

Development Aid Projects
and Migration Across
Mobility Margins

Narcisse Cha'ngom

GODAD

Development Aid projects and Migration Across Mobility Margins

Narcisse Cha'ngom*

October 7, 2026

Abstract

This paper estimates the causal effect of local exposure to development aid on migration across several margins: internal versus international, and regular versus irregular international. Using Senegal as a setting, I combine geo-referenced project-level aid with census-based bilateral flows from two waves, one of which uniquely records emigrants' legal status at departure. For identification, I predict where projects land from districts' vote shares for the incumbent at the preceding election, then let predicted aid arrive over time according to donor-sector disbursement profiles in other low- and lower-middle-income countries. The exogeneity comes from this externally-determined timing, which is orthogonal to local migration shocks. Doubling aid per capita reduces overall migration by 9.4 percent, driven by a 9.8 percent decline in internal migration, with no effect on regular international migration. Aid also curbs irregular migration by roughly 9.7 percent, concentrated among low-skilled young men along Europe-bound corridors. I provide evidence for an opportunity-cost channel: aid raises local self-employment among prime-age workers.

Keywords: Migration, Internal, International, Development aid, Employment.

JEL codes: F22, F35, R2, O55

*I thank Jesús Fernández-Huertas Moraga, Albrecht Glitz, André Groeger, Furio Camillo Rosati, Frédéric Docquier, and Joël Machado for their helpful comments. I am also grateful to participants at the 10th Conference on Understanding Voluntary and Forced Migration in France, the 2024 Barcelona Summer Forum (BSF) in Spain, and the 2023 Winter School on Inequalities, Social Norms and Migration in Italy for valuable feedback. Correspondence: Narcisse Cha'ngom ([\[narcisse.changom@liser.lu\]](mailto:narcisse.changom@liser.lu) (<mailto:narcisse.changom@liser.lu>)).

1 Introduction

Development aid is increasingly used as a policy instrument to address the “*root causes*” of migration ([European-Commission, 2015](#), [National-Security-Council, 2021](#)). Yet whether aid reduces migration is theoretically ambiguous. By improving incomes, public services, and employment opportunities at origin, aid may raise the opportunity cost of leaving and reduce migration. At the same time, by relaxing liquidity constraints, expanding access to information about donor countries, aid may also enable migration, especially costly long-distance moves. This ambiguity is particularly important because different mobility margins respond to development incentives in different ways. Internal migration, and international migration differ in cost. Also, regular international, and irregular international migration differ in cost, risk, liquidity requirements, and policy relevance. Aggregating these margins may therefore obscure how development interventions shape mobility decisions. Distinguishing between them is thus essential for understanding how development aid interacts with migration.

This paper takes up that task by estimating the causal effect of local development aid projects on migration across distinct mobility margins. Using Senegal as an empirical setting, I combine geo-referenced project-level aid data ([Bomprezzi et al., 2025](#)), with census-based bilateral migration flows to examine, within a unified framework, how local exposure to development aid projects affects internal and international migration. I then further distinguish between regular and irregular international migration to investigate how each margin responds to local exposure to aid projects.

Senegal is unusually well suited to this question. The country combines substantial internal mobility, concentrated international migration corridors, and rare census-based information on irregular migration. The 2013 and 2023 *full population censuses* allow me to construct harmonized five-year migration flows across Senegalese districts and from Senegalese districts to foreign destination countries. The 2023 census further identifies emigrants abroad without regular legal status at departure, allowing for a direct analysis of irregular migration, a margin that is central to policy debates but rarely observed in standard data.

While Senegal provides unusually rich data for studying these margins, its migration patterns are not exceptional within Sub-Saharan Africa. As shown in [section 4.5](#), Senegal’s internal migration rates, age profile of mobility, international destination structure, and irregular-route composition all fall within the range observed for comparable SSA countries. The Senegalese case therefore offers a relevant setting for studying how local development aid impact mobility across margins, with lessons that extend beyond the country itself.

Estimating the effect of aid on migration is challenging because aid is not randomly allocated. Projects may be targeted to districts with weak economic conditions, high out-migration pressure, or adverse local shocks. To address this concern, I implement a placement-adjusted synthetic disbursement IV design. Aid tends to flow toward districts that are politically aligned with the sitting president at the time of commitment, but the pace at which committed projects actually disburse is dictated by donor routines far beyond any single recipient. I exploit this by predicting where projects land from districts' pre-commitment vote shares for the incumbent, then letting that predicted aid flow over time according to the disbursement profile of comparable donor-sector projects outside Senegal and the wider WAEMU region. The instrument is local aid projects whose placement is politically driven and whose timing is externally set, plausibly orthogonal to local migration shocks conditional on district and time fixed effects. The estimating framework is gravity-like and includes origin–destination fixed effects and destination–time fixed effects, absorbing time-invariant corridor frictions and destination-specific shocks. I further check the robustness of my identification strategy by employing an alternative design. I implement standard shift–share instrumental-variables (SSIV) design. The instrument interacts predetermined district exposure to donor–sector cells with changes in donor–sector disbursements to non- West Africa Economic and Monetary Union (WAEMU) Low- and Lower-Middle Income countries, excluding Senegal. The identifying variation therefore comes from external donor–sector shocks interacted with historical local aid portfolios, rather than from contemporaneous migration conditions in Senegal.

The results show that local development aid reshapes migration in highly margin-specific ways. A doubling of local aid per capita reduces five-year migration flows by about 9.4 percent in the pooled sample. This aggregate decline is driven primarily by internal migration, which falls by roughly 9.8 percent. By contrast, overall international migration is statistically unaffected. While regular international migration also appears unaffected by exposure to local aid projects, irregular international migration, however, declines substantially. A doubling of aid per capita reduces irregular migration by approximately 9.7 percent. These findings indicate that aid does not uniformly deter migration. Rather, it reduces mobility along margins that are more directly exposed to local economic opportunities or involve high risks, while leaving regular international migration largely unchanged.

I then examine the mechanisms behind these effects. The evidence supports an opportunity-cost channel operating through local labor markets. Using individual-level census data, I show that a 10 percent increase in local project aid per capita raises the probability of employment by about 0.13 percentage points, with gains concentrated in

self-employment and among prime-age workers. These are precisely the groups for whom local employment opportunities are most likely to affect migration decisions. In contrast, I find no evidence for an information channel through which bilateral donor presence facilitates migration to donor countries. Linking aid projects to their respective donors, I estimate whether donor-specific aid exposure increases migration to the corresponding donor country and find uniformly null effects. Together, the evidence suggests that aid affects migration primarily by improving local outside options, rather than by transmitting destination-specific information.

Although seemingly disparate, these margin-specific responses follow a single logic. Development aid increases the value of staying in the country of origin by improving local labor market opportunities. This shifts the migration decision against leaving, but not uniformly across all margins. Internal migration, the cheapest and most easily substituted form of mobility, decreases the most because, for these migrants, a local opportunity closely substitutes for a move. Irregular international migration, undertaken at high financial and physical cost by those with the weakest outside options, also falls sharply because the value of the local alternative is weighed against the risk of a clandestine journey. Regular international migration, by contrast, is governed less by conditions at origin than by rationing at the destination, such as visa quotas, family reunification eligibility, and admission policy, so improvements in local opportunities leave it largely unchanged. Thus, aid compresses the margin it can reach through the outside option while leaving the destination-side constrained margin unchanged.

This paper contributes to three strands of literature. First, it speaks to the literature on development aid and international migration. Existing cross-country studies document mixed effects of aid on emigration, with estimates varying by aid type, time horizon, and empirical specification (Belloc, 2015, Berthélemy et al., 2009, Dreher et al., 2019a, Gamso and Yuldashev, 2018a, Lanati and Thiele, 2018a,b, Marchal et al., 2025, Murat, 2020). These studies are informative about aggregate aid–migration relationships but generally rely on national-level aid and migration measures, which mask substantial sub-national heterogeneity. This paper complements them by exploiting geo-referenced project-level aid exposure and origin–destination internal and international migration flows within a single country. Second, the paper relates to work emphasizing bilateral information and network channels linking aid and migration (Belloc, 2015, Marchal et al., 2025, Menard and Gary, 2018). These studies suggest that donor-recipient relationships may facilitate migration by transmitting destination-specific information or strengthening bilateral ties. I test this mechanism directly by linking projects to their donor countries and estimating whether donor-specific aid exposure increases migration to the corresponding destination. I find no evidence for this channel. Third, the paper contributes to the emerging sub-

national literature using geo-coded development projects to study migration responses. Closest to this paper, [Fuchs et al. \(2023\)](#) show that World Bank projects affect migration aspirations and later regular international migration, while [Lanati et al. \(2023\)](#) study internal migration responses to local aid projects. I extend this literature by jointly studying internal, regular international, and irregular international migration within a unified origin–destination framework. This distinction is central because these margins differ sharply in cost, risk, liquidity requirements, and policy relevance. In doing so, the paper helps reconcile mixed findings in the literature: aid may reduce some forms of migration, leave others unchanged, and do so through local labor-market improvements rather than through bilateral information effects.

The remainder of the paper is organized as follows. Section 2 positions the paper relative to the existing literature. Section 3 describes the migration and development aid data and the construction of local exposure to project aid measures. Section 4 provides some descriptive statistics. Section 5 outlines the empirical framework and identification strategy. Section 6 presents the main results on the effects of local exposure to development aid on internal and international migration. Section 7 investigates how irregular migration responds to local aid projects. Section 8 investigates the mechanisms underlying the migration responses, with a focus on local labor-market and information channels, and section 9 concludes.

2 Related literature

Existing empirical work on aid-migration nexus mirrors the conceptual ambiguity discussed above. A first empirical strand examines the aid-migration nexus using cross-country data. Several studies find heterogeneous effects of aid on emigration, depending on the type of aid, time horizon, and empirical specification ([Belloc, 2015](#), [Berthélemy et al., 2009](#), [Dreher et al., 2019b](#), [Gamso and Yuldashev, 2018a,b](#), [Lanati and Thiele, 2018a,b](#), [Marchal et al., 2025](#), [Murat, 2020](#)). Some studies find that aid is associated with reduced emigration, especially when it improves public services ([Lanati and Thiele, 2018b](#)), governance ([Gamso and Yuldashev, 2018b](#)), or targets rural development ([Gamso and Yuldashev, 2018a](#)). Moreover, foreign aid has been shown to reduce refugee flows to donor countries ([Dreher et al., 2019b](#), [Murat, 2020](#)). Other studies argue that foreign aid can increase emigration by relaxing liquidity constraints, raising aspirations, or strengthening bilateral ties between donors and recipients, particularly through information and network channels. Cross-country evidence shows that aid and migration may act as complements rather than substitutes ([Berthélemy et al., 2009](#), [Marchal et al., 2025](#)), while related work for Sub-Saharan Africa documents a positive

association between development assistance and outward migration (Belloc, 2015). However, much of this literature relies on national-level aggregates of aid and migrant stocks, which mask substantial sub-national heterogeneity.

