Stop That Train, I'm Starving
Access to Outside Labor Markets and Rural Living Standards

Mohamed Coulibaly*
HEC Montréal

February 10, 2026

Abstract

I exploit the abrupt closure of rural rail stations after Côte d'Ivoire's 2011 post-presidential-election crisis to study how losing access to outside labor markets affects rural living standards. Using nationally representative microdata, I find sizable declines in per-capita spending and consumption in rural areas exposed to rail closures relative to nearby controls. These declines coincide with a labor reallocation from higher-paying nonfarm employment into contributing-family work. I rationalize these patterns using a Roy model with heterogeneous nonfarm returns and sectoral mobility costs. The estimated nonfarm returns reveal strong heterogeneity, implying that similar transport investments can generate different welfare outcomes depending on which workers gain access to nonfarm jobs. Counterfactual simulations show that spatially targeted reopenings of a subset of rural rail stations are economically effective as they can achieve nearly the same welfare gains as full network restoration. These findings help inform infrastructure policy when budget constraints limit expansion.

Keywords: Rural transport infrastructure, Labor market access, Nonfarm employment, Sectoral mobility costs, Marginal treatment effects, Côte d'Ivoire
JEL Classification: O18, R42, O14, C21

*Department of Applied Economics, HEC Montréal, 3000 chemin de la Côte-Sainte-Catherine, Montréal, QC H3T 2A7, Canada. Email: mohamed.coulibaly@hec.ca.

1 Introduction

In low income rural economies, nonfarm employment opportunities are relatively scarce within villages but typically concentrated in nearby towns.[1] For rural workers, access to these towns often determines whether they remain in low productivity agriculture or transition into better-paying nonfarm work. Transport infrastructure that connects villages to nearby towns can therefore integrate rural communities into outside labor markets. Yet empirical evidence shows that such infrastructure-induced integration does not systematically translate into improvements in living standards, with sharply different welfare outcomes documented across settings (Asher and Novosad, 2020; Brooks and Donovan, 2020).[2]

This paper asks how reducing rural areas' access to outside labor markets affects living standards, and why similar forms of rural connectivity can generate different welfare outcomes. In particular, I study whether these welfare effects depend primarily on the scale of labor reallocation or the composition of workers who gain access to nonfarm jobs. To answer this question, I exploit the abrupt closure of rural rail stations in southern Côte d'Ivoire following the 2011 post-presidential-election crisis. Prior to the crisis, a publicly subsidized rural rail service connected small farming communities to nearby towns and allowed passengers to transport small quantities of merchandise, facilitating participation in nonfarm self-employment and petty commerce. After the crisis, this rural service was permanently discontinued, increasing the cost of reaching town-based employment for affected communities. The sudden and centralized nature of these closures provides a natural experiment to study how losing access to outside labor markets affects household living standards and labor allocation.

I estimate the causal effects of reducing rural areas' access to outside labor markets using a difference-in-differences design that exploits spatial and temporal variation in exposure to rural rail closures. The analysis combines nationally representative household survey data collected before and after 2011, with detailed information on spending, demographics, and labor outcomes. Rural subprefectures are linked to their nearest pre-closure rail station using geospatial coordinates, allowing me to compare outcomes in (treatment) areas that lost rail access after 2011 to those in nearby rural (control) areas that were never served by rail.

This paper delivers two central reduced-form results. First, I find that the closure of rural rail stations caused a 16 percent decline in household per-capita spending and consumption in treated areas relative to nearby control areas. Using pre-closure data, I show that per-capita

[1] These activities typically include petty commerce, retail trade, street vending, or small scale services that require regular access to nearby towns.

[2] Asher and Novosad (2020) show that rural road construction in India induces substantial labor reallocation out of agriculture without improving living standards, whereas Brooks and Donovan (2020) find large income gains from improved access to outside labor markets from footbridge constructions in rural Nicaragua.

spending and consumption evolved similarly in treatment and control areas prior to the closure, supporting the credibility of the estimated effects. Second, I find that the closure of rural rail stations induced a reallocation of labor away from nonfarm activities and into contributing family work, a less productive segment of the agricultural sector. Nonfarm employment contracted by roughly one quarter, while farm employment and other labor market margins, including labor force participation and hours worked, remained largely unchanged. Because nonfarm workers earn substantially more per hour than contributing family workers in the data, this reallocation provides a natural mechanism linking the loss of rail access to the observed decline in household living standards. These reduced-form results are not driven by selective migration: I find no evidence of differential population movements across treatment and control areas, consistent with high migration costs in low-income rural settings.[3]

Taken together, the reduced-form evidence shows that reducing rural areas' access to outside labor markets led to sizable losses in living standards and a reallocation of workers away from higher-paying nonfarm activities. However, these reduced-form results alone cannot distinguish whether the spending decline primarily reflects the number of workers who exit nonfarm employment or the types of workers whose access to nonfarm jobs is reduced. Distinguishing between these channels is essential for interpreting the welfare effects of rural transport connectivity and motivates the analysis that follows.

I rationalize the two central reduced-form results using a Roy model of sectoral choice in which workers choose between farm and nonfarm employment based on their potential returns and the cost of moving into nonfarm jobs. In the model, rail access lowers sectoral mobility costs rather than directly affecting productivity. As a result, workers with high potential nonfarm returns may remain in agriculture when mobility costs are high, while workers with lower returns may be more responsive to changes in sectoral mobility costs. The model provides a framework for linking changes in local access to nonfarm employment opportunities to aggregate welfare outcomes.

I estimate nonfarm returns using an instrumental variable (IV) strategy that exploits, as an instrument, the unexpected increase in distance to the nearest rail station following the closure of rural rail stations. The IV estimates imply high nonfarm returns for workers whose sector choice is affected by the instrument, consistent with high sectoral mobility costs. By construction, however, the IV estimates are local to compliers and do not characterize returns for all workers. I therefore extend the analysis using the marginal treatment effect framework to recover how nonfarm returns vary with unobserved costs to nonfarm employment (Heckman and Vytlacil, 2005). The estimated marginal nonfarm returns rise with these unobserved costs, revealing

[3] See Bryan et al. (2014) and Morten and Oliveira (2024) for evidence on migration costs in rural economies.

strong reverse selection: workers with the highest potential gains from nonfarm employment are also those least likely to work in the nonfarm sector. This pattern implies that improving access to nonfarm employment for workers with high unobserved costs and high potential returns can generate much larger welfare gains in terms of per-capita spending.

Finally, I use the estimated marginal nonfarm returns to simulate counterfactual rail reopening policies. Reopening all closed stations at their original locations raises average per-capita spending, but much of the resulting labor reallocation involves workers with relatively low nonfarm returns. I then consider spatially targeted reopenings. The key intuition is that different stations connect to different worker populations. Reopening a station that serves many constrained high-return workers generates large welfare gains, while reopening stations that primarily serve low-return workers yields much smaller benefits. My counterfactual results show that reopening only half the stations can generate nearly all (more than 95 percent) or very little (less than 40 percent) of the welfare gains from full network restoration, depending entirely on which workers regain access to nonfarm employment. These results show that the welfare effects of rural connectivity depend fundamentally on who gains access to nonfarm employment opportunities, not on the scale of infrastructure expansion alone.

This paper contributes to the literature on rural transport infrastructure and labor market integration by explaining why similar connectivity shocks can generate divergent welfare outcomes across settings. Recent evidence documents substantial labor reallocation following improvements in rural connectivity, yet reaches contrasting conclusions about welfare. While Asher and Novosad (2020) find that new rural roads in India induce large shifts out of agriculture with limited effects on living standards, Brooks and Donovan (2020) show that improved access to outside labor markets in rural Nicaragua generates sizable income gains. I show that these differences arise not from the scale of infrastructure investment or the extent of labor reallocation, but from the composition of beneficiaries. Because different transport nodes serve different worker populations, identical aggregate investments can yield very different welfare outcomes. Using spatially targeted counterfactual reopenings of rural rail stations, I show that reopening only a subset of stations can generate nearly all or very little of the welfare gains from full network restoration, depending entirely on which locations regain access. These findings are informative for infrastructure policy design in settings where fiscal and administrative constraints limit the scope for universal network expansion.

The paper also contributes to the literature on structural transformation by showing that sectoral reallocation out of agriculture does not systematically improve living standards.[4] Evi-

[4]Most work on structural transformation emphasizes rural–urban migration as a pathway out of low-productivity agriculture. Recent research highlights the role of mobility costs, rather than selection alone, in shaping sectoral allocation and productivity differences across agricultural and non-agricultural sectors (Bryan et

dence from rural road construction and urban-rural migration shows that lowering transport or migration costs can induce substantial reallocation out of agriculture without generating significant gains in living standards (Asher and Novosad, 2020; Lagakos et al., 2020). I study the reverse setting and document a case of adverse structural change: rural rail station closures push workers out of higher-return nonfarm activities and into lower-productivity contributing family work.[5] By allowing for heterogeneity in both nonfarm returns and sectoral mobility costs, my selection model clarifies why structural transformation sometimes fails to raise living standards. When workers can sort efficiently across sectors, labor reallocation has limited effects on living standards. By contrast, when mobility costs prevent high-return workers from entering nonfarm activities, reducing these constraints generates large gains. These results highlight the central role of worker composition in linking labor reallocation to changes in living standards.

The remainder of the paper proceeds as follows. Section 2 describes the institutional context and the natural experiment. Section 3 presents the data and empirical strategy. Section 4 reports the reduced-form results. Section 5 develops the model and presents the estimation results. Section 6 analyzes spatial targeting and counterfactual policies. Section 7 concludes.

2 Background and Natural Experiment

2.1 Context

Côte d'Ivoire's railway, linking the Abidjan seaport to landlocked Burkina Faso, was built during the colonial period primarily to export raw materials. After independence, the railway remained publicly managed before being privatized in the mid-1990s. Today, the line operates as a binational corridor connecting Abidjan to Ouagadougou, comprising a profitable freight system and a passenger service that historically included rural mixed trains. These mixed train services, carrying passengers and their merchandise, were historically used to serve low-traffic corridors where running separate lines was not economically viable (Burns, 2024).

Freight component as a corridor system. Since privatization, freight has been the financial backbone of Côte d'Ivoire's railway, contributing nearly 80 percent of operator revenue (Dagnogo et al., 2012). The system mainly supports international transit, carrying cement, fertilizer, and containers northward, and cotton, livestock, and agricultural goods southward, through a few major urban stations (Abidjan, Bouaké, Ferkessédougou, and Ouangolodougou).

al., 2014; Lagakos et al., 2020; Gai et al., 2025).

[5] See also McMillan et al. (2014) for direct evidence that labor reallocation into lower-productivity activities can generate growth-reducing (negative) structural change.

Freight traffic largely bypasses rural stations and plays little direct role in integrating small villages (Konan, 2021).

Passenger services and the role of mixed rail services. Historically, Côte d'Ivoire operated two passenger services: the Express, serving major cities, and the Omnibus, a slower mixed rail service critical to rural connectivity. The Express functioned like an intercity train with few urban stops, whereas the Omnibus, government-subsidized and primarily serving southern rural areas, stopped at nearly every station and carried both passengers and their small agricultural goods. It enabled small traders to travel along with their small quantities of products such as rice, bananas, okra, charcoal, and palm kernels from remote villages to urban markets (Dagnogo et al., 2012). This dual role made the Omnibus an important mode of transport linking rural areas to nearby towns and outside labor markets.

Throughout the paper, I refer to these mixed services as "rural rail services" to distinguish them from the more profitable freight and intercity lines.

2.2 The 2011 Crisis and Rail Station Closures

After two decades of strong post-independence growth, Côte d'Ivoire entered a period of economic decline and political instability marked by a 1999 military coup and a 2002 conflict that divided the country between a rebel-controlled north and a government-controlled south (Soumahoro, 2017). This de facto partition lasted until the disputed 2010 presidential election, which triggered a brief post-presidential-election crisis from December 2010 to April 2011, ended by an unexpected change in government. Violence during the crisis was highly localized: three regions, the economic capital Abidjan (49.6 percent of all deaths), the Cavally region (30.3 percent), and the city of Duékoué (11.2 percent), accounted for more than 90 percent of total fatalities (Léon and Dosso, 2020).[6]

The crisis led to a nationwide closure of rail operations in early 2011. Freight and intercity passenger services resumed later that year, but rural rail services along the southern corridor remained closed. Of the 28 passenger stations operating before the crisis, only 10 reopened, all located in urban centers. The remaining 18 stations, including 16 in southern rural regions, remained closed (Figure A.1).

These closures were abrupt, exogenous to local economic conditions, and geographically concentrated. They were not driven by declining demand but became possible following the unexpected change in government. The sudden and localized closure of rural rail stations re-

[6]Below, these three areas are excluded from the analysis sample.

duced access to rail-based passenger transport in previously connected rural areas.

The closure left a major gap in rural transport access. While road networks have expanded in urban areas, rural regions continue to face severe transport constraints. For instance, in Côte d'Ivoire, fewer than one in four rural residents live within two kilometers of an all-season road (Mikou et al., 2019). In these settings, rural rail services often represent the only reliable connectivity between rural areas and outside labor markets. Their closure can therefore significantly reduce rural areas' access to outside labor markets where alternative transport options are scarce.

Further details on the historical background of the 2011 post-presidential-election crisis are provided in the Online Appendix.

3 Data and Empirical Design

3.1 Data and Key Outcomes

Data sources. I use two main sources of data to analyze the effects of rural rail station closures. The first is a list of all 28 rail stations that were operational in Côte d'Ivoire before the 2011 post-presidential-election crisis. I match these stations with GPS coordinates extracted from OpenStreetMap (see Figure A.1).

The second source is three waves of the Enquête Niveau de Vie (ENV), Côte d'Ivoire's national household budget survey, conducted in 2002, 2008, and 2015. These surveys are representative at the national and subnational (region) level. They follow similar sampling and questionnaire design across years. In rural areas, the ENV covers 5,819 households in 2002, 6,000 in 2008, and 7,115 in 2015, representing 32,464, 29,998, and 26,488 individuals, respectively. Because the surveys are repeated cross-sections, I aggregate household and individual data at the subprefecture level for my reduced-form analysis in Section 4 below.[7] This allows me to construct a (pseudo) panel of subprefectures, unbalanced across years.

To measure geographic exposure to rail services, I geocode the survey locations and compute distances to rail stations using GPS coordinates. Then, I assign each household to a subprefecture, and compute the distance to all 28 rail stations using Vincenty (1975)' formula. I retain the minimum distance as the distance to the nearest station for each household, separately by year.

[7] Subprefectures are the lowest administrative level in Côte d'Ivoire during my study period.

Main outcomes. I focus on four sets of variables to measure household living standards, labor market outcomes, and basic demographics.

1. *Household per-capita spending.* I construct this measure from the raw survey data to ensure consistency across the three survey waves. It aggregates household expenditures on education, health, clothing, personal care, transportation, and housing. I then divide the total by household size to obtain per-capita spending.

2. *Household per-capita consumption.* This measure is provided by the national statistical office and is already included in the surveys. It adds to the spending measure the value of goods consumed from the household's own production.

3. *Demographics.* I track the age and gender structure of individuals and household heads to verify the composition of my sample over time.

