there. Recent work on factorial designs shows that when a program assigns more than one treatment, inference on either one is more reliable in the specification that includes both treatments and their interaction (Muralidharan et al., 2025). Appendix Table C2 reports this fully interacted specification, and the estimated effect of market access is essentially unchanged.

4 Reduced Form Results

4.1 The Effects of Market Access on Local GDP per Capita

I find strong evidence that higher initial market access leads to worse aggregate output per capita. Table 4 shows that across a variety of specifications I find that increased initial market access causes villages to have lower GDP per capita between 2012 and 2019. Not only are the effects of initial market access negative, but they are also economically significant; moving from the 25th to 75th percentile of market access translates to a decrease in GDP per capita of slightly over 6%. This explains nearly 20% of the interquartile spread in GDP per capita that I observe in my villages. The table also reports the robustness measures described above for three separate specifications. In the final column, my preferred specification with the full control set, I find strong evidence that my results are robust and not likely explained by omitted variable bias: the fraction of times that the breakdown point from Diegert et al. (2025) is larger than the ρ_k values for my local controls is over 85%, higher than the same fraction Diegert et al. (2025) report in their empirical application where they conclude the results are robust. The comparison across specifications provides more direct evidence that my preferred empirical strategy mitigates concerns around bias: the results are less robust (though still quite robust) when I do not include district controls or the expected market access control.

Appendix Table C1 shows that these results survive approaches to inference that are more demanding than the baseline on both margins. When I replace the island fixed effects with province fixed effects, the estimate is unchanged whether standard errors are clustered at the village, regency, district, or province-by-cohort level. Retaining the island fixed effects of the baseline while clustering at the regency level yields the least precise estimate, significant at the ten percent level, and the point estimate remains similar in magnitude across island, province, province-by-Repelita, and province-by-cohort fixed effects, ranging from -0.0026 to -0.0044. The province-by-cohort specification is demanding, as the median cell contains only two villages and over a third of cells contain one, yet the estimate remains significant at the one percent level. A randomization inference exercise makes the same point without any assumption on the error structure, as 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. Because the program administers two treatments, Appendix Table C2 additionally reports the specification that enters market access and initial settlement size jointly with their interaction (Section 3.3). The negative effect of market access is concentrated among smaller settlements, a pattern I return to with the model in Section 8.

To ensure that these results are not an artifact of the particular dataset in use, I also estimate the effect of initial market access on an estimate of GDP per capita that separately combines the effect of market access on population and nightlights and uses estimated conversions found in Henderson et al. (2012), Hu and Yao (2022), and Beyer et al. (2022) to back out an implied measure of GDP per capita and report the results in Table 5. While less precisely estimated, I find even more negative impacts of connectivity on long run aggregate output using these measures. Finally, using publicly available videos that document the conditions of a limited subset of Transmigration villages and their inhabitants, I construct a measure of village material well-being using a large vision-language model that analyzes each video according to a defined rubric and I find similarly large negative effects of initial market access. Appendix E provides more detail on this process and presents the results.

Table 4: Effects of Market Access on GDP per Capita

	ln(GDP per Cap)	ln(GDP per Cap)	ln(GDP per Cap)	ln(GDP per Cap)
ln(Market Access)	-0.0060***	-0.0059***	-0.0057***	-0.0058***
	(0.0022)	(0.0019)	(0.0019)	(0.0020)
Expected MA Control	No	No	No	Yes
Local Controls	Yes	Yes	Yes	Yes
District Controls	No	No	Yes	Yes
Cohort FE	No	Yes	Yes	Yes
Island FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Unique Villages	1152	1152	1152	1152
DMP Fraction	0.714	0.643	0.643	0.857

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 GDP per capita according to Equation 3. Each column represents an alternative specification, as described in the table rows. Additionally, the last three columns include the robustness results described in Section 3.3.

Table 5: Effects of Market Access on GDP per Capita, Alternative Measures

	ln(GDP per Cap), HSW (2012)	ln(GDP per Cap), Hu and Yao (2022)	ln(GDP per Cap), BHY (2022)
ln(Market Access)	-0.0070	-0.0154**	-0.0200***
	(0.0063)	(0.0063)	(0.0062)
Expected MA 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	1142	1142	1142

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 GDP per capita according to Equation 3, where each column constructs an estimate of GDP per capita based on observed levels of night lights and population density according to a different methodology. Column 1 uses the estimated long run elasticity between night lights and GDP from Henderson et al. (2012). Columns 2 and 3 do the same but following Hu and Yao (2022) and Beyer et al. (2022), respectively.

4.2 The Effects of Market Access on Local Productivity

I now turn to corroborating the negative effects of market access on GDP per capita by estimating the effect of market access on various local measures of productivity. Table 6 reports the effects of initial market access on proxies for both agricultural and non-agricultural productivity. I find that, outside of agriculture, the ratio of formal employees to employers is smaller in high market access locations, suggesting that the typical firm is smaller. In the agricultural sector, I find that higher market access leads to lower rates of mechanization (as measured by tractors per capita and tractors per agricultural worker). The direct yield measures point the same way, though less precisely. Villages with higher market access produce less yield per hectare for the top crops and for rice, with these estimates significant at the ten percent level, and the palm oil estimate is negative but imprecise. These patterns are unchanged when I control for the agricultural employment share. Taken together, the proxies consistently point toward lower productivity in more connected villages. Appendix Table C3 further reports additional results that directly measure individual wages in a subset of my villages. I find qualitative evidence that individual wages are lower in high market access villages, which supports the results on village level measures of productivity.

Table 6: Effects of 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)	-0.6994**	-0.0501**	-0.0555**	-0.0095*	-0.0097*	-0.0073
	(0.2858)	(0.0203)	(0.0247)	(0.0053)	(0.0050)	(0.0217)
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	1105	1030	1030	974	583	296

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 multiple measures of productivity according to Equation 3. Column 1 uses data from the 2000 and 2010 population census. Columns 2-3 combine data on tractor usage from the 2000 PODES and population and agricultural employment data from the 2000 population census. Columns 4-6 use data from the 2000 PODES.

4.3 Potential Mechanisms

I find robust, economically meaningful negative effects of initial market access on village GDP per capita and productivity. Given the surprising nature of these findings, I next turn to examining potential mechanisms motivated by theory and previous work that could explain these results.

4.3.1 The Role of Sectoral Shifts

One potential explanation for the negative impacts of increased connectivity on output could be that increased exposure to competition in more connected locations leads to the flight of productive industry and more specialization in agricultural activities. If more connected places were led to specialize in agriculture due to the increased market access for their output and the competition from connected, more industrial, locations, then I would find that increased market access leads to increased agricultural employment shares compared to less connected locations. These results would be consistent with theoretical arguments in Matsuyama (1992) or empirical results in Baum-Snow et al. (2020) or Faber (2014)[8]. However, Table 7 reveals that I find the opposite effect; increased market access leads to more structural transformation out of agriculture. Not only does increased market access lead to less growth, but it simultaneously spurs structural transformation. The fact that increased structural transformation coincides with lower output per capita and productivity is all the more concerning, given the general focus on promoting structural transformation as a means to promote growth in developing countries. As I will show in Section 6, this combination of more structural transformation and lower productivity is exactly the pattern predicted by a model in which agricultural TFP depends on the scale of the agricultural labor force.

[8]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.

Table 7: Effects of Market Access on Structural Transformation

	Agriculture Share	Change in Ag Share	Processing Share	Services Share	Trade Share
ln(Market Access)	-0.0038***	-0.0027***	0.0013***	0.0017***	0.0009***
	(0.0008)	(0.0007)	(0.0003)	(0.0004)	(0.0002)
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	1152	1147	1152	1152	1152

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 sectoral employment according to Equation 3. Columns 1-5 use data from the 2000 and 2010 population census. Column 2 takes the difference in agricultural employment share between 2010 and 2000.

4.3.2 Firm Entry and Size

The expected effects of increased connectivity on firm entry and size are not immediately clear. While the increased competition and declines in output would suggest that increased market access leads to fewer firms that are then larger due to selection forces, the fact that more connected villages are situated around more demand–simply as a function of population–could encourage firms to enter. I find empirical evidence that increased connectivity in my setting spurs firm entry (Table 8). Across a variety of sectors and firm sizes, I see that increased market access increases the number of firms that are located in that village. I also find statistically significant increases in the share of individuals that report operating either an informal or formal firm in the census in higher market access locations. I also find that, within the villages with medium and large firms, increased market access is associated with smaller average firm size. This is consistent with entry-driven selection, in which more entry means the marginal entrant is smaller, though I note that this estimate conditions on the endogenous presence of such firms and other forces could produce the same sign.

These results also point away from a competition-based channel of the type identified by Faber (2014). If increased market access primarily exposed transmigration villages to greater competition from more productive external firms, I would expect to see firm exit and a decline in the number of local firms. Instead, I find that more connected villages have more firms, not fewer. While firm counts cannot rule out compositional shifts within the firm population, the fact that I do not find net exit coupled with the price effects discussed in Section 4 does not support a competition-based channel that can explain these negative effects.

Table 8: Effects of Market Access on Firm Employment Dynamics

	Employers per 1000	Total Number of Firms	ln(Employment)	ln(Employment)
ln(Market Access)	1.7717***	0.1331*	-0.0289**	-0.0244**
	(0.5354)	(0.0785)	(0.0130)	(0.0102)
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	No	No	No
Unique Villages	1152	944	128	128

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 firm dynamics according to Equation 3. Column 1 uses data from the 2000 and 2010 economic census. Column 2 combines data from the 2003 PODES and the 2016 economic census. Columns 3-4 use firm level data from the 2016 economic census.

4.3.3 Migration and Commuting

Given the central role of migration in determining the spatial distribution of economic activity, it is important to assess whether population movements can account for the differences in outcomes between more and less connected locations. If initially connected villages offered their residents better opportunities elsewhere, part of the negative effect on GDP per capita could reflect people or production leaving the village rather than lower resident incomes. I measure migration two ways. Because transmigration villages were created for individuals from the inner islands, the number and share of residents in the 2000 Census born on the inner islands proxy for permanent out-migration. In addition, the 2003 PODES records whether residents have left the village to work elsewhere, a direct, if coarse, indicator of out-migration that complements the birthplace stocks. I find little to no relationship between initial market access and these measures of migration. The share of inner-island-born residents and the out-migration indicator show no relationship, and the level of inner-island-born residents shows only a small negative association, significant at the ten percent level (Table 9).

Commuting is the harder channel to measure, as there is no available dataset where I can observe the commuting behavior of individuals at a granular enough level to match to transmigration villages. The concern with commuting would be that some portion of the output generated by commuters in a location would be recorded outside the bounds of a village, and so this could potentially make high-commuting villages appear poorer than they really are. I argue that differences in commuting patterns are highly unlikely to explain the negative effects of market access on GDP per capita for two reasons. First, commuting in rural areas is uncommon. Using data from the 2005 Intercensal Population Survey, I find that the majority of regencies have rural commuting rates of the working population below 2%. The increase in commuting rates in high market access locations would need to be implausibly high to fully explain the negative effects I find. Second, and most directly, commuting can depress measured local production but not other outcomes that I also measure. I find that market access decreases the wages of village residents, regardless of where they work (Table C3), the material-welfare index constructed from village footage (Appendix E), and the nightlight-based measures of GDP per capita (Table 5). All three of these measures should increase with market access if the negative effect of market access on GDP per capita were entirely explained by commuting differences attributing production outside the village boundaries. Taken together, this evidence makes it unlikely that the negative effects of market access simply reflect output that is produced by residents, and therefore recorded, outside of the village.

Importantly, these results do not rule out that individuals in more connected villages travel for work more often. In fact, the model I develop in Section 6 suggests this could be the case, since rural services are delivered in person and urban demand for these services is concentrated in the villages closest to urban areas. To the extent that workers in more connected villages reallocate toward work that requires travel, this reallocation is itself the structural transformation documented above, and it lowers village incomes through the agricultural scale externality that I estimate in Section 5.

Table 9: Effects of Market Access on Migration

	ln(Inner Islanders)	Share of Inner Islanders	Village Out-Migration (0/1)
ln(Market Access)	-0.0113*	-0.0006	0.0006
	(0.0058)	(0.0011)	(0.0013)
Expected MA 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	No
Unique Villages	1136	1147	944

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 migration according to Equation 3. Columns 1-2 use data from the 2000 and 2010 population census. Column 3 uses data from the 2003 PODES.

4.3.4 Amenities and Compensating Differentials

An alternative explanation for the decrease in GDP per capita and productivity stemming from increased market access lies in the role of a location's amenity value. It could be the case that individuals in high market access locations are not worse off from a welfare perspective if the amenity value of high market access locations is higher than lower market access locations. In order to assess the plausibility of this explanation, I estimate the effects of initial market access on a public goods index that is designed to capture many important facets of the amenities of a location. Table 10 reports the estimated effect of market access on this index and a disaggregated version of it. I find no evidence that increased market access leads to increased amenities in my context. Furthermore, the only significant effects of market access on my various public goods are for outcomes that directly capture aspects of market access itself–road quality. The fact that I am able to detect significant improvements in public goods associated with market access but not any other measure of amenities highlights that my measure is able to capture important differences across villages, and despite this fact I do not find that market access improves the measured amenities of a location. The public goods index cannot cover every dimension households might value, but the estimates show no support for a compensating amenity channel on their own.

Table 10: Effects of Market Access on Amenities

	Public Goods Index - Full	Public Goods Index - Excl Roads	Public Goods Index - Roads Only
ln(Market Access)	0.0023	0.0005	0.0106***
	(0.0020)	(0.0022)	(0.0035)
Expected MA 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	1147	1147	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 public goods and amenities according to Equation 3. Columns 1-3 use data from the 2000 PODES.

4.3.5 Price Effects and Expenditure

I next show that the decreased level of GDP per capita is not offset by declines in the price level that households in high market access locations face. Using detailed consumption data, I estimate the effects of market access on prices and quantities for food consumption. I construct two separate indices that weight goods differently. The index indicated by Ind constructs a price level by summing the log unit price of all goods weighted by the household specific expenditure shares, while the index indicated by Ref instead uses product weights from the average expenditure share of rural individuals on the same island that are *not* living in transmigration villages. Table 11 shows the results for both price and quantity levels by both weighting schemes. I find consistent, statistically significant evidence that increased initial market access actually increases the price level that households face. The mixed effect on quantities is less conclusive, because the consumption module primarily captures market purchases rather than consumption of food from own-production. This means that the higher quantities consumed that I observe in more connected villages likely reflect substitution from own-produced toward purchased food as households leave agriculture, rather than higher total food consumption. Importantly, this does not affect the unit-price results, which are the primary focus for understanding how prices respond to market access. The fact that households face higher prices in high market access locations indicates that the lower measured production per capita is 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 increased market access in this context.

Table 11: Effects of Market Access on Price Levels

	Price Index^{Ind}	Price Index^{Ref}	Quantity Index^{Ind}	Quantity Index^{Ref}	ln(Food Expenditure)
ln(Market Access)	0.0135***	0.0106***	0.0060	0.0133***	0.0284***
	(0.0034)	(0.0031)	(0.0048)	(0.0044)	(0.0088)
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	Yes	Yes	Yes	Yes
Unique Villages	84	84	84	84	84
Unique Households	1234	1234	1234	1234	1234

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 price levels according to Equation 3. Columns 1-5 use data from the 2002 and 2011 SUSENAS household food consumption module.

I also examine the effects of market access on housing expenditure, which represents the largest non-food consumption category in my setting. I find that increased market access leads households to spend substantially more on housing, and the point estimates attribute most of this increase to the price of housing rather than the quantity consumed, though the two components are individually imprecise (Table C4). Land-use records 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 (Table C5 in Appendix F). Converting land between uses appears difficult in this context, so structural transformation operates through labor rather than land. These patterns motivate the model's treatment of land in Section 6: land works only in agriculture, and non-agricultural activity absorbs labor, not land.

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 average firm size, no sizable migration response, no improvements in amenities, and higher local prices. The standard channels (agricultural specialization, competition, labor drain, compensating differentials) find no support in the data as explanations for the negative effects of connectivity in this setting. The combination of accelerated structural transformation and declining agricultural productivity points to an agricultural scale externality. The remainder of the paper identifies this elasticity and quantifies its consequences: Section 5 estimates the structural elasticity of agricultural TFP with respect to the scale of the agricultural workforce, Section 6 embeds it in a spatial general equilibrium model of the full economy and derives the threshold at which it reverses the gains from connectivity, Sections 7 and 8 take the model to the data and show the curse emerges as an untargeted prediction, and

Section 9 uses the estimated model to evaluate policy.