A second strand focuses on bilateral relationships between donors and recipient countries. This literature emphasizes potential information and network channels, whereby aid flows may transmit destination-specific information or strengthen political and economic ties that facilitate migration (Belloc, 2015, Marchal et al., 2025, Menard and Gary, 2018). While informative, these studies typically treat migration as a single international margin and abstract from internal mobility or differences between regular and irregular migration.

More recently, research has begun to exploit geo-referenced project-level data and sub-national variation. Fuchs et al. (2023) combine geo-coded World Bank projects with global survey data on migration to show that aid reduces migration aspirations in the short run but increases regular international migration in the longer run, consistent with mobility transition dynamics. Lanati et al. (2023) focus on internal migration and document that local aid projects can alter domestic mobility patterns by operating as a pull factor in location where projects are implemented. Together, these contributions establish that development aid can affect migration when measured at the local level. At the same time, existing sub-national evidence typically examines individual margins in isolation and does not jointly evaluate how aid reshapes internal, regular international, and irregular migration within a unified empirical framework.

Distinguishing mobility margins is economically important. Internal migration is quantitatively dominant worldwide and often serves as an adjustment mechanism to local shocks (Marchiori et al., 2012, UNDP, 2009). Regular and irregular international migration differ substantially in cost, risk, and liquidity requirements. Development interventions that improve local labor markets may reduce internal mobility while having different or even opposite effects on international migration. Aggregating these margins may therefore obscure important reallocation effects.

3 Data

This paper combines two rounds of uniquely rich Senegalese population census data with geo-coded information on development aid projects from Geo-coded Official Development Assistance Dataset (GODAD) to construct district-level measures of dyadic migration flows (internal and international) and local exposure to aid projects.

3.1 Migration

I use the 2013 and 2023 Senegalese population censuses to measure migration over two harmonized five-year windows: 2009–2013 and 2019–2023. The censuses record individuals’ current location of residence and their location of residence five years earlier. In addition, households report whether any member emigrated abroad during the previous five years, together with the year of departure and destination country. These features allow the construction of consistent origin–destination dyadic migration flows, district-to-district (internal) and district-to-abroad (international). Migration flows are aggregated to Senegal’s 45 districts (ADM2-level).¹ This level of aggregation is chosen to match the geographical details of the development aid data.

Internal and international migration flows are constructed using a common accounting framework. Let $\tau \in \{2013, 2023\}$ index the census year and λ index individuals. Denote by $o_{\lambda, \tau-5}$ the individual’s location of residence five years prior to the census and by $d_{\lambda, \tau}$ the location of residence at the census date. Locations may correspond either to Senegalese districts or to foreign destination countries.

I define the bilateral five-year migration flow from origin o to destination d as

$$M_{od, \tau: \tau-5} \equiv \sum_{\lambda=1}^{N_{\tau}} 1\{o_{\lambda, \tau-5} = o, d_{\lambda, \tau} = d, d \neq o\}, \quad (1)$$

where N_{τ} denotes the population in census year τ .

The bilateral five-year migration flow from origin o to destination d is defined as the number of individuals whose residence changed from o to d over the previous five years. Throughout the analysis, migration is measured net of stayers, focusing on origin–destination pairs with $d \neq o$. This definition applies uniformly to internal and international migration; the distinction arises solely from the set of possible destinations.

Internal migration. For internal migration, both o and d correspond to Senegalese districts. Flows are obtained by comparing district of residence at time t with district of residence at $t - 5$. Aggregating individual transitions yields a district-by-district origin–destination matrix for each census round.

International migration. For international migration, the origin o is a Senegalese district and the destination d is a foreign country. The census records household-reported

¹It is worth noting that two years prior the 2023 census, a 46th district was created in Senegal, *Keur-Massar*, but mainly consisted in splitting the district of *Pikine* into two separate districts. To keep consistent boundary across the two census rounds, I recoded Keur-Massar as Pikine such has to remain with 45 districts in 2023 as well.

emigrants who left Senegal during the previous five years, including their year of departure, destination country, and socio-demographic characteristics at the time of departure.

3.2 Development aid projects

Local exposure to development aid project is measured using the Geo-coded Official Development Assistance Dataset ([Bomprezzi et al., 2025](#)),² which provides sub-national information on aid project locations, commitments, disbursements, sectors, donors, and implementation periods. The dataset covers 3,867 projects implemented across Senegal between 1989 and 2023, with total disbursements of approximately \$7.9 billion overall. Projects are classified into four OECD’s DAC sectors: social infrastructures and services, economic infrastructure and services, production sector, and other (residual category). For each census wave, I link migration outcomes to development projects implemented in the five years preceding the migration observation window. Specifically, migration observed in 2009–2013 is matched to projects implemented between 2004 and 2008, while migration observed in 2019–2023 is matched to projects implemented between 2014 and 2018. This alignment ensures that project exposure systematically predates migration decisions while maintaining a consistent five-year look-back window across census waves and that the estimated effect will capture the medium/long run responses. This yields a total of 165 projects implemented between 2004 and 2008 across 468 localities, and 319 projects implemented between 2014 and 2018 across 638 localities (Figure [E.6](#)).

4 Descriptive statistics

4.1 Migration: evolution, composition and spatial distribution

Table [1](#) provides a comprehensive snapshot of five-year migration rates and flows in Senegal between the 2008–2013 and 2018–2023 extracted from the 2013 and 2023 census waves, disaggregated by migration type and key demographic characteristics. Overall mobility declined markedly over the decade: the share of individuals changing residence over a five-year window fell from 9.4 percent to 7.2 percent, driven primarily by a contraction in internal migration (from 6.8 to 5.3 percent), while international migration, smaller in magnitude, also decreased from 2.7 to 1.9 percent. Despite lower rates, total migration flows increased, reflecting strong population growth (38.5 percent between 2013 and 2023): aggregate migrant stocks rose by 13.5 percent, with particularly pronounced growth in internal flows (15.6 percent). Migration is sharply positively selected on

²Geo-coded Official Development Assistance Dataset (GODAD).

education. In both periods, high-skilled individuals exhibit migration rates more than twice those of low-skilled individuals, and the education gradient is especially steep for international migration where high-skilled are four times more likely to migrate than their non high-skilled counterparts. Age profiles display the familiar inverted-U pattern, with mobility peaking among young adults aged 20–29 and declining steadily thereafter, while international migration is overwhelmingly concentrated among prime-age cohorts. Finally, migration is strongly gendered: men are substantially more likely to migrate than women, a gap that is almost entirely explained by international migration, whereas internal migration rates are similar across genders. Taken together, Table 1 highlights declining migration propensities alongside expanding migrant flows, and documents strong selection by education, age, and gender that is stable across census rounds.

Figure E.7 complements the aggregate patterns documented in Table 1 by illustrating the spatial distribution of migration rates across Senegalese districts and how this geography has evolved over time. It further highlights systematic differences between internal and international mobility. Panels E.7a and E.7b show that internal migration exhibits a stable and relatively diffuse geography over time. Although internal migration rates decline on average between 2009–2013 and 2018–2023, high-mobility districts remain concentrated in similar western, central, and selected southern areas, while large parts of the interior consistently display low internal mobility. The aggregate contraction thus reflects a broad-based reduction rather than a spatial reallocation of internal migration. Panels E.7c and E.7d reveal a more concentrated and equally persistent spatial structure for international migration. Districts with high international emigration rates in 2009–2013, particularly along the eastern and north-eastern corridor and in other peripheral areas, remain dominant origins in 2019–2023, despite lower overall intensities. Compared to internal migration, international migration is sharply clustered in a small number of districts, underscoring the role of localized migration networks, historical corridors, and geographic proximity to international routes (Bocquier et al., 2023). Taken together, the decline in both internal and international migration rates operates primarily as a level shift, reflecting faster population growth relative to the expansion of migration flows, rather than a spatial reallocation of mobility, leaving the underlying hierarchy of migration origins largely intact and reinforcing the view that internal and international migration are governed by distinct and persistent spatial mechanisms.

Overall, three salient facts stand out: (i) migration in Senegal is predominantly internal; (ii) international migration is highly selective along education, age, and gender dimensions, and (iii) despite declining mobility between 2013 and 2023, the cross-sectional structure of migration remains largely unchanged.

Table 1: Five Years migration rates and flows across various dimensions

	Migration rates (%)						Migration flows (× 1,000)					
	Both types		Internal		Int'l		Both types		Internal		Int'l	
	2013	2023	2013	2023	2013	2023	2013	2023	2013	2023	2013	2023
Panel A: Overall migration												
All	9.4	7.2	6.8	5.3	2.7	1.9	517.7	587.6	371.7	429.8	146.0	157.8
Δ	-23.4%		-22.1%		-29.6%		+13.5%		+15.6%		+8.1%	
Panel B: By education level												
HS	23.8	18.8	14.0	11.5	9.9	7.4	41.0	66.7	24.0	40.5	16.9	26.1
MS	13.1	13.1	10.2	9.3	2.8	3.8	73.9	61.4	58.0	43.5	15.9	17.9
LS	8.5	6.3	6.1	4.7	2.4	1.6	402.9	459.5	289.7	345.8	113.2	113.8
Panel C: By age group												
15-19	10.2	7.4	8.9	6.6	1.3	0.8	109.8	122.1	96.0	109.2	13.8	13.0
20-24	11.5	8.5	8.6	6.3	2.9	2.1	112.0	111.7	83.7	83.4	28.3	28.2
25-29	11.4	8.5	7.3	5.8	4.1	2.7	92.6	100.6	59.4	68.3	33.2	32.4
30-34	10.5	7.8	6.4	5.2	4.1	2.6	67.8	78.7	41.2	52.6	26.6	26.2
35-39	8.6	7.2	5.4	4.7	3.2	2.5	46.2	57.9	29.1	37.7	17.2	20.2
40-44	7.5	6.4	4.9	4.2	2.6	2.2	30.6	39.5	19.8	25.9	10.7	13.6
45-49	6.3	5.5	4.6	3.7	1.7	1.8	24.5	29.9	17.7	20.2	6.7	9.7
50-54	5.9	5.1	4.0	3.2	1.9	1.9	16.4	19.6	11.1	12.3	5.2	7.3
55-59	4.9	4.5	3.8	3.4	1.1	1.1	12.0	16.8	9.3	12.9	2.6	4.0
60-64	4.2	4.3	3.0	2.9	1.2	1.4	5.9	10.6	4.2	7.3	1.7	3.4
Panel D: By gender												
Women	7.5	6.0	6.7	5.4	0.8	0.6	214.1	248.5	190.1	222.9	24.1	25.5
Men	11.4	8.5	6.8	5.2	4.6	3.3	303.6	339.1	181.6	206.8	122.0	132.3

Notes: This table reports five-year migration rates (in percent) and corresponding migration flows (in thousands) constructed from the 2013 and 2023 population censuses of Senegal. Migration rates are defined as the number of individuals who changed their place of usual residence during the five years preceding each census relative to the initial population of the district of residence (i.e., the population 5 years earlier). Internal migration refers to moves across districts within Senegal, while international migration captures moves from districts in Senegal to a foreign country. “Both types” aggregates internal and international migration. Education groups are defined as high-skilled (HS) i.e. individuals with at least one year of tertiary education, middle-skilled (MS) i.e., individuals with Upper secondary up to high school diploma, and low-skilled (LS) reflects the residual category, based on the highest level of completed schooling at the time of the census. Age groups are defined at census date minus 5. Panel A reports aggregate migration outcomes, while Panels B–D dis-aggregate by education, age, and gender, respectively. The row labeled Δ reports the percentage change between 2013 and 2023.