4. *Labor outcomes.* I analyze both extensive and intensive margins of labor supply for the working-age population (15–64 years old). At the extensive margin, I use an indicator for whether an individual worked during the past 12 months. At the intensive margin, I use total hours worked per year. I also classify employment into three mutually exclusive sectors: nonfarm, farm, and contributing family work.

More details on the construction of these variables are provided in the Online Appendix.

3.2 Identification Strategy

Sample Selection I restrict the analysis to rural southern Côte d'Ivoire, where all rural rail stations under consideration operated before 2011 and where 16 of the 18 closures occurred after the 2011 post-presidential-election crisis. This geographic focus ensures that the analysis captures the main variation from the natural experiment while avoiding potential confounding from long-standing north–south differences.

I also exclude the three areas most directly affected by the 2011 post-presidential-election crisis: the economic capital city of Abidjan, the city of Duékoué, and the Cavally region. These zones accounted for over 90% of conflict-related casualties during the 2011 post-presidential-election crisis (Léon and Dosso, 2020). Excluding these areas helps isolate the effects of rail station closures from those driven by short-term violence. The final sample includes only rural subprefectures in southern Côte d'Ivoire, covering: 19,500 individuals (3,521 households) in 2002, 17,622 individuals (3,480 households) in 2008, and 13,198 individuals (3,947 households) in 2015.

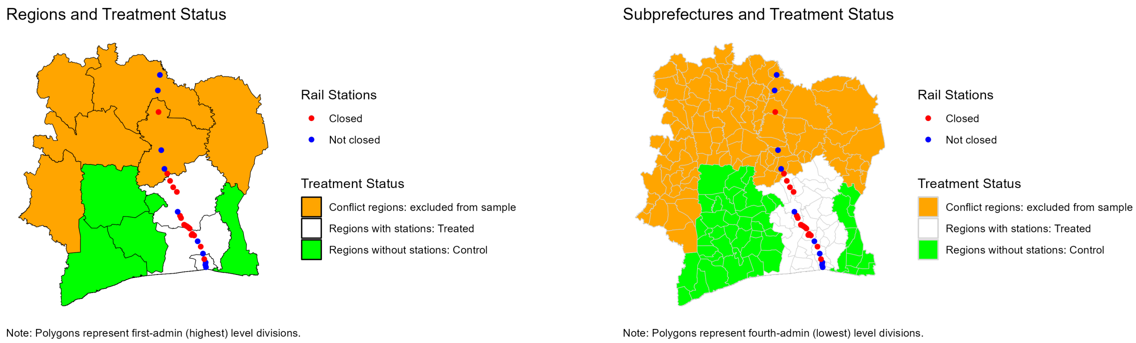

(a) Baseline regions

(b) Baseline subprefectures

Figure 1: Spatial Coverage of Treatment and Control Areas in Côte d'Ivoire

Note: These figures present maps of Côte d'Ivoire distinguishing treatment and control areas. The former are shown in white, while the latter are shown in green, both located in the southern regions. Orange-colored areas represent the northern regions, which are excluded from the study sample.

Figure 1a and Figure 1b illustrate the spatial coverage of the study. Treatment subprefectures (in white) correspond to rural areas in regions previously traversed by the rail line, while control subprefectures (in green) are adjacent rural areas in the southern regions that were never served by rail. Excluded areas (in orange) primarily correspond to northern regions.[8]

Research Design. To identify the causal impact of rural rail closures, I use a difference-in-differences (DiD) approach that combines spatial and temporal variations. The temporal variation arises from the 2011 post-presidential-electoral crisis, which caused the permanent closures of rural rail stations in rural southern Côte d'Ivoire. The spatial variation comes from whether a subprefecture was located in a southern region historically traversed by the rail line before 2011. The treatment group includes all rural subprefectures in southern regions in Côte d'Ivoire that were traversed by the rail line and had stations closed after 2011. The control group includes rural subprefectures in adjacent southern regions that were never served by rail.[9]

This empirical design exploits sharp, plausibly exogenous variation in rural areas' access to outside labor markets. The 2011 closures were abrupt, localized, and unrelated to pre-existing demand, as they resulted following the unexpected change in governemnt (see the Section 2).

Identification relies on parallel trends in outcomes between treatment and control subprefectures absent the closures. This is supported by the unexpected nature of the closure shock and

[8] The excluded areas also include the three most crisis-affected zones.

[9] The treated regions are: Abidjan (excluding the economic capital), Lagunes, and Lacs. The controls are: Bas-Sassandra, Comoé, Gôh-Djiboua, Sassandra-Marahoué, and Yamoussoukro.

by pre-crisis balance between the two groups, as shown in Section 4. Additionally, the causal interpretation of my results relies on the fact that my considered control group, not directly exposed, is plausibly unaffected by rail station closures.

The formal regression specification is:

$$Y_{S,t} = \beta_0 + \beta_1(post \times treat) + \lambda_S + \lambda_t + \epsilon_{S,t} \tag{3.1}$$

where $Y_{S,t}$ denotes the subprefecture-level mean of the outcome variables (e.g., log spending, sectoral choices) in subprefecture S and year t. $post$ equals 1 for 2015 (after the closures) and 0 for 2002–2008; $treat$ equals 1 for treatment subprefectures. λ_S and λ_t are subprefecture and year fixed effects, respectively. Standard errors are clustered at the subprefecture level (Bertrand et al., 2004), and all regressions use survey weights.

Because all treated subprefectures were exposed in the same year, the design is a standard DiD rather than a staggered adoption one, and the two-way fixed-effects estimator yields the conventional treatment-on-the-treated effect (Wooldridge, 2021).

To test for parallel pre-trends, I estimate the same equation using only pre-treatment years (2002 and 2008), with $post = 1$ for 2008.

3.3 Descriptive Statistics

Table A.1 summarizes key characteristics of the sample. Treated subprefectures tend to have slightly older household heads and a higher share of female-headed households throughout my study period. Access to credit is limited in my sample with fewer than 25% of household heads report having any form of credit in 2008 or 2015. Fewer than one in five household heads identify as migrants in 2002 and 2008, confirming limited migration in rural southern Côte d'Ivoire.

Both treatment and control groups show increasing (log) spending and consumption over time. However, growth was slower in the treatment group post-2011, where consumption and spending levels, initially higher before 2011, fell behind control areas by 2015. In terms of labor outcomes, employment rates remain broadly similar, but sectoral composition diverged. Before 2011, treated subprefectures had a higher share of nonfarm workers and fewer contributing family workers. After 2011, the opposite occurred: a visible shift back toward contributing family work in treated areas.

4 Reduced-form results

4.1 Effects on Household Living Standards

This subsection examines the main impacts of rural station closures on household per capita spending and consumption.

Table A.2 reports the baseline household-level pre-trend estimates (Panel A) and the DiD estimates (Panel B). Prior to the 2011 natural experiment of rural rail station closures, rural subprefectures served by rail followed similar trends as those never served. There is no evidence of systematic pre-closure differences. Pre-trend estimates for per capita spending and consumption are small and insignificant, supporting the parallel-trend assumption prior to the closures.

Following the closures, average household per capita spending declined by about 16 percent. Per capita consumption fell by a similar 15.7 percent in treated subprefectures relative to control areas. These are sizable and economically meaningful effects, indicating that rural rail station closures substantially reduced household living standards. The magnitude of these effects is substantial. They are comparable, though opposite in sign, to the income gains documented following improvements in rural connectivity from footbridge constructions in rural Nicaragua (Brooks and Donovan, 2020).

Robustness. I assess robustness using alternative spatial definitions of exposure.

Spatial heterogeneity. Because rural rail stations are highly localized, I first test whether the estimated effects are spatially concentrated near the closed lines. I divide the treatment group into two subgroups: those located closer than the median distance to the nearest closed station and those farther away. Both subgroups are compared against the same control group. This corresponds to the same regression specification as in Equation (3.1), except that two treatment variables are included instead of one. The results, reported in Table A.3, show that the negative effects of rural rail station closures are indeed locally concentrated in areas closest to the former rail stations, with negligible and statistically insignificant effects in more distant rural subprefectures. These findings suggest that the control group provides a valid counterfactual. They are located even farther from the rail line than the more distant treated subprefectures.

Redefining treatment by distance. The baseline treatment definition, whether a subprefecture was located in a region served by rural rail stations, might be imperfect. Some control subprefectures may lie closer to a closed station than some treated ones, and individuals can move freely across subprefectures. To address this, I redefine treatment status using distance to the nearest rail station. Specifically, I estimate effects using three distance thresholds: 80 km, 100 km, and 120 km. Households within these thresholds of a closed station are considered treated. These alternative definitions mitigate potential misclassification in the baseline setup. They also provide a more geographically precise measure of exposure to the closures. While the first test captures within-treatment spatial variation, the second redefines treatment boundaries altogether.

Accounting for treatment intensity. Exposure to rail services varies among treated households: those living closer to a station likely relied more on it than those farther away. If so, treatment intensity varies continuously with distance. To account for this, I follow Callaway et al. (2024). They show that even with heterogeneous treatment intensities ('doses'), one can identify average effects under parallel trends by dose. I apply their framework, defining treatment intensity by proximity to a station and using households beyond 120 km as the control group.

Figures A.2a and A.2b present these robustness results. Across all specifications, the pattern remains consistent: while treated and control groups evolve similarly before 2011, household spending and consumption decline significantly in treated areas after the closures.

4.2 Labor Market Responses

I also examine the effects of rural rail station closures on labor supply and sectoral allocation among working-age adults. Household spending and consumption depend partly on the labor activities of working-age members. Examining these outcomes sheds light on the potential mechanisms behind the main results.

Panel A of Table A.4 reports the pre-trend estimates between treatment and control subprefectures. Labor supply and sectoral choices among working-age adults evolved in parallel between 2002 and 2008, indicating no systematic differences prior to the closures. These negligible pre-trends are consistent with those observed for household consumption and spending in the pre-treatment period.

Panel B of Table A.4 presents the post-treatment estimates, showing that while overall labor supply appears unaffected, the composition of employment shifts significantly across sectors. Specifically, the closures led to a reallocation of labor from nonfarm to contributing family work, with virtually no effect on farm activities. Nonfarm employment declines by about nine percentage points, a magnitude comparable in size but opposite in sign to the reallocation of labor out of agriculture documented by Asher and Novosad (2020) following rural road construction in India.

The null effect on farm work likely reflects that farming depends on access to land. This is unlikely to be influenced by the presence or absence of a rural rail station. In contrast, rural rail services provide vital connectivity to towns and cities, facilitating nonfarm self-employment opportunities. When transport links disappear, the nonfarm sector contracts. Many displaced workers shift into contributing-family jobs, which are typically unpaid. Consistent with this, only 2 percent of nonfarm workers and 12 percent of farm workers receive no compensation. Among contributing-family workers, the figure rises to 42 percent. Nonfarm workers earn, on average, more than nine times the hourly earnings of contributing-family workers (Figure A.3). This reallocation pattern is consistent with an adverse structural change (McMillan et al., 2014). This shift from paid nonfarm to unpaid family work thus provides a plausible mechanism underlying the observed decline in household spending and consumption.

4.3 Other Outcomes: Population and Demographic Composition

I further examine whether the results might reflect selective migration or demographic shifts. If individuals moved out of treated subprefectures after the closures, or if household composition changed, the spending and consumption estimates could be biased.

To assess this possibility, I estimate DiD regressions using demographic and population-level outcomes: (i) the average age of household heads, (ii) the share of female-headed households, and (iii) total population size, using survey weights to approximate population counts. The surveys are probability-based, with household weights equal to the inverse of each household's selection probability. These weights allow the sample to represent the reference population, producing estimates that are representative at reasonably aggregated levels.

Table A.5 reports no significant post-2011 differences between treatment and control subprefectures along these dimensions. This suggests that the observed decline in spending is unlikely to be driven by selective migration or demographic changes. There is little evidence of population movement. This pattern aligns with prior research showing that high mobility costs in low-income rural settings constrain migration (Bryan et al., 2014; Morten and Oliveira, 2024; Gai et al., 2025).

4.4 Summary of reduced-form results

In summary, the reduced-form evidence indicates that rural rail station closures led to sizable declines in household living standards, with per capita spending and consumption falling by about 16 percent in affected areas. These effects are robust to pre-trend tests, alternative treatment definitions based on distance, and specifications accounting for treatment intensity. The analysis of demographic and population outcomes shows no evidence of selective migration or compositional change, reinforcing the causal interpretation of the results. Instead, the decline in spending coincided with adjustments in the labor market: the closures induced a reallocation of work from higher-paid, less agriculture-intensive nonfarm activities toward unpaid, more agriculture-intensive contributing-family work, consistent with an adverse structural change (McMillan et al., 2014). Taken together, these findings suggest that the spending decline was primarily driven by reduced access to better-paying nonfarm employment. Population shifts or aggregate labor contraction played little role.

At the same time, the reduced-form evidence does not reveal whether the magnitude of the spending decline reflects the number of workers who exited nonfarm employment or the types of workers whose access to nonfarm jobs was curtailed. This distinction matters because substantial labor reallocation need not translate into welfare gains, as shown by Asher and Novosad (2020), who document large shifts out of agriculture following rural road construction in India with limited effects on living standards. Distinguishing between these channels is therefore essential for interpreting the welfare effects of rural connectivity. The next section develops a model that allows me to quantify these margins and evaluate counterfactual reopening policies.

5 Model and Estimation

The reduced-form evidence showed sizable declines in per capita spending after rural rail station closures, alongside a shift from higher-paid nonfarm to lower-paid contributing-family work. A natural mechanism is that sectoral mobility costs limit access to nonfarm jobs and closures raise these costs.

I formalize this mechanism using a Roy model of sectoral choice with heterogeneous returns and mobility costs, which I use to recover the distribution of marginal nonfarm returns and evaluate welfare under counterfactual rail policies.

5.1 Roy Model with Sectoral Mobility Costs

Intuition: A Simple Roy Model of Costly Selection. To build intuition, consider a simple Roy model where workers choose between farm and nonfarm sectors by comparing relative nonfarm returns with the cost of accessing those jobs. Both returns and costs vary across individuals, so some who would earn more in nonfarm still do not enter because costs exceed gains.

Figure 2 illustrates this costly selection. Worker A, with modest returns but low costs, selects into nonfarm, while Worker B, with high returns but high costs, does not. This reverse selection pattern rationalizes the reduced-form reallocation away from nonfarm work after the closures.

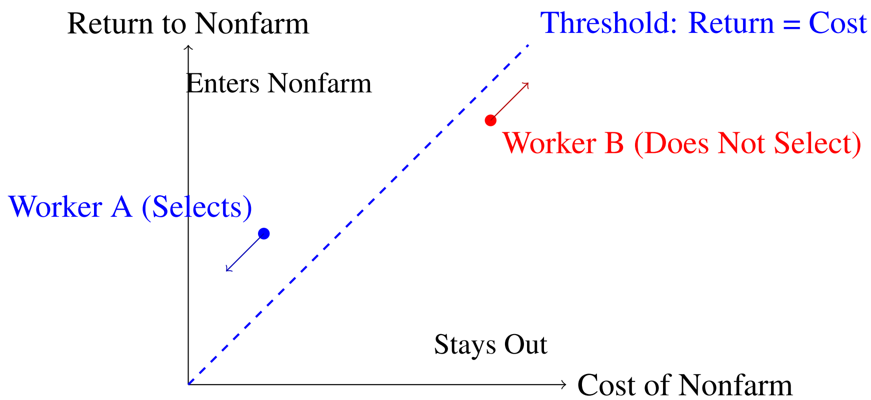

Figure 2: A Simple Roy Model of Costly Selection into Nonfarm

Generalized Roy Model: Selection Framework. I consider a two-sector Roy model for working-age employed adults who choose farm ($D = 0$) or nonfarm ($D = 1$).[10] Each sector yields a potential (log) spending outcome $Y_d(X, U_d)$, where $d \in \{0, 1\}$. Here, X is a vector of observed characteristics, and U_d represents unobserved individual heterogeneity in sector-specific productivity. Working in the nonfarm sector is contingent to an individual-specific cost $C(Z, U)$, where Z is an observed cost shifter and U an unobserved component of the cost.