5 The Agricultural Externality: Direct Evidence

The program design also allows me to directly identify the elasticity of agricultural TFP with respect to the scale of the agricultural labor force in this context, using exogenous variation in initial settlement size as an instrument for current agricultural employment. The model in Section 6 embeds this elasticity as the central novel ingredient and quantifies its general-equilibrium consequences. Across six specifications I find point estimates of the production-function coefficient that support the claim that there exist economically meaningful scale externalities in my setting. Following the best practices outlined by Andrews et al. (2019), I rely on weak-instrument-robust inference. In these data, the Anderson and Rubin (1949) confidence sets are bounded in every specification and exclude a zero production-function coefficient in all six. The implied lower bound on the externality elasticity φ at the standard agricultural labor share $\gamma_a = 0.304$ is positive in five of the six specifications. The IV evidence therefore supports the existence of the externality under conservative inference.

5.1 Empirical Framework

Let agricultural production in village v for crop c be

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

where L_v^{ag} is village agricultural labor, H_{vc} is crop-c area planted, and γ_a is the private agricultural labor share. Village-level productivity depends on the scale of agricultural labor through an externality:

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

Substituting into the production function and taking logs yields the estimating equation:

$$\ln Q_{vc} = \underbrace{\delta_c}_{\ln \bar{A}_c} + (\gamma_a + \varphi)\ln L_v^{ag} + \gamma_H \ln H_{vc}. \tag{6}$$

The coefficient $\beta \equiv \gamma_a + \varphi$ on $\ln L_v^{ag}$ is the object of interest. With an external estimate of γ_a, the externality elasticity is identified as $\varphi = \beta - \gamma_a$.

Because crop-specific labor is unobserved in the PODES data, the production function specifies the private labor input using total village agricultural labor. Identifying the parameter $\gamma_a + \varphi$ using total village labor requires the assumption that the share of the village workforce allocated to a specific crop does not systematically covary with the instrument. Three factors mitigate this potential concern. First, crop fixed effects absorb invariant differences in labor intensity across crops. Second, the specification directly controls for crop-specific planted area, which absorbs differences in crop scale that are likely to be associated with the allocation of labor across crops. I treat crop-specific planted area as conditionally exogenous and control for it directly. Under this assumption the coefficient on agricultural labor measures the response of production to the scale of the agricultural workforce holding cultivated area fixed. Section 5.4 examines the related concern that initial settlement size may shift cultivated area. I find no statistically significant relationship between initial settlement size and planted area, and the estimated labor coefficient remains similarly large when planted area is omitted from the 2SLS. Finally, to ensure the instrument itself did not inadvertently shift crop composition, I regress the number of individuals initially placed on the share of the village working in agricultural cash crops versus staple crops in the Population Census and find insignificant effects. While I can only split individuals based on crop *type* rather than a specific crop, these null effects suggest that initial settlement size is not shifting the more granular agricultural labor share toward different crops. Together, these features provide reassurance that the instrument affects crop output by isolating the scale of the agricultural workforce, rather than by shifting its composition.

5.2 Data

The estimation uses village-by-crop microdata from the 2003 wave of PODES. For each village and crop, I observe production quantity Q_{vc}, area planted H_{vc} (hectares), and yield $y_{vc} = Q_{vc}/H_{vc}$. Both yield and production are winsorized. The sample consists of transmigration villages from the main analysis sample. Agricultural labor L_v^{ag} is the number of individuals working in agriculture in village v, treated as contemporaneous with the PODES 2003 outcomes. An important detail here is that I observe production quantities directly, so I do not need to worry about how prices might vary across locations.

5.3 Instrument: Initial Settlement Size

The main challenge is that L_v^{ag} is endogenous: sorting, local fundamentals, and equilibrium responses can jointly determine agricultural employment and production. To address this I instrument for $\ln L_v^{ag}$ using $\ln Z_v$, the number of individuals initially settled by the government in transmigration village v. This variable is determined well before I observe crop production information in 2003.

The instrument is relevant because larger initial settlements generate larger agricultural workforces decades later. Figure 6 displays a binscatter of agricultural employment against initial settlement size, showing a clear positive relationship. Exogeneity is supported by the balance tests presented in Table 12. Conditional on island and cohort fixed effects, initial settlement size is orthogonal to pre-determined location characteristics. The qualitative evidence on the chaotic nature of the program further supports the claim that settlement sizes were not systematically assigned based on location quality.

Figure 6: First Stage: Ag Working Population vs. Initial Settlement Population

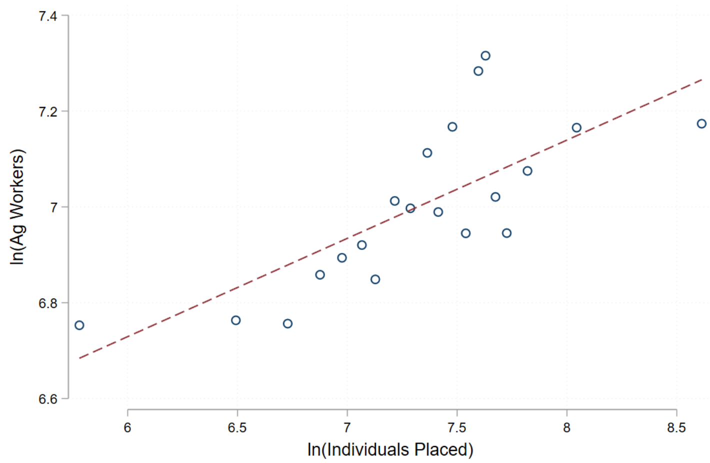

Note: Binscatter of log agricultural working population against log initial settlement population, controlling for island and cohort fixed effects.

5.4 Threats to Exclusion and Controls

The key identifying assumption is that initial settlement size affects agricultural outcomes only through the scale of the agricultural labor force it generates. Even if initial settlement size is balanced on initial location characteristics, the exclusion restriction could fail if settlement size shifts agricultural production through other channels.

I first provide direct evidence on the balance of initial settlement size with respect to observable village characteristics. Table 12 reports coefficients from regressions of each standardized characteristic on the log of individuals initially placed, controlling for island and cohort fixed effects. The results show that initial settlement size in this sample is uncorrelated with village-level characteristics. Importantly, initial settlement size is also uncorrelated with both market access and expected market access, confirming that the instrument captures variation in village scale that is orthogonal to the connectivity variation I exploit in the previous analysis. 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 I control for these district-level measures directly in my specifications.

A related concern is that initial settlement size may affect agricultural production by changing the amount of land cultivated, rather than only through the scale of the agricultural workforce. In the baseline specification, I treat crop-specific planted area as an exogenous production input and control for it directly. Consistent with the assumption that the instrument does not meaningfully shift this input, I do not find a statistically significant relationship between planted area and the number of individuals initially settled in a location.[9] If initial settlement size does not affect planted area, including or excluding planted area should not change the population IV estimand for agricultural labor. Consistent with this idea, the same 2SLS strategy that omits planted area as a control produces a significant, slightly smaller coefficient of 1.07, compared with 1.39 when planted area is included. The two estimates are not statistically different, and even the direction of the difference is reassuring. If initial settlement size increased cultivated area through its effect on agricultural labor, omitting land would tend to raise the estimated labor coefficient, since the land channel would be loaded onto the labor effect. Instead, the coefficient falls from 1.39 to 1.07 when planted area is omitted. Given that the two estimates are statistically indistinguishable, I do not interpret this difference too strongly, but its direction provides no indication that an omitted positive land response is inflating the baseline estimate.

As a sensitivity exercise, I also allow the positive point estimate on planted area to represent a genuine response operating through agricultural employment. Let π_H denote the effect of log initial settlement size on log planted area and π_L the corresponding first-stage effect on log agricultural labor. If planted area responds to the instrument only through agricultural labor, the implied elasticity of planted area with respect to agricultural labor is π_H/π_L. The IV specification that omits planted area then identifies

$$\beta^{\text{no } H} = \gamma_a + \varphi + \gamma_H \frac{\pi_H}{\pi_L}. \tag{7}$$

Using $\hat{\pi}_H = 0.118$, $\hat{\pi}_L = 0.200$, and the private constant-returns agricultural technology used in the quantitative model, $\gamma_a = 0.304$ and $\gamma_H = 1 - \gamma_a = 0.696$, the estimate $\hat{\beta}^{\text{no } H} = 1.07$ implies

$$\varphi = 1.07 - 0.304 - 0.696 \left(\frac{0.118}{0.200} \right) \simeq 0.36. \tag{8}$$

Thus, even if the entire positive point estimate of the relationship between initial settlement size and planted area is interpreted as an endogenous land response generated by additional agricultural labor, the implied scale externality remains economically large and positive.

[9] A regression of log planted area on log individuals initially placed, using the same controls and fixed effects, yields a positive but statistically insignificant coefficient of 0.118. In contrast, the first-stage coefficient of initial settlement size on log agricultural employment is 0.200 and statistically significant.

Table 12: 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 district (kecamatan) level, are reported in parentheses. The omnibus F-statistic tests joint significance across all dependent variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Beyond the balance evidence, I examine several specific threats to exclusion. Larger settlements may develop in different ways that could affect agricultural productivity beyond the size of the sector itself. Table 13 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 above address. While it is impossible to definitively rule out violations of the exclusion restriction, the patterns I can examine do not point toward one.

Table 13: 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 presented in Section 5.5. All columns except column 2 use data from the 2003 PODES wave. Column 2 uses my constructed measure of market access, recentered by the expected market access measure described in Section 3.

5.5 2SLS Specification

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

$$\ln Q_{vc} = \beta \widehat{\ln L_v^{ag}} + \gamma_H \ln H_{vc} + \Gamma X_v + \delta_c + \delta_{\text{island}} + \delta_{\text{cohort}} + \varepsilon_{vc}, \tag{9}$$

where the endogenous regressor $\ln L_v^{ag}$ is instrumented by $\ln Z_v$, δ_c are crop fixed effects, and X_v includes baseline village covariates and agro-ecological controls. Standard errors are clustered at the village level. The coefficient β estimates $\gamma_a + \varphi$. With an Indonesia-specific external estimate of $\gamma_a = 0.304$ (Fuglie, 2010), I can infer $\varphi = \beta - \gamma_a$.

5.6 Results

Table 14 provides consistent evidence for economically meaningful external scale elasticities in the agricultural sector. Across all six specifications, point estimates of $\beta = \gamma_a + \varphi$ are positive and economically large, ranging from 1.04 to 1.39. The effective F-statistic outlined by Montiel Olea and Pflueger (2013) falls below the threshold where conventional Wald confidence intervals perform well. I therefore turn to weak-IV-robust inference as suggested by Andrews et al. (2019).

Table 14: Estimation of the Agricultural Technology 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**
	(0.583)	(0.705)	(0.598)	(0.703)	(0.697)	(0.575)
Effective F (Olea-Pflueger)	7.32†	6.31†	7.87†	6.23†	6.04†	7.32†
AR 95% CI for β	[0.45, 4.81]	[0.52, 6.91]	[0.26, 4.43]	[0.50, 6.98]	[0.48, 7.17]	[0.42, 4.72]
Implied φ ($\gamma_a = 0.304$)	0.902	1.096	0.748	1.077	1.060	0.872
Implied φ lower bound (AR, $\gamma_a = 0.304$)	0.141	0.212	-0.045	0.191	0.177	0.117
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 X Isle + Cohort	Crop X Cohort + Isle	Crop + Isle + Cohort
Observations	1487	1487	2261	1487	1487	1487
Villages (clusters)	602	602	612	602	602	602

Notes. 2SLS estimates with cluster-robust standard errors at the village level. $\ln(L_v^{ag})$ is instrumented by ln(individuals initially placed). **Effective** F is the Montiel Olea & Pflueger (2013) cluster-robust first-stage statistic; † flags values below 16.38, the threshold above which conventional Wald inference using $|t| > 1.96$ is approximately valid. **AR 95% CI** is the Anderson and Rubin (1949) confidence set. Implied externality $\varphi = \widehat{\beta} - \gamma_a$ at $\gamma_a = 0.304$ (Fuglie 2010, livestock-adjusted labor share). The *lower bound* on φ is computed as (AR lower bound on β) $-\gamma_a$ and reflects the smallest value of the externality elasticity not rejected by the data at the 5% level under weak-IV-robust inference.

5.7 Weak-IV-Robust Inference

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, will serve as the primary outcome of interest in this section. AR confidence sets remain valid under weak identification. Two features of the sets are worth describing. First, the AR upper bounds are wider than the Wald upper bounds (4.43 to 7.17 across the six specifications, versus Wald upper bounds in the 2.2–2.7 range), reflecting the additional range of β_0 values the data cannot reject once weak identification is taken into account. Second, the AR lower bounds sit *above* the Wald lower bounds in every specification, generally between 0.42 to 0.51. The two intervals answer different questions, and in these data the weak-identification-robust one turns out to be more, not less, favorable to a positive coefficient at its lower end. The implication I rely on is that the data reject $\beta = 0$ in all six specifications and reject $\beta \leq \gamma_a = 0.304$ in five of six under weak-IV-robust inference.

The implied lower bound on the externality elasticity φ at the standard agricultural labor share $\gamma_a = 0.304$ is reported in Table 14 as well. The AR-implied lower bound on φ is positive in all but one case. The value the model adopts, $\varphi = 0.70$, lies below every point estimate of $\varphi = \beta - \gamma_a$ and inside every AR confidence set. Section 7.2 discusses this choice, and Appendix K.1 additionally validates the estimator itself. Using data generated by the calibrated model, where the true coefficient is known, the same 2SLS recovers the model's labor coefficient of $\gamma_a + \varphi$ to within two to four percent, and recovers γ_a when the externality is shut off. Section 6 embeds the elasticity in a spatial model, shows in closed form when it reverses the gains from connectivity, and quantifies its consequences for real GDP, structural transformation, and welfare across the connectivity gradient.

5.8 Interpreting the Magnitude

Before incorporating φ into the spatial model, it is worth discussing the magnitude of the estimated externality I obtain. The externality is large, and it 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. This finding complicates the notion of there being "surplus labor" in the agricultural sector in developing countries. Two points reconcile this finding with the modern evidence for surplus labor. First, my elasticity is a long-run object, while recent empirical evidence on surplus labor is from short-run experimental variation (Breza et al., 2021). Second, the existing evidence measures the private return to agricultural labor, while my estimates capture the social return.

In terms of the magnitude, the natural benchmark is Bartelme et al. (2025)'s estimates of sector-level economies of scale across manufacturing industries. They find scale elasticities averaging roughly 0.15, and the externality I estimate is roughly four to five times larger than the average here. A potential reason for the discrepancy is that my variation is coming from relatively small movements in the size of villages, while other work often looks at variation at the country-by-sector level. The average population of transmigrant villages is around 2000 individuals, and it seems plausible that scale externalities are non-linear and stronger at lower initial levels. Consistent with this idea, I find noisy evidence that the externality I estimate is larger when examining variation in settlement size among smaller villages. 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. Finally, the model does adopt a conservative value for φ, and I report model results across a wide range of values for φ (discussed more in the next sections).

6 Model

The empirical results establish three facts. Villages with higher initial market access have lower GDP per capita and 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 and quantify whether the estimated externality is large enough to reverse the standard gains from market access. The model is also the basis for the policy experiments of Section 9, since the presence of the externality changes who benefits from connectivity and how alternative policy instruments compare.

I build the model in three steps. Section 6.1 describes the two-sector structure that every location carries, which combines the agricultural scale externality estimated in Section 5 with a differentiated local non-agricultural sector. Section 6.2 then adds enough structure to analytically characterize the model. In a simplified economy of a single village trading with one city, the contest between the externality and the gains from proximity can be written intuitively and solved in closed form, and Proposition 1 gives the externality threshold at which the effect of connectivity changes sign. Sections 6.3 through 6.7 then write out the full quantitative spatial general equilibrium model that I calibrate and use for the counterfactual analysis. Section 7 provides details on the calibration process. Importantly, the income gradient itself is never targeted, and Section 8 shows that the model is able to reproduce it.