4.2 Irregular international emigration from Senegal

The 2023 Senegalese population census provides a unique opportunity to document irregular international emigration directly and at a fine geographic scale.³ Focusing on the working-age population aged 15–64 over the 2019–2023 period, census households report that more than one in five international emigrants from Senegal departed without regular legal status. This corresponds to 21.7 percent of international emigrants, a share that is quantitatively substantial and rarely observed in standard migration data. A closer look at the demographic profile of Senegalese international out-migrants reveals striking patterns.

Irregular migration is large but is negatively selected on education. Irregular migration is concentrated overwhelmingly among low-skilled individuals, who account for 90.7 percent of all irregular migrants while representing 71.3 percent of total international emigrants. The irregularity rate among low-skilled migrants (27.2 percent) is more than seven times that of the high-skilled (3.6 percent). This strong education gradient is consistent with an opportunity cost channel: high-skilled individuals have access to regular visa pathways and stronger local labor market alternatives, whereas low-skilled workers face both fewer legal migration options and weaker employment prospects at origin.


Irregular migration is overwhelmingly a male phenomenon. Men account for 94.8 percent of all irregular emigrants, compared to 80.9 percent of regular emigrants. The irregularity rate for men reaches 24.1 percent, versus only 6.9 percent for women, a gap that is remarkably stable across age groups (Panel 1a of Figure 1). Panel 1b further indicates that, among the low-skilled men, the irregularity rate is 29 percent against 12.9 percent for low-skilled women. Among high- and medium-skilled migrants, women have irregularity rates below 1 percent, reflecting near-exclusive reliance on legal channels (family reunification, student visas).

Young adults, especially young men, drive irregular flows. The 25–34 age group alone accounts for 41.2 percent of all irregular migrants. Together with the 15–24, the 15–34 age group accounts for 71.6 percent of total irregular migration among the working age population. Young men aged 15–34 represent 68 percent of all irregular emigrants, with an irregularity rate of 28.1 percent. Moreover, The irregularity rate declines monotonically with age, from 29.9 percent among men aged 15–24 to 11.4 percent among men aged 60–64, whereas they remain nearly flat for women. This is consistent with the idea that

³Emigrants residing abroad without regular legal status at the time of departure.

younger cohorts face the largest liquidity constraints and the highest returns to migration relation to local outside options.

Figure 1: Demographic profile of irregular migration from Senegal



Notes: Panel 1a shows irregularity rates by five-year age group and gender, where the irregularity rate is defined as the ratio of irregular migrants to total migrants (regular + irregular) × 100. Panel 1b presents irregularity rates disaggregated by educational attainment across three education categories: low-skilled (no formal education or primary only), medium-skilled (secondary education), and high-skilled (tertiary education). Panel 1c displays the top 15 destinations ranked by number of irregular migrants; Libya is highlighted as a transit corridor rather than a final destination, as it serves as the principal entry point to the Central Mediterranean route toward Southern Europe. All figures are computed from the 2023 Senegalese Population Census.

Spain and Italy are the primary irregular gateways to Europe. Europe absorbs 59 percent of all irregular Senegalese emigrants, making it the dominant destination region. As depicted by Panel 1c of Figure 1, within Europe, Spain alone receives 28.4 percent, the largest among all major destinations. This reflects route from Senegal’s northwest coast to the Canary Islands, a well documented clandestine corridor. Italy, the second largest European irregular destination, records an irregularity rate of 21.9 percent, consistent with Central Mediterranean arrivals. By contrast, France, the second largest destination overall, shows a much lower irregularity rate of 12.1 percent, reflecting its large established diaspora and the predominance of family reunification flows.

Libya signals transit-driven irregularity and within Africa irregular migration is substantial. Libya stands out with an irregularity rate of 80.4 percent, the highest among all destinations (Panel 1c). The near total absence of regular status for Senegalese nationals in Libya is consistent with their documented role as transit node on the Central Mediterranean route. The fact that the census-measured irregularity rate is highest precisely where the clandestine transit is known to be concentrated corroborates the validity of the self-reported measure. Moreover, Sub-Saharan African destinations account for 27.7 percent of all irregular Senegalese emigrants, with an irregularity rate of 23.8 percent, comparable to European rate. Mauritania, the Gambia, Côte d’Ivoire and Gabon together account for 19.7 percent of all irregular emigrants. This within continent irregular mobility rarely features in policy debates focused Europe- and North-America-bound flows but represents a numerically comparable phenomenon.

Validation of irregular migration data. A natural concern is whether self-reported legal status in a population census can credibly measure irregular migration. I address this directly by validating the census-based counts against independent UNHCR asylum data and Frontex records of illegal border crossings. Figure E.3 plots the log number of Senegalese irregular migrants recorded in the 2023 Census against, respectively, the log number of new asylum applications (Panel E.3a), the log number of asylum seekers (Panel E.3b) filed by Senegalese nationals at each destination, and the log number of illegal border crossings by migration route (Panel E.3c). Across all three benchmarks, the census-based counts are strongly and positively correlated with the external measures, and the leading destinations and corridors line up as expected: Spain, Italy, and France dominate the asylum-based panels, while the Western African (to Spain) and Central Mediterranean (to Italy) routes account for the bulk of recorded crossings. The route-based comparison in Panel E.3c is especially tight, suggesting that the census captures the same underlying spatial structure of irregular flows reflected in administrative border

crossings data. This consistency across conceptually distinct sources (self-reported census status, asylum filings, and physical illegal border crossings) provides reassurance that the 2023 Census meaningfully measures irregular migration.

4.3 Aid projects: evolution, composition and spatial distribution

Figure E.8 documents a pronounced increase in aggregate aid disbursements over time, with particularly rapid growth in the later years (Panel E.8a). The sectoral composition of aid projects is heavily skewed toward social infrastructure, which accounts for 56 percent of total disbursements, followed by economic infrastructure at 29 percent (Panel E.8b). Aid projects in Senegal are characterized by long implementation horizons. Focusing on closed projects, Panel E.9a of Figure E.9 depicts the average project duration which is about 7.6 years, and this duration is remarkably stable across project start cohorts (Panel E.9b). Disbursements are also spread over time rather than concentrated at approval. Panel E.9a of Figure E.9 plots the share of total disbursements by event time, measured as “years since project start”, and shows that a substantial fraction of payments occurs several years after project initiation.

Table D.7 further summarizes the evolution of local development aid commitments and disbursements in Senegal over four five-year windows between 2004 and 2022, disaggregated by donor type and sector. Total disbursement increased sharply over time, rising from USD 0.7 billion in 2004–2008 to nearly USD 1.4 billion in 2018–2022, with multilateral donors accounting for the bulk of financing in all periods. Although disbursement to commitment gaps vary substantially across periods, donors, and sectors, these gaps should not be interpreted as evidence of systematic withholding of funds. The table reports disbursements observed within fixed five-year windows, while the average duration of completed projects is 7.6 years (Figure E.9, Panel E.9a), implying that a share of committed funds is mechanically disbursed outside the window defined. Consistent with this timing mismatch, disbursement shares are relatively low in some periods—most notably in 2008–2012, yet exceed 100 percent in others. Sectoral patterns mirror this dynamic, with economic and social infrastructure projects displaying particularly uneven disbursement profiles over time. Overall, the table highlights a sustained scaling-up of aid commitments alongside substantial intertemporal smoothing of disbursements, underscoring the importance of accounting for project duration when interpreting aid flows in short analytical windows.

Taken together, the aid data exhibit three salient features: (i) aggregate aid volumes increase sharply over time; (ii) disbursements are concentrated in social and economic infrastructure; and (iii) project implementation and payment profiles are long and

gradual. These features motivate an empirical measure of aid exposure based on realized disbursements and explicitly accounting for time since project initiation when relating local aid to migration outcomes measured over five-year census windows.

4.4 Spatial distribution of migration and aid projects


Focusing on the most recent period, Figure 2 reveals a meaningful, though spatially selective, alignment between the geography of development aid projects and migration. Consistent with the 2019–2023 patterns documented in Table 1 and Figure E.7, international migration and irregular migration are highly concentrated in eastern and peripheral districts. Panel 2d shows that these same districts tend to receive relatively high levels of aid per capita, particularly along the eastern border corridor, suggesting that development projects are disproportionately implemented in areas that are major origins of international and irregular migrants. In contrast, internal migration in 2019–2023 remains more spatially diffuse, with prominent origin districts located in the north-western and western parts of the country; these areas also exhibit comparatively high aid intensity, indicating a second axis of alignment between aid allocation and internal mobility pressures. While this correspondence is not uniform, some high-aid districts exhibit low migration rates, and vice versa, the overall pattern suggests that development aid is not randomly distributed across space but instead partially overlaps with entrenched migration corridors. Taken together, the figure points to a non-trivial spatial coincidence between exposure to aid projects and both regular and irregular migration in the recent period, providing suggestive descriptive evidence that the spatial allocation of aid projects is closely linked to the geography of migration pressures.

4.5 External validity

Before turning to the empirical analysis, I assess whether Senegalese migration patterns are representative enough for conclusions to generalize to other Sub-Saharan African countries. I treat internal and international migration separately, since the available comparison data differ across the two margins.

Internal migration. Census-based rates from IPUMS allow direct comparison across several Sub-Saharan African (SSA) countries. Panel E.1b of Figure E.1 shows that Senegal lies well within this range: its internal migration rate (6.8 percent) is close to the weighted cross-country average (5.9 percent), and its age profile follows the same inverted-U shape, peaking among young adults.