Workers decide whether to work in the nonfarm sector by comparing their potential returns across sectors, net of costs. The selection rule is therefore:

$$D = 1\{\, Y_1(X, U_1) - Y_0(X, U_0) \geq C(Z, U) \,\}. \tag{5.1}$$

Under this rule, both the return to nonfarm work and the cost of sectoral mobility are heterogeneous across workers. Some workers may have high potential productivity in nonfarm activities but still choose to remain in the farm sector if their mobility costs exceed their expected gain. When $D = 1$, Y_1 is observed, and when $D = 0$, Y_0 is observed. The realized outcome is thus: $Y = D \cdot Y_1 + (1 - D) \cdot Y_0$.

I make the following standard assumptions (Heckman and Vytlacil, 2005):

Assumption 1

A.1 $Y_d(X, U_d, z) = Y_d(X, U_d)$ *for all* $z \in \mathrm{supp}(Z)$ *(exclusion restriction).*

A.2 $Z \mid X \perp (Y_1, Y_0, U)$ *(exogeneity of the cost shifter).*

A.3 *The distribution of* U *is absolutely continuous with respect to the Lebesgue measure.*

[10] I aggregate farm and contributing-family work into a single farm sector. I provide descriptive evidence supporting this grouping in the next subsection.

A.4 $\mathbb{E}[Y_d] < \infty, \quad d \in \{0, 1\}$ *(finite moments).*

Assumption **A.1** imposes that variation in the cost shifter affects outcomes only through the cost of working in the nonfarm sector, not directly through productivity. Assumption **A.2** ensures that the cost shifter is exogenous conditional on observables, justifying its use as an instrument. The remaining assumptions, **A.3** and **A.4**, are standard technical conditions ensuring well-defined parameters and continuous heterogeneity.

These assumptions justify using Z as a valid instrument and allow linking rail policy shocks to workers' sectoral choices in a structurally consistent way.[11]

5.2 Estimating Mean Nonfarm Returns: IV Evidence

I begin with a linear IV specification to estimate average returns to nonfarm work in the model using working-age employed adults in treated regions in 2015, consistent with the reduced-form sample. I adopt the following baseline specification:

$$Y = \alpha + \beta D + \varepsilon, \tag{5.2}$$

where Y denotes (log) per capita spending, and D is an indicator for working in the nonfarm sector. Here, $\alpha = \mathbb{E}[Y_0]$, $\beta = Y_1 - Y_0$ represents the individual return to nonfarm work, and $\varepsilon = Y_0 - \mathbb{E}[Y_0]$. Because sectoral choice is endogenous, workers with lower mobility costs are more likely to work in the nonfarm sector, OLS estimates of β are likely biased. An external source of variation in D is thus required to identify causal returns to nonfarm employment.

Instrument: Cost Shifter from Rural Rail Closures. I use as an instrument the closure-induced variation in the cost of accessing nonfarm employment:

$$Z = \Delta \log Dist_S \equiv \log(\text{Dist}_{S,\text{close}}) - \log(\text{Dist}_{S,\text{base}}),$$

where $Dist_S$ measures the distance from subprefecture S to its nearest rail station. This variable captures how much farther a subprefecture moved from rail access due to rural rail closures.

Identification of the IV estimand $\widehat{\beta}^{IV}$ relies on the standard conditions of *exogeneity*, *exclusion*, and *monotonicity* (Imbens and Angrist, 1994; Heckman et al., 2006), which follow from the structural model under Assumptions **A.1**-**A.4** (Vytlacil, 2002). Under these assumptions, the IV estimator recovers the mean return to nonfarm work for the subgroup of *compliers*,

[11] Assumptions **A.1**-**A.4** are common in the heterogeneous treatment-effect literature (Heckman and Vytlacil, 2005, 2007; Mogstad and Torgovitsky, 2018).

workers whose sectoral choice is affected by the instrument. The instrument plausibly satisfies these requirements. In my context, conditional on initial distance, variation in Z is driven by station closures rather than sorting or local shocks, and Section 4 shows no migration responses. Monotonicity is natural: higher distance should only lower nonfarm participation for everyone.

Results. Table A.6 reports summary statistics for the estimation sample. Farmers and contributing-family workers are far more likely than nonfarm workers to live in the same household. Their per capita spending levels are likewise much more similar. Both occupations are also heavily agricultural, with over 70% of employment in that sector (Figure A.3). These patterns support treating farm and contributing-family work as a single "farm" category in the two-sector model.

I assess whether the instrument behaves as expected. Regressing Z on observable workers' characteristics (Table A.7) shows no systematic differences, with coefficients globally insignificant,[12] confirming that Z is unrelated to pre-existing worker attributes. Placebo first-stage and reduced-form regressions in 2015 control regions, where rail lines never existed, likewise show no relationship between Z and outcomes (Table A.8). These results confirm the instrument's validity and that control regions in Section 4 were plausibly unaffected by rail closures.

IV Estimates. Table A.9 reports OLS, first-stage, reduced-form, and IV estimates, controlling for initial distance to the nearest rail station, demographics, and their interactions. OLS estimates show a positive association between nonfarm employment and household spending, likely reflecting selection bias. The first stage confirms that greater exposure to rail closures reduces nonfarm employment, consistent with the DiD evidence of reallocation away from nonfarm work. The reduced-form results likewise show lower per-capita spending in more affected areas.

The IV estimate of β identifies the mean nonfarm return for workers whose sectoral choice is affected by the closure-induced increase in distance. The IV estimate exceeds OLS, consistent with negative selection: workers who enter nonfarm at low mobility costs have lower gains than workers whose entry is induced by the instrument.[13]

5.3 Marginal Treatment Effects and Selection

Because returns to nonfarm employment are heterogeneous, the linear IV estimates in the previous subsection identify only an average effect for compliers, those whose sectoral choice responds to the instrument, rather than the full distribution of nonfarm returns. With a contin-

[12]Except "never attended school," marginally significant at the 10% level.

[13]This pattern mirrors Gai et al. (2025), who find similarly large IV returns to migration, pointing to significant mobility costs among compliers.

uous instrument, like in my setting, the linear IV estimand represents a weighted average of marginal treatment effects (MTEs) rather than a single, policy-invariant local average treatment effect (LATE) (Blandhol et al., 2022; Alvarez and Toneto, 2024; Heckman and Urzua, 2010). To recover this underlying heterogeneity and evaluate counterfactual rail policies, I adopt the MTE framework of Heckman and Vytlacil (2005), which models the distribution of marginal nonfarm returns and allows aggregation under alternative policy scenarios such as rural rail closures or reopenings.

Definition and Identification of the MTE. Following Vytlacil (2002), the selection rule in Equation 5.1 can be expressed as a single-crossing condition:[14]

$$D = 1\{\, U_D \leq p(X, Z) \,\},$$

where $p(X, Z)$ denotes the propensity score or probability of working in the nonfarm sector, and $U_D \sim \text{Unif}[0, 1]$ represents the unobserved cost (resistance) to nonfarm employment.

Under this formulation, the Marginal Treatment Effect (MTE) is defined as:

$$\text{MTE}(u) = \mathbb{E}\left[Y_1 - Y_0 \mid U_D = u\right],$$

which measures the mean nonfarm return for a marginal worker whose unobserved cost to nonfarm employment is $u \in [0, 1]$. Low values of u correspond to individuals who are more likely to work in the nonfarm sector, while higher values indicate those less likely to work in nonfarm. The MTE thus traces out the entire distribution of expected nonfarm returns across workers with varying unobserved costs.

The Marginal Treatment Effect (MTE) is identified via the local-IV estimand, that is, as the derivative of the observed outcome with respect to the propensity score (e.g., Heckman and Vytlacil, 2007). Identification of the MTE requires that the propensity score spans the unit interval and that the outcome varies smoothly with it. I estimate the MTE curve following Cornelissen et al. (2016) using a two-stage procedure.[15] In the first stage, I estimate the propensity score $p(X, Z)$ using a probit model. In the second stage, I estimate the outcome equation as a flexible function of the propensity score and compute the MTE as its derivative. For robustness, I report results under both a quadratic polynomial and a standard normal specification for the second

[14] For simplicity, I assume the cost function $C(Z, U)$ is separable in its observed and unobserved components: $C(Z, U) = C(Z) + U$. Rewriting the selection equation is possible under assumptions **A.1**-**A.4**.

[15] The approach of Cornelissen et al. (2016) assumes that potential outcomes are additively separable in observed and unobserved components conditional on the unobserved costs to treatment, implying a partially linear outcome equation where nonlinearity enters only through a flexible function of the propensity score. Further details are provided in Appendix B of that paper.

stage.

Results. Figure A.4 presents the distribution of estimated propensity scores from the first-stage probit regression. The estimates display sufficient variation, covering almost the entire $[0, 1]$ interval with support from 0 to 0.91. Consequently, fewer than 10% of MTE values require extrapolation beyond the observed support.

Figures A.5a and A.5b present the estimated MTE curves with 90% confidence intervals. The MTE is upward sloping in the unobserved cost, mirroring the pattern in the illustrative example of Figure 2. These findings are robust across both polynomial and normal specifications, which yield similar patterns and magnitudes of heterogeneity in nonfarm returns.

The upward slope of the MTE indicates that workers who would gain most from moving into nonfarm employment are the least likely to do so. This pattern, consistent with constrained comparative advantage or reverse selection on returns (Cornelissen et al., 2018), suggests that mobility costs might distort sectoral allocation: some individuals with high potential nonfarm returns remain in low-return farm activities because of these costs.

The evidence of reverse selection further suggests that mobility costs can inhibit structural transformation by preventing high-return workers from shifting into more productive nonfarm sectors. This interpretation aligns with the reduced-form evidence of adverse structural change following rail station closures, pushing workers away from more productive nonfarm employment. This empirical link motivates the analysis in the next section, which evaluates the welfare effects of reducing such costs through rail reopenings.

6 Spatial Targeting and Counterfactual Policies

The estimated MTEs imply large unrealized gains: the highest-return workers are least likely to enter nonfarm employment. I quantify the welfare gains from reversing the closure shock through counterfactual rail reopenings, and show how spatial targeting changes what the same budget achieves.

6.1 Benefits of Counterfactual Rail Reopenings

From Rail Policy to Costs of Sectoral Mobility. I map rail policy to the cost shifter via

$$Z = \log Dist_{S,\text{rail}} - \log Dist_{S,\text{base}},$$

where $Dist_{S,\text{base}}$ denotes the no-closure benchmark distance from subprefecture S to the nearest rail station, and $Dist_{S,\text{rail}}$ is the distance under a given policy scenario.

Under the counterfactual of full network restoration, $Dist_{S,\text{rail}} = Dist_{S,\text{base}}$, so that $Z = 0$ for everyone. Because closures primarily altered sectoral sorting rather than migration (Section 4), I hold residence fixed across counterfactuals.

In the model, distance shifts sectoral mobility costs and therefore the propensity to work in nonfarm. Consistent with this mechanism, nonfarm employment falls in areas exposed to rail closures (Section 4) and the IV first stage is negative (Section 5).

Policy Invariance and the PRTE Framework. To evaluate the average benefits of rail re-openings, I adopt the policy invariance framework of Heckman and Vytlacil (2005). The key identifying condition is that the joint distribution of (Y_1, Y_0, U_D) given X remains invariant across rail policies:[16]

Assumption 1 ***A.5*** *$(Y_1, Y_0, U_D) \mid X$ is invariant to rail policy changes.*

Assumption **A.5** implies that rail policy affects outcomes only through sectoral selection, not through changes in sector-specific productivity. Given this invariance, the estimated MTE curve can be used to simulate how average per capita spending would respond to alternative rail policies. Hence, the average policy benefit can be expressed as the *Policy-Relevant Treatment Effect* (PRTE; Heckman and Vytlacil 2001, 2005; Carneiro et al. 2010):[17]

$$\underbrace{\mathbb{E}[Y \mid \text{rail policy}] - \mathbb{E}[Y \mid \text{baseline}]}_{\text{Average Benefit}} = \underbrace{\int \text{MTE}(u)\, w(u)\, du}_{\text{PRTE}}, \tag{6.1}$$

where $w(u) = F_{P^*,X}(U_D \mid X) - F_{P,X}(U_D \mid X)$ measures the policy-induced change in the share of individuals entering nonfarm work at each level of unobserved cost u.[18] Intuitively, $w(u)$ shows which part of the unobserved-cost distribution the policy moves, so the PRTE aggregates gains for the workers who become marginal under that policy.

Estimated benefits from full network restoration. Figure A.6a shows the distribution of policy weights, $w(u)$, for the counterfactual re-opening of all rural rail stations. The policy weights are concentrated among low-cost individuals (low u) with lower expected nonfarm

[16] The distribution of X is also assumed invariant to policy changes.

[17] I use the unnormalized definition of the PRTE following Heckman and Vytlacil (2001).

[18] $F_{P^*,X}(U_D \mid X)$ and $F_{P,X}(U_D \mid X)$ are the cumulative distributions of the estimated propensity score under the counterfactual and baseline rail systems, respectively.

returns. Thus, workers with the highest potential nonfarm returns (high u) remain constrained even under full re-opening.

Table A.10 reports the PRTE estimates. Average per-capita spending rises by just over 5 percent. The gain is limited because the policy shifts entry mostly among low u workers with relatively low expected nonfarm returns, while many high u workers remain constrained by other frictions.

Benchmark (composition holding scale fixed). To benchmark the role of composition, I also compute a reweighted counterfactual that preserves the same aggregate increase in nonfarm participation but shifts nonfarm entry toward high-cost workers. This benchmark is not a literal rail policy; it isolates the role of composition holding the overall scale of reallocation fixed (Figure A.6b). This targeted counterfactual yields substantially larger benefits, raising average per capita spending by more than 9 percent.

These results reveal a broader mechanism: identical aggregate increases in nonfarm participation can produce sharply different benefits depending on which margin of selection the policy activates. For instance, when returns to nonfarm work and sectoral mobility costs are heterogeneous, policies that ease constraints for high-cost, high-return workers reallocate individuals with the greatest potential productivity gains, generating far larger benefits per worker moved. The findings are robust across both polynomial and normal specifications of the MTE, which yield similar magnitudes of per capita spending gains (Table A.10).

6.2 Spatially Targeted Reopenings

The counterfactual simulations above showed that fully restoring rural rail service would raise per-capita spending by about 5 percent by reducing sectoral mobility costs and enabling labor reallocation toward more productive nonfarm work. In practice, however, reopening every closed station might not be feasible. Budget and administrative constraints might limit the number of reopenings, shifting the policy question from *how much* to rebuild to *where* to rebuild. When resources are scarce, the planner must decide which locations to prioritize to maximize welfare.