6.1 Setup

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 different goods and 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 described in Section 2. 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 sector-specific travel costs, 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. The difference in trade costs is what drives the labor reallocation. Section 6.5 describes how each friction operates in detail.

Labor is immobile across locations, matching the empirical setting of Section 4 where migration does not respond to market access. 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. This reflects the land-use evidence of Section 4, where cultivated land adjusts in response to market access but land devoted to housing and firms does not expand to accommodate non-agricultural activity.

Agriculture is the same technology at every location, rural or urban. 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,

$$A_{a,\ell} = \bar{A}_{a,\ell}\, Z_{a,\ell}^{\varphi}, \qquad \varphi > 0, \tag{10}$$

external to the individual farmer and therefore uncompensated in factor markets. Substituting (10) into the production function,

$$Y_{a,\ell} = \bar{A}_{a,\ell}\, Z_{a,\ell}^{\gamma_a+\varphi}\, H_{\ell}^{1-\gamma_a}, \tag{11}$$

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 (10) is the same object estimated in Section 5, whose estimating equation is the log of (11) 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 of Section 6.6.

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,

$$Y_{k,\ell} = A_{k,\ell}\, Z_{k,\ell}, \qquad k = n \text{ (rural)},\ k = u \text{ (urban)}, \tag{12}$$

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

These ingredients already outline how connectivity can lower income. Because services are more costly to trade over space, proximity to a city concentrates urban demand on a nearby village's services. That demand raises the private returns to non-agricultural work and pulls workers out of agriculture. Then, the externality converts the reallocation into lower agricultural productivity and lower income. However, this model still contains the standard gains from trade, since the same proximity lets the village sell its services and its food at better prices and buy city goods more cheaply. Which force dominates depends on the size of φ. The next subsection characterizes this contest exactly.

6.2 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 H contains the derivation, and Appendix G 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,

$$\mathcal{E} \equiv s_a\,\varphi\,\frac{u}{1-u}\,\eta, \tag{13}$$

which I call the externality drag. A one 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 subsection, derived in Appendix H, 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 H, equation (C15)). 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. If sector-specific trade costs were symmetric, there would be no sectoral reallocation in response to the connectivity shock. At the measured values in

Section 7 this condition holds easily. The sections that follow determine the quantitative strength of each mechanism.

6.3 The Quantitative Spatial General Equilibrium Model

The full model embeds the structure of Section 6.1 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; Section 8 examines structurally whether any productivity gradient hides behind the connectivity gradient. Every other location's fundamentals are recovered from its observed outcomes (Section 7.3).

6.4 Households

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

$$U_\ell = c_{a,\ell}^{s_{a,\ell}}\, Q_{n,\ell}^{s_{n,\ell}}\, Q_{u,\ell}^{s_{u,\ell}}, \qquad s_{a,\ell} + s_{n,\ell} + s_{u,\ell} = 1, \tag{14}$$

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,

$$Q_{n,\ell} = \Big[\sum_{o \in \text{rural}} \omega_o^{1/\sigma_n}\, q_{n,\ell o}^{\frac{\sigma_n - 1}{\sigma_n}} \Big]^{\frac{\sigma_n}{\sigma_n - 1}}, \qquad \sigma_n > 1, \tag{15}$$

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 $\omega = 1$. 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), but assigning every location one variety instead leaves the estimated model's moments essentially unchanged (Appendix K.4).

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

$$P_\ell = p_{a,\ell}^{s_{a,\ell}}\; P_{n,\ell}^{s_{n,\ell}}\; P_{u,\ell}^{s_{u,\ell}}, \tag{16}$$

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 (16) 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

$$\pi_{a,\ell} = \frac{T_a w_{a,\ell}^\theta}{T_a w_{a,\ell}^\theta + T_k w_{k,\ell}^\theta}, \qquad Z_{j,\ell} = L_\ell\, \Gamma\!\left(1 - \tfrac{1}{\theta}\right) T_j\, w_{j,\ell}^{\theta-1} \left(T_a w_{a,\ell}^\theta + T_k w_{k,\ell}^\theta\right)^{\frac{1}{\theta}-1}, \tag{17}$$

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.

6.5 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, $p_{a,\ell} = \tau_\ell^{-\kappa_a}$. 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[10]. Because cities are urban and therefore sit at $\tau_\ell = 1$,[11] 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.

According to the production function in the model, 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 in a similar way to agricultural goods, and person-embodied town services, which behave similarly to the rural service good. To account for this, 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. Section 7.2 provides details on the measurement of all three frictions, and Appendix K.8 quantifies how the results change when either

[10] 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. However, the accounting audit of Appendix J.4 shows the resulting error is small where the importers cluster.

[11] Recall, τ_ℓ is defined as a location's distance to an urban location.

friction is set to the other's value.

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. Ruling this out is consistent with the setting, as commuting is rare in rural Indonesia and the cost of regularly traveling to urban areas is high (Section 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

$$p_{n,v}Y_{n,v} = \omega_v\, p_{n,v}^{1-\sigma_n}\Big[\beta_n \sum_{\ell\in\text{urban}} \tau_{\ell v}^{-\kappa_n(\sigma_n-1)}\,\frac{I_\ell}{\Phi_\ell} \;+\; \alpha_n \sum_{\ell\in\text{rural}} \tau_{\ell v}^{-\kappa_n(\sigma_n-1)}\,\frac{I_\ell}{\Phi_\ell}\Big], \tag{18}$$

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 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.

6.6 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 (17), (ii) the agricultural wage satisfies its first-order condition (11), and (iii) the non-farm labor market clears, service revenue (18) 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.

Because the externality leaves effective returns to scale in agriculture slightly above constant at the measured parameter values (Section 7), uniqueness cannot be guaranteed in general. The reported equilibrium is solved by a continuation in φ approach. 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 then verify that solving the model without this continuation still converges to the identical equilibrium at every location and perturbations do not alter final equilibrium. Additionally, the calibration of Section 7 recovers non-program fundamentals by inverting the model on observed outcomes, so re-solving the full equilibrium forward from those recovered fundamentals should return the observed data. It does, with a mean absolute error of 0.6 percent in incomes across rural locations and 0.6 percent across urban ones.

6.7 The Mechanism in the Full Model

In the full model the causal chain runs in four links, each tied to a reduced-form pattern from Section 4. Demand for a village's service variety decays exponentially with bilateral distance at rate $\kappa_n(\sigma_n - 1)$. Villages near cities and other central locations therefore capture a disproportionate share of spending on rural services. That demand raises the service wage relative to the agricultural wage and, through (17), workers reallocate out of agriculture, producing the structural transformation effect. The reallocation lowers effective agricultural labor and, through the externality (10), agricultural TFP, explaining the lower agricultural productivity in more connected villages. Because food prices do not fall, and in fact increase, with connectivity, the productivity loss translates to lower factor incomes in agriculture. A large enough magnitude here explains the negative effect of initial market access on GDP per capita. The price index (16) is approximately flat because the rising farm-gate food price and the falling city-good index offset at the measured frictions and consumption shares. Against all of this stand the standard gains, since connected villages buy city goods cheaper and sell services dearer. Sections 7 and 8 determine which force wins in the full economy.

7 Bringing the Model to the Data

Most of the model's inputs are measured directly. Following the standard practice for models of this type, I invert the model to recover the fundamentals that reproduce the observed income per capita and agricultural employment share of every non-program location so that the world reproduces the data exactly by construction. The program villages are held out from this process. In line with the identification argument and balance table in Section 3, they carry *common* fundamentals, and their outcomes are solved from the model. The rest of this section is devoted to describing the measurement of relevant parameters, the inversion, and the model calibration; Table 15 collects the parameter map.

7.1 Measured Geography, Endowments, and Outcomes

The trade-cost matrix is the travel-time network of Section 2, evaluated between all locations. The variety weights $\omega_o = N_o/\bar{s}$ of Section 6.4 use populations from the settlement layer underlying the market-access measure. Transmigration villages enter with a common scale, every program village is assigned the mean recorded labor and land endowment[12]. All other land comes from the agricultural census. Because that source records household landholdings, it misses land devoted to

[12]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. However, 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. This method also isolates the effects of geography itself on my outcomes of interest, which is what the model seeks to measure.

estate, plantation, and fishery agriculture. This results in a small number of populated locations recording near-zero household land despite substantial populations and agricultural employment shares. I winsorize non-program land-per-worker from below at 0.02 hectares per person. While these locations are not an issue for the inversion process itself, the winsorization process is helpful for the counterfactual analysis. Near-zero land would make these locations extremely responsive to shocks and appear as spuriously cheap service supplies. Appendix I.4 details the measurement gap.

The outcomes the inversion consumes are observed at the location level. Income per capita comes from village-level GDP for every location in the economy, normalized to the measured rural income anchor. Agricultural employment shares come from the 2010 census wherever possible to match locations, and from the 2000 census elsewhere.[13] Urban locations are inverted on the same two outcomes as rural ones. Of note, their agricultural shares are far from zero (0.26 at the median administrative city, 0.36 for functionally urban locations), which is why the two-sector city block is a measurement requirement rather than simply a modeling choice.

The decision on how to classify locations as either urban or rural does not have a clear answer. The administrative classification produces an implausibly low urban share for Indonesia as it assigns many peri-urban and urban areas on the outskirts of cities as rural. The urban share according to the administrative classification is 44%, roughly 15 percentage points lower than aggregate statistics suggest. My baseline classification therefore assigns to the urban side every location whose centroid falls inside a satellite-defined functional urban area[14], in addition to the administratively urban set. This definition more accurately reflects the split between urban and rural Indonesia, with an urban share just under 65%. Appendix K.5 reports the estimated point under alternative classifications. The results are similar across different definitions and qualitatively unchanged even under the administrative classification, where the same patterns appear at smaller magnitudes because peri-urban locations, treated as rural suppliers, absorb part of the urban service demand before it reaches the countryside.

7.2 Externally Set and 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 5. 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.[15] Section 9 reports every experiment at $\varphi = 0$ as well, and Appendix K.12

[13] The 2000 values are adjusted multiplicatively by province-level ratios of worker-weighted mean shares in 2010 to 2000, measured on the twenty-eight thousand villages present in both census waves.

[14] Data from the GHS-FUA layer

[15] Additionally, this decision keeps the model inside the range the solver can represent easily, with the stress tests of Appendix K.12 solving the economy up to $\varphi = 0.77$, near the computational boundary discussed there. Beyond that point solving the model becomes more computationally difficult, and adds little to the intuition of the relevant mechanisms.

spans a wide range of externality values. Because the 2SLS instruments for agricultural *headcount* while the model uses effective labor, I demonstrate that mapping my estimated φ to the model is nearly exact in this setting. Appendix K.1 runs the paper's exact estimator on data generated by the calibrated model, where the coefficient is known, at the recorded settlement sizes and recovers an estimate of $\gamma_a + \varphi$ that is within four percent of the true value. Then, in the model with the externality shut off I recover the true estimate of γ_a. 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 this effect I find in Section 4.

Trade frictions. I measure separate trade frictions for services, food, and the city good. Appendix I.2 provides the full details on the data and estimating equations. Services in the model require a person to travel, so I measure their trade cost using the time and monetary cost of making that trip. The 1993 IFLS community survey provides detailed information on travel times and fares to a variety of destinations, local wages, and doctors' consultation fees for 312 communities. I use these data in two ways to try and estimate κ_n. First, I compare the time and cost of a round trip to a destination with the value of a day of work. For a typical community, a trip to the regency capital costs about half the value of a day of work, while a same-day trip to the province capital is too costly for most villages. Second, I compare the cost of traveling to receive an urban service with the posted price of a doctor's consultation. The two approaches give very similar estimates of the service trade-cost exponent, 0.94 and 0.89, respectively. I therefore set $\kappa_n = 0.90$. These same trip costs explain why commuting is rare in rural Indonesia, as regular travel to urban work is uneconomic for most villages (Section 4).

I measure the food friction using the same survey and a similar method. The 1993 IFLS also contains information on staple food prices in each community. Therefore, I compare round-trip travel costs with the value of a fifty-kilogram load of rice to arrive at an upper bound on κ_a of roughly 0.10–0.12. This is likely to overstate the true cost of shipping food because commercial shipments are larger and cheaper than passenger fares, which is what the data captures. I also estimate the relationship between staple prices and travel time and find a small, precisely estimated null relationship. Together, these estimates point to a relatively small food friction, and I set $\kappa_a = 0.05$. Appendix K.7 shows that the results remain disciplined when this value is varied between one-half and twice the baseline.

Because I lack the data to measure the city-good trade friction in the same way as κ_a and κ_n, I construct the city-good friction from the service and food frictions themselves. The city good combines physical goods, which I assume can be shipped at the food friction, with services such as formal health and education, which require people to travel, which I assume can be shipped at the village service friction. In the rural SUSENAS expenditure data, physical goods account for 82 percent of this bundle and person-embodied services account for the remaining 18 percent. Under the Cobb–Douglas aggregation used in the model, the corresponding trade-cost exponent is therefore $\kappa_u = 0.82\kappa_a + 0.18\kappa_n = 0.203$. The resulting gap between the service friction and the goods frictions captures the intuitive feature of the data, that moving people is substantially more costly than moving goods.

Demand. I measure rural expenditure shares using outer-island rural households in the 2002 SUSENAS. Mapping the survey categories into the three goods in the model requires one adjustment. Many purchases that appear as local services in the survey, such as prepared food and small-scale retail, combine local service labor with purchased food inputs. I therefore express expenditure shares in value-added terms. Using Economic Census cost data, I assign a fraction $\ell = 0.48$ of these purchases to local service labor and the remainder to food inputs. Appendix I.1 describes this crosswalk in detail.

One part of this construction cannot be measured from the data directly. The SUSENAS records details on what households purchase, but not whether the service was sourced within the village or from outside it. Because the value-added crosswalk depends on whether these purchases are local, the final rural expenditure shares also depend on this value. I denote the local sourcing share by s^{own} and estimate it internally in Section 7.3 using the observed relationship between market access and structural transformation. Intuitively, a higher s^{own} shifts expenditure from services toward food in the value-added accounting framework, because for locally sourced purchases only the labor component is counted as service demand, while the purchased ingredients are assigned to food. At the estimated value, $s^{\text{own}} = 0.532$, the shares used in the model are $(\alpha_a, \alpha_n, \alpha_u) =$ (0.667, 0.103, 0.230). Thus, the survey determines the overall pattern of household spending, while the model estimates the one sourcing margin that the survey does not observe.

The urban expenditure shares are measured differently. Rather than estimating a separate set of shares from urban households, I use aggregate market-clearing conditions to pin down the values. Given observed incomes and the rural service expenditure share α_n, total aggregate spending on rural services must equal total service revenue. I therefore select a value of β_n that allows for this identity to hold in the model. The same accounting for city goods determines β_u, and the food share β_a follows from simply adding up. This procedure also gives a useful check on the calibration. The model implies an urban food share of $\beta_a = 0.595$. Applying the same value-added crosswalk directly to urban SUSENAS households and adjusting for the measured Engel relationship gives 0.560, roughly the same value. Appendix K.6 shows that remeasuring the demand block using different survey years and samples produces very similar results.

Elasticities. At the top level I use Cobb–Douglas preferences across food, rural services, and the city good. Appendix I.6 tests this restriction directly in the SUSENAS data. Expenditure shares on city goods do not decrease as their delivered price rises with isolation from urban locations. The implied substitution elasticity is close to one, and a substantially more price-elastic specification is rejected according to this exercise. For the sorting elasticity θ I do not have a direct measurement of my own. There is, however, an estimate from this same context that I can impose. Bryan and Morten (2019) recover a choice elasticity of 12.5 (s.e. 1.36) governing how Indonesian workers reallocate across alternatives in the same national data environment. While their margin (migration across locations) is not exactly mine (choice of sector within a village), it is the closest measured elasticity the mechanism runs on, and imposing it removes a free parameter that would otherwise have to be fitted. I therefore impose $\theta = 12.5$. Two exercises show the results do not lean on the specific value. Appendix K.3 frees θ in a two-parameter variant and recovers a value just beyond the upper end of the Bryan and Morten (2019) confidence interval, with every prediction essentially unchanged, and the simple model of Appendix G sweeps θ widely and finds the curse's sign preserved throughout. To the best of my knowledge there is no accepted estimate of a rural service substitution elasticity σ_n, but some external estimates provide a rough range of values that this parameter could take. Applications of the Gervais and Jensen (2019) service-to-manufacturing ratio to standard trade elasticities put substitution across service-industry varieties near 2.8[16], while substitution across individual service establishments within a city is about 9 (Couture, 2016). A village's service bundle is likely something roughly between an industry composite and a single establishment. I set $\sigma_n = 5$ and equal to the city-variety elasticity σ_u, so that a single substitution value governs service trade throughout the model. There are limited concerns that this value is driving any important results; the composition moment of Section 7.3 is close to flat in σ_n near this range, so the data would not distinguish values within it, and Appendix K.13 sweeps the city-good elasticity σ_u and finds similar results throughout the range.