Figure 2: Spatial distribution of migration and development aid projects



Notes: This figure overlays the spatial distribution of migration outcomes and development aid projects across Senegalese districts. Panels 2a–2c report district-level five-year migration measures for the 2019–2023 period, constructed from the 2023 population census. Panel 2a shows internal migration rates, Panel 2b shows international migration rates, and Panel 2c reports the share of international migrants who migrated irregularly over the same period relative to the countrywide number of irregular migrants. Panel 2d maps average development aid (project aid) per capita at the district level over 2014–2018, aggregated across all projects and donors and normalized by district working age population (15–64). Differences across panels reflect variation in intensities.

International migration. No comparable census-based international emigration flows exist for other SSA countries, so I use the five-year bilateral estimates of [Abel and Cohen \(2022\)](#). I first confirm these estimates reproduce Senegal’s census patterns: Panel E.2a of Figure E.2 shows that destination-region distribution matches across both windows (2005–2010 estimates against the 2008–2013 census; 2015–2020 against 2018–2023). I then extract emigration flows by World Bank destination region for the remaining SSA

origins and benchmark Senegal against the SSA average. Panel E.1d shows it again falls within the SSA range.

Irregular migration. Using Frontex data on irregular border crossings by route, Panel E.1a shows that the route composition of Senegalese nationals broadly matches that of other SSA origins.

Together, these checks indicate that Senegal is broadly representative of SSA mobility along the internal, international, and irregular margins, supporting the external relevance of the empirical findings.

5 Empirical Strategy

5.1 Baseline Specification

The estimating framework is a gravity-style model in which bilateral five-year migration flows from origin district o to destination d are related to lagged local aid exposure at the origin. Let $\tau \in \{2013, 2023\}$ index census rounds and $\mathcal{D} \equiv \mathcal{D}^{int} \cup \mathcal{D}^{intl}$ denote the universal destination set, where $\mathcal{D}^{int}$ includes the 45 Senegalese districts and $\mathcal{D}^{intl}$ the 50 retained destination countries. The baseline estimating equation is:

$$\ln M_{od,\tau:\tau-5} = \delta_{od} + \delta_{d\tau} + \eta \ln(\text{Aidpc}_{o,\tau-5:\tau-9}) + \kappa \ln N_{od,\tau-5}^{tw} + X'_{o,\tau-5} \beta + \varepsilon_{od\tau} \quad (2)$$

Origin-destination fixed effects δ_{od} absorb all time-invariant corridor-specific drivers, including distance, and pre-existing migration networks, for both internal and international corridors. Destination-time fixed effects $\delta_{d\tau}$ absorb all pull-side shocks: economic or policy shifts in Dakar, France, Spain, or Italy affect estimated coefficients only insofar as they vary differentially across sending districts, which the origin-level variation in aid precludes. The variable of interest, $\text{Aidpc}_{o,\tau-5:\tau-9}$, is total disbursements per capita received by origin district o over the five years immediately preceding the migration window, ensuring that aid exposure is causally prior to migration decisions. The dyadic network term $N_{od,\tau-5}^{tw}$ captures pre-existing corridor-specific connections: for internal corridors it is the stock of origin-born residents who arrived in the destination district 5 years prior each census round; for international corridors it is the stock of returnees from that destination country, the most informative available proxy for bilateral information and support networks in the Senegalese census. Origin-time controls $X_{o,\tau-5}$ include climatic shocks (SPEI), conflict-related deaths, population size, and urbanization rates. Standard errors are clustered at the origin-time level.

To allow aid to have heterogeneous effects by migration type, I augment equation (2) with an interaction term:

$$\begin{aligned} \ln M_{od,\tau:\tau-5} &= \delta_{od} + \delta_{d\tau} + \eta_1 \ln(\text{Aidpc}_{o,\tau-5:\tau-9}) + \eta_2 \ln(\text{Aidpc}_{o,\tau-5:\tau-9}) \times \mathbf{1}\{d \in \mathcal{D}^{intl}\} \\ &\quad + \kappa \ln N_{od,\tau-5}^{tw} + X'_{o,\tau-5} \beta + \varepsilon_{od\tau}, \end{aligned} \tag{3}$$

where η_1 identifies the elasticity of internal migration with respect to local aid exposure, while $\eta_1 + \eta_2$ identifies the corresponding elasticity for international migration.

5.2 Identification

Estimating η by Ordinary Least Squares (OLS), or by Pseudo-Maximum Likelihood (PPML) to accommodate the large share of zero flows that characterizes international migration corridors, may not identify the causal effect of development aid projects on migration, for two distinct reasons: (i) the endogenous placement of aid projects across space, and (ii) the endogenous disbursement timing of projects once placed. First, aid projects are likely to be disproportionately allocated to districts with weaker economic conditions, precisely those that generate higher out-migration or rural exodus ([European Commission, 2015](#)). Second, the disbursement timing of projects already allocated to specific locations may respond to contemporaneous local shocks ([Kraay, 2012, 2014](#)) that speed up or delay payments while also affecting migration decisions. The placement channel biases the naïve estimates toward zero, attenuating the true causal impact of aid on migration; the timing channel introduces an additional bias of ambiguous sign. The instrumental-variables strategy below addresses both.

To overcome these threats, I exploit two sources of variation, each of which handles one source of endogeneity and which are jointly plausibly exogenous to local migration outcomes. First, I address the endogenous spatial allocation of aid by predicting project placement from district-level vote shares for the incumbent at the time of commitment, with district fixed effects absorbing the cross-sectional level of political alignment so that identification rests on the politically-predicted component of placement. Second, I address the endogeneity of disbursement timing using synthetic disbursement profiles estimated from projects in non-WAEMU (West African Economic and Monetary Union) low- and lower-middle-income countries; these externally estimated profiles generate predicted disbursement flows orthogonal to contemporaneous local shocks affecting migration. Combined, the two ingredients isolate exogenous differences in the spatial allocation and temporal realization of aid, supporting a causal interpretation of the estimated coefficients under the exclusion restrictions discussed below.

Formally, the instrument is defined as:

$$z_{o,\tau} = \frac{1}{L_{o,\tau}} \cdot \sum_{t \in \tau} \left(\sum_s \sum_{k=0}^K \widehat{Com}_{o,s,t-k} \cdot \varphi(s,k) \right), \quad (4)$$

where $\widehat{Com}_{o,s,t-k}$ is the predicted commitment in district o in year $t - k$ using vote shares, which feeds into year- t disbursements according to the sector-specific profile $\varphi(s,k)$ estimated from comparable projects in non-WAEMU (West African Economic and Monetary Union) low- and lower-middle-income countries. The profile $\varphi(s,k)$ gives the average fraction of a project's total disbursements in sector s released k years after commitment. The window τ denotes the commitment period preceding five-year migration window in each census round (2004–2008 and 2014–2018), over which the resulting district-year series are aggregated, and $L_{o,\tau}$ is the district's working-age population, used to express the instrument in per capita terms. The full construction is detailed in Section A.

The concern with using vote shares in favor of the incumbent is that politically aligned districts may differ in patronage exposure, economic structure, or ethnic composition, and these characteristics could shape migration through channels unrelated to aid. To the extent that these are persistent district traits, district fixed effects absorb them. Identification rests on within-district changes in alignment across commitment windows, not on cross-sectional comparisons of politically loyal and disloyal districts. The exclusion restriction would fail only if changes in a district's vote share were correlated with contemporaneous migration shocks, conditional on district and time fixed effects, through a non-aid channel, a much more demanding condition. Two features make this implausible. First, the variation is generated on the supply side of aid, through the incumbent's allocation of visible public goods to politically valuable districts (Jablonski, 2014, Masaki, 2018), rather than through migration-relevant local conditions. Second, the placement signal is combined with a disbursement profile estimated outside Senegal and the whole WAEMU region, so the timing through which aid could affect migration is set by donor routines in unrelated countries. A local migration shock coinciding with political realignment would therefore still have to operate through this externally fixed schedule to bias the estimate.

Importantly, I also use an alternative shift-share instrument (see Section B) in which the shares are constructed from predetermined district-level donor-sector exposure and the shifts from donor-sector disbursement shocks to non-WAEMU low- and lower-middle-income countries, and I check that it delivers similar results. The two instruments share a common external-timing logic, both draw their identifying time variation from donor disbursement behavior outside Senegal and the WAEMU region, but differ in

how they model the cross-sectional placement of aid: the baseline instrument predicts placement from politically-driven commitments, whereas the alternative relies on districts’ predetermined historical donor portfolios. Agreement between the two therefore shows that the results do not hinge on any single way of modeling where aid lands, holding the external-timing component similar across both designs.

6 Main results

6.1 Benchmark results

Table 2 reports benchmark estimates of the migration response to local exposure to development aid projects. Column (1) uses the full sample, while columns (2) and (3) separately examine internal and international migration responses. Panel **A** presents ordinary least squares (OLS) estimates, Panel **B** reports Poisson Pseudo-Maximum Likelihood (PPML) estimates, and Panel **C** reports PPML instrumental-variables estimates, implemented via a control-function approach. Across specifications, higher local exposure to aid projects is associated with lower migration flows. In Panel **A**, a doubling of aid per capita is associated with about a 1.5 percent reduction in five-year migration flows in the full sample and a 2.4 percent reduction in internal migration, while the estimated effect on international migration is small and statistically indistinguishable from zero. Panel **B** shows that these patterns are robust to a PPML specification that accounts for the count nature of migration flows and the presence of zeros. Panel **C** reports the IV estimates. A doubling of aid per capita causally reduces five-year migration flows by 9.4 percent in the full sample and by 9.8 percent for internal migration, while the estimated effect on international migration remains small and statistically indistinguishable from zero. The larger IV magnitudes relative to both OLS and PPML suggest that endogenous project placement and timing attenuate the naïve estimates, consistent with aid being disproportionately allocated to districts experiencing adverse shocks that would otherwise increase migration.

6.2 Robustness checks

To assess the stability of the results, I conduct a wide range of robustness checks, including using an alternative instrument, excluding small corridors, excluding dominant destinations, and explicitly considering the intensive versus extensive margins.