Previous results in Section 5 indicated substantial heterogeneity in nonfarm gains across workers, with the largest benefits accruing to those least likely to work in nonfarm sectors. This heterogeneity implies scope for spatial targeting: by reopening stations in areas where constrained high-return workers are concentrated, welfare gains can be amplified. I operationalize this idea by examining how aggregate benefits varies across all feasible spatial allocations of a fixed number of reopenings. Specifically, I consider the case where only half of the closed rural stations can be reopened and ask which should be prioritized to maximize welfare.

Formally, the problem can be expressed as choosing the reopening allocation A^* that maximizes expected welfare:

$$A^* = \arg \max_{A \in \mathcal{A}_{8/16}} \mathbb{E}[Y \mid A],$$

where $\mathcal{A}_{8/16}$ denotes the 12,870 possible combinations of reopening eight of the sixteen closed stations. Under the maintained identification and policy-invariance assumptions, this Empirical Welfare Maximization (EWM) is equivalent to maximizing the policy-relevant treatment effect, PRTE(A), which summarizes the aggregate gains induced by each spatial configuration of reopenings. I compute PRTE(A) for each of the 12,870 feasible allocations and compare the implied welfare gains. Intuitively, the optimal allocation relaxes mobility constraints for the most constrained yet high-return workers, generating the largest welfare-improving reallocation from farm to nonfarm activities. A full derivation of the welfare–maximization framework in terms of PRTE and estimation algorithm is provided in the Online Appendix.

Results. The simulated welfare effects reveal large variation across alternative reopening allocations. Figure A.7 shows the distribution of estimated PRTE(A) values for all 12,870 feasible reopening allocations. Welfare gains differ markedly depending on where stations are reopened. The empirical welfare–maximizing allocation yields an average spending increase of 5.07 percent, nearly matching full restoration (5.28 percent), whereas the least efficient allocation produces only 1.98 percent. Reopening the same number of stations can therefore generate more than twice the welfare gain depending solely on their spatial allocation.

Figure 3 illustrates the spatial configurations of the optimal, worst, and a randomly chosen allocation. In the worst case, reopened stations cluster in the central corridor, leaving peripheral areas disconnected. By contrast, the optimal allocation distributes reopenings more evenly across space.

To understand the mechanism behind these welfare differences, Figure A.8 compares the policy weights $w^A(u)$ under the best and worst allocations. The efficient allocation reallocates a much larger share of constrained, high-return workers from farm to nonfarm sectors, generating stronger welfare effects. The share of workers in the upper half of the expected nonfarm return distribution who switch sectors is more than three times larger under the optimal allocation. This composition effect captures more efficient structural transformation: targeted infrastructure investments enable higher-return workers to move into more productive activities, raising aggregate welfare through improved sectoral allocation.

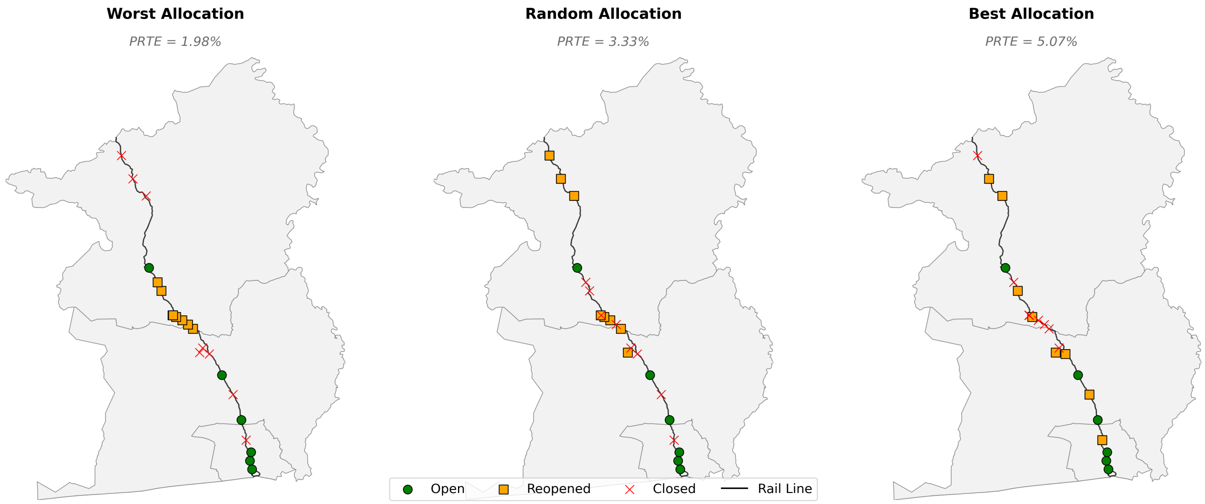

Figure 3: Worst, Random, and Optimal Station Reopening Configurations

Note: The map shows alternative reopening configurations of eight among the sixteen closed rural stations. The optimal (worst) allocation maximizes (minimizes) the estimated PRTE across all 12,870 possible combinations, consistent with the empirical welfare maximization framework in **??**. Green circles indicate stations that were never closed; yellow squares mark stations reopened under the given allocation; and red crosses denote stations that remain closed. Peripheral areas that remain disconnected under the worst allocation gain access under the optimal allocation, which distributes reopenings more evenly across space.

Net Gains and Benefit–Cost Ratios (BCR). I translate the PRTE into aggregate monetary gains and compare them to annual operating costs for rural service. Because stations are pre-existing, the relevant cost is operations rather than construction.[19] The BCR is defined as:

$$\text{BCR} = \frac{\text{Aggregate Benefit}}{\text{Total Operating Cost}}.$$

I evaluate four counterfactual scenarios already considered earlier: the worst, random, and optimal allocations of half the stations, and the full network restoration. The optimal half-allocation, generating more than 95% of the full-restoration benefit at roughly half the cost, achieves an almost twofold increase in the BCR compared to the full-restoration counterfactual (Table A.11).

Together, these results show that spatial targeting can substantially enhance the returns to rural transport infrastructure by directing resources toward locations where constrained, high-return workers are concentrated. By improving the match between high-return, constrained workers and access to outside labor markets, targeted reopenings magnify welfare gains even

[19] Further details about the data sources and cost assumptions for the BCR calculation are provided in Appendix C.

under tight budget constraints.

7 Conclusion

Access to nearby towns often determines whether workers in isolated rural areas can work in nonfarm employment. Yet the evidence on whether improving rural connectivity raises living standards is mixed.

This paper studies the reasons behind these contrasting findings. Using the abrupt closure of rural rail stations in southern Côte d'Ivoire after the 2011 post-presidential-election crisis as a natural experiment, I find large losses in living standards: household per-capita spending and consumption fell by about 16 percent in exposed rural areas relative to nearby controls. The closures also led to a reallocation of labor out of nonfarm work and into lower-paid contributing-family work, with little evidence of migration or changes in overall labor supply.

I rationalize these patterns using a Roy model in which workers face heterogeneous returns to nonfarm work and heterogeneous sectoral mobility costs. Estimating marginal nonfarm returns from the model reveals strong reverse selection: workers with the highest potential gains from nonfarm employment are also the least likely to work in that sector. This heterogeneity helps explain why similar transport infrastructure investments can generate different welfare outcomes across settings. Using spatially targeted counterfactual rail reopenings, I show that investments of similar scale can yield very different welfare gains, depending entirely on which workers regain access to nonfarm employment.

Although the analysis relies on data from Côte d'Ivoire, the mechanisms documented here extend to other contexts. Evidence from the model helps reconcile findings in the structural-transformation literature showing that movement out of agriculture does not necessarily translate into large welfare improvements (Asher and Novosad, 2020; Lagakos et al., 2020). My results highlight a key mechanism: welfare gains depend less on how many workers move and more on who is able to move. By shifting attention from the scale of labor reallocation to the composition of workers who reallocate, the paper offers a framework for interpreting when and why sectoral labor reallocation translates into meaningful welfare gains.

From a policy perspective, these findings are informative for the design of rural transport investments in settings where fiscal and administrative constraints limit the scope for universal network expansion. By highlighting the role of worker heterogeneity, the analysis suggests that targeting connectivity toward locations serving constrained, high-return workers can substantially enhance the welfare impact of infrastructure spending.

A Appendix Tables

Table A.1: Summary Statistics

	2002		2008		2015	
	Treatment	Control	Treatment	Control	Treatment	Control
Household head characteristics						
Age (years)	46.2	43.3	44.7	42.3	42.7	38.8
	15.990	*15.080*	*15.886*	*14.019*	*15.304*	*13.656*
Female (%)	20.6	11.4	20.0	9.0	24.0	13.9
	0.404	*0.318*	*0.400*	*0.286*	*0.427*	*0.346*
Credit access (%)	–		19.9	21.0	12.8	10.7
	–		*0.400*	*0.408*	*0.334*	*0.310*
Migrant (internal/external, %)	4.0	4.8	16.4	15.5	–	
	0.197	*0.213*	*0.371*	*0.362*	–	
Household outcomes (log)						
Consumption pc (XOF)	12.51	12.38	12.60	12.50	12.76	12.76
	0.684	*0.727*	*0.735*	*0.701*	*0.782*	*0.748*
Spending pc (XOF)	11.90	11.83	12.06	12.01	12.22	12.28
	0.759	*0.821*	*0.933*	*0.840*	*0.971*	*0.934*
Working-age adults outcomes						
Working (%)	71.9	71.9	73.9	71.2	65.9	64.1
	0.450	*0.450*	*0.439*	*0.453*	*0.474*	*0.480*
Non-farm workers (%)	19.6	16.0	27.6	22.6	27.0	30.6
	0.397	*0.366*	*0.447*	*0.418*	*0.444*	*0.461*
Farm workers (%)	45.8	43.0	51.4	45.1	44.0	41.0
	0.498	*0.495*	*0.500*	*0.498*	*0.497*	*0.492*
Contributing family workers (%)	34.3	40.6	18.4	29.7	18.8	18.9
	0.475	*0.491*	*0.387*	*0.457*	*0.391*	*0.392*
Sample size						
Households	980	2,561	1,000	2,480	1,534	2,406
Working-age adults (15–64)	2,862	7,289	2,488	6,976	2,761	4,562
Subprefecture	23	29	27	41	30	41
Minimum # of households	20	20	20	20	12	12
Average # of households	43	88	37	60	50	57

Notes: Means are reported with standard deviations in italics below each estimate. "Treatment" refers to rural households located in southern regions historically served by rail stations before 2011; "Control" refers to rural households in southern regions that were never served by rail. All values are survey-weighted. Monetary figures are expressed in XOF (CFA Franc). Credit access data were not collected in 2002, and migration data were not collected in 2015.

Table A.2: Effects of Rural Rail Station Closures on per capita Spending and Consumption

	(1) Spending pc (log)	(2) Consumption pc (log)
Panel A: Pre-treatment (2002 vs. 2008)		
c.treatment × year=2008	-0.017	0.022
	(0.1303)	(0.1113)
year=2008	0.157^{**}	0.068
	(0.0648)	(0.0646)
Constant	11.883^{***}	12.472^{***}
	(0.0306)	(0.0268)
Observations	88	88
R-squared	0.126	0.049
Number of subprefectures	44	44
Subprefecture FE	YES	YES
Treated average in 2008	12.05	12.62
Panel B: Full period (2002, 2008, 2015)		
c.treatment × year=2015	-0.161^{*}	-0.157^{*}
	(0.0894)	(0.0846)
year=2008	0.155^{***}	0.072
	(0.0564)	(0.0490)
year=2015	0.487^{***}	0.396^{***}
	(0.0619)	(0.0527)
Constant	11.867^{***}	12.455^{***}
	(0.0341)	(0.0275)
Observations	189	189
R-squared	0.371	0.335
Number of subprefectures	75	75
Subprefecture FE	YES	YES
Treated average in 2015	12.21	12.76

Notes: Each column reports estimates from the baseline Difference-in-Differences specification in Equation 3.1. The dependent variables are the natural logarithms of household per capita spending and per capita consumption, respectively. The variable *post* equals 1 for 2015 (after the rural rail station closures) and 0 for 2002–2008 for Panel B, and it equals 1 for 2008 and 0 for 2002 for the pre-trend panel A. *treat* equals 1 for subprefectures located in southern regions historically traversed by the rail line and affected by rural station closures (Abidjan–Lagunes–Lacs), and 0 for rural subprefectures in adjacent southern regions never served by rail (Bas-Sassandra, Comoé, Gôh-Djiboua, Sassandra-Marahoué, and Yamoussoukro). All regressions include subprefecture and survey-year fixed effects, and are weighted using household sampling weights. Standard errors, clustered at the subprefecture level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.3: Spatial Heterogeneity: Below vs. Above Median Distance to Nearest Rail Stations

	(1) Spending pc (log)	(2) Consumption pc (log)
Panel A: Pre-treatment (2002 vs. 2008)		
Below median × year=2008	-0.068	-0.062
	(0.2097)	(0.1701)
Above median × year=2008	0.039	0.115
	(0.1099)	(0.0942)
year=2008	0.157^{**}	0.068
	(0.0652)	(0.0649)
Constant	11.883^{***}	12.472^{***}
	(0.0306)	(0.0266)
Observations	88	88
R-squared	0.132	0.075
Number of subprefectures	44	44
Subprefecture FE	YES	YES
Below median treated average in 2008	12.11	12.69
Above median treated average in 2008	11.98	12.53
Panel B: Full period (2002, 2008, 2015)		
Below median × year=2015	-0.269^{***}	-0.275^{***}
	(0.0984)	(0.1035)
Above median × year=2015	-0.047	-0.033
	(0.1203)	(0.1028)
year=2008	0.156^{***}	0.073
	(0.0567)	(0.0492)
year=2015	0.487^{***}	0.397^{***}
	(0.0621)	(0.0529)
Constant	11.865^{***}	12.453^{***}
	(0.0336)	(0.0265)
Observations	189	189
R-squared	0.387	0.358
Number of subprefectures	75	75
Subprefecture FE	YES	YES
Below median treated average in 2015	12.19	12.72
Above median treated average in 2015	12.24	12.79

Notes: Each column reports estimates from an adapted version of the Difference-in-Differences specification in Equation 3.1, where two treatment indicators are included to distinguish treated subprefectures located below and above the median distance to the nearest rail station. The dependent variables are the natural logarithms of household per capita spending and per capita consumption, respectively. In Panel A (*Pre-treatment*), the variable *post* equals 1 for 2008 and 0 for 2002; in Panel B (*Full period*), *post* equals 1 for 2015 (after the rural rail station closures) and 0 for 2002–2008. "Below median" and "Above median" thus capture differential post-closure effects among treated subprefectures that were, respectively, closer to or farther from the nearest rail station prior to 2011. All regressions include subprefecture and survey-year fixed effects and are weighted using household sampling weights. Standard errors, clustered at the subprefecture level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.4: Potential Mechanisms: Labor Market Outcomes