Table 15: 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 5 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. K.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. I.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 7.3)			
s_{own}	0.533	moment	structural-transformation slope b_{ST} on the reduced-form measure, full sample, joint equilibrium
T_a	5.4×10^{-23}	inner	transmigration agricultural employment level (0.648)

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

7.3 Inversion and Internal Calibration

Inverting the non-program economy. I begin by recovering location-specific fundamentals for every non-program location. For each rural location, I choose the agricultural productivity and service-sector fundamental that make the model reproduce its observed income per capita and agricultural employment share. These two observed outcomes are enough to recover the two fundamentals for each location. Urban locations are treated in the same way, using an agricultural productivity and

[16] Their own direct estimates are higher, around 5 to 6

the city-good fundamental. Given the measured parameters and demand shares from Section 7.2, this construction matches the observed income and sectoral employment of every location that is not part of the transmigration program. Appendix J.3 provides the algebra and computational details.

Transmigration villages. The Transmigration villages are deliberately excluded from this inversion. Rather than choosing village-specific fundamentals to match their outcomes, I give them common agricultural and service fundamentals and solve for their outcomes in the model. The common agricultural sorting scale T_a is chosen so that the model reproduces the program-wide agricultural employment share of 0.648. Under the common endowment convention of Section 7.1, geography is therefore what generates differences across Transmigration villages in the model. This is important for the exercise that follows: the relationship between market access and income is not produced by choosing village-specific fundamentals to fit those outcomes.

Internal calibration. One parameter remains to be estimated inside the model. As discussed above, SUSENAS does not fully reveal whether some household purchases are sourced within the village or from outside it. I denote this local sourcing share by 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 s^{own} using this relationship. The targeted moment is the relationship between market access and structural transformation among the Transmigration villages. Specifically, I use the coefficient b_{ST} from a regression of the non-agricultural employment share on the same reduced-form measure of market access used in the empirical analysis. In the data this coefficient is 0.0053 (s.e. 0.0009).[17] For each candidate value of s^{own}, I update the corresponding expenditure shares, re-invert the non-program economy, and solve the full joint equilibrium. I then estimate the relationship between market access and non-agricultural employment shares in the model. At each iteration I also adjust T_a to continue matching the program-wide agricultural employment share. I choose s^{own} so that the model's structural transformation slope matches its empirical counterpart. This procedure gives $s^{\text{own}} = 0.532$. The estimate also falls comfortably within the range of values suggested by the SUSENAS data. The observed purchase categories imply that s^{own} could plausibly lie between about 0.471 and 0.904, so the calibrated value of 0.532 is within the range suggested by the survey.

[17] Both the data and model regressions include island and cohort fixed effects, with standard errors clustered by village in the data.

7.4 Fit

Table 16: 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)	

Notes: 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 J.4). Estimation on the functional-urban inverted world with the city-good price index from the inverted urban block.

The calibration uses two moments. The common agricultural sorting scale T_a is chosen to match the program-wide agricultural employment share, while s^{own} is chosen to match the relationship between market access and structural transformation. Because the system is exactly identified, both moments are matched exactly. The estimated sourcing share, $s^{\text{own}} = 0.532$, lies comfortably within the range allowed by the expenditure data.

There is also a useful untargeted check on this estimate. At the estimated sourcing share, the value-added crosswalk implies an own-village share of service consumption of about 0.35. The model independently produces an equilibrium own-village share of about 0.30. These objects need not coincide, since the first comes from the expenditure accounting and the second is an equilibrium outcome of the spatial model, but they are close. I now turn to characterizing how well the model matches the untargeted relationship between market access and income.

8 Model Results

This section reports the estimated model's ability to generate a negative relationship between market access and income. While the relationship between market access and structural transformation is targeted, there is no guarantee that I recover (i) a negative relationship between market access and income and (ii) that the strength of this relationship is of a similar magnitude as what I observe in the data.

8.1 The Income Gradient

The model is able to consistently generate the negative effects of market access on village income per capita. I find that the model is able to explain between roughly 30–130% 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. Table 17 reports the comparison between data and model for the full sample of transmigration villages, but my discussion in this section will focus on the results when I ignore the most isolated 5% of my sample.

Using the same reduced form measure of market access as in the results from Section 4, I find that the model is able to explain just over 80% of the relationship between market access and GDP per capita in the transmigration villages, while using the model-consistent measure explains 50% of the effect. Figure 7 shows the relationship between data and model together, while Figure 8 demonstrates the impact that excluding the most isolated villages has on the estimates. The premium earned by the most remote villages is a distinct phenomenon the model does not capture, and Section 8.2 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 17: Income gradients: model against data

Slope of ln real GDP pc	Sample	Data (s.e.)	Model	
estimation measure	full	−0.0191 (0.0050)	−0.0056	untargeted
estimation measure	excl. bottom 5%	−0.0142 (0.0050)	−0.0071	untargeted
estimation measure	excl. bottom decile	−0.0091 (0.0049)	−0.0087	untargeted
reduced-form measure	full	−0.0066 (0.0016)	−0.0028	untargeted
reduced-form measure	excl. bottom 5%	−0.0047 (0.0017)	−0.0038	untargeted
reduced-form measure	excl. bottom decile	−0.0030 (0.0018)	−0.0040	untargeted
externality off ($\varphi = 0$)	full	—	+0.0070	untargeted

Notes: Regressions include Island and cohort fixed effects, district-clustered standard errors. Model slopes are computed in equilibrium. The $\varphi = 0$ row solves the full economy with the externality removed and the productivity fields re-inverted to reproduce the same observed baseline and with the transmigration agricultural sorting scale re-pinned to match the observed mean program agricultural employment share. "Estimation measure" refers to the measure of market access computed within the model equilibrium, while "reduced-form measure" corresponds with the exact measure used in Section 4.

Figure 7: Income and Market Access in the Data and the Model

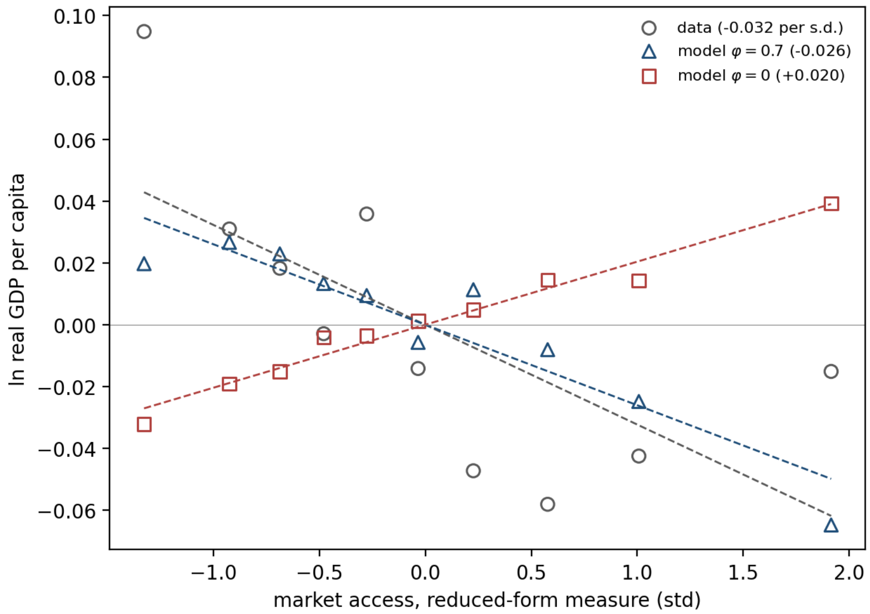

Note: I plot the relationship for transmigration villages between income and market access in the data and the model equilibria using the same method. The lowest 5% of transmigration villages in market access are excluded. For both data and model, I plot ten quantile bins of the island-and-cohort fixed-effect residuals for the 1,089 transmigration villages, with the fitted slope per standard deviation of log access and its district-clustered standard error in the legend. I include the model results both with and without the externality.

Importantly, shutting off the externality reverses the effect, and Figure 7 highlights this result by including 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 results from the model show that the externality is not one contributor to the curse among several, but rather the primary driver of the negative relationship between market access and GDP per capita.

Figure 8: Income-Access Relationship as Isolated Villages Are Excluded

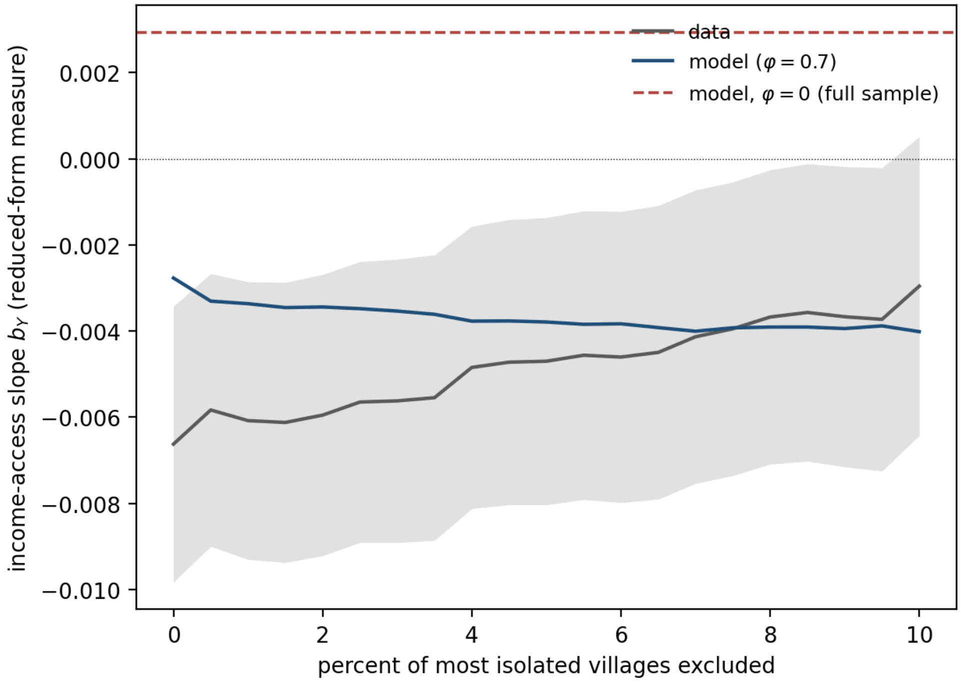

Note: I plot the relationship for transmigration villages between income and market access in the data and the model as the lowest market access villages are excluded, from zero to ten percent. I also plot the 95% confidence intervals for the data results. For reference, I include the full transmigration sample estimated relationship between market access and income in the model with $\varphi = 0$ as a horizontal line in red.

The model also makes an untargeted prediction about how the curse varies with settlement size. In the data, the fully interacted specification of Appendix Table C2 shows that the negative effect of market access weakens with initial settlement size. Solving the model at each village's recorded settlement size produces the same configuration (Appendix Table C13). The effect of access at the mean settlement size is nearly identical in the model and the data (-0.0042 against -0.0044), and the interaction is positive in both, though the model understates the attenuation, at about 0.002 per log point of settlement size against 0.008 in the data. The direct return to settlement size is strongly positive in the model but a noisy zero in the data, which is expected, as the model treats initial size as the permanent workforce while in the data five decades of population change separate the two. The model also explains the attenuation. Larger settlements are more agricultural, because agriculture absorbs labor at nearly constant returns while the demand for a village's services is limited by its access to buyers, and the curse operates on the non-agricultural margin, so more agricultural villages are less exposed to it. No moment involving settlement size is targeted in the calibration.

8.2 Structural Balance

I next turn to an alternative exercise to provide additional evidence that the effects I am identifying are the result of market access and not selection. The balance of observable characteristics with respect to market access in Section 3 makes the case that there are no omitted or unobserved factors that are correlated with initial market access and decrease income. While the balance across a wide range of controls provides convincing empirical support to this case, the model permits a structural analog to the same test, conditional on the model's structure. Inverting each transmigration village with the same process described in Section 7.3 to reproduce their observed employment share and income exactly allows me to recover the fundamentals for each location that exactly replicate the observed relationship between market access and income, by construction. I can then regress the recovered fundamentals on market access to ask whether the estimated model needs a relationship between fundamentals and market access in order to fit the data exactly. It does not. Figure 9 shows that I find no relationship between a transmigration village's baseline agricultural productivity and market access, and I find a similar null effect for the service fundamental. This result is the structural equivalent of the balance tests in Section 3, and I again find no evidence of any unobserved factor driving the relationship between market access and income. More convincingly, this null relationship is only a property of the model *with* the externality. The same exercise carried out in the $\varphi = 0$ world uncovers a strong negative relationship between market access and baseline agricultural productivity (also shown in Figure 9). 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, this exercise provides evidence that excluding an externality in the agricultural sector actually produces a world that is at odds with the data. Leaning on the success of the model in reproducing an untargeted relationship between market access and income in the sample of transmigration villages where I have identified a similar effect, I now conduct various counterfactual exercises to highlight the policy relevance of this channel and to estimate the potential welfare gains from various policy interventions.

Figure 9: Recovered Agricultural Productivity and Market Access

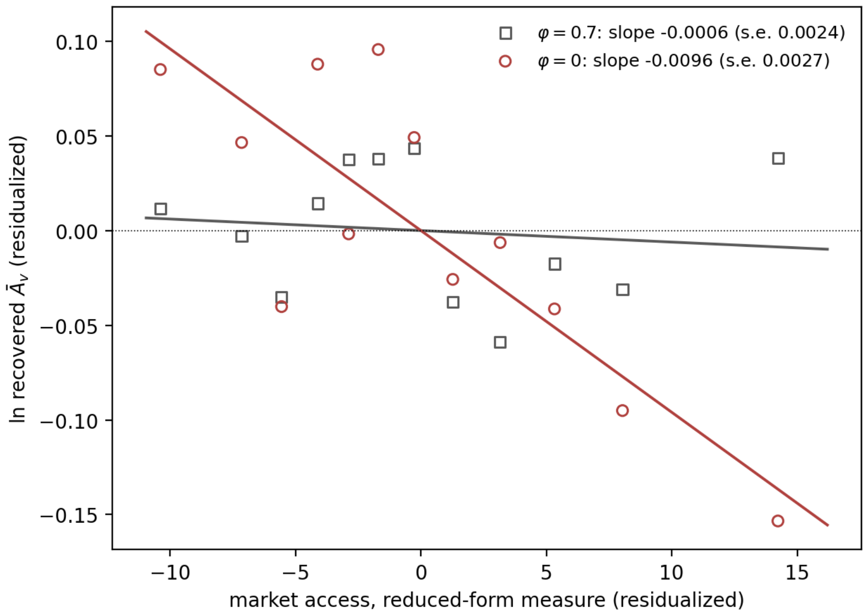

Note: I plot the relationship between the recovered program-village agricultural productivity and market access, residualized on island and cohort fixed effects. I include twelve quantile bins and the linear fit, slope and standard errors clustered at the district level in the legend. I plot this relationship both with the externality ($\varphi = 0.7$) and without ($\varphi = 0$).

9 Counterfactuals

This section uses the estimated model to evaluate various policy-relevant counterfactuals. Every experiment re-solves the full general equilibrium under the counterfactual shock, so the baseline is the observed economy whose behavior has been detailed above. Village and city incomes, occupational sorting, service prices, and the urban variety block all respond endogenously, food is the numeraire, and the national food market clears by Walras' law. Fundamentals are held fixed at their baseline inverted values throughout, with transmigration villages keeping their common fundamentals. For each experiment I also solve for the estimated effects when $\varphi = 0$ to quantify how the externality changes the results.