Robustness to an alternative instrument. A natural concern is that the estimated effect of aid on migration may depend on the specific construction of the predicted-

Table 2: Migration responses to development aid

	Full sample	Internal	International
	(1)	(2)	(3)
Panel A. OLS			
Log Aid per capita	-0.023** (0.010)	-0.038* (0.020)	-0.013 (0.015)
Panel B. PPML			
Log Aid per capita	-0.026** (0.010)	-0.023** (0.011)	-0.023* (0.013)
Panel C. PPML-IV			
Log Aid per capita	-0.143*** (0.039)	-0.149*** (0.044)	0.005 (0.046)
Residuals	0.125*** (0.040)	0.132*** (0.044)	-0.032 (0.050)
Origin $\times$ Destination FE	✓	✓	✓
Destination $\times$ Time FE	✓	✓	✓
Origin-Destination-Time controls	✓	✓	✓
Origin-Time controls	✓	✓	✓
Mean dep. (baseline)	161	352	48
K. Paap F-Stat	11.10	10.79	11.58
N	10,644	3,944	6,700

Notes: The dependent variable is the five-year bilateral migration flow from origin district o to destination d over period $\tau:\tau-5$. Column (1) pools all destinations, while columns (2) and (3) separately estimate internal (within-country) and international migration flows. The main regressor is log aid per capita received by the origin district, averaged over the preceding five-year period $\tau-5:\tau-9$. Panel A reports ordinary least squares estimates, Panel B reports the Poisson Pseudo Maximum Likelihood (PPML) estimates, and Panel C reports instrumental-variables estimates. All specifications include origin–destination fixed effects and destination–time fixed effects, as well as time-varying origin-level controls. Each specification controls for origin district number of non-migrants (stayers), climatic shocks, conflict related deaths, and urbanization rates. Standard errors are reported in parentheses and are clustered at the origin–time level. The Kleibergen–Paap F-statistic reports first-stage instrument strength. The mean dependent variable refers to the baseline five-year migration flow (2009-2013). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

disbursement instrument. To assess this, I re-estimate the baseline specification using a standard shift-share instrument that combines pre-period district-level donor-sector exposure shares with donor-sector disbursement shocks observed in non-WAEMU low- and lower-middle-income countries (i.e., outside the eight-member WAEMU bloc, which includes Senegal). Unlike the baseline instrument, which exploits politically predicted project placement interacted with external disbursement profiles, this alternative design relies exclusively on external donor-sector funding shocks and historical local exposure. The two instruments therefore draw on distinct sources of identifying variation.

Table D.5 reports the results. The shift-share IV estimate (column 4) is negative and statistically significant at the one percent level. The coefficient of -0.066 implies that doubling local aid per capita reduces five-year migration flows by approximately 4.5 percent. The instrument remains strong, with a Kleibergen-Paap first-stage F -statistic of 16.56. Although the point estimate is smaller than the corresponding predicted-disbursement estimate (-0.143 in column 3), both instruments yield economically meaningful and statistically significant reductions in migration. The similarity in sign and significance across two conceptually distinct identification strategies strengthens the conclusion that higher local aid exposure reduces migration.

The positive and significant control-function residual (0.058) further suggests that aid is not randomly allocated across space. Instead, aid appears to be disproportionately directed toward districts facing adverse conditions that would otherwise increase out-migration, implying that conventional OLS and PPML estimates are biased toward zero. This pattern is consistent with endogenous aid placement attenuating the estimated effect of aid on migration.

Robustness to the exclusion of small corridors. A natural concern is that the estimates may be driven by origin-destination corridors with very few migrants at baseline, where small absolute changes in flows translate into disproportionately large proportional changes. A corridor starting from three migrants that gains ten in the following period records a 333 percent increase, whereas a corridor of one hundred migrants gaining the same ten experiences only a 10 percent increase. Because such thin corridors mechanically generate noisy and high-variance growth rates, they may exert undue influence on the estimated migration response. To address this, I follow [Shrestha et al. \(2026\)](#) and restrict the sample to “regular” corridors, defined as origin-destination pairs with at least ten migrants observed in each five-year window. Table D.8 reports the results. The estimates closely mirror the baseline: higher local exposure to aid projects reduces overall and internal migration, while leaving international migration unchanged, confirming that the findings are driven by variation in active corridors rather than by

sparse, high-variance cells.

Robustness to the dominant destination. Another concern is that the estimated effects may be driven by the disproportionate role of the capital region, Dakar, comprising the districts of Dakar, Pikine, Rufisque, and Guédiawaye, which together account for approximately 32 percent of internal migration in Senegal. To address this, I exclude those districts from the set of destination locations. Table D.9 reports the corresponding estimates. The results closely mirror the baseline findings. Quantitatively, a doubling of local aid per capita reduces overall migration in corridors excluding the capital region by 8.6 percent and internal migration by 9 percent, while the estimated effect on international migration remains statistically indistinguishable from zero.

Intensive and extensive margins. A non-negligible share of observations are zero: 55 percent overall, with sharp heterogeneity across types (2 percent for internal but 68 percent for international migration). The prevalence of zeros in international corridors raises the concern that the estimated effect is driven primarily by the extensive margin. Following Chen and Roth (2024) and Shrestha et al. (2026), I assess the intensive margin by re-estimating the response while progressively restricting the sample to corridors with increasingly large minimum migrant counts (Table D.6). Excluding zero-migration corridors (column 2) raises the OLS coefficient substantially (Panel A), consistent with OLS being ill-suited to data with many zeros, whereas the PPML and PPML-IV estimates (Panels B and C) are essentially unchanged. The qualitative conclusions hold throughout, with effects becoming slightly larger once attention is restricted to active corridors. This pattern indicates that project aid reduces out-migration primarily along established routes, where migration decisions operate on an active margin (Munshi and Rosenzweig, 2015). The stability of the PPML and PPML-IV estimates confirms that the main results are not driven by the large number of zeros, reinforcing the results. The OLS estimates, by contrast, are more sensitive to the prevalence of zeros.

6.3 Heterogeneity analysis

I examine how the migration-reducing effect of local exposure to aid projects varies across age, skill, gender, and migration margin. As in the baseline, OLS and PPML estimates are small and clustered near zero across all subgroups, while the IV estimates are consistently larger in magnitude. The discussion therefore centers on the IV results.

Skill and gender. Figure E.12 reports the effect by education level, pooling both genders (Panel A) and separately for women (Panel B) and men (Panel C). In the pooled

sample, the deterrent effect is present at all skill levels, with IV estimates of roughly -0.15 for the high-skilled, -0.17 for the medium-skilled, and -0.16 for the low-skilled. The gender split reveals a clearer pattern: the effect is more pronounced for women than for men at every skill level, and for women it is strongest among the low-skilled, where the IV estimate reaches roughly -0.22 . For men the estimates are more uniform across skill groups and somewhat smaller in absolute terms. This is consistent with the gender gap documented in the baseline heterogeneity, and suggests that low-skilled women are the group most responsive to local exposure to development aid.

Age and migration margin. Figure E.13 dis-aggregates the effect by age at migration, separately for total mobility, internal migration, and international migration. For pooled flows (Panel A), the IV estimates are negative and broadly similar across age bands, ranging from roughly -0.10 for the 25–34 group to about -0.13 to -0.16 for the youngest and oldest groups, with no clear monotonic gradient. Panels B and C show that this effect is driven almost entirely by internal migration: the internal estimates (Panel B) are uniformly negative and sizable across all age groups, whereas the international estimates (Panel C) are close to zero and statistically indistinguishable from zero in every cohort. Local exposure to development aid thus deters movement within Senegal across the entire working-age population, while leaving international migration largely unaffected.

Taken together, the heterogeneity analysis points to a deterrent effect that operates primarily on internal mobility, is stable across age, and is somewhat stronger for women, particularly low-skilled women, than for men.

7 Aid projects and irregular international migration

While the benchmark results show that development aid has no robust average effect on international migration, this aggregate masks the margin that motivates the policies in the first place. The instruments behind the “root causes” agenda: the *EU Emergency Trust Fund for Africa*, and the *U.S. root-causes strategy*, for instance, are aimed not at migration in general but at *irregular* migration specifically. Whether aid reaches that target is hard to document based on the analysis of average international flows, since irregular migration is rarely observed directly in standard migration data, and almost never separated from regular flows. This section takes the policy target seriously and asks directly whether local exposure to development aid projects reduces irregular migration.

I draw on a feature of the 2023 Senegalese full-population census introduced in Section 4.2. For the first time, it identifies emigrants residing abroad without regular status at departure. Using this information, I construct origin district-to-destination

country flows of irregular emigration over 2019–2023 and estimate a gravity-style specification relating these flows to district-level project aid:

$$\ln M_{od}^{2019-2023,I} = \delta_{rD} + \delta_d + \beta \ln Aidpc_{o,2014-2018} + X'_{o,2018}\lambda + V'_{od}\lambda + \varepsilon_{od}, \quad (5)$$

where δ_{rD} and δ_d are higher level origin-region $\times$ destination group (neighbors, other Africa, Europe, North America, and the Rest of the World), and destination country fixed effects, and $X_{o,2018}$ is the vector of origin district level controls, V_{od} is the vector of origin-destination controls that includes geographical distance and migrants' network.

Table 3: Migration responses to development aid by migration type

	Full sample	Regular	Irregular
	(1)	(2)	(3)
Panel A. PPML			
Log Aid per capita	-0.036 (0.026)	-0.020 (0.026)	-0.085*** (0.031)
Panel B. PPML-IV			
Log Aid per capita	-0.059* (0.036)	-0.020 (0.039)	-0.147*** (0.037)
Residuals	0.050 (0.058)	-0.001 (0.067)	0.135* (0.073)
Origin region $\times$ Destination group FE	✓	✓	✓
Destination FE	✓	✓	✓
Origin-destination controls	✓	✓	✓
Origin controls	✓	✓	✓
Mean dep. (baseline)	49	39	11
K. Paap F-Stat	17.97	17.97	17.97
N	3,348	3,315	3,082

Notes: This table reports estimates of the effect of exposure to development project on migration flows by legality status. Panel A reports PPML estimates, and Panel B instrumental-variables (IV) estimates. The dependent variable is the five year migration flows from origin district o to destination country d . Standard errors clustered at the origin district level are reported in parentheses. The Kleibergen–Paap F-statistic reports first-stage instrument strength. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3 examines whether the migration response to local aid exposure differs by the legal status of international migrants. The results reveal a clear asymmetry: while neither total nor regular migration responds significantly to aid exposure, irregular migration does. In the PPML specification (Panel A), the coefficient on log aid per capita is small and statistically insignificant for total and regular flows, but negative and significant for irregular migration (-0.085). The IV estimates (Panel B) amplify this pattern, consistent with attenuation arising from the endogenous placement of aid across districts. The IV coefficients remain imprecisely estimated for total and regular flows, whereas the effect on irregular migration is larger in magnitude and statistically significant (-0.147). The IV coefficient of -0.147 implies that doubling local aid per capita reduces irregular migration by about 9.7 percent.