	(1) Working	(2) Hours	(3) Earnings	(4) Nonfarm	(5) Farm	(6) Family work
Panel A: Pre-treatment (2002 vs. 2008)						
c.treatment × year=2008	0.065	0.013	0.295	0.032	0.033	-0.062
	(0.0580)	(0.0660)	(0.2848)	(0.0688)	(0.0767)	(0.0657)
year = 2008	-0.028	-0.112^{***}	0.612^{***}	0.040	0.040	-0.108^{**}
	(0.0387)	(0.0357)	(0.1813)	(0.0380)	(0.0461)	(0.0432)
Constant	0.722^{***}			0.208^{***}	0.442^{***}	0.344^{***}
	(0.0144)			(0.0164)	(0.0186)	(0.0163)
Observations	88	88	88	88	88	88
R-squared	0.029			0.065	0.051	0.299
Number of subprefectures	44	44	44	44	44	44
Subprefecture FE	YES	YES	YES	YES	YES	YES
Treated average in 2008	0.74	40.03	332.34	0.31	0.50	0.16
Panel B: Full period (2002, 2008, 2015)						
c.treatment × year=2015	-0.004	-0.036	-0.238	-0.094^{**}	0.006	0.099^{***}
	(0.0396)	(0.0396)	(0.2903)	(0.0409)	(0.0397)	(0.0347)
year = 2008	0.012	-0.126^{***}	0.766^{***}	0.055^{*}	0.042	-0.116^{***}
	(0.0282)	(0.0293)	(0.1395)	(0.0302)	(0.0338)	(0.0300)
year = 2015	-0.036	-0.200^{***}	1.411^{***}	0.156^{***}	-0.033	-0.217^{***}
	(0.0303)	(0.0268)	(0.2393)	(0.0335)	(0.0303)	(0.0297)
Constant	0.708^{***}			0.199^{***}	0.453^{***}	0.341^{***}
	(0.0197)			(0.0181)	(0.0192)	(0.0190)
Observations	189	185	185	189	189	189
R-squared	0.045			0.164	0.066	0.318
Number of subprefectures	75	71	71	75	75	75
Subprefecture FE	YES	YES	YES	YES	YES	YES
Treated average in 2015	0.68	35.09	430.26	0.29	0.43	0.18

Notes: Each column reports estimates from the baseline Difference-in-Differences specification in Equation 3.1, applied to labor market outcomes for the working-age population (15–64 years old). The dependent variables are: the share of individuals working (column 1), average weekly hours worked (column 2), average hourly earnings in local currency units (column 3), and the shares of workers employed in nonfarm, farm, and contributing-family work (columns 4–6, respectively). In Panel A (*Pre-treatment*), the variable *post* equals 1 for 2008 and 0 for 2002; in Panel B (*Full period*), *post* equals 1 for 2015 (after the rural rail station closures) and 0 for 2002–2008. *treat* equals 1 for subprefectures located in southern regions historically traversed by the rail line and affected by rural station closures (Abidjan–Lagunes–Lacs), and 0 for rural subprefectures in adjacent southern regions never served by rail (Bas-Sassandra, Comoé, Gôh-Djiboua, Sassandra-Marahoué, and Yamoussoukro). All regressions include subprefecture and survey-year fixed effects and are weighted using household sampling weights. Standard errors, clustered at the subprefecture level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.5: Population and Demographic Characteristics: Pre- and Post-Treatment

	(1) Population	(2) Population (log)	(3) Age	(4) Female
Panel A: Pre-treatment (2002 vs. 2008)				
c.treatment × year=2008	0.190	0.023	0.000	0.037**
	(0.2182)	(0.2462)	(0.0351)	(0.0154)
year = 2008	−0.225	0.129	−0.011	−0.026***
	(0.1686)	(0.1849)	(0.0213)	(0.0095)
Constant		11.166***		0.498***
		(0.0632)		(0.0038)
Observations	88	88	88	88
R-squared		0.028		0.158
Number of subprefectures	44	44	44	44
Subprefecture FE	YES	YES	YES	YES
Treated average in 2008	76,142.91	11.08	23.52	0.51
Panel B: Full period (2002, 2008, 2015)				
c.treatment × year=2015	0.099	0.081	−0.003	0.004
	(0.1716)	(0.1979)	(0.0272)	(0.0138)
year = 2008	−0.107	0.187	−0.007	−0.012
	(0.1323)	(0.1204)	(0.0167)	(0.0079)
year = 2015	−0.326*	−0.093	0.061***	−0.023**
	(0.1804)	(0.1758)	(0.0188)	(0.0095)
Constant		11.094***		0.497***
		(0.0812)		(0.0051)
Observations	185	189	185	189
R-squared		0.045		0.058
Number of subprefectures	71	75	71	75
Subprefecture FE	YES	YES	YES	YES
Treated average in 2015	65,950.01	10.87	25.31	0.49

Notes: Each column reports estimates from the baseline Difference-in-Differences specification in Equation 3.1, applied to demographic outcomes at the subprefecture level. The dependent variables are total population (proxied by the sum of survey weights), the natural logarithm of population, average individual age, and the share of females, respectively. In Panel A (*Pre-treatment*), the variable *post* equals 1 for 2008 and 0 for 2002; in Panel B (*Full period*), *post* equals 1 for 2015 (after the rural rail station closures) and 0 for 2002–2008. *treat* equals 1 for subprefectures located in southern regions historically traversed by the rail line and affected by rural station closures (Abidjan–Lagunes–Lacs), and 0 for rural subprefectures in adjacent southern regions never served by rail (Bas-Sassandra, Comoé, Gôh-Djiboua, Sassandra-Marahoué, and Yamoussoukro). All regressions include subprefecture and survey-year fixed effects and are weighted using household sampling weights. Standard errors, clustered at the subprefecture level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.6: Descriptive Statistics (Working-Age Adults, Treatment Group, 2015)

	Nonfarm		Farm		Contributing family work	
	Mean	SD	Mean	SD	Mean	SD
Assortative matching (share within column)						
Nonfarm			0.108	0.311	0.101	0.302
Farm	0.190	0.393			0.640	0.481
Family work	0.064	0.246	0.236	0.425		
Working-age characteristics						
Age	36.0	11.148	40.1	11.939	30.6	11.157
Female	0.46	0.499	0.24	0.427	0.63	0.483
log spending pc	12.5	0.946	12.1	0.886	11.9	0.776
Sample size						
N	424		832		370	

Notes: Means and Standard Deviations (SD) reported for working-age adults in the treatment group in 2015. Estimates are survey-weighted. "Assortative matching" are shares within each employment-type column.

Table A.7: Balance Test for Treatment Group in 2015: Correlation with Instrument

	(1) Balance test
Age of person	0.000
	(0.0007)
Individual is female	-0.005
	(0.0195)
Never been to school	0.064*
	(0.0339)
log(Household size)	0.026
	(0.0229)
Number of rooms	0.009
	(0.0146)
Dwelling is apartment	0.022
	(0.0628)
Constant	0.161**
	(0.0686)
Observations	1,804
R-squared	0.012

Notes: Each coefficient reports the estimated correlation between the instrument $Z = \Delta \log(\text{Dist}_S)$—the change in log distance from a subprefecture to its nearest rail station following the 2011 closures—and observable individual or household characteristics in the 2015 treated regions. Estimates are obtained from simple cross-sectional OLS regressions weighted by survey sampling weights. A lack of statistically significant coefficients indicates that the instrument is uncorrelated with pre-existing socioeconomic characteristics, supporting its exogeneity. Standard errors, clustered at the Primary Sampling Unit level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.8: Placebo Estimates for Control Group in 2015

	(1) First Stage	(2) Reduced Form
Instrument: log difference in distance	-0.117	1.657
	(0.7650)	(1.3725)
Observations	2,898	2,898
R-squared	0.040	0.230

Notes: Each column reports placebo estimates from the linear IV specification applied to individuals residing in control subprefectures in 2015; areas that were never served by rail stations. The dependent variable is log per capita spending, and the endogenous regressor is the choice of working in the nonfarm sector. The instrument is the survey-weighted mean log difference in distance to the nearest rail station before and after the 2011 closures, computed at the subprefecture level. All regressions include controls for initial (2008) log distance to the nearest rail station and its square, individual characteristics (age, gender, education), and household characteristics (log household size, number of rooms, apartment dwelling), as well as interactions of each control with initial log distance. Estimates are weighted using survey sampling weights. Standard errors, clustered at the Primary Sampling Unit level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.9: IV Estimates: Effect of Non-Farm Employment on Per Capita Spending in 2015

	(1) OLS	(2) First Stage	(3) Reduced Form	(4) IV
Non-farm employment share	0.319***			1.289**
	(0.0621)			(0.6197)
Instrument: log difference in distance		-0.145***	-0.187**	
		(0.0358)	(0.0853)	
Observations	1,804	1,804	1,804	1,804
R-squared	0.322	0.098	0.305	0.118

Notes: Each column reports estimates from the linear instrumental-variable specification described in Section 5, estimated for working-age adults residing in treated areas in 2015. The dependent variable is log per capita spending, and the endogenous regressor is the choice of working in the nonfarm sector. The instrument is the survey-weighted mean log difference in distance to the nearest rail station before and after the 2011 closures, computed at the subprefecture level. All regressions include controls for initial (2008) log distance to the nearest rail station and its square, individual characteristics (age, gender, education), and household characteristics (log household size, number of rooms, apartment dwelling), as well as interactions of each control with initial log distance. Estimates are weighted using survey sampling weights. Standard errors, clustered at the Primary Sampling Unit level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.10: Average Benefit of Reopening Rural Railway Stations

	Actual re-openings			**Counterfactual:reverse targeting**		
Model	**Estimate**	**90% LB**	**90% UB**	**Estimate**	**90% LB**	**90% UB**
Normal model	0.0527	0.0267	0.3033	0.1173	0.0119	0.4197
Polynomial model	0.0528	0.0283	0.3045	0.0945	0.0228	0.3978

Notes: This table presents PRTE estimates of re-opening all rural railway stations in treated areas in 2015. Confidence intervals are bootstrapped using a Bayesian bootstrap with 500 replications, clustered at the Primary Sampling Unit level. Estimates are based on the PRTE formula provided in Section 6.

Table A.11: Comparative Welfare Gains and Benefit–Cost Ratios Across Reopening Scenarios

	(1) Worst Allocation	(2) Random Allocation	(3) Optimal Allocation	(4) Full Reopening
Budget Share (relative to full)	0.5	0.5	0.5	1.0
Average Welfare Gain (%)	1.98	3.33	5.07	5.28
[90% CI]	[0.26; 2.34]	[0.13; 3.58]	[2.73; 30.25]	[2.83; 30.44]
Benefit–Cost Ratio (BCR)	1.14	1.91	2.90	1.51
[90% CI]	[0.15; 1.34]	[0.07; 2.05]	[1.56; 17.34]	[0.81; 8.73]
Relative Efficiency (vs. Full)	0.75	1.26	1.92	1.00

Note: This table compares welfare gains and benefit–cost ratios (BCRs) across four reopening scenarios: the worst, random, and optimal allocations of half the closed stations, and full network restoration. Each entry reports the average per-capita spending gain relative to the baseline, that is, the PRTE estimates using the polynomial specification. 90% confidence intervals, shown in brackets, are computed using Bayesian cluster-robust bootstrap standard errors with 500 repetitions. BCRs are defined as the ratio of welfare gain to total reopening operating cost, while Relative Efficiency measures each allocation's BCR relative to that of the full reopening scenario.

B Appendix Figures

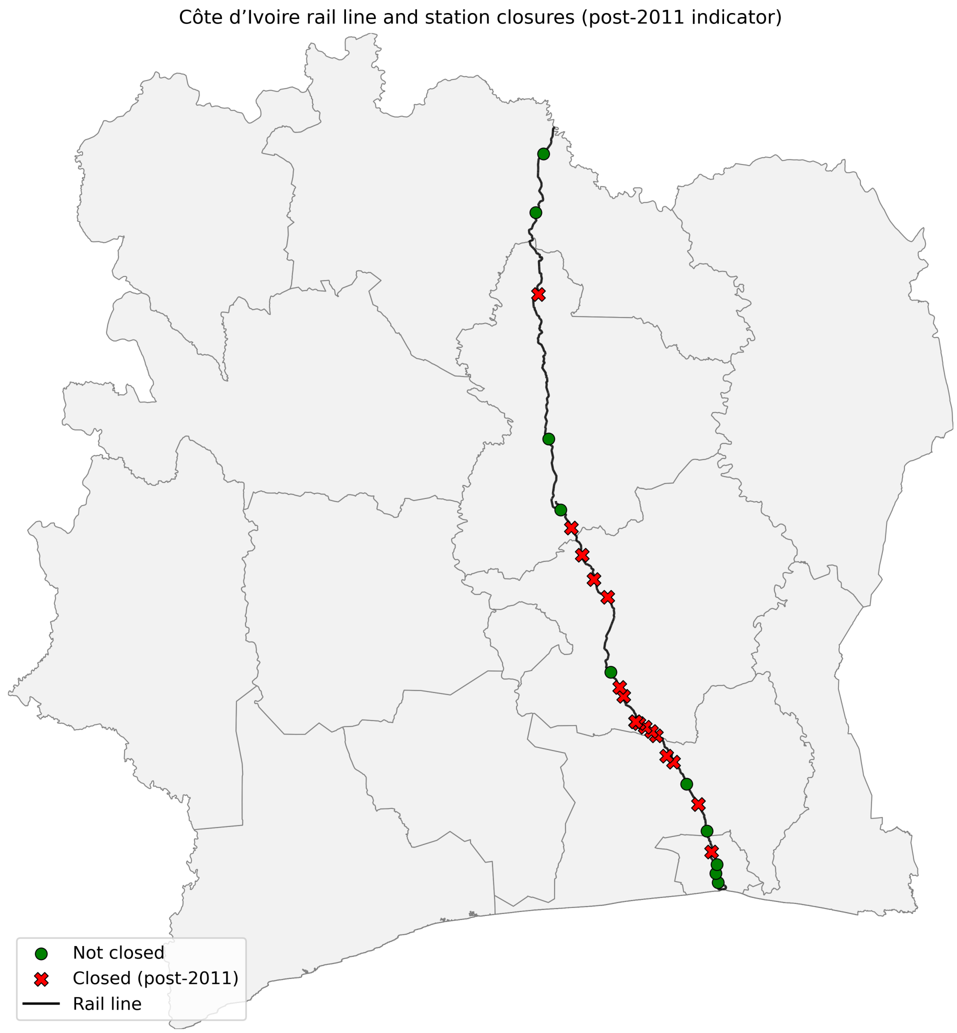

Figure A.1: Rural Rail Stations Before and After 2011

Note: These figures show the number of rail stations that were open before and after the post-presidential-election crisis. 18 of all existing rail stations were closed following the crisis, with 16 of these closures occurring in southern Côte d'Ivoire.
Source: Own computation (see Dagnogo et al. (2012) for more details)

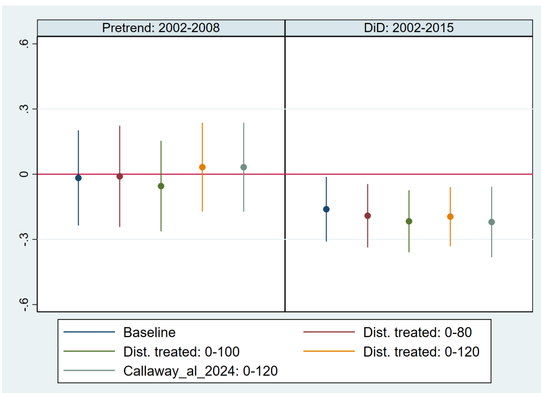

(a) Spending per capita

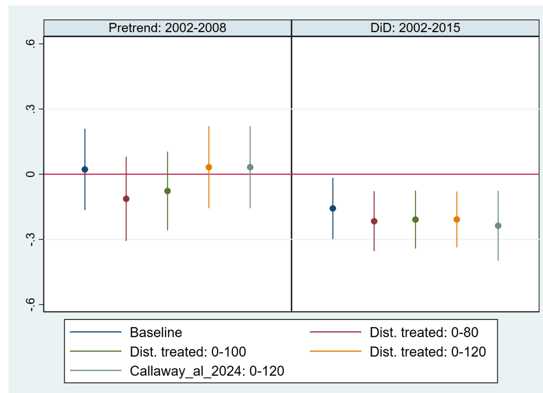

(b) Consumption per capita

Figure A.2: Robustness of DiD Estimates to Alternative Distance-Based Treatment Definitions and Continuous Treatment Intensity

Note: Each figure shows robustness checks for the baseline DiD estimates reported in Section 4. Estimates correspond to alternative definitions of exposure to rural rail station closures. The specifications labeled "Dist. treated: 0–80," "0–100," and "0–120" redefine treatment by distance to the nearest closed station, while "Callaway_al_2024: 0–120" implements the continuous-intensity estimator of Callaway et al. (2024), where exposure varies with proximity. Vertical bars represent 90% confidence intervals. Results are survey-weighted, and standard errors are clustered at the subprefecture level.