The section comes in two parts. Sections 9.1 and 9.2 use the model exactly as estimated and ask cross-sectional questions, a road built to a single program village and the anatomy of which villages such a road hurts. The remaining experiments evaluate aggregate programs, and for those the model needs one addition. Under homothetic demand the composition of national spending is fixed, so aggregate structural transformation is ruled out by construction and the aggregate cost the externality attaches to it cannot appear. Section 9.3 extends the model with a measured non-homothetic demand system, and Sections 9.4 through 9.7 then evaluate the whole road network and its cost, agricultural subsidies against roads, rural agricultural productivity growth, and the value of the program itself, each with and without the externality.

9.1 A Road to a Program Village

The first experiment improves a single transmigration village's travel times by ten percent and nothing else. I conduct this exercise for each of the 1,147 villages and then aggregate the results. In line with the empirical and model predictions, I find that, with or without the externality, improved connectivity leads to reallocation out of agriculture and into rural services. However, 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 sixty percent of villages made worse off, while without the externality every village is made better off by the improved connectivity. Figure 10 plots the distribution of effects for both income and agricultural employment. These results highlight the fact that observing structural transformation itself is not a sufficient condition for improved welfare.

The decomposition of the effects is also presented in Table 18. In both economies villages benefit from improved connectivity via cheaper access to goods, but the increase in external demand for services is the primary driver of the overall effect of the policy. Increased demand for services pulls workers out of agriculture regardless of whether there is an externality in the agricultural sector. However, the erosion of agricultural productivity that comes with this when there is an externality in agriculture outweighs the potential gains from this and all other positive channels.

Table 18: 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%	

Notes: Population-weighted means across the 1,147 single-village experiments. Each experiment cuts one village's excess travel times by ten 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 10: Effects of a Road Built to a Single Program Village

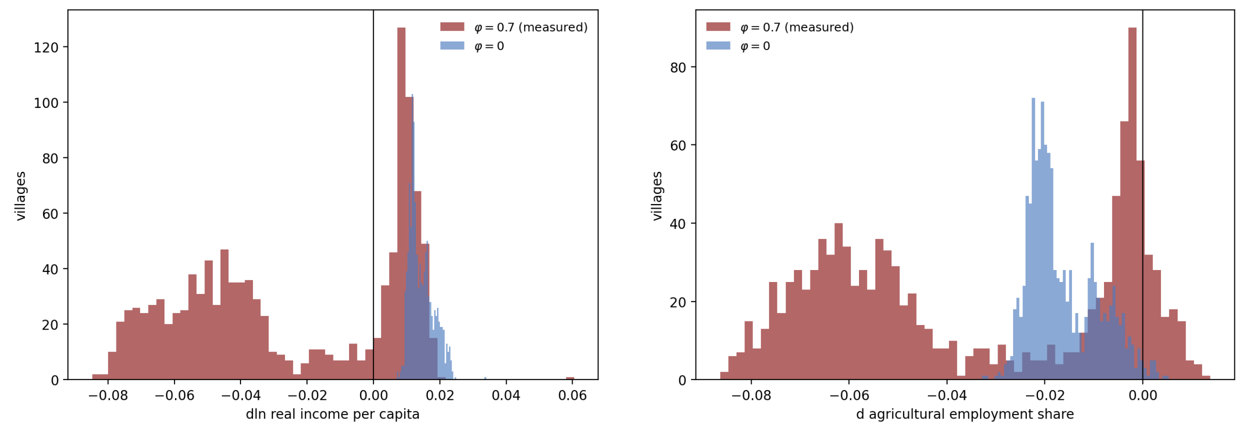

Note: I plot the distribution of effects across the 1,147 separate experiments, each of which cuts one village's travel times by ten percent and re-solves that village against the unchanged rest of the economy. The left panel plots the change in log real income per capita and the right panel the change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

This experiment is the model's answer to the additional evidence on the effects of village roads. Asher and Novosad (2020)'s estimates for India's rural road construction program find 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 itself. 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 ten percent at once, the program-wide version of the one-village experiment, 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.

9.2 The Targeted Road by Initial Specialization

This experiment also allows me to examine the heterogeneous responses to improvements in connectivity. Figure 11 plots the change in real income from the connectivity shock against each village's agricultural employment share before the shock. Without the externality every village gains between one and two percent.[18] The model reveals more interesting heterogeneity when the

[18] 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.

externality is included. While villages that were almost entirely agricultural still gain from a connectivity shock, every village with an agricultural share below 0.7 loses. The negative relationship between the change in income and the non-agricultural employment share is predicted in Proposition 1. 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 outlined in Proposition 1, while a village with a significant share of workers in both agriculture and non-agriculture is more responsive to the connectivity shock; it experiences a larger decrease in the agricultural employment share, and a more significant loss in productivity due to the externality.

Figure 11: Road Effects by Initial Agricultural Employment Share

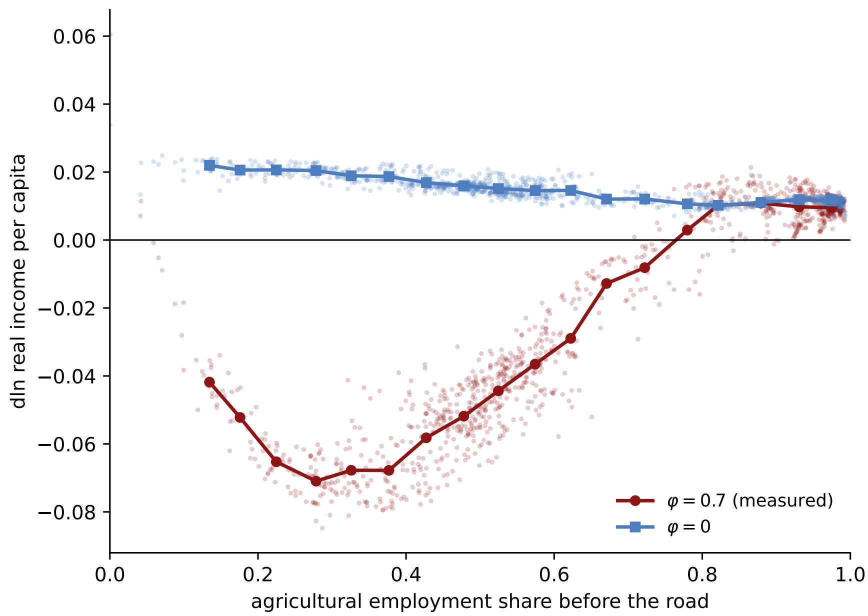

Note: I plot the change in log real income per capita from the single-village road experiment against the village's agricultural employment share before the shock. Each point is one of the 1,147 experiments, and the lines connect means within bins of width 0.05 in the pre-shock share, using bins that contain at least eight villages. I plot both the model with the externality ($\varphi = 0.7$) and the model without it ($\varphi = 0$).

9.3 Extending Demand for Aggregate Counterfactuals

The experiments so far move activity across space. In order to examine what the model predicts for larger shocks and policies, I need to extend the model slightly. The remaining experiments improve the entire road network, subsidize agricultural earnings, raise rural agricultural productivity, and remove the transmigration program. Shocks of this size move the aggregate agricultural employment share in the data. The model described in Section 6 features homothetic demand, meaning that the composition of national spending does not change in response to shocks, so even large shocks cannot generate *aggregate* structural transformation. This has two consequences. First, it shuts down a response that these policies should produce. Second, the agricultural externality attaches a cost to that response, so the homothetic setting cannot capture what the externality implies for their aggregate welfare effects.

I therefore extend the model with the non-homothetic demand system of Boppart (2014). Fan et al. (2023) use the same system for India. The PIGL structure I impose means that as households grow richer they shift spending away from food and toward rural services and city goods. Budget shares take the form

$$s_{g,\ell} = \omega_g + \nu_g \, x_\ell^{-\varepsilon}, \tag{19}$$

where x_ℓ is real expenditure per capita, ω_g is the budget share a household approaches as it grows rich, and the loading ν_g is positive for food and negative for services and city goods. Differentiating gives the response the counterfactuals run on,

$$\frac{\partial s_{g,\ell}}{\partial \ln x_\ell} = -\varepsilon \left(s_{g,\ell} - \omega_g \right), \tag{20}$$

so the single parameter ε governs how strongly spending recomposes as income rises, and the response fades as a location's budget approaches ω_g. Homothetic demand is the case $\nu_g = 0$, where every share is constant and (20) is zero.

I estimate the Engel curves and curvature parameter with the estimator used in Fan et al. (2023). My estimated ε of 0.34 is nearly identical to Fan et al. (2023)'s estimate for India (0.33). Appendix I.7 writes out the preferences and the estimation details. Because the system is anchored to each location's observed baseline budget, the calibration and every result of Section 8 is unchanged. Here the income margin operates only when a counterfactual moves real incomes for a location. Real income is measured exactly as in the main model, with income deflated at each location's baseline budget shares, which Appendix I.7 shows is still the correct first-order welfare metric.

In this context, this extension does almost nothing on its own. At $\varphi = 0$, switching demand systems changes the aggregate gains from the road improvement and rural productivity growth experiment by 1.2% and 0.8%, respectively, since recomposing demand is close to costless when there is no externality. Anything the income margin contributes is therefore due to the interaction of the externality and non-homotheticity. Each experiment is solved under both demand systems and at both values of φ, so the tables report what the externality changes in both demand systems.

9.4 A Uniform National Road Improvement

A uniform ten-percent cut in all travel times in Indonesia makes every type of location better off, with or without the externality; the program villages, the rest of the countryside, and the cities all gain. What the externality changes is the size of the aggregate gain. Table 19 reports the results from the counterfactual. National real income rises by 0.73 percent with the externality in place, against 0.86 percent without it, so the externality reduces the aggregate gains by roughly fifteen percent.

The aggregate gain is smaller with the externality because of the interaction between the externality and non-homothetic demand. As real incomes rise in response to the positive effects of the aggregate road improvement, households in both urban and rural locations respond by spending a smaller share of their budget on food. The relative decrease in aggregate demand for agricultural output pulls workers out of agriculture, which lowers agricultural TFP through the externality. Without the externality that reallocation does not lower TFP, and the aggregate gain is essentially unchanged.

Comparing the aggregate gains from the counterfactual is where the extension of Section 9.3 is highlighted. Under homothetic demand the same experiment generates an aggregate income increase of +0.83 percent with the externality against +0.85 percent without. With homothetic demand, the conclusion would be that the externality has no bite outside of how it shapes the cross sectional distribution of income in the economy. With non-homothetic demand, the aggregate gains of the road are roughly fifteen percent lower with the externality than without it. Nearly all of the difference operates through the combination of the income response and the externality, and almost none of it through the reallocation across space that the homothetic model already captures. The reason is the one given in Section 9.3. With fixed budget shares the aggregate agricultural spending share is the same at every equilibrium, so there is no aggregate reallocation for the externality to act on, and the cross-sectional reshuffle that remains nets out at the national level.

The gains above are gross, but maintaining transport infrastructure quality requires constant investment. Appendix L.2 works out that this sort of infrastructure improvement would cost at least 0.68 percent of GDP a year. The improvement results in an increase in aggregate income of 0.73 percent, so the improvement is just worth the cost. However, that is the favorable case, and if raising speeds instead requires other investment, the benefits come to well under half of what the work costs. Published benefit-cost ratios for road maintenance in Indonesia are higher than these (Gertler et al., 2024), though they are not measuring quite the same thing. They estimate in the short-run what individual districts gain when their roads improve, while the figure I obtain is from a full general equilibrium treatment of an improvement in the entire travel network.

Table 19: A uniform ten-percent travel-time reduction

	$\varphi = 0.7$ (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%

Notes: Population-weighted changes in log real income, in percent, from a uniform ten-percent cut in all travel times. 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 L.2. The $\varphi = 0$ columns re-solve the identical experiments in the re-inverted externality-free world.

9.5 Agricultural Subsidies versus Roads

Given the introduction of an externality into the model, a natural question is what the effects are of policies that try to target the externality directly. Agricultural support programs are commonplace in developing countries, and my model provides a theoretical justification for how policies that subsidize agriculture can be welfare-improving. I carry out two different counterfactuals that feature a subsidy on agricultural earnings, and both are financed *within* each subsidized village by a lump-sum tax on its households. First, I calculate the optimal subsidy[19] to maximize transmigration villages' income and apply it to each one. Second, I estimate the effects of an agricultural earnings subsidy at the same fiscal cost as the road improvement in Section 9.4 and spread that across all rural villages.

Table 20 presents the results of both experiments, with a comparison to the national road improvement included. The fully corrective subsidy to program villages is highly effective at raising transmigration village incomes, with villages increasing their real income by 24.2 percent. Additionally, because the externality has been correctly priced in this experiment, I find that after the subsidy the relationship between market access and income is positive and essentially the same as the estimated relationship at $\varphi = 0$ in Table 17.

While the full subsidy to transmigration villages raises their own income, the policy itself results in small aggregate losses (−0.06 percent of national income). As the subsidized villages expand agriculture, they push nearby villages out of agriculture and also reduce their supply of the non-agricultural variety of their village. A national planner has no reason to correct the externality for these 1,147 villages and nowhere else, so the next experiment applies a smaller subsidy more broadly and examines its impact.

The subsidy becomes valuable in aggregate when it is applied to every rural village. I use

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

the estimated annual budget of the road program of Section 9.4, roughly 0.68 percent of GDP, and use that as the total outlay of the rural agricultural earnings subsidy program. Because the subsidy affects labor earnings directly, and only in rural locations, the subsidy rate is almost exactly 8 percent. Note this is small compared to the earnings subsidy of 230 percent in the previous experiment on program villages. The subsidy raises national income by 0.20 percent, with the gains coming from rural locations while urban locations are made slightly worse off. The gains to rural locations are several times what the road program generates. Once the cost of the road program, which is not recouped via taxes, is taken into account, the subsidy delivers roughly four times the national gains.

The gains from the subsidy do not survive without the externality. The $\varphi = 0$ column of Table 20 shows the results from the identical subsidy when the externality is shut off. Without the externality, the subsidy simply introduces a distortion in the economy and generates a small loss in national income. These results suggest a new interpretation of agricultural support programs that are common across developing countries. While these programs are usually defended as redistribution and criticized as distortions, in settings where agriculture carries scale economies they can also increase aggregate income.

Table 20: 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%	—

Notes: All entries are percent of national income per year. The road row comes from the uniform ten-percent network improvement described in Section 9.4. 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 that are associated with infrastructure projects. 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.

9.6 Agricultural Productivity Growth

The agricultural externality I find also has important implications for the effects from other drivers of aggregate structural transformation. I simulate the effects of an increase in agricultural productivity by raising the agricultural TFP of every rural location by five percent. This is the closest the model comes to the force usually credited with producing aggregate structural transformation. Here the externality has large aggregate implications. Table 21 reports the effects of rural agricultural productivity growth both with and without the externality.

Table 21: 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

Notes: Population-weighted changes in log real income, in percent, and the change in national real income, from a five-percent rise in the agricultural TFP of every rural location. 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.

I find that the externality significantly dampens the aggregate gains from agricultural productivity growth. Rural agricultural productivity growth raises national income by 0.74 percent with the externality compared to 1.21 percent without the externality, roughly forty percent less. There are two reasons for the gap in aggregate gains. Rural agricultural productivity growth shifts more agricultural activity from urban locations to rural areas. Urban agricultural employment falls by almost 2 percentage points after rural areas become more productive. Because of the externality, this erodes the cities' agricultural TFP and income. This channel operates regardless of demand system; with homothetic demand the gains from this experiment are still roughly 25 percent smaller due to the externality. The rest of the difference is explained by the interaction of non-homothetic demand with the externality. Because income increases due to the positive productivity shock, spending shifts away from food and toward non-agricultural products. The decrease in relative demand for food translates into more structural transformation, which reduces the agricultural productivity gains from the shock. Of note, without the externality the choice between demand systems has little impact on the aggregate gains from this policy (1.21 against 1.20).