7.1 Robustness checks

I assessed the robustness of the findings on irregular migration using three complementary tests: an alternative shift-share instrumental variables (SSIV) test, a restriction to regular origin-destination corridors with at least ten migrants, and a leave-one-destination-group-out analysis. The results, reported in Tables D.10 and D.11 and Figure E.4, confirm that the estimated deterrent effect of local aid on irregular migration is not due to a particular identification strategy, sample composition, or destination structure. Replacing the baseline predicted-disbursement instrument with a canonical donor-sector shift-share instrument yields similar estimates in terms of sign and magnitude (Table D.10). This indicates that the findings are not driven by the specific construction of the instrument. Restricting the analysis to regular corridors (Table D.11) leaves the main conclusions unchanged: Aid has no detectable effect on regular migration, yet it significantly reduces irregular migration. This suggests that the estimates are not driven by noisy, low-volume corridors. Finally, the leave-one-destination-group-out analysis (Figure E.4) shows that the estimates remain stable when excluding any destination group other than Europe. The effect disappears only when the Europe-bound corridor is omitted, confirming that the deterrent impact of local aid operates almost entirely through migration routes to Europe, remaining robust to the exclusion of all other destination groups.

7.2 Heterogeneity in the irregular-migration response

The average deterrent effect of local exposure to aid projects on irregular migration masks substantial heterogeneity across destinations and demographic groups. Table D.12 shows that the response is concentrated entirely in Europe-bound corridors. While the pooled IV estimate implies that doubling local aid per capita reduces irregular emigration

by roughly 9.7 percent, the effect reaches about 13.8 percent for migration to Europe and is statistically indistinguishable from zero for neighboring countries, other African destinations, North America, and the rest of the world. This pattern is consistent with the descriptive evidence that Europe absorbs the majority of irregular Senegalese emigrants and that high-risk routes through West Africa and the Mediterranean account for most irregular departures. The results therefore suggest that improved local opportunities primarily deter migration along the corridors where migration costs and risks are highest.

Figure E.5 further shows that the deterrent effect is concentrated among the demographic groups most likely to migrate irregularly. The response is strongest among young adults and declines steadily with age (Panel E.5a), mirroring the age profile of irregular migration documented in the descriptive analysis. Splitting the estimates by gender (Panel E.5b) reveals that the effect is driven almost entirely by men, while the coefficients for women are small and statistically indistinguishable from zero across all age groups. Similarly, the education heterogeneity (Panel E.5c) indicates that the response is concentrated among low-skilled individuals, whereas the estimates for higher-skilled groups are much smaller and generally insignificant. Taken together, these findings indicate that local aid primarily affects irregular migration decisions among young, low-skilled men, the group with the highest baseline propensity to migrate irregularly and the most limited access to regular migration channels.

8 Mechanisms

Standard migration models predict two countervailing forces: a constraints channel, whereby higher local incomes or better information relax financial and informational barriers to migration and raise mobility; and an opportunity-cost channel, whereby improved local employment prospects increase the returns to staying and reduce migration incentives. I now test which channel is operative and which dominates.

Information channel. A natural concern is that bilateral aid does not affect migration only through local labor-market conditions but through an *information channel*, whereby donor-funded projects convey destination-specific information and lower the informational frictions facing potential migrants. If this mechanism were operative, migration responses would track donor-specific allocation patterns rather than origin conditions. I test this directly by relating bilateral, donor-specific aid exposure to origin–destination migration flows, instrumenting aid with a leave-one-out predicted-allocation strategy following Shrestha et al. (2026) and Bryan and Morten

(2019) that isolates donor-driven variation in where projects are placed. Across OLS, PPML, and IV specifications, the estimated effect of bilateral aid on migration is small and statistically indistinguishable from zero, and this holds even for the education groups for whom information frictions should be most salient. Results reported in Table C.3 provide no evidence that aid operates through a donor-specific information channel. Full details of the design and results are reported in Appendix C.1.

Opportunity-cost/development channel. A second mechanism operates through local labor markets: to the extent that aid-funded projects expand local employment opportunities, they raise the returns to remaining and may deter migration precisely where they succeed in strengthening local economies. A growing literature documents that development interventions targeting economic infrastructure, productive capacity, and human capital generate sizable and persistent local employment effects (Bandiera et al., 2017, Dinkelman, 2011, Donaldson and Hornbeck, 2016, Duflo, 2001, Faber, 2014), motivating a direct test of whether the migration response to aid runs through improvements in local employment conditions.

I test this channel by relating individual employment status to lagged district-level aid exposure across four Senegalese census rounds (1988, 2002, 2013, 2023), as specified in equation (18), restricting the sample to stayers for whom local labor-market exposure is well defined. Because aid is plausibly targeted toward districts with weaker baseline conditions, I instrument local aid intensity with a synthetic-disbursement instrument (equation (23)) in the spirit of Kraay (2012, 2014), which combines each district’s commitments, fixed at approval, with a typical sector-specific disbursement profile that allocates committed funds across years from the project start. The profile is estimated from the disbursement life-cycle of Senegalese projects of the same sector, so that the instrument isolates variation driven by the mechanical timing of disbursements rather than by contemporaneous local conditions. The instrument is a strong predictor of realized aid (Panel B of Table C.4), with a Kleibergen–Paap F -statistic of 15.72.

The IV estimates (Panel C of Table C.4) confirm and sharpen the OLS results in Panel A, showing a positive and significant effect of local exposure to project aid on overall employment: a 10 percent increase in local project aid per capita raises the probability of being employed by roughly 0.13 percentage points. This aggregate effect masks a compositional shift across employment forms. Aid significantly raises self-employment and reduces unpaid work: the same 10 percent increase leads to a 0.52 percentage point rise in self-employment and a 0.31 percentage point fall in unpaid employment, while wage employment does not respond positively. Aid-funded projects therefore expand local labor demand mainly through self-employment and own-account activity rather

than formal wage jobs, consistent with broader evidence that interventions in developing settings generate employment gains largely through micro-enterprise creation and the rise of income generating activities.

Dis-aggregating by age (Figure E.15), these gains are concentrated among prime-age workers, peaking around ages 20–29 and declining steadily to reach zero by the 60–64 group; dis-aggregating by skill (Figure E.16), they fall almost entirely on the low-skilled, while high- and medium-skilled estimates are indistinguishable from zero. Because low-skilled prime-age workers are precisely the group with the highest baseline propensity to migrate irregularly, the employment response is largest on the margin most relevant to the headline results, corroborating their opportunity-cost interpretation. Full specification, identification details, and robustness are reported in Appendix C.2.

9 Concluding remarks

This paper asked whether local development aid alters migration decisions, and importantly, whether it does so uniformly across the very different margins along which people move. Using Senegal as a setting that combines unusually rich census-based migration data with geo-referenced project-level aid, and an identification strategy that separates the politically-driven *placement* of projects from their externally-determined disbursement *timing*, I show that aid reshapes mobility, but it does so selectively, and aggregating across margins can lead to misleading conclusions.

Three findings stand out. First, local exposure to development projects deters migration, but the effect is concentrated where it is least visible in standard migration data: a doubling of aid per capita reduces overall migration by roughly 9.4 percent, almost entirely through a contraction of internal mobility (−9.8 percent), while leaving total international migration unchanged. Second, this aggregate null on international migration conceals the margin that motivates the “root causes” agenda in the first place. When regular and irregular flows are separated, made possible by the uniquely rich 2023 census’s direct measurement of the legal status at departure, aid leaves regular migration unaffected but reduces irregular migration by about 9.7 percent, an effect that operates almost entirely through Europe-bound corridors and falls on young, low-skilled men: precisely the group with the highest baseline propensity to migrate irregularly and the most constrained access to legal alternatives. Third, the mechanism is an opportunity-cost channel grounded in local labor markets rather than an information channel running through donor ties. Aid raises local employment, concentrated in self-employment among prime-age low-skilled workers, while donor-specific exposure to

aid does nothing to redirect migration toward donor countries. Aid deters migration by improving the value of staying, not by transmitting destination-specific information.

Taken together, these results help reconcile a literature that has long disagreed about whether aid substitutes for or complements migration. The disagreement, this paper suggests, is partly a result of aggregation. Aid can simultaneously suppress internal mobility, leave regular international migration unaffected, and curb irregular migration, outcomes that average to ambiguity when margins are pooled, but that follow a coherent logic once the heterogeneity in cost, risk, and liquidity requirements across margins is taken seriously. The same intervention that anchors a prime-age worker to a local labor market does little to alter the calculus of a family-reunification migrant on a regular visa.

The findings carry a measured policy implication. The architecture of the EU Emergency Trust Fund for Africa and analogous “root causes” instruments rests on the premise that development spending can reduce irregular departures ([European-Commission, 2015](#), [National-Security-Council, 2021](#)). The evidence here is consistent with that premise, but only narrowly. The deterrent effect is real, statistically robust across two distinct identification strategies, and concentrated on exactly the high-risk, Europe-bound flows these policies target.

Several limitations frame these conclusions and point toward future work. The analysis rests on two census waves, so identification exploits medium-run variation and cannot speak to the longer-horizon mobility-transition dynamics through which aid might eventually *raise* international migration as incomes grow ([Clemens, 2014](#), [Clemens and Mendola, 2024](#), [Dao et al., 2018](#)). While the results survive an extensive battery of robustness checks, an alternative shift–share instrument, corridor restrictions, dominant-destination exclusions, and leave-one-group-out tests, the external validity exercise, though reassuring, establishes representativeness rather than transportability. Extending this framework to other settings with comparable census-based measurement of irregular status, tracing the dynamics of the employment response over a longer panel, and distinguishing the productive content of aid projects more finely would each sharpen the picture. The broader methodological contribution, that census data can credibly measure irregular migration, validated here against independent UNHCR and Frontex benchmarks, suggests that the margin most central to policy debate, and most rarely observed, need not remain a black box.

References

- Abel, G. J. and Cohen, J. E. (2022). Bilateral international migration flow estimates updated and refined by sex. *Scientific Data*, 9(1):173.
- Bandiera, O., Burgess, R., Das, N., Gulesci, S., Rasul, I., and Sulaiman, M. (2017). Labor markets and poverty in village economies. *The Quarterly Journal of Economics*, 132(2):811–870.
- Belloc, F. (2015). International economic assistance and migration: The case of sub-saharan countries. *International Migration*, 53(1):187–201.
- Berthélemy, J.-C., Beuran, M., and Maurel, M. (2009). Aid and migration: Substitutes or complements? *World Development*, 37(10):1589–1599.
- Bocquier, P., Cha’ngom, N., Machado, J., and Docquier, F. (2023). The within-country distribution of brain drain and brain gain effects: A case study on senegal. *Manuscript*.
- Bomprezzi, P., Dreher, A., Fuchs, A., Hailer, T., Kammerlander, A., Kaplan, L., Marchesi, S., Masi, T., Perlik, K., and Robert, C. (2025). Wedded to prosperity? Spousal favoritism in foreign aid and regional development. *CEPR Discussion Paper 18878*.
- Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213.
- Bryan, G. and Morten, M. (2019). The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy*, 127(5):2229–2268.
- Chen, J. and Roth, J. (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, 139(2):891–936.
- Clemens, M. A. (2014). Does development reduce migration? In *International handbook on migration and economic development*. Edward Elgar Publishing.
- Clemens, M. A. and Mendola, M. (2024). Migration from developing countries: Selection, income elasticity, and simpson’s paradox. *Journal of Development Economics*, 171:103359.
- Dao, T. H., Docquier, F., Parsons, C., and Peri, G. (2018). Migration and development: Dissecting the anatomy of the mobility transition. *Journal of Development Economics*, 132:88–101.