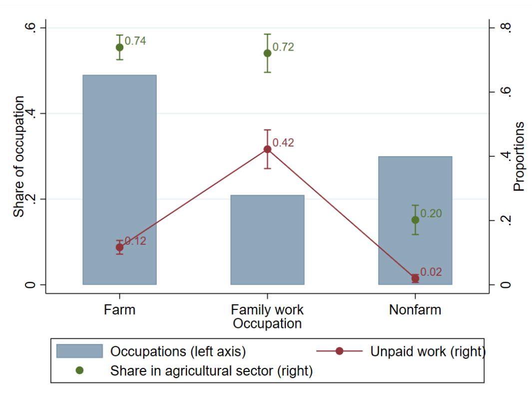

(a) Occupations and non-paying workers

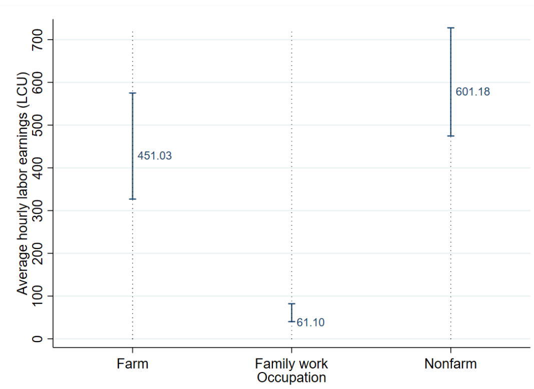

(b) Average hourly wage per occupation

Figure A.3: Stylized Facts: Nonfarm, Farm, and Contributing Family Work

Note: The left panel shows the share of workers across nonfarm, farm, and contributing-family occupations, together with the share of individuals receiving no payment within each group. The right panel reports average hourly earnings by occupation. All estimates are survey-weighted and based on working-age adults in treated regions in 2015.

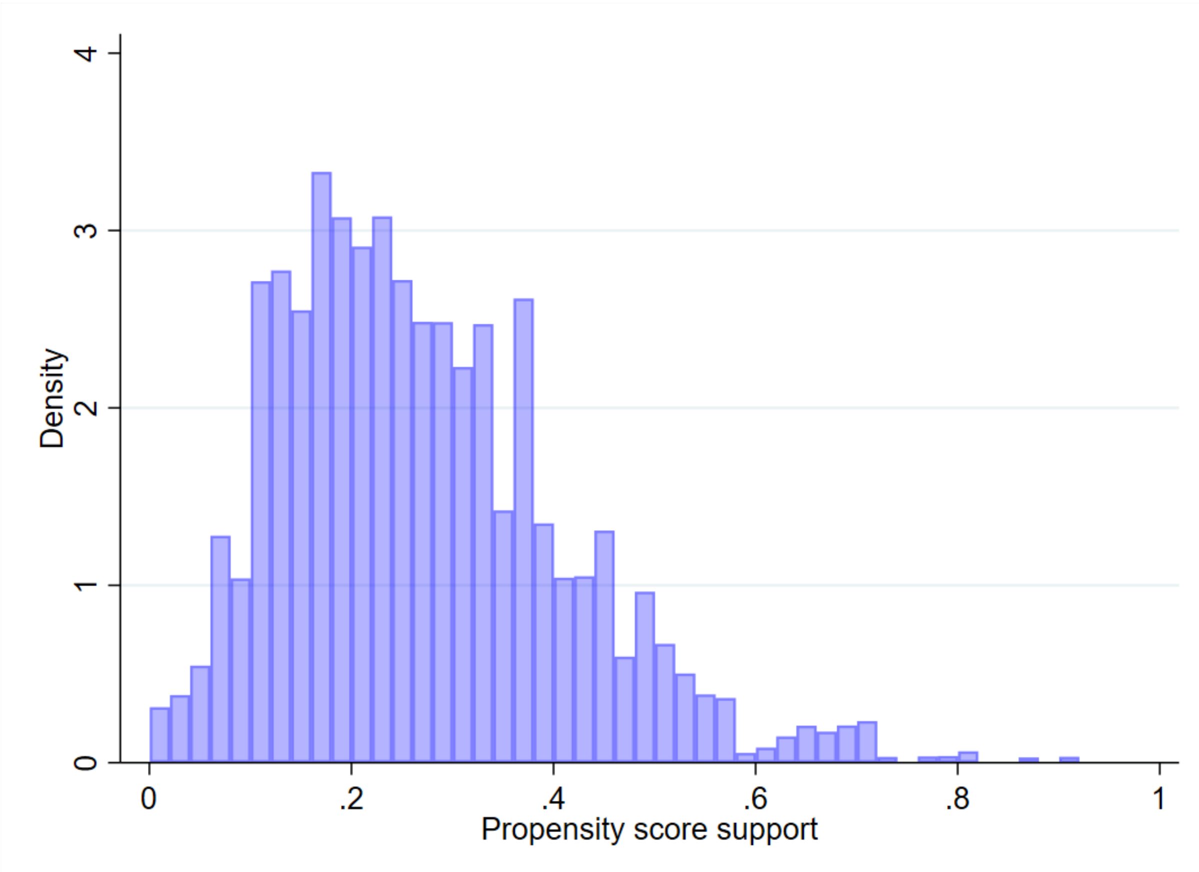

Figure A.4: Distribution of Estimated Propensity Scores

Note: The figure plots the distribution of the estimated propensity scores from the first-stage probit regression used in the MTE estimation. The sample includes all working-age adults with employment in treated regions in 2015. The estimated propensity scores exhibit substantial variation across individuals, with support ranging from 0 to 0.91, covering nearly the full $[0, 1]$ interval. Fewer than 10% of the MTE values therefore require extrapolation beyond the observed support.

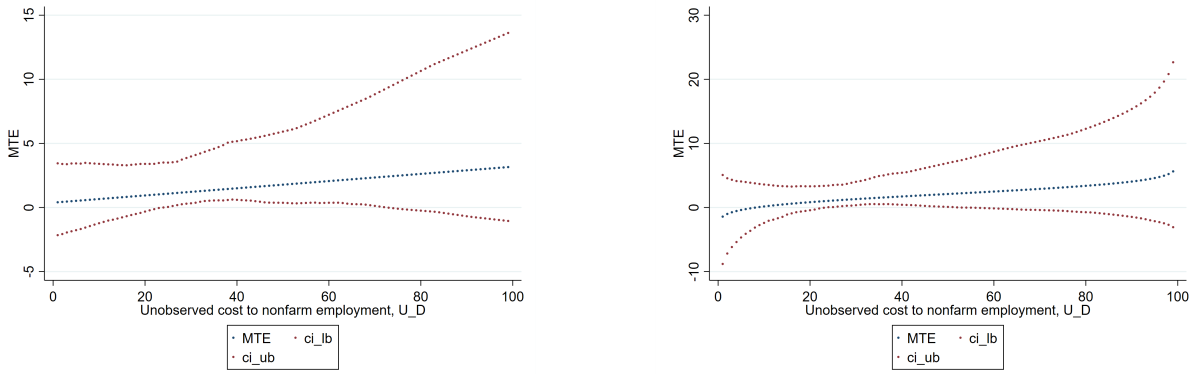

(a) MTE curve, polynomial model

(b) MTE curve, normal model

Figure A.5: MTE Curves for (Log) Per Capita Spending

Note: The plots depict the MTE curves for (log) per capita spending and evaluated at mean values of the covariates. The MTE is estimated based on the local IV approach (Heckman and Vytlacil, 2007), as described in Section 5. The 95 percent confidence interval is based on Bayesian cluster-robust bootstrapped standard errors with 500 repetitions. Standard errors are clustered at the Primary Sampling Unit of the survey data.

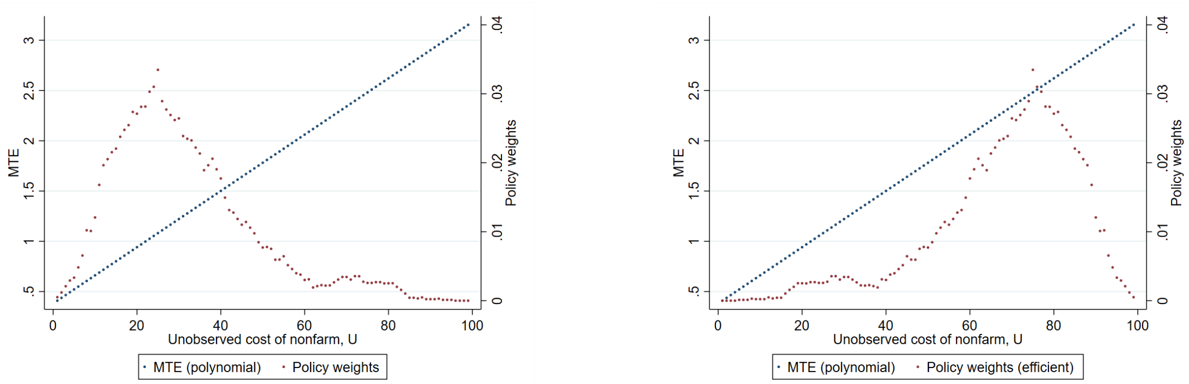

(a) Actual counterfactual of re-opening

(b) Counterfactual with reverse targeting

Figure A.6: Policy Weights, $w(u)$, from Reopening Rural Rail Stations

Note: The plots depict the distribution of the policy weights, $w(u)$, associated with reopening all rural rail stations that were previously closed. The policy weights are estimated as the difference in the cumulative distribution function (CDF) of the propensity score, evaluated when $Z = 0$ and when Z equals its original value. The left-hand figure shows the actual policy weights, while the right-hand figure shows the reverse policy weights.

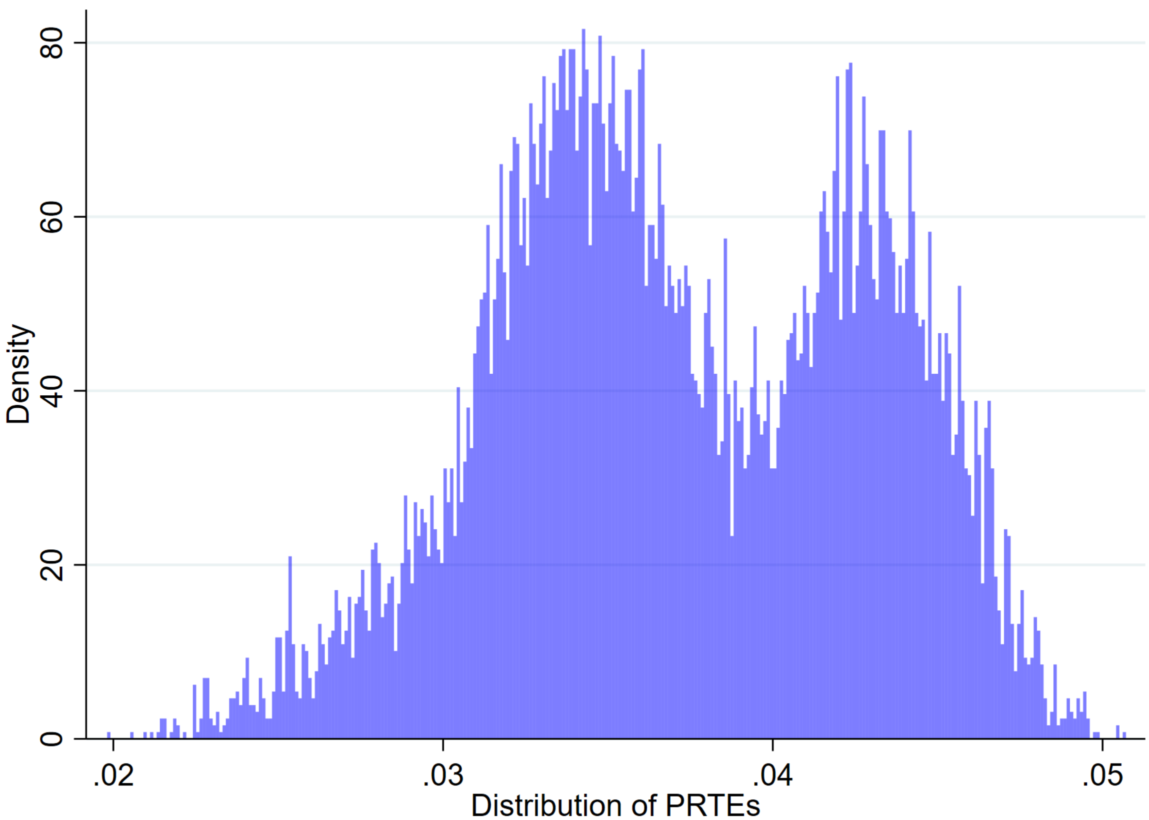

Figure A.7: Distribution of Estimated PRTE Across All 12,870 Reopening Allocations

Note: The figure plots the distribution of simulated Policy-Relevant Treatment Effects (PRTEs) for all 12,870 possible reopening allocations of 8 out of 16 closed rural rail stations. Each PRTE measures the expected per capita spending gain relative to the closure baseline. The dispersion across allocations reflects the variation in welfare gains arising from different spatial configurations of station reopenings.