9.7 The Value of the Program

The most basic counterfactual is a world in which the transmigration program never ran. I remove the 1,147 transmigration villages and return their 2.3 million settlers to the Java/Bali countryside, spread across its rural villages in proportion to population (an eleven percent increase in the inner-island rural population). Table 22 reports the results. The largest effect depends on land access itself, which was a central motivation for the program. On the outer islands, settlers cultivate roughly half a hectare per person, several times what the same households would hold in Java or Bali. Because agricultural income rises steeply with land per worker, returning them cuts their real income by 59 percent. The program more than doubled the settlers' income and raised national income by about 1.4 percent, and both numbers are nearly identical in the $\varphi = 0$ world because they come from moving people onto land rather than from anything related to agricultural scale. These two figures should be read as upper bounds. The inversion matches each inner-island village's observed income, so the gap a settler faces between Java or Bali and the outer islands reflects everything that differs between those regions, not only the land they would farm. Part of what these numbers credit to land access is really that wider difference. The incidence result below does not depend on it, since it compares the same inner-island households across the two externality values.

Here the externality determines the incidence of the program. Returning 2.3 million individuals to rural areas of Java and Bali helps or hurts the households already there depending entirely on whether agricultural scale matters. Without the externality the extra workers simply crowd fixed land and incumbent real income falls 6.5 percent. With the externality, the more populated countryside is more productive and incumbent income rises 0.8 percent, despite the negative effect of increased land crowding. An evaluation of the program that ignores the externality therefore gets the incidence on the sending region wrong. There are two important caveats to this experiment. First, I do not have any data on where transmigrants came from, so I am not able to reallocate migrants back to their original locations. Second and most importantly, the results rely on my estimated agricultural scale externality translating to the more-crowded inner islands. It is possible that the externality I estimate is actually non-linear, and I am capturing a particularly steep part of that relationship among smaller villages. If this were the case, then this would have important implications for the incidence I estimate here in this experiment, as the gains in productivity from a larger agricultural workforce might not outweigh the land-crowding effect that reduces rural incumbent income.

Table 22: No transmigration program

No transmigration program	$\varphi = 0.7$	$\varphi = 0$
National real income	-1.4	-1.5
Transmigrants' own real income	-59	-58
Java/Bali rural incumbents	$+0.8$	-6.5
Outer islands, total real income	-4.1	-3.5
Inner islands, total real income	$+1.6$	$+0.8$

Notes: The counterfactual removes the 1,147 transmigration locations and redistributes their population across Java and Bali rural villages in proportion to population, holding national population fixed. Entries are percent changes in real income relative to the observed economy with the program in place. "Transmigrants" is the change for the resettled population, evaluated at the inner-rural incomes they would have earned. "Java/Bali rural incumbents" is the within-region change for the population already living there. The $\varphi = 0$ column re-solves in the re-inverted externality-free world.

10 Conclusion

I leverage exogenous variation in the initial market access of Indonesian Transmigration villages to study the long-run effects of connectivity on measures of economic development. I find that increased initial market access causes villages to have lower productivity and estimated GDP per capita as measured contemporaneously. I present evidence that these results are not driven by agricultural specialization, competition, differential migration, compensating differentials, or price effects. Instead, the evidence points to a mechanism rooted in agriculture: more connected villages undergo more structural transformation out of agriculture, which erodes agricultural TFP through a scale externality. I directly estimate this externality using program variation in settlement size and find a positive scale externality. A quantitative spatial general equilibrium model of the full Indonesian economy with the estimated externality reproduces the negative effect of market access on income among program villages as an untargeted prediction.

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 often make that village poorer. Additionally, I find that not modeling the agricultural externality leads to overly optimistic estimates of the gains from different development policies. Once the externality is taken into account, the aggregate gains from improving the whole travel network and from rural agricultural productivity growth are both significantly smaller.

More generally, this paper makes the point that structural transformation itself is not sufficient for generating growth. 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.

Baragwanath, Kathryn, Gordon Hanson, Amit Khandelwal, Chen Liu, and Hogeun Park (2024) "Using Satellite Imagery to Detect the Impacts of New Highways: An Application to India," *Working Paper*.

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*.

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. 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*.

Desmet, Klaus, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg (2018) "The Geography of Development," *Journal of Political Economy*, 126 (3), 903–983, 10.1086/697084.

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).

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).

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.

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).

Muralidharan, Karthik, Mauricio Romero, and Kaspar Wüthrich (2025) "Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments," *The Review of Economics and Statistics*, 107 (3), 589–604.

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*.

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).

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).

Appendix for
"The Curse of Connectivity"

A Direct Qualitative Evidence of Program Implementation and Data Table

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 periods (called Repelitas) over this period, I document each period separately as well:

A.1 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)

A.2 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."

A.3 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% relative to Repelita III[20]. 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

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

by itself rule out subtler channels, such as provincial governments choosing which land to submit for settlement; the empirical design of Section 3, 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.

B Data Appendix: Sources and Construction

This appendix provides detailed descriptions of the data sources summarized in Table A1 and used throughout the analysis. Construction of the pre-transmigration market access measure is described in the main text and in Appendix C.

B.1 Local Measures of GDP per Capita

The outcome of first order importance involves measures of per capita output. Due to the difficulty in obtaining measures of local output at granular geographic level, my primary measure of local GDP per capita comes from high-resolution estimates of GDP from Rossi-Hansberg and Zhang (2025). The measure combines a diverse set of inputs beyond simply population and nightlights, and performs significantly better at predicting local GDP than previous measures. As Figure 5a shows, I find significant variation in this measure across transmigration villages. Additionally, I provide an alternative measure of local aggregate output per capita by combining measures of nightlights from Li et al. (2020)[21] and population from the Global Human Settlement Layer. While this measure is subject to some of the potential issues raised in Rossi-Hansberg and Zhang (2025), it represents a more transparent estimate of aggregate output. 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 5, 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 5d 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.

[21] Converted to estimates of GDP following the elasticities in Beyer et al. (2022), Hu and Yao (2022), and Henderson et al. (2012)

Table A1: 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 5)
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)

Additionally, I complement my measures of GDP per capita and productivity with individual level survey data on wages from multiple waves of the National Socioeconomic Survey (SUSENAS) between 2000-2007. The SUSENAS survey randomly samples multiple individuals within random villages across many districts in Indonesia each year. While there is no guarantee that a Transmigration village will be included in the sample villages in a given year, by pooling multiple waves of the survey I am able to observe wages of individuals working in formal employment in roughly one third of my villages at some point in time. I then adjust the monthly and hourly wage responses to convert all wages to their 2007 equivalents. While the fact that I only observe wages in a smaller fraction of my transmigration villages and only of those that work in formal employment poses some challenges related to statistical power and the interpretation of the causal effects, it provides important contextual backing to the more abstract measures of GDP per capita and proxies for local productivity.

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 5. 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 seven 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

In order to measure the effects of initial market access on the sectoral employment composition I use data from two waves of the Indonesia Population Census. In both 2000 and 2010, I have access to a complete census of individuals in the country. I am able to observe every individual's demographic information, education, employment status, and sector of employment. I then use this information to construct village-level measures of sectoral employment shares. Figures 5b and 5c plot the distribution of agricultural and formal employment for the transmigration villages in 2000. While most villages are still primarily agricultural, there exists a good deal of variation within transmigration villages. Similarly, the most common employment type in transmigration villages is informal employment. I supplement these census measures with village level data on sectoral employment from the 1993 wave of the village level potential survey (PODES).

B.5 Firm Dynamics

In addition to the effects of market access on individuals and their employment decisions, I am able to estimate the effects of market access on firm characteristics at the village level. I create village-level measures of formal firm presence and employment using the listing of all medium and large enterprises from the 2016 Economic Census. This data contains the universe of medium and large enterprises across Indonesia, as well as a large sample of micro and small enterprises in each urban location, and collects information on the owner of the firm, detailed sector level classifications, and employment counts. I collapse this listing down to the village level to construct an indicator for if a village has a medium or large firm, as well as the average employment of individuals in medium and large firms within a village. While only 13% of transmigration villages have at least one of these firms, these measures still provide important details about the nature of firms within villages.

I also am able to provide additional village-level information about firm presence from PODES waves. These detailed surveys are administered to village heads and cover a variety of questions about life within their village. From these surveys I construct measures of the number and density of small industries within a village and the number of shops in a village. These measures provide a more accurate picture of the smaller scale production and trade that is present within transmigration villages.

B.6 Local Prices and Expenditure

In order to better understand local variation in prices and expenditure patterns in transmigration villages, I leverage data from two waves (2002 and 2011) of the SUSENAS that contained detailed consumption modules. In both years I am able to construct measures of the price level in transmigration villages for food expenditure[22]. Additionally, in 2002 I can construct measures of total expenditure on other non-food categories, and the price per square meter estimates of housing value. Because this data only comes from two waves of the SUSENAS, I have a relatively limited sample where I can measure price and expenditure behavior; I observe the food price index for 84 villages and 1234 households and non-food expenditure for 70 villages and 1086 households. While other waves of the SUSENAS contain more coarse measures of total consumption expenditure, they do not separate out prices and quantities, and so I am unable to construct a measure of prices for this larger sample.

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

B.7 Amenities

Given the importance of a location's amenity value in determining the distribution of individuals across space, I measure each transmigration village's provision of public goods across a wide range of indicators. I use measures of public goods that are provided by both local village and regional governments and include measures of health care availability and quality, school presence and staffing, the provision of safe drinking water, waste collection and other measures of sanitation, and road maintenance and quality. I then create a public goods index by standardizing each of these 18 measures and taking the average.

B.8 Migration and Commuting

Because the census also asks each individual their birth location, I construct the two village level migration variables used in Section 4: the log number and the share of current residents who were born on the inner islands (Java and Bali), the sending region of the transmigration program. Perhaps not surprisingly given the nature of transmigration villages, these inner-island-born shares are often large. Finally, I measure the degree of village-level commuting using information from the 2003 PODES that includes questions on whether individuals within the village commute outside the village for work. Although the survey does not record the number of individuals that regularly commute, I am still able to assess the prevalence of commuting at the village level.

B.9 Land Use and Development

In order to measure dynamics around land use at the village level, I leverage detailed information on land use from the 2000 and 2003 waves of the PODES. These surveys record information on the amount of land within a village that is devoted to agriculture, housing, and firms. Crucially, the 2003 PODES also measures *changes* in land use over the last 3 years. This allows me to determine if villages with increased connectivity have different allocations in land use, in levels and changes.

B.10 Time-Invariant Characteristics

In order to characterize the time-invariant characteristics of transmigration villages *before* they were settled, I use a rich set of geographic variables that characterize a village's soil quality, terrain, and location. This information will allow me to demonstrate how location-specific traits of transmigration villages do not systematically vary with the initial market access of the village. These variables are the same as those used in Bazzi et al. (2016). I also identify and classify transmigration villages using data from the Transmigration Census. The census contains information on settled transmigration villages and includes their location, the year they were settled, and how many individuals were relocated to the village. This information is crucial in identifying transmigration villages, as well as how long that village has existed. The number of individuals initially settled in each village also serves as the instrument for the agricultural externality estimation in Section 5.

To complement these measures of the initial conditions of transmigration villages, I also use measures of the characteristics of the broader districts transmigration villages were placed in that are roughly contemporaneous to the transmigration program. Using the 1980 Population Census, I build a rich set of district-level variables that include the sectoral employment composition, formal employment rates, housing quality, agricultural potential yield, education and literacy, and asset ownership rates.

C Map Digitization Process

See main text for the location of the maps, and Figure A1[23] 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. Additionally, trails and other roads are the same color red, but dashed lines. 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.

In order 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[24] 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 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.

To estimate the travel times between origin o and all destinations d I implemented a fast

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

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

Figure A1: Northern Sumatra, Example Map

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. Table A2 displays the travel cost parameters.

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.

Figure A2: 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		

D Empirical Strategy Simulation

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[25]. 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[26]. 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 B1 displays an example draw of this simulation exercise.

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 exactly as in the empirical section as the province-level average market access of new villages. 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 B2 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

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

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

Figure B1: Map of Simulated Space

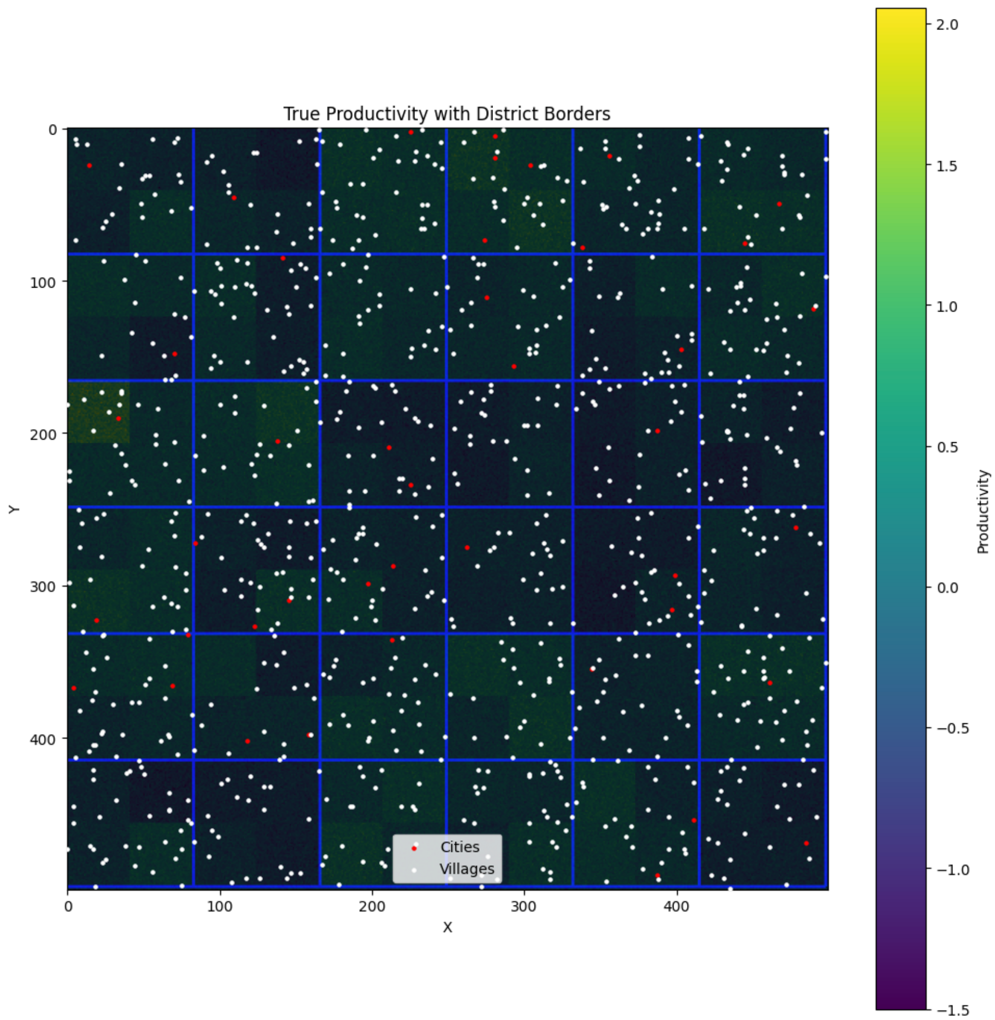

Figure B2: Simulation Results

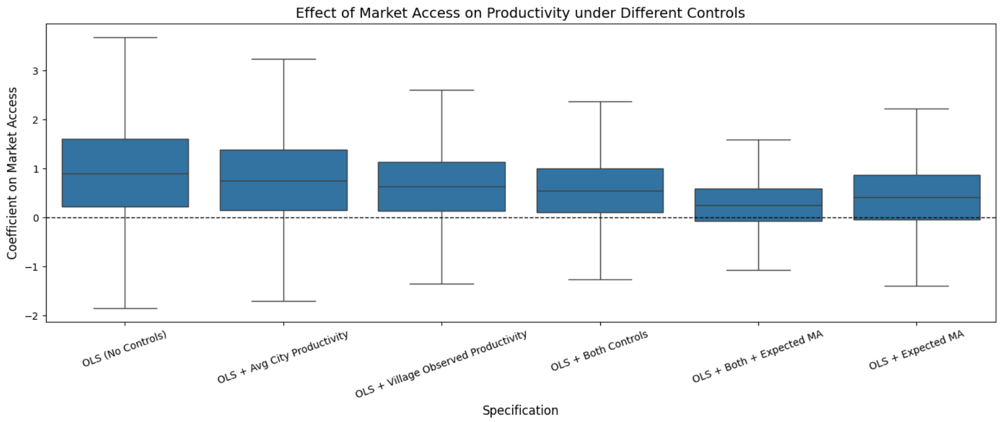

despite the fact that new villages are placed randomly. Importantly, I also find evidence that in this type of assignment process, each of my different controls helps address this threat to identification. In the specification that most closely matches my preferred estimation strategy in the true data (OLS + Both + Expected MA), the interquartile range of my estimates contains 0.