- Dinkelman, T. (2011). The effects of rural electrification on employment: New evidence from south africa. *American Economic Review*, 101(7):3078–3108.
- Donaldson, D. and Hornbeck, R. (2016). Railroads and american economic growth: A “market access” approach. *The Quarterly Journal of Economics*, 131(2):799–858.
- Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., and Tierney, M. J. (2019a). African leaders and the geography of China’s foreign assistance. *Journal of Development Economics*, 140:44–71.
- Dreher, A., Fuchs, A., and Langlotz, S. (2019b). The effects of foreign aid on refugee flows. *European Economic Review*, 112:127–147.
- Duflo, E. (2001). Schooling and labor market consequences of school construction in indonesia: Evidence from an unusual policy experiment. *American economic review*, 91(4):795–813.
- European-Commission (2015). Governance and procedure: Eu trust fund for africa for stability and addressing root causes of irregular migration and displaced persons. *European Union*.
- Faber, B. (2014). Trade integration, market size, and industrialization: evidence from china’s national trunk highway system. *Review of Economic Studies*, 81(3):1046–1070.
- Fuchs, A., Gröger, A., Heidland, T., and Wellner, L. (2023). The effect of foreign aid on migration: Global micro evidence from world bank projects. Technical report, Kiel Working Paper.
- Gamso, J. and Yuldashev, F. (2018a). Does rural development aid reduce international migration? *World Development*, 110:268–282.
- Gamso, J. and Yuldashev, F. (2018b). Targeted foreign aid and international migration: Is development-promotion an effective immigration policy? *International Studies Quarterly*, 62(4):809–820.
- Jablonski, R. S. (2014). How aid targets votes: the impact of electoral incentives on foreign aid distribution. *World Politics*, 66(2):293–330.
- Kraay, A. (2012). How large is the government spending multiplier? evidence from world bank lending. *The Quarterly Journal of Economics*, 127(2):829–887.

- Kraay, A. (2014). Government spending multipliers in developing countries: evidence from lending by official creditors. *American Economic Journal: Macroeconomics*, 6(4):170–208.
- Lanati, M., Sanfilippo, M., and Santi, F. (2023). Aid and internal migration in malawi. *World Development*, 162:106134.
- Lanati, M. and Thiele, R. (2018a). Foreign assistance and migration choices: Disentangling the channels. *Economics letters*, 172:148–151.
- Lanati, M. and Thiele, R. (2018b). The impact of foreign aid on migration revisited. *World development*, 111:59–74.
- Marchal, L., Naiditch, C., and Simsek, B. (2025). Aid as a migration policy tool. *The World Economy*.
- Marchiori, L., Maystadt, J.-F., and Schumacher, I. (2012). The impact of weather anomalies on migration in sub-saharan africa. *Journal of Environmental Economics and Management*, 63(3):355–374.
- Masaki, T. (2018). The political economy of aid allocation in africa: Evidence from zambia. *African Studies Review*, 61(1):55–82.
- Menard, A. and Gary, A. (2018). Aid, trade and migration: How do bilateral flows interact? *The World Economy*, 41(2):431–456.
- Munshi, K. and Rosenzweig, M. (2015). Insiders and outsiders: local ethnic politics and public goods provision. Technical report, National Bureau of Economic Research.
- Murat, M. (2020). Foreign aid, bilateral asylum immigration and development. *Journal of Population Economics*, 33(1):79–114.
- National-Security-Council (2021). U.s. strategy for addressing the root causes of migration in central america. *Washington, DC: White House.*, [Available at](#):4.
- Saasa, O. S. (2010). Political economy of budget support in zambia. *Lusaka: Premier Consult.*
- Shrestha, M., Kaila, H., Cha’ngom, N., and Adhikari, S. (2026). Do bilateral labor agreements increase migration? global evidence from 1960 to 2020. *Journal of Development Economics*, 179:103673.

UNDP (2009). Human development report 2009. overcoming barriers: Human mobility and development. *Overcoming Barriers: Human Mobility and Development (October 5, 2009)*. UNDP-HDRO Human Development Reports.

Appendices

A Instrument Construction

This section describes the construction of the predicted-disbursement instrument for local aid exposure. The instrument isolates the component of aid that flows to a district for politically motivated allocation reasons, combined with an externally estimated disbursement-timing profile, so that neither the cross-sectional placement of projects nor the timing of their disbursements reflects local economic conditions in Senegal.

A.1 Predicted Project Placement

A key feature of aid allocation is that recipient governments often retain substantial influence over where donor-funded projects are implemented. In many cases, donors delegate project-site selection to government ministries, while weak bureaucratic capacity further concentrates decision-making power within the executive branch (Jablonski, 2014, Saasa, 2010). Bilateral donors likewise typically rely on project requests submitted by recipient governments, which often specify target locations (Masaki, 2018). As a result, political leaders have considerable scope to shape the spatial distribution of aid. Since development projects are visible public goods that can generate electoral support, vote-maximizing incumbents have incentives to direct projects toward areas that are politically valuable. This political-allocation mechanism provides a rationale for using vote shares for the incumbent as a predictor of project placement.

To isolate the component of project allocation driven by political considerations, I predict project commitments using variation in districts' political proximity to the incumbent leader at the time of project commitment. Specifically, I estimate

$$\begin{aligned} \ln Com_{i,o,s,n,t} = & \beta_0 + \beta_1 \ln Vote_{o,t} + \beta_2 (\ln Vote_{o,t} \times State_i) \\ & + \beta_3 (\ln Vote_{o,t} \times Post_t) + \beta_4 (\ln Vote_{o,t} \times Pre_t) \\ & + \gamma_1 Lon_i + \gamma_2 Lat_i + \delta_o + \delta_t + \delta_s + \delta_n + \varepsilon_{i,o,s,n,t}, \end{aligned} \quad (6)$$

where $Com_{i,o,s,n,t}$ denotes the commitment amount of project i implemented in district o , sector s , by donor n , in year t . The variable $Vote_{o,t}$ is the presidential vote share obtained by the incumbent leader in district o , $State_i$ is an indicator equal to one when the implementing agency is a government entity, and Pre_t and $Post_t$ indicate whether the commitment occurred before or after a presidential election. The specification further includes district, year, sector, and donor fixed effects (δ_o , δ_t , δ_s , and δ_n), as well as geographic controls for project longitude and latitude.

The fitted value from equation (6), $\widehat{Com}_{i,o,s,n,t}$, captures the component of project commitments explained by political alignment with the ruling coalition, net of donor, sector, location, and time-specific factors.

A.2 Synthetic Disbursement Profile

To address the potential endogeneity of disbursement timing, I construct donor-sector specific synthetic disbursement profiles using project-level data from low- and lower-middle-income countries, excluding the WAEMU (West African Economic and Monetary Union) countries to which Senegal belongs. For each donor-sector s and project age k (years since project start), the profile $\varphi(s, k)$ measures the average fraction of total project disbursements released at age k :

$$\varphi(s, k) = \frac{\sum_c \sum_{p \in s} Disb_{c,p,k}}{\sum_c \sum_{p \in s} \sum_{j=0}^K Disb_{c,p,j}}, \quad \sum_{k=0}^K \varphi(s, k) = 1, \quad (7)$$

where $Disb_{c,p,k}$ denotes the amount disbursed to project p in country c during its k -th year of implementation. By construction, $\varphi(s, k)$ summarizes the typical implementation trajectory of projects within a sector and is estimated entirely from countries outside Senegal and WAEMU.

A.3 Constructing Predicted Aid Flows

I first aggregate predicted commitments to the district $\times$ sector $\times$ commitment-year level. I then distribute each predicted commitment over future years according to the corresponding synthetic disbursement profile. Predicted disbursements received by district o in calendar year t are therefore given by

$$\widehat{D}_{o,t} = \sum_s \sum_{k=0}^K \widehat{Com}_{o,s,t-k} \varphi(s, k), \quad (8)$$

where commitments made in year $t - k$ contribute to current disbursements according to the sector-specific share $\varphi(s, k)$. The resulting district-year series is then aggregated to the two time windows of interest. Let τ denote a migration window. Calendar years are mapped into migration windows as

$$\tau = \begin{cases} 2013, & \text{if } t \in [2004, 2008], \\ 2023, & \text{if } t \in [2014, 2018]. \end{cases} \quad (9)$$

For each district and migration window, the instrument is defined as cumulative predicted aid per worker:

$$z_{o,\tau} = \frac{\sum_{t \in \tau} \widehat{D}_{o,t}}{L_{o,\tau}}, \quad (10)$$

where $L_{o,\tau}$ denotes the district labor force.

Table A.1: Vote shares and the spatial placement of aid projects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Vote share	2.343** (0.974)	2.345** (0.974)	1.846* (1.009)	2.725*** (0.943)	2.291** (0.982)	2.291** (0.982)	2.292** (0.983)
Vote share $\times$ Government agency		-0.585 (1.130)				-0.513 (1.135)	-0.520 (1.127)
Vote share $\times$ Pre			2.978 (1.830)		3.413* (1.814)	3.409* (1.816)	3.401* (1.812)
Vote share $\times$ Post				-1.468 (1.626)	-1.988 (1.631)	-1.982 (1.628)	-1.977 (1.625)
Geographical longitude							0.067 (0.532)
Geographical latitude							0.313 (0.767)
District FE	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
Sector FE	✓	✓	✓	✓	✓	✓	✓
Donor FE	✓	✓	✓	✓	✓	✓	✓
Adj. R^2	0.189	0.189	0.190	0.190	0.190	0.190	0.190
Observations	4,008	4,008	4,008	4,008	4,008	4,008	4,008

Notes: This table reports estimates from the project-placement regression in equation (6). The dependent variable is the log commitment amount of an individual aid project, observed at the project $\times$ district $\times$ sector $\times$ donor $\times$ year level. *Vote share* is the presidential vote share obtained by the incumbent leader in the district. *State* is an indicator equal to one when the implementing agency is a government entity; *Pre* and *Post* indicate commitments made in the years before and after a presidential election, respectively. All specifications include district, year, sector, and donor fixed effects, and columns (7) adds geographic controls for project longitude and latitude. The fitted values from column (7) are used to construct the predicted component of project commitments, $\widehat{Com}$, underlying the instrument. Standard errors, clustered at the district level, are reported in parentheses. R^2 and the number of observations are reported at the foot of the table. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Alternative Instrument: standard Shift–Share

As a robustness check on the predicted-disbursement instrument, I construct a standard shift-share (Bartik) instrument that combines predetermined within-Senegal exposure shares (district level) with externally determined donor-level shifters. The appeal of this design is that it relies on a transparent source of identifying variation: each district’s historical reliance on particular donors, interacted with the subsequent expansion or

contraction of those donors' aid envelopes in countries other than Senegal. The two ingredients are constructed from non-overlapping samples, so that neither the shares nor the shifters reflect contemporaneous economic conditions in the Senegalese districts.