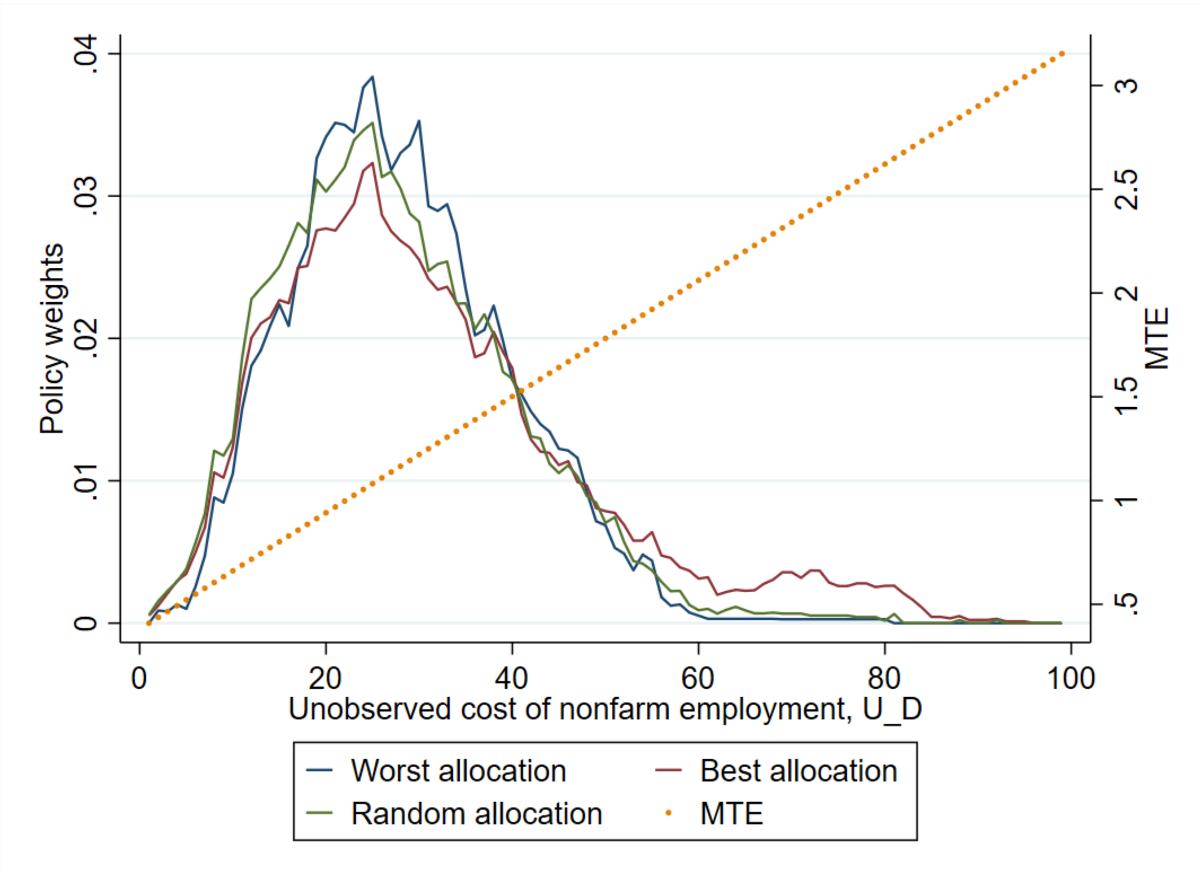

Figure A.8: Policy Weights, $w^A(u)$, for Worst and Optimal Reopening Allocations

Note: The figure compares the distribution of policy weights, $w^A(u)$, across the unobserved cost distribution for the optimal and worst reopening allocations shown in Figure 3. Each $w^A(u)$ measures the share of workers at a given unobserved mobility cost u who are induced to enter nonfarm employment under reopening allocation A. Higher values of u correspond to workers with higher mobility costs and higher potential nonfarm returns. The optimal allocation disproportionately reconnects these more constrained high-return workers, while the worst allocation primarily benefits workers with lower costs and lower expected returns. This pattern explains why spatial targeting yields large welfare heterogeneity across reopening configurations.

C Cost Benefit Analysis Data

C.1 Aggregate Benefits

Aggregate benefits are computed as the product of the average simulated welfare gain in per capita spending and the size of the affected rural population. The monetary gain in per capita spending is obtained by multiplying the estimated Policy-Relevant Treatment Effect (PRTE) by the average baseline level of per capita spending in the treated areas in 2015. This measure captures the total increase in per capita spending that would result from reopening rural rail stations.

To estimate the population exposed to the policy, I draw on population data from the 2014 General Census of Population in Côte d'Ivoire for the three affected regions. I then scale these totals by the rural population share, as estimated from the nationally representative household survey data, to obtain the rural population in 2014. Finally, I project this figure to 2015 using official population growth rates, yielding an estimate of the affected rural population during the 2015.

C.2 Operating Costs

Because the policy counterfactual concerns reopening previously built but inactive stations, the relevant costs are operating rather than construction costs. I recover operating costs using variation in total rail subsidies before and after the closures, normalized by the number of stations. Direct data on rail subsidies for Côte d'Ivoire are not publicly available; however, I use data from Burkina Faso as a proxy for three reasons.

First, the rail line under study spans both countries and is operated by the same concessionaire under a joint bilateral agreement. Second, both countries experienced station closures during the same 2009–2011 period. Third, detailed annual data on rail operating subsidies are available for Burkina Faso surrounding the closure years. Using these data, I estimate the average annual operating cost per rural station, which serves as the benchmark cost for Côte d'Ivoire (see Figure A.9).

Figure A.9: Government Operating Subsidies to Rural Rail Stations

- Operating costs are estimated from Burkina Faso and used as a proxy for Côte d'Ivoire:
  - The considered rail line spans both countries and is operated by the **same concessionaire** under joint agreements.
  - Both countries experienced rural station closures during the same period (2009–2011).
  - Detailed rail operating subsidy data are available for Burkina Faso (BFA) and provide a credible proxy for Côte d'Ivoire (CIV).

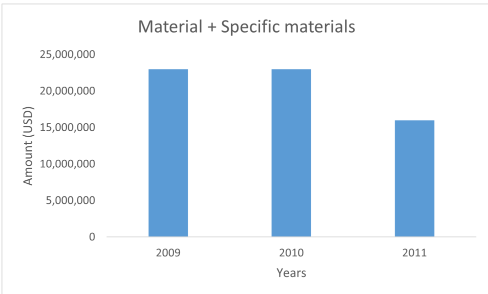

- Between 2009–2011, **four rural stations** were closed in Burkina Faso.
- ⇒ **Estimated operating subsidy per rural station: 1.75 million USD/year.**

Notes: The figure reports annual government operating subsidies to the rail network in Burkina Faso, obtained from the World Bank BOOST database. These data are used to proxy the operating costs of rural railway stations in Côte d'Ivoire, as both countries share the same concessionaire and experienced parallel station closures during 2009–2011.

Acknowledgments

I express my gratitude to Deico Coviello, Fabian Lange, and Ismael Mourifié for their guidance and continued support. I am also thankful to Jean-François Gauthier, Matteo Bobba, and Gilles Duranton for their insightful suggestions. I thank Souleymane Soumahoro, Sidi Mohamed Sawadogo, Fansa Koné, and N'dri Guillaume for their helpful comments. I further thank various seminar and conference participants for helpful comments. I acknowledge financial support from the Fonds de recherche du Québec—Société et culture (FRQSC) and the Programme de bourses d'excellence pour étudiants étrangers (PBEEE).

References

Alvarez, Luis AF and Rodrigo Toneto, "The interpretation of 2SLS with a continuous instrument: A weighted LATE representation," *Economics Letters*, 2024, *237*, 111658.

Asher, Sam and Paul Novosad, "Rural roads and local economic development," *American economic review*, 2020, *110* (3), 797–823.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How much should we trust differences-in-differences estimates?," *The Quarterly journal of economics*, 2004, *119* (1), 249–275.

Blandhol, Christine, John Bonney, Magne Mogstad, and Alexander Torgovitsky, "When is TSLS actually late?," Technical Report, National Bureau of Economic Research Cambridge, MA 2022.

Brooks, Wyatt and Kevin Donovan, "Eliminating uncertainty in market access: The impact of new bridges in rural Nicaragua," *Econometrica*, 2020, *88* (5), 1965–1997.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak, "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh," *Econometrica*, 2014, *82* (5), 1671–1748.

Burns, Adam, "The Mixed Train: A Unique Aspect of Railroading," https://www.american-rails.com/mixed.html 2024. Accessed: 2024-11-30.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna, "Difference-in-differences with a continuous treatment," Technical Report, National Bureau of Economic Research 2024.

Carneiro, Pedro, James J Heckman, and Edward Vytlacil, "Evaluating marginal policy changes and the average effect of treatment for individuals at the margin," *Econometrica*, 2010, *78* (1), 377–394.

Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg, "From LATE to MTE: Alternative methods for the evaluation of policy interventions," *Labour Economics*, 2016, *41*, 47–60.

_ , _ , _ , and _ , "Who benefits from universal child care? Estimating marginal returns to early child care attendance," *Journal of Political Economy*, 2018, *126* (6), 2356–2409.

Dagnogo, Foussata, Olivier Ninot, and Jean-Louis Chaléard, "Le chemin de fer Abidjan-Niger: la vocation d'une infrastructure en question," *EchoGéo*, 2012, (20).

Gai, Qingen, Naijia Guo, Bingjing Li, Qinghua Shi, and Xiaodong Zhu, "Rural Pensions, Labor Reallocation, and Aggregate Income: An Empirical and Quantitative Analysis of China," *Econometrica*, 2025, *93* (5), 1663–1696.

Heckman, James J and Edward J Vytlacil, "Econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments," *Handbook of econometrics*, 2007, *6*, 4875–5143.

_ and Edward Vytlacil, "Policy-relevant treatment effects," *American Economic Review*, 2001, *91* (2), 107–111.

_ and _ , "Structural equations, treatment effects, and econometric policy evaluation 1," *Econometrica*, 2005, *73* (3), 669–738.

_ and Sergio Urzua, "Comparing IV with structural models: What simple IV can and cannot identify," *Journal of Econometrics*, 2010, *156* (1), 27–37.

_ , _ , and Edward Vytlacil, "Understanding instrumental variables in models with essential heterogeneity," *The review of economics and statistics*, 2006, *88* (3), 389–432.

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, *62* (2), 467–475.

Konan, Suzanne Aya, "LA GARE SITARAIL, PÔLE DE DISTRIBUTION DES PRODUITS DÉRIVÉS DU MANIOC VERS LE BURKINA-FASO THE SITARAIL, DISTIBUTION

CENTER FOR CASSAVA-DERIVED PRODUCTS TO BURKINA FASO," *Journal de Géographie Rurale Appliquée et Développement*, 2021, *1* (03), 79–90.

Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E Waugh, "Migration costs and observational returns to migration in the developing world," *Journal of Monetary Economics*, 2020, *113*, 138–154.

Léon, Florian and Ibrahima Dosso, "Civil conflict and firm recovery: Evidence from post-electoral crisis in Côte d'Ivoire," Technical Report, FERDI Working Paper 2020.

McMillan, Margaret, Dani Rodrik, and Íñigo Verduzco-Gallo, "Globalization, structural change, and productivity growth, with an update on Africa," *World development*, 2014, *63*, 11–32.

Mikou, Mehdi, Julie Rozenberg, Elco Eduard Koks, Charles James Edward Fox, and Tatiana Peralta Quiros, "Assessing rural accessibility and rural roads investment needs using open source data," *World Bank Policy Research Working Paper*, 2019, (8746).

Mogstad, Magne and Alexander Torgovitsky, "Identification and extrapolation of causal effects with instrumental variables," *Annual Review of Economics*, 2018, *10*, 577–613.

Morten, Melanie and Jaqueline Oliveira, "The effects of roads on trade and migration: Evidence from a planned capital city," *American Economic Journal: Applied Economics*, 2024, *16* (2), 389–421.

Soumahoro, Souleymane, "Export taxes and consumption: Evidence from Côte d'Ivoire's de facto partition," *Economic Development and Cultural Change*, 2017, *65* (3), 425–456.

Vincenty, Thaddeus, "Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations," *Survey review*, 1975, *23* (176), 88–93.

Vytlacil, Edward, "Independence, monotonicity, and latent index models: An equivalence result," *Econometrica*, 2002, *70* (1), 331–341.

Wooldridge, Jeffrey M, "Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators," *Available at SSRN 3906345*, 2021.

Online Appendix to "Stop That Train, I'm Starving"

Mohamed Coulibaly

February 10, 2026

D Appendix: Data and Variable Construction

This appendix provides additional details on the construction of the variables and datasets used in the paper. I describe the aggregation of household- and individual-level data to the subprefecture level, the construction of key outcome variables, and the steps taken to ensure comparability across survey waves. All procedures are implemented consistently across the three waves of the *Enquête Niveau de Vie* (ENV), for the years 2002, 2008, and 2015.

D.1 Aggregation and Unit of Analysis

Most reduced-form analyses are conducted at the subprefecture level. Subprefectures represent the lowest administrative units in Côte d'Ivoire during the study period and allow for consistent aggregation of socioeconomic outcomes across years. I aggregate household-level variables and individual-level variables separately. For household outcomes, I retain only subprefectures with at least 12 sampled households in a given survey year. For labor market outcomes, I require at least 12 sampled working-age adults (aged 15–64). This restriction ensures sufficient within-cell variation and reduces measurement noise in small subprefectures. The resulting dataset is an unbalanced panel of subprefectures across the three survey waves.

D.2 Construction of Key Variables

Household Per Capita Spending

To maintain consistency across survey years, I construct total household expenditure directly from the raw sections of the ENV. Each spending category is computed following the original questionnaire structure and recall periods, which vary across sections. I winsorize all expenditure components to limit the influence of extreme outliers before aggregation.

The construction follows these main steps:

1. **Education spending:** annual expenditures on school fees, supplies, uniforms, transportation, canteen meals, extracurricular activities, and tutoring.
2. **Health spending:** all medical expenses (consultations, medication, hospitalization, and other health-related costs) over the past three months, annualized by multiplying by four.
3. **Clothing spending:** reported annual expenditures on garments, tailoring, footwear, jewelry, and hairdressing.
4. **Personal care and leisure:** spending on hygiene products and domestic help, annualized where necessary.
5. **Transport and communication:** expenditures on routine and exceptional transport, fuel, vehicle repairs, and trips, adjusted to annual equivalents based on reported frequency.
6. **Housing:** rent, home repair, water, electricity, and cooking fuel costs. Water, electricity, and fuel expenditures are converted to annual terms using reported payment frequencies.
7. **Food:** separated into food consumed at home and food consumed outside the home, both annualized based on reported purchase frequency.

Each component is winsorized at the upper tail (typically the 99.9th percentile for all expenditure categories expect health expenditures at the 99.5th percentile) and then collapsed to the household level. Total household expenditure is the sum of all categories:

$$\begin{aligned}\text{Total household expenditure} &= \text{Education} + \text{Health} + \text{Clothing} + \text{Personal care} \\ &\quad + \text{Transport} + \text{Housing} + \text{Food at home} + \text{Food outside home}.\end{aligned}$$

I then compute per capita spending by dividing total household expenditure by household size. This measure serves as a consistent indicator of household outcomes across survey waves.

Labor Market Outcomes

Labor market outcomes are derived from the individual questionnaire, specifically question *Eb.3* in the 2015 survey wave ("Quelle est votre catégorie socio-professionnelle à l'obtention de ce travail et actuellement?"). This question identifies detailed occupational categories ranging from civil servants and private employees to self-employed and agricultural workers.