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 3 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.

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 4 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 seven 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 four 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 five 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 seven 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 B1 have slightly smaller samples than the composite.

For the regression reported in Section E.5, the two composite measures and the seven 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 seven 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 seven 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 B1 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. Reassuringly, 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 4 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 B1: 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

Notes. 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 seven 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 C1: Robustness of the Market Access Estimate to Inference and Assignment Strata

	ln_gdpc	ln_gdpc	ln_gdpc	ln_gdpc	ln_gdpc	ln_gdpc	ln_gdpc
ln(Market Access)	-0.0031***	-0.0031***	-0.0031***	-0.0031***	-0.0044*	-0.0026***	-0.0035***
	(0.0009)	(0.0007)	(0.0010)	(0.0008)	(0.0024)	(0.0008)	(0.0008)
Expected MA Control	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
Geography FE	Province	Province	Province	Province	Island	Prov x Repelita	Prov x Cohort
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SE Clustering	Kabupaten	Village	Kecamatan	Prov x Cohort	Kabupaten	Kabupaten	Kabupaten
Permutation p (within-cell)	0.000						
Unique Villages	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 more conservative approaches to inference and under fixed effects that absorb the assignment strata. The main specification of Table 4 uses island fixed effects and clusters standard errors at the village level. Columns 1-4 instead use province fixed effects and share an identical specification and sample, differing only in the level at which standard errors are clustered. Columns 5-7 cluster standard errors at the kabupaten level and vary the fixed effects, with column 5 retaining the island fixed effects of the main specification. 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 exact conditional on as-if random placement within cells. No reassignment does, so the p-value is below 0.001. The sample contains 290 province-by-cohort cells with a median of two 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 7 and are unaffected by the permutation reassignments.

Table C2: Fully Interacted Specification: Market Access and Settlement Size

	ln_gdpc
ln(Market Access), recenter	-0.0044***
	(0.0017)
ln(Individuals Placed), recenter	0.0083
	(0.0208)
ln(Market Access), recenter × ln(Individuals Placed), recenter	0.0083***
	(0.0026)
Expected MA Control	Yes
Local Controls	Yes
District Controls	Yes
Cohort FE	Yes
Island FE	Yes
Year FE	Yes
Unique Villages	890

Standard errors in parentheses

Standard errors clustered at the village level

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

Note: This table reports the specification of Table 4 with both of the program's treatments entered jointly. Market access and initial settlement size are in logs and centered at their sample means, so the market access coefficient is the effect at the mean settlement size and the settlement size coefficient is the effect at the mean level of access. Standard errors are clustered at the village level.

Table C3: Effects of MA 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

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

Table C4: Effects of MA 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 3. Columns 1-4 use data from the 2002 SUSENAS household expenditure module.

Table C5: 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 3. 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 4.

G The Simple Model

This appendix presents a simplified version of the quantitative model of Section 6. 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 H).

G.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,

$$Y_{a,v} = \bar{A}_a \, Z_{a,v}^{\gamma_a + \varphi} \, H_v^{1-\gamma_a}, \qquad Y_{n,v} = A_n \, Z_{n,v}^{g_n}, \tag{C1}$$

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.

G.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 service share $\beta_n = 0.10$ 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).

G.3 Baseline Results

Figure C1 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 C1: Simple Model: Outcomes by Distance from the City

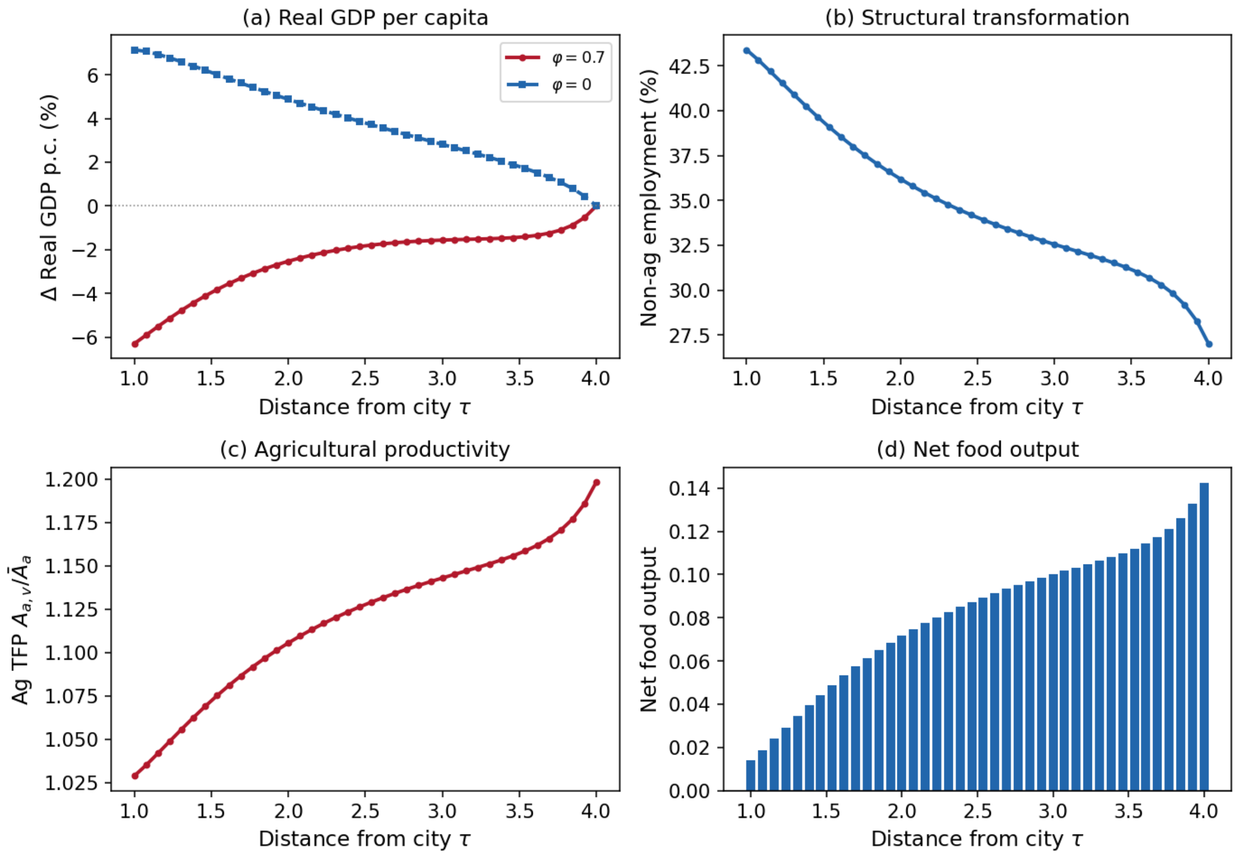

Note: I plot four outcomes for the forty villages of the simple model against their distance from the city. Panel (a) plots real income per capita relative to the most distant village, at the calibrated $\varphi = 0.70$ and at $\varphi = 0$. Panel (b) plots the non-agricultural employment share, panel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

G.4 Stress Tests

Table C6 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 C2 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 C6: 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

Notes: Each row re-solves the simple model at $\varphi = 0.70$ under one deviation from the baseline calibration of Section G.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 C2: Simple Model: Real-Income Profiles Across Stress-Test Variants

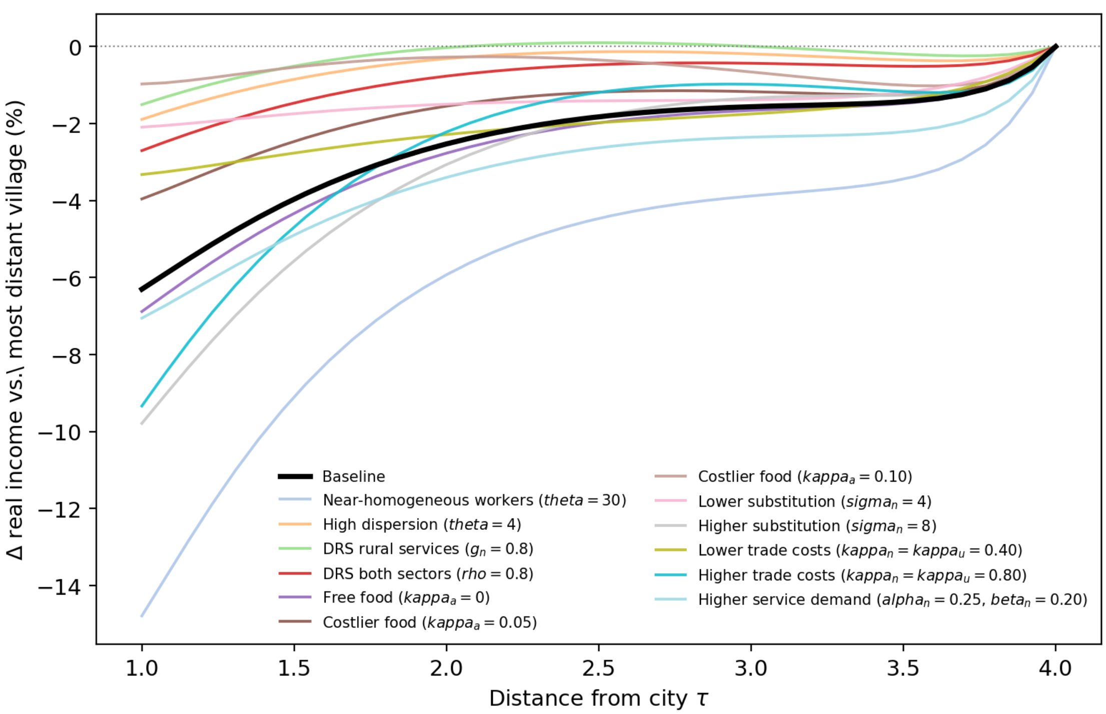

Note: Each curve plots real income relative to the most distant village against distance from the city, solved at $\varphi = 0.70$ under one deviation from the baseline calibration of Section G.2. The baseline is plotted as the heavy black line. The corresponding summary statistics are reported in Table C6.

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, if all sectors had the same trade friction, there would be no difference in structural transformation across the villages. 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 (Section K.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 5 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 G.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 6–9. 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.

H Derivation of the Sign-Reversal Threshold

This appendix derives the closed-form connectivity threshold stated in Proposition 1 of Section 6.2. The environment is the simple model of Appendix G 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.[27]

H.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,

$$Y_a = \bar{A}_a \, Z_a^{\gamma_a+\varphi} H^{1-\gamma_a}, \qquad Y_n = A_n \, Z_n^{\gamma_n}, \qquad \gamma_n \in (0,1], \tag{C2}$$

where $\gamma_n = 1$ is the constant-returns baseline and $\gamma_n < 1$ the decreasing-returns variant of Appendix G.[28] 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 H.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:

$$Z_a = 1 - u, \qquad Z_n = \frac{k}{k-1}\,\underline{z}\, u^{\frac{k-1}{k}}. \tag{C3}$$

[27] 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 C6 confirms the crossing in the Fréchet simple model.

[28] 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,

$$\frac{(1-u)^B}{u^A} = \frac{\tilde{C}}{\rho}, \qquad A = \frac{1+(1-\gamma_n)(k-1)}{k}, \qquad B = 1-\gamma_a-\varphi, \tag{C4}$$

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. Log-differentiating (C4) gives the structural transformation elasticity

$$\eta \equiv \frac{d\ln u}{d\ln \rho} = \frac{1}{A + B\,\dfrac{u}{1-u}} > 0, \tag{C5}$$

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:

$$d\ln I = s_a\, d\ln p_a + s_n\, d\ln p_n \;-\; \mathcal{E}\; d\ln\rho, \qquad \mathcal{E} \equiv s_a\,\varphi\,\frac{u}{1-u}\,\eta, \tag{C6}$$

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 one 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.

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

$$d\ln V = (s_a - \alpha_a)\, d\ln p_a + (s_n - \alpha_n)\, d\ln p_n - \alpha_u\, d\ln p_u^d \;-\; \mathcal{E}\; d\ln\rho. \tag{C7}$$

Each price enters with the village's net exposure to that good.[29] Net food exports, net service exports, and city-good imports sum to zero by the budget identity. At $\varphi = 0$ equation (C7) is the textbook terms-of-trade calculus; the externality contributes the single extra term $-\mathcal{E}\, d\ln\rho$.

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

H.2 The Three Frictions

Now place a continuum of such villages at distances τ from a single city, the arrangement of Appendix G. 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

$$D_c(\tau, p_n) = D\, \tau^{\kappa_n(1-\sigma_n)}\, p_n^{-\sigma_n}, \qquad D \equiv X_n/\Phi, \tag{C8}$$

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,

$$Y_n = \alpha_n \frac{I}{p_n} + D_c(\tau, p_n), \tag{C9}$$

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

Log-differentiating (C9) 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 H.1 (supply $\varepsilon_S \equiv d \ln Y_n / d \ln \rho = \gamma_n \frac{k-1}{k} \eta$ and the income response (C6), with $d \ln p_a / d \ln \tau = -\kappa_a$ substituted throughout), the service-price and relative-price gradients solve out to

$$\frac{d \ln \rho}{d \ln \tau} = -\frac{\theta_c \left[\kappa_n(\sigma_n - 1) - \kappa_a \sigma_n\right]}{\varepsilon_S + \theta_c\, \sigma_n + \theta_l\, (s_a + \mathcal{E})}, \qquad \frac{d \ln p_n}{d \ln \tau} = \frac{d \ln \rho}{d \ln \tau} - \kappa_a. \tag{C10}$$

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

$$\kappa_n \frac{\sigma_n - 1}{\sigma_n} \; > \; \kappa_a : \tag{C11}$$

The service friction shrinks city demand at rate $\kappa_n(\sigma_n - 1)$ rather than $\kappa_n \sigma_n$, because shipping 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 6.5 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 4: villages closer to the city do more of their work in services and have lower agricultural productivity.

H.3 The Welfare Gradient and Its Decomposition

Substituting the three price gradients into (C7) decomposes the welfare gradient into four named channels:

$$\frac{d\ln V}{d\ln\tau} = \underbrace{(s_n-\alpha_n)\,\frac{d\ln p_n}{d\ln\tau}}_{\substack{\text{service terms of trade}\\ (<0:\text{ proximity gain})}} \underbrace{-\,(s_a-\alpha_a)\,\kappa_a}_{\substack{\text{farm-gate food price}\\ (<0:\text{ proximity gain})}} \underbrace{-\,\alpha_u\,\kappa_u}_{\substack{\text{delivered city goods}\\ (<0:\text{ proximity gain})}} \underbrace{-\,\mathcal{E}\,\frac{d\ln\rho}{d\ln\tau}}_{\substack{\text{externality drag}\\ (>0:\text{ proximity curse})}}\;. \tag{C12}$$

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:

$$\frac{d\ln V}{d\ln\tau} = \big[(s_n-\alpha_n)-\mathcal{E}\big]\,\frac{d\ln\rho}{d\ln\tau} \;-\; \alpha_u\,(\kappa_a+\kappa_u). \tag{C13}$$

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 (C10). 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.

H.4 Proof of Proposition 1

Under condition (C11) the relative-price gradient is strictly negative, so (C13) turns into a condition where real income increases with distance if and only if

$$\underbrace{\mathcal{E}\left|\frac{d\ln\rho}{d\ln\tau}\right|}_{\text{externality loss}} > \underbrace{(s_n-\alpha_n)\left|\frac{d\ln\rho}{d\ln\tau}\right|}_{\text{terms-of-trade gain}} + \underbrace{\alpha_u\,(\kappa_a+\kappa_u)}_{\text{market-access gain}}. \tag{C14}$$

Condition (C14) is a threshold in φ. Both $\mathcal{E}$ and the denominator of (C10) 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 (C14) 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. 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

$$\varphi^* = \frac{S\left[(1-\gamma_a) + A\,\dfrac{1-u}{u}\right] \;+\; b\,\gamma_n\,\dfrac{k-1}{k}\,\dfrac{1-u}{u}}{S + s_a\,(1-b\,\theta_l)}, \tag{C15}$$

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 (C15) reduces, using $s_a + s_n = 1$, to

$$\varphi^*\big|_{\kappa_a=\kappa_u=0} = \frac{(s_n-\alpha_n)(1-u)}{s_a\,u\,\eta} = \frac{s_n-\alpha_n}{1-\alpha_n}\left[(1-\gamma_a) + A\,\frac{1-u}{u}\right], \tag{C16}$$

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.