Predetermined Exposure Shares. The exposure shares capture how strongly each district relied on a given donor before the start of the analysis period. Using project-level data for Senegal, I retain only commitments approved on or before 2004, which precede the first migration window. For each project I take a single commitment per district, attributed to its approval year, and I aggregate commitments to the district $\times$ donor $\times$ sector level. Because the pre-2005 project data are sparse at the district $\times$ sector $\times$ donor level, with many donor cells empty or supported by only a handful of projects, I group donors into a limited set of major bilateral and multilateral donors (Austria, Belgium, Germany, Spain, France, Italy, and the World Bank Group), with all remaining donors collapsed into a residual “Other” category. This grouping ensures that the exposure shares are estimated from a sufficient number of underlying projects rather than from isolated observations.

Let $Commit_{o,n,s}^{\leq 2004}$ denote pre-2005 cumulative commitments to district o from donor n in sector s . The baseline share is

$$\omega_{o,n,s} = \frac{Commit_{o,n,s}^{\leq 2004}}{\sum_{n'} Commit_{o,n',s}^{\leq 2004}}, \quad \sum_j \omega_{o,n,s} = 1, \quad (11)$$

so that $\omega_{o,n,s}$ measures donor n 's share of district o 's historical aid portfolio within sector s . Because these shares are fixed at their pre-2005 values, they are predetermined with respect to aid flows and migration outcomes realized over the two analysis windows.

B.0.1 External Donor Shifters

The shifters measure the growth of each donor's aid envelope in a sample that excludes Senegal and its monetary union, removing any mechanical link between the shifter and local Senegalese conditions. I draw on project-level data for low- and lower-middle-income countries (LLMICs), excluding all WAEMU (West African Economic and Monetary Union) members, the bloc to which Senegal belongs. I apply the same project- and donor-level filters used in the Senegal sample.

For each donor (after the same grouping into major donors plus an “Other” category), sector, and window τ , the shifter $g_{j,s,\tau}^{-\text{WAEMU}}$ is the total external commitment volume, where commitments are attributed to a window according to project start year:

$$\tau = \begin{cases} 2013, & \text{if } start \in [2004, 2008], \\ 2023, & \text{if } start \in [2014, 2018]. \end{cases} \quad (12)$$

Forming the Instrument. The shift-share instrument for district o in window τ is the inner product of the predetermined shares and the external shifters, summed across donors within each sector and then aggregated:

$$z_{o,\tau} = \frac{\sum_s \sum_n \omega_{o,n,s} g_{j,s,\tau}^{-\text{WAEMU}}}{L_{o,\tau}}. \quad (13)$$

A district is therefore predicted to receive larger aid flows when donors on which it historically relied expand their overall aid envelopes elsewhere.

Identifying Assumptions. Identification in this design rests on the exogeneity of the shifters rather than the shares, following the logic of [Borusyak et al. \(2022\)](#). Even if a district’s historical donor portfolio is correlated with persistent local characteristics that also shape migration, the instrument remains valid as long as the donor-level shifters are as-good-as-randomly assigned conditional on their expected values, so that the realized fluctuations in each donor’s aid envelope are uncorrelated with the unobserved determinants of Senegalese migration.

Two features of the construction make the required shift exogeneity plausible. First, the shifters are computed entirely outside Senegal and the WAEMU bloc, from the broader pool of low- and lower-middle-income aid recipients. Each donor’s window-specific envelope is therefore driven by its global programming priorities, domestic budget cycles, and commitments to unrelated recipients, none of which respond to labor-market or migration shocks in a given Senegalese district. The leave-WAEMU-out construction further removes any common regional component that might link the shifter to conditions in Senegal through shared shocks to neighboring economies. Second, the variation that identifies the second stage is differential: because districts held different pre-2005 donor portfolios, a common movement in a donor’s envelope is projected onto districts in proportion to their predetermined exposure, generating district-level variation in predicted aid that is plausibly orthogonal to local shocks. The exclusion restriction requires only that these donor-wide funding swings affect migration through their effect on local aid receipts, and not through some other channel correlated with a district’s historical donor mix.

I treat this instrument as a robustness check rather than the headline specification because the effective concentration of aid across sectors and donors within districts is high, leaving relatively few donors with meaningful weight in any district. This concentration places a heavy burden on a small number of shifters and makes fine-grained exposure shares more demanding to defend than the predicted-disbursement instrument used in the main analysis, even though the underlying source of identifying variation: exogenous, externally determined donor funding cycles, is shared across the two designs.

B.1 First stage

The first-stage specification is

$$\ln(Aidpc_{o,\tau:\tau-5}) = \alpha_o + \alpha_\tau + \pi z_{o,\tau:\tau-5} + X'_{o,\tau-5}\Gamma + \zeta \ln N_{od,\tau-5}^{tw} + u_{od\tau}, \quad (14)$$

where $Aidpc_{o,\tau:\tau-5}$ denotes cumulative aid per capita in district o over the five-year period preceding migration window τ , $X_{o,\tau-5}$ is a vector of predetermined district characteristics, and $N_{od,\tau-5}^{tw}$ captures lagged migrant network effects at the origin–destination level. District and time fixed effects (α_o and α_τ) are included throughout.

The regression is estimated on the stacked origin–destination–period dataset because migrant networks vary at the corridor level. Importantly, however, both the endogenous regressor and the instrument vary only at the district-period level. Consequently, identification comes exclusively from differences across districts in exposure to predicted aid flows over time, while corridor-level networks enter solely as control variables.

C Mechanisms

C.1 Information channel

This section tests whether bilateral development aid affects migration primarily through an *information channel*, whereby aid projects convey destination-specific information and reduce informational frictions faced by potential migrants. The identifying implication of this mechanism is that migration responses should be driven by donor-specific aid allocation patterns rather than by local economic conditions in origin districts.

Let $M_{od,\tau:\tau-5}$ denote five-year migration flows from origin district o to destination country d at between $\tau - 5$ and τ . The baseline estimating equation is:

$$\log M_{od,\tau:\tau-5} = \beta Aid_{od,\tau-5:\tau-9} + \mu_{o\tau} + \mu_{d\tau} + \mu_{od} + \varepsilon_{od,\tau}, \quad (15)$$

where $Aid_{od,\tau-5:\tau-9}$ is a dummy variable that takes the value 1 if the d implemented aid projects in o between $\tau - 5$ and $\tau - 9$. $\mu_{o\tau}$, $\mu_{d\tau}$, and μ_{od} denote origin-year, destination-year, and origin-destination fixed effects, respectively. Equation (15) is estimated using OLS and Poisson pseudo-maximum likelihood (PPML), with PPML accommodating zero migration flows and heteroskedasticity.

To isolate variation in bilateral aid that plausibly operates through information transmission rather than local conditions, I instrument $Aid_{od\tau}$ using a leave-one-out strategy following [Shrestha et al. \(2026\)](#) and [Bryan and Morten \(2019\)](#). The instrument is constructed as the predicted probability that district o receives aid from donor d ,

Table B.2: First stage results

	Origin-destination-year		Origin-year	
	(1)	(2)	(3)	(4)
Placement-adjusted synthetic disbursement	0.702*** (0.211)		0.652*** (0.206)	
Shift-Share IV (SSIV)		-4.480*** (1.101)		-4.621*** (1.304)
Origin $\times$ Destination FE	✓	✓	-	-
Origin-Destination-Time controls	✓	✓	-	-
Destination $\times$ Time FE	✓	✓	-	-
Origin FE	-	-	✓	✓
Time FE	-	-	✓	✓
Origin-Time controls	✓	✓	✓	✓
K. Paap F-Stat	11.10	16.56	10.11	12.73
N	10,644	9,708	90	80

Notes: This table reports first-stage estimates of local aid exposure on the excluded instruments. The dependent variable is cumulative predicted aid per worker in the origin district. Columns (1)–(2) are estimated at the origin $\times$ destination $\times$ year level, with origin $\times$ destination, destination $\times$ time, and origin fixed effects; columns (3)–(4) collapse to the origin $\times$ year level, with time fixed effects and origin-time controls. *Placement-adjusted synthetic disbursement* is the predicted-disbursement instrument used in the main analysis, constructed from politically predicted project commitments distributed through externally estimated, sector-specific disbursement profiles. *Shift-Share IV (SSIV)* is the canonical Bartik instrument formed by interacting pre-2005 district $\times$ donor exposure shares with leave-WAEMU-out donor funding shifters. *K. Paap F-Stat* is the Kleibergen–Paap rank Wald F -statistic for weak identification. Standard errors, clustered at the origin-district level, are reported in parentheses.

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

computed using donor d 's aid allocations to all other Senegalese districts, excluding district o . Formally,

$$Z_{od,\tau} = \Pr(\text{Aid}_{od\tau} > 0 \mid \text{Aid}_{-od\tau}), \quad (16)$$

where $\text{Aid}_{-o,d,\tau}$ denotes aid from donor d to all districts other than o . The first-stage relationship is given by:

$$\text{Aid}_{od\tau} = \pi Z_{od\tau} + \mu_{o\tau} + \mu_{d\tau} + \mu_{od} + u_{od\tau}. \quad (17)$$

This identification strategy directly targets the information channel. Conditional on fixed effects, variation in $Z_{od\tau}$ reflects donor-level allocation behavior rather than shocks to district o . Such variation can affect migration only if aid projects transmit information about destination-specific opportunities.

Table C.3 reports estimates of the effect of bilateral development project aid on migration flows. Across OLS and PPML specifications (Panels A and B), the coefficients on bilateral aid are not statistically different from zero. These results suggest that there is no evidence that bilateral aid increases migration through an information channel. Panel C reports instrumental-variables estimates based on a leave-one-out strategy that isolates donor-driven variation in aid allocations. The IV estimates are likewise statistically indistinguishable from zero, and the Kleibergen–Paap first-stage statistics indicate strong instrument