I construct three mutually exclusive employment categories based on occupation codes and payment type:

1. **Nonfarm employment (*total_emp_NF*):** includes wage or self-employed work in non-agricultural activities.
   - Casual nonfarm workers: occupations 1–13 or 16, paid irregularly.
   - Regular nonfarm workers: same occupations, paid monthly.
   - Self-employed nonfarm workers: independent workers and employers (codes 14–15).
2. **Farm employment (*total_emp_F*):** includes individuals engaged in agricultural activities for pay or self-employment.
   - Casual farm workers: paid agricultural laborers (codes 18 or 20).
   - Self-employed farmers: independent agricultural producers (code 17).
3. **Contributing family work (*total_emp_FA*):** unpaid family workers (code 19), classified as casual or regular depending on payment frequency.

Each variable is coded at the individual level for those reporting any work activity during the reference year and averaged within subprefectures for the reduced-form analysis. The resulting measures capture the share of working-age adults employed in each sector.

D.3 Sample Restrictions and Weighting

All monetary variables are expressed in nominal CFA francs of the survey year. I apply the survey sampling weights provided in the ENV to maintain representativeness at the national level prior to aggregation. For subprefecture-level averages, I compute weighted means using household or individual weights as appropriate.

Individuals with missing occupational codes or employment status are omitted from the construction of labor variables. These individuals are not employed in the first place, so their exclusions do not affect representativeness.

D.4 Consistency Across Survey Waves

All variables are constructed using identical procedures across the 2002, 2008, and 2015 surveys. The ENV questionnaires maintain a consistent structure, though minor differences in item coding or recall periods are harmonized consistently based on the approach described above. Aggregation, winsorization, and labeling steps are fully reproducible and applied identically in each wave.

E Appendix: 2011 Post-Presidential-Election Crisis's Historical Background

Côte d'Ivoire, one of West Africa's fastest-growing economies after independence, experienced sustained growth averaging over 7 percent annually for two decades. Its heavy reliance on primary commodities, particularly cocoa and coffee, which accounted for more than half of exports by 2000, made it highly vulnerable to terms-of-trade shocks. A prolonged downturn in the 1980s, driven by falling export prices, eventually led to a currency devaluation in 1993.

The following decade was marked by growing political instability. A military coup in 1999 and a subsequent conflict in 2002 divided the country into a rebel-controlled north and a government-controlled south (Soumahoro, 2017). Under the Linas–Marcoussis peace agreement of 2003, this de facto partition persisted for nearly a decade, until the contentious presidential election of 2010.

The 2010 election was held in two rounds and produced unexpected results. In the first round (October 2010), the incumbent president led with 38.3 percent of the vote, followed by the main opposition candidate with 32.1 percent. In the second round (November 2010), the opposition candidate formed a broad coalition with the third-largest political party. When results were announced, the country's institutions split: the Independent Electoral Commission declared the opposition the winner (54.1 percent), while the Constitutional Council declared the incumbent the winner (51.45 percent).

With both candidates claiming victory, Côte d'Ivoire descended into a brief but intense post-presidential-election crisis between December 2010 and April 2011. Violence was highly localized: three areas, the economic capital Abidjan (49.6 percent of all deaths), the Cavally region (30.3 percent), and the city of Duékoué (11.2 percent), accounted for more than 90 percent of total fatalities (Léon and Dosso, 2020). These three areas are excluded from my analysis sample.

The conflict caused short-term disruption across major economic sectors. During the 2011 crisis, GDP contracted by 5.4 percent, but growth rebounded quickly, averaging over 5 percent annually from 2012 until the COVID period. Land borders were closed, and the national rail network was fully suspended between January and April 2011. Freight operations resumed in late April and intercity passenger services in June, but rural rail stations along the southern corridor remained closed.

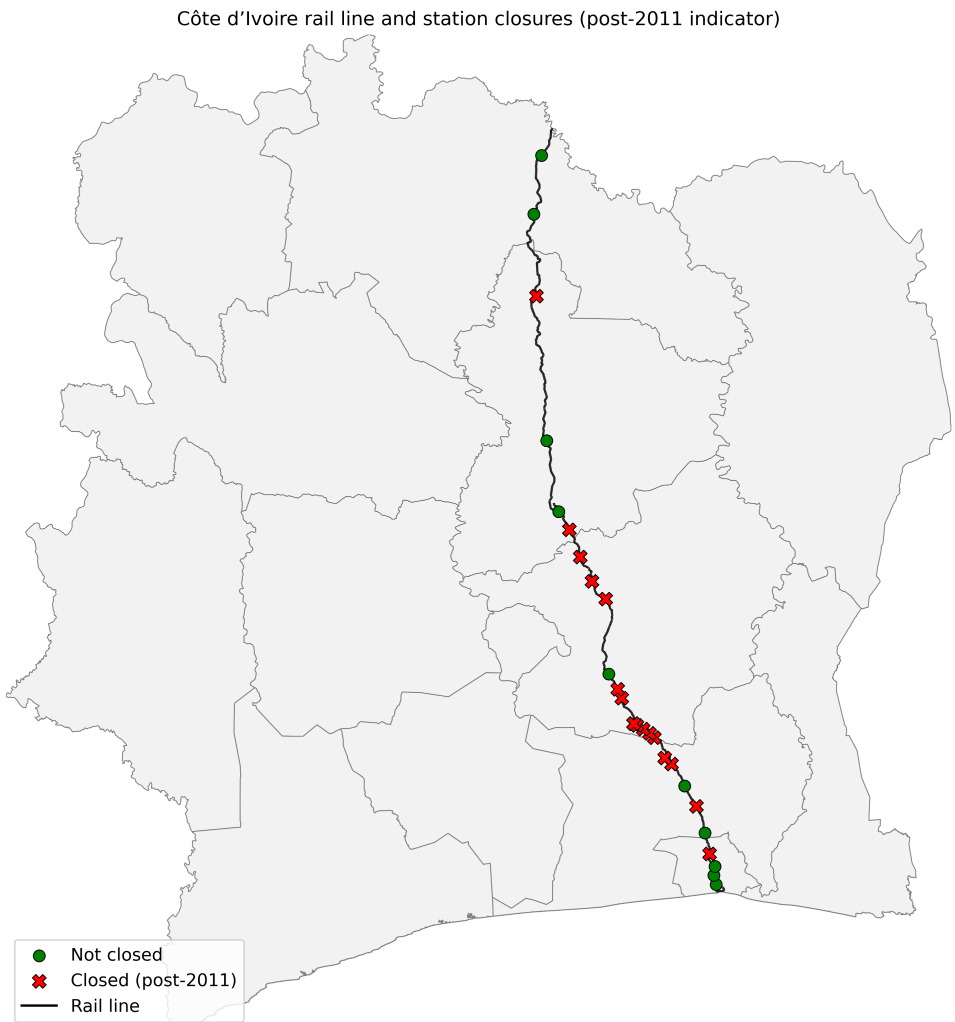

Figure 1: Rural Rail Stations Before and After 2011

Note: These figures show the number of rail stations that were open before and after the post-presidential-election crisis. 18 of all existing rail stations were closed following the crisis, with 16 of these closures occurring in southern Côte d'Ivoire.
Source: Own computation (see Dagnogo et al. (2012) for more details)

Figure 1 illustrates the geographic contraction of rail access: all rural stations in southern Côte d'Ivoire closed permanently after the crisis. Of the 28 passenger stations operating before 2011, only 10 reopened, all in major urban centers. The remaining 18 stations, 16 of them in southern rural regions, remained closed.

The closures were abrupt, exogenous to local economic conditions, and geographically concentrated. They were not driven by declining demand but became possible following the sudden and unexpected change in government. This makes the 2011 crisis a natural experiment for studying how the loss of rural rail access affected household living standards and labor allocation in previously connected communities.

F Appendix: Spatial Targeting Via Empirical Welfare–Maximization Framework

This appendix derives the empirical welfare-maximization framework used to evaluate spatially targeted counterfactual rail-reopening allocations. The objective is to characterize how alternative spatial reopening allocations affect aggregate welfare and to derive the corresponding optimal allocation rule.

Intuition. Reopening a rail station in a given location reduces sectoral mobility costs for a specific group of nearby workers. Because these workers differ in both their potential nonfarm returns and their costs of switching sectors, the welfare impact of reopening a given station depends on who gains access. Allocations that connect areas with a high concentration of constrained yet high-return workers generate the largest aggregate gains. The optimal allocation thus relaxes mobility constraints where they are most binding and productive, accelerating structural transformation in the rural economy.

Setup and Notation. Let each subprefecture S be characterized by a distance to the nearest active rail station under allocation A, denoted by $Dist_{S,A}$. Define the change in log distance relative to the baseline (pre-closure) network as:

$$Z_A = \log Dist_{S,A} - \log Dist_{S,base}.$$

This variable, compared to the initial rail closure-induced shock to rail access ($\log Dist_{S,close} - \log Dist_{S,base}$), improves rail access by reducing the closure-induced shock through reopenings under allocation A. Note that $Z_A \leq Z_{\text{close}}$ almost surely. That is, compared to the scenario with all closed rural rail stations, reopenings under allocation A would reduce sectoral mobility costs, thereby increasing nonfarm participation and improving welfare.

Under Assumptions **A.1-A.5** (defined in the paper), expected welfare under allocation A satisfies:

$$\mathbb{E}[Y \mid A] = \mathbb{E}\left[D(A) \cdot Y_1 + (1 - D(A)) \cdot Y_0\right]$$

where $D(A)$ indicates nonfarm participation under allocation A. Hence, alternative reopening allocations affect welfare only through their impact on the induced sectoral allocation $D(A)$.

Empirical Welfare–Maximization Problem. The policymaker's problem is to select the reopening allocation A^* that maximizes expected welfare:

$$A^* = \arg \max_{A \in \mathcal{A}_{8/16}} \mathbb{E}[Y \mid A],$$

where $\mathcal{A}_{8/16}$ denotes the set of 12,870 possible reopening configurations that involve reopening eight of the sixteen closed rural stations. The welfare criterion is utilitarian, emphasizing average living standards rather than distributional concerns.

Because the baseline welfare level $\mathbb{E}[Y \mid \text{baseline}]$ is constant across allocations,[1] maximizing expected welfare is equivalent to maximizing the welfare gain relative to the baseline:

$$\max_{A \in \mathcal{A}_{8/16}} \{\mathbb{E}[Y \mid A] - \mathbb{E}[Y \mid \text{baseline}]\}.$$

Under Assumptions **A.1–A.5**, this welfare difference corresponds to the Policy-Relevant Treatment Effect (PRTE), that is:

$$\text{PRTE}(A) = \int \text{MTE}(u)\, w^A(u)\, du$$

where $\text{MTE}(u)$ is the marginal treatment effect as a function of unobserved costs u, and $w^A(u)$ represent the policy weights capturing how each allocation changes the distribution of workers across sectors for different levels of unobserved costs (resistance) to nonfarm employment.

[1] It is finite per assumption **A.4**.

Hence, the empirical welfare–maximization problem can be rewritten as

$$A^* = \arg \max_{A \in \mathcal{A}_{8/16}} \text{PRTE}(A).$$

This representation links the planner's decision directly to heterogeneity in returns and mobility costs, derived from microdata rather than from a calibrated general-equilibrium structure.

Implementation. To implement the empirical welfare maximization problem described above, I compute the policy relevant treatment effect, PRTE(A), for every feasible reopening allocation $A \in \mathcal{A}_{8/16}$. Each allocation corresponds to reopening 8 of the 16 closed rural stations. The procedure for computing the PRTE under each allocation follows Algorithm 1.

Algorithm 1 ***Computing the PRTE for Each Reopening Allocation***

1. ***Initialize the set of allocations.*** *Construct the set $\mathcal{A}_{8/16}$ containing all 12,870 possible combinations of reopening 8 out of 16 closed stations.*

2. ***Compute updated distances.*** *For each allocation $A \in \mathcal{A}_{8/16}$, compute the distance from each subprefecture S to the nearest active rail station, denoted $Dist_{S,A}$.*

3. ***Update cost shifters.*** *For each subprefecture, update the policy-induced cost shifter as*
$$Z_{S,A} = \log(Dist_{S,A}) - \log(Dist_{S,base}).$$

4. ***Recover policy weights.*** *Using the updated $Z_{S,A}$, recover the policy weights $w^A(u)$ across the unobserved mobility cost distribution, following the expression in the previous subsection.*

5. ***Compute welfare effects.*** *For each allocation A, compute the Policy-Relevant Treatment Effect:*
$$PRTE(A) = \int MTE(u) \cdot w^A(u)\, du.$$

6. ***Store and rank results.*** *Record the estimated $PRTE(A)$ for each allocation and rank them by the resulting welfare gains, identifying the allocation A^* that maximizes $PRTE(A)$.*

The resulting distribution of PRTE(A) values across all 12,870 allocations summarizes the potential welfare gains achievable under alternative spatial-targeting strategies.

Relation to Existing Frameworks. Unlike standard empirical welfare-maximization frameworks such as Kitagawa and Tetenov (2018), where assignment is exogenous, the current setup accounts for endogenous sectoral participation driven by spatially varying mobility costs. The resulting welfare problem is most closely related to the extensions in Sasaki and Ura (2024) and Liu (2022), which address policy evaluation in the presence of endogenous selection and heterogeneous returns.

Discussion. The empirical welfare–maximization approach used here should be interpreted as a partial-equilibrium exercise. It captures the first-order effects of reallocating labor in response to improved transport access while abstracting from general-equilibrium adjustments such as large-scale migration or or intersectoral linkages. Evidence from the reduced-form analysis in the paper suggests the migration assumption is reasonable, supported by limited migration responses observed following the larger shock of all rural rail station closures. This is also supported by existing literature on limited migration in low-income rural contexts (Bryan et al., 2014; Morten and Oliveira, 2024).

Summary. This appendix establishes the empirical foundation for the spatial-targeting exercise in the paper. By mapping each feasible reopening configuration to its implied policy weights and corresponding PRTE, this framework quantifies how the spatial allocation of limited infrastructure investments shapes aggregate welfare and identifies the configuration that maximizes it.

References

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak, "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh," *Econometrica*, 2014, *82* (5), 1671–1748.

Dagnogo, Foussata, Olivier Ninot, and Jean-Louis Chaléard, "Le chemin de fer Abidjan-Niger: la vocation d'une infrastructure en question," *EchoGéo*, 2012, (20).

Kitagawa, Toru and Aleksey Tetenov, "Who should be treated? empirical welfare maximization methods for treatment choice," *Econometrica*, 2018, *86* (2), 591–616.

Léon, Florian and Ibrahima Dosso, "Civil conflict and firm recovery: Evidence from post-electoral crisis in Côte d'Ivoire," Technical Report, FERDI Working Paper 2020.

Liu, Yan, "Policy learning under endogeneity using instrumental variables," *arXiv preprint arXiv:2206.09883*, 2022.

Morten, Melanie and Jaqueline Oliveira, "The effects of roads on trade and migration: Evidence from a planned capital city," *American Economic Journal: Applied Economics*, 2024, *16* (2), 389–421.

Sasaki, Yuya and Takuya Ura, "Welfare analysis via marginal treatment effects," *Econometric Theory*, 2024, pp. 1–24.

Soumahoro, Souleymane, "Export taxes and consumption: Evidence from Côte d'Ivoire's de facto partition," *Economic Development and Cultural Change*, 2017, *65* (3), 425–456.