H.5 Numerical Confirmation

The analytical gradients are confirmed against full numerical solutions of the clearing condition (C9) 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 (C10) and (C13) 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 and measured frictions the welfare gradient crosses zero at $\varphi = 0.376$, and the closed form (C15), evaluated at that equilibrium's own shares, returns the same value; shutting off κ_a and κ_u reproduces the frictionless threshold (C16) (about 0.16).

I Measurement of Model Inputs

This appendix describes how I construct each measured input to the model of Section 6. I describe the measure of 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 extension of Section 9.3.

I.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. I develop a crosswalk in order to translate spending shares to value-added measures. In order to most closely calibrate rural households to match the characteristics of the transmigration villages, I use the 2002 SUSENAS consumption module and restrict my sample to rural households in the outer islands. I am able to divide 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 non-food intermediates such as fuel, transport, and utilities. Pricing those at the village-gate food price rather than at delivered city-good prices strengthens the negative effect of market access on income slightly, since food gets cheaper with distance while city goods get more expensive. Appendix K.9 bounds how much this matters. Repricing the measured non-food portion at delivered city-good prices attenuates the income gradient by under ten percent, and moving ℓ changes the results very little.

Naively assigning all service consumption that is 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 K 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 at a given period of time.

The model does use external estimates for the urban budget shares. The aggregate accounting identities imply them (Section 7.2). This makes them an overidentifying check against the survey. 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 I.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 three percentage points of the accounting-implied $\beta_a = 0.595$.

I.2 Trade Frictions

I attempt to measure the three sector-specific trade frictions in a variety of ways. I measure the food and rural service friction from a combination of an accounting exercise and the relationship between prices and travel times, while the city-good friction is constructed to be a combination of the other two frictions.

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 nine destinations.[30] I also observe local wages and the prices of 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 C7 reports both.

[30] These destinations are: the nearest bus terminal, market, public telephone, post office, health facility, and bank, and the kecamatan, kabupaten, and province capitals.

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

$$\text{cost share} = \underbrace{2t/8}_{\text{time}} + \underbrace{2f/w}_{\text{fare}}, \tag{C17}$$

with the workday set at eight 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 six 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 kabupaten capital consumes about forty 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 six-day stay leaves 38 percent. A doctor 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. Since there is no spatial price gradient to discipline κ_n, that agreement is what does it instead. I take $\kappa_n = 0.90$.

Two things bound how much weight the exact figure carries. The trip cost is a fixed fare plus a time cost linear in t rather than a power law, so the fitted exponent moves with the convention, running to between 1.6 and 2.1 on feasible day trips alone and returning 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. And 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 C7: 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

Notes: 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 six 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 C8).

The upper side comes from freight. The same round trip that delivers a day of service labor can instead carry a fifty-kilogram rice load, and the round-trip passenger fare over the load's value at the measured rice price is 4 percent to the kabupaten 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 kabupaten capital and 0.12 to the province capital. Both are upper bounds, for two reasons. The fifty-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 kabupaten gives gradients of -0.006 (s.e. 0.051) to the nearest market, -0.035 (s.e. 0.022) to the kabupaten 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 K.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 $\kappa_u = 0.82\times0.05+0.18\times0.90 = 0.203$ (Table C8). It ranges roughly 0.18 to 0.22 across the service-definition and network-utility scenarios.

Building it this way has three consequences. It uses only frictions that have running measurement code behind them. 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 carries no σ_u dependence. 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 0.21 is the asymmetry the mechanism runs on, and Appendix K.8 shows what breaks when either friction is set to the other's value.

Table C8: 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	$s_g\kappa_a + s_s\kappa_n$	$\kappa_u = 0.21$

Notes: The freight accounting prices the round-trip passenger fare against a fifty-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.

I.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 7.1.

I.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 shows up 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 twelve 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. 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.

I.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 ten 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 regency median imputed to the remainder, and the functionally urban locations of Section 7.1 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 J.4).

I.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 I.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, and pronouncedly 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) buckets 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, regency fixed effects, and survey weights. Table C9 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 bucket 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 C9: 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

Notes: 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 kabupaten 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.

I.7 The Engel Slopes and the Extended Demand System

This section describes the demand system used in the extension of Section 9.3. 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 kabupaten-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 C10 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 C10: 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	

Notes: 2002 SUSENAS rural households; desa fixed effects, survey weights, kabupaten-clustered standard errors. The value-added row applies the crosswalk transform of equation (C18).

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 I.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

$$b_n = f\, b_n^{\text{raw}}, \qquad b_a = b_a^{\text{raw}} + (1-f)\, b_n^{\text{raw}}, \qquad b_u = b_u^{\text{raw}}, \tag{C18}$$

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 I.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 three 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 kabupaten-clustered standard errors. Under the PIGL form this recovers ε as the negative of the food-share elasticity. 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 C11). 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 C11: 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)	

Notes: 2002 SUSENAS rural households, desa fixed effects, survey weights, kabupaten-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

$$V_\ell(p,e) = \frac{1}{\varepsilon}\left(\frac{e}{A_\ell(p)}\right)^{\varepsilon} - \sum_g \nu_g \ln p_{g,\ell}, \qquad A_\ell(p) = \prod_g p_{g,\ell}^{\omega_{g,\ell}}, \tag{C19}$$

by Roy's identity, with the composite CES price indices of Section 6.4 standing in for the service and city-good prices by two-stage budgeting. Two properties of this class are why I use it rather than some other 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). Both are established in Boppart (2014) and used this way by Fan et al. (2023) and Peters et al. (2026). The share equations are

$$s_{g,\ell} = \omega_{g,\ell} + \nu_g\, x_\ell^{-\varepsilon}, \qquad \textstyle\sum_g \omega_{g,\ell} = 1, \quad \sum_g \nu_g = 0, \tag{C20}$$

with x_ℓ real expenditure per capita. The logarithmic price term in (C19) is the $\sigma_c \to 1$ case of Boppart (2014)'s price aggregator, the case the price test of Appendix I.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 two-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 three to five 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 food-share elasticity rises with income, from 0.20 at the rural tenth percentile to 0.37 at the ninetieth, rather than staying constant as the power form implies. 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 counterfactual-sized income move and separate only at the program-removal shock of Section 9.7, where the predicted food-share response is $+0.17$ under the operative form, $+0.15$ under PIGLOG, and $+0.12$ under the unrestricted exponent, which is the band reported there. I carry the Boppart (2014) form because it is the standard in this literature, because its estimator delivers the cross-country comparison above, and because its curvature keeps the shares globally bounded.

Welfare. With (C19) 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 (C19) 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

$$\Delta_\ell^{EV} = \frac{1}{\varepsilon} \ln\Big[e^{\varepsilon\, \Delta x_\ell} - \varepsilon \sum_g \big(\bar{s}_{g,\ell} - \omega_{g,\ell}\big) \Delta \ln p_{g,\ell}\Big], \tag{C21}$$

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 (C21) 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 $\varepsilon\, \Delta x_\ell \sum_g (\bar{s}_{g,\ell} - \omega_{g,\ell}) \Delta \ln p_{g,\ell}$, which at the estimated ε and the price movements these counterfactuals generate is on the order of two thousandths of a percentage point, against reported effects of tenths of a percent.

J Computation

This appendix documents how the 60,466-location equilibrium is represented, solved, selected, and audited, and how the estimation wraps the solver.

J.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 $69{,}225 \times 60{,}466$ 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. A natural economy would be sparsity: truncate each buyer's sourcing at a travel-time cutoff and bound the discarded tail. That economy 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 eight hours one-way, population mass at distance grows faster than the $\tau^{-\kappa_n(\sigma_n-1)}$ kernel decays, leaving median tail shares above one percent and, for isolated buyers, near one hundred 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 one to two 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.

J.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 50–90 iterations, about eleven minutes.

Selection follows the protocol of Section 6.6: continuation in φ from the externality-free equilibrium, in thirteen 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 current calibrated point both return the identical equilibrium, maximum log-wage discrepancy 4×10^{-8} for the perturbation and 3×10^{-7} for the cold start, with zero of 60,466 locations differing beyond 10^{-6}. 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.

J.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 J.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. 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. (The two-parameter variant of Appendix K.3 adds an inner root that sets θ on the composition moment before the outer root closes the income moment.)

J.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 six gigabytes) and applied through single-precision chunked products; the whole solve holds in about 45 gigabytes and converges in 30–70 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, with the net-importing remainder (concentrated among the most transformed, best-connected locations, including a third of the program villages) mispriced by at most the arbitrage band $2\kappa_a \ln \tau$, under four percent where they cluster.

K Quantitative Robustness

This appendix collects the robustness of the quantitative results to every input that could reasonably be measured differently: the measure of the externality from the model vs the data, 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 (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.

K.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 5 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 C12 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: the estimator recovers the structural coefficient to within two to four percent in both worlds, and does not manufacture an externality where none exists.

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 5. 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 5.4, 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 C12: 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]		

Notes: 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 5 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.

K.2 The Size Interaction in the Model

Table C13 reports the model counterpart of the fully interacted specification of Table C2, alongside the data. Each transmigration village is solved in the per-village battery frame at its recorded settlement size, with the sorting scale held at its baseline value, and the model's real income per capita is regressed on centered log market access, centered log settlement size, and their interaction, with island and cohort fixed effects. The memo row reports the complementary experiment, the change in a village's real income from the ten-percent targeted road of Section 9.1, regressed on log settlement size within the same fixed effects.

Table C13: The fully interacted specification, data against model

	Data	Model	
		H at program mean	H at recorded allocation
ln(Market Access)	-0.0044	-0.0042	-0.0056
	(0.0017)	(0.0006)	(0.0022)
ln(Settlement Size)	0.0083	0.1122	0.2083
	(0.0208)	(0.0056)	(0.0206)
ln(Market Access) × ln(Settlement Size)	0.0083	0.0021	-0.0003
	(0.0026)	(0.0007)	(0.0027)
Memo: one-village road effect on ln(Settlement Size)		0.0020	
		(0.0013)	

Notes: The data column reports the fully interacted specification of Table C2, estimated on the 890 villages with recorded settlement sizes, with standard errors clustered at the village level. The model columns solve the 1,147 transmigration villages at their recorded settlement sizes at the calibrated $\varphi = 0.70$. The first model column holds land at the program mean, so that settlement size varies population alone. The second assigns each village its recorded land allocation. Market access and settlement size are in logs and centered in all columns. Model standard errors describe dispersion across the model cross-section rather than sampling uncertainty. The memo row reports the change in a village's real income from the ten-percent targeted road of Section 9.1, regressed on log settlement size within island and cohort fixed effects.

K.3 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 C14, 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 held throughout as the sourcing share climbs from 0.67 with θ. 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 five 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 Bryan and Morten's 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 C14: 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

Notes: 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.

K.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 C14: 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 K.3 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.

K.5 The Urban Classification

Section 7.1 classifies as urban the administratively urban locations plus every location inside a satellite-defined functional urban area. Table C15 reports the estimated point under alternative classifications.

Table C15: 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

Notes: Model slopes are computed at $(\theta, \sigma_n) = (12.5, 5.76)$. 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 $\kappa_u = 0.21$ across classifications. Slopes are reported in the estimation frame, and the baseline row's joint-equilibrium counterparts appear in Table 17. "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 thirty minutes of a city. Data comparisons are as in Table 17.

Every functional definition of urban, satellite-based, outcome-based, or widened by a commuting buffer, delivers the same slopes to within a few percent: nothing 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.

K.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.

K.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 C14). 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 the toy curse entirely (Appendix G), which is why κ_a is measured and its measurement double-checked against the null staple-price slope.

K.8 The Friction Asymmetry

Section 6.5 states that the gap between the person-embodied service friction ($\kappa_n = 0.90$) and the goods friction ($\kappa_u = 0.21$) is central to the results; Panel B of Table C14 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.

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

The Economic Census cost structures put the bundle labor share in a 0.42–0.54 band, 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 C14) moves the sourcing estimate as the crosswalked shares shift (s^*_{own} from 0.60 to 0.76) but leaves every reported moment, targeted and untargeted, 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 ten percent, the relevant cost-of-living weight being about 0.06.

K.10 Outer Land and the Extrapolated Externality

In the inverted-world architecture these two inputs are null directions at the baseline. The outer land floor of Appendix I.4 and the externality's extrapolation to non-program agriculture 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 Appendix K.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.

K.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 J.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 four percent, within the arbitrage band.

K.12 The Road Across Externality Values

Table C16 solves the whole economy across the externality, re-pinning the transmigration villages' common agricultural scale T_a by bisection at every φ so the baseline reproduces the observed program agricultural employment share throughout; the non-program productivity fields re-invert analytically, so every φ reproduces the same observed baseline at every non-program location. The ladder is solved under homothetic demand, which isolates the relative-access margin. Three patterns hold across the range.

Table C16: The road across the externality

φ	access gradient b_Y	uniform road (national)	program villages
0	+0.0070	+0.85	+0.95
0.1	+0.0063	+0.85	+0.92
0.2	+0.0055	+0.85	+0.90
0.3	+0.0044	+0.85	+0.87
0.4	+0.0031	+0.85	+0.83
0.70	−0.0056	+0.83	+0.64
0.74	−0.0080	+0.83	+0.65

Notes: Each row 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 access gradient b_Y is the model's real-income slope on the estimation measure. The uniform road column reports the ten-percent network cut of Section 9.4, in percent of national income, and the program villages column reports the same shock's effect on the transmigration villages.

First, the cross-sectional access gradient reverses at $\varphi \approx 0.5$: identical villages placed on the measured geography come out richer where access is higher below that value and poorer above it, and the calibrated $\varphi = 0.70$ sits well 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 Pareto case once the measured frictions are included (Appendix H), and 0.33 to 0.65 across the stress-test variants of the Fréchet simple model (Appendix G), all below the calibrated 0.70.

Second, the relative-access margin's aggregate is flat in φ: under homothetic demand the uniform ten-percent road raises national income by about eight tenths of a 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 9.3 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.

Third, what the externality moves is incidence: the program villages' share of the uniform gains falls monotonically with φ, from +0.95 percent at $\varphi = 0$ to +0.64 at the calibrated value, 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 joint degree in labor and land 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 C17: 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

Notes: Each row re-solves the ten-percent uniform road of Section 9.4 under homothetic demand. The baseline row reproduces the homothetic rows of Table 19 and the $\varphi = 0.70$ row of Table C16. 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 I.6 rejects. Results under non-homothetic demand are reported in Table 19.

K.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 ($\kappa_u = 0.21$, Appendix I.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 C17 reports the ten-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 (whose leg alone runs +0.59, +0.53, +0.48 percent as σ_u rises). No sign changes anywhere on the dial.

K.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 C17). At the rejected $\sigma_c = 2$ the ten-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 I.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 9.3; 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.

K.15 The Service-Price Slope

The spatial slope of village service prices (+0.006, estimated from a specification that strips the delivered-goods and composition content that inflates the broad non-food index to +0.01) is a held-out check in the general-equilibrium calibration (Section 7.3), 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 4 is suggestive 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.

L Additional Counterfactuals

This appendix holds two exercises that support Section 9 without belonging in it: a connectivity experiment that lowers the cost of distance itself rather than travel times, and the costing of the road network that Section 9.4 summarizes.

L.1 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 ten 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 as opposed to what asphalt does. Table C18 reports the results.

Table C18: 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%

Notes: Population-weighted changes in log real income, in percent, from cutting each friction parameter by ten 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 L.2. Cutting both 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 while enriching everyone else. The countryside and the cities gain 0.2 to 1.0 percent while the transmigration villages alone lose, 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 six percent: when distance stops shielding anyone, the curse shrinks by pulling the protected down rather than lifting the exposed up.

L.2 The Cost of the Road Network

The gains reported in Section 9.4 are annual flows, while roads are stocks that depreciate. Pricing the ten-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 ten-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 eight 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 ten to twenty percent; I read the ten-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 540,000 km × \$93,000 per overlay, spread over the eight-year cycle, is 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 (EINRIP: roughly

$850,000 per kilometer 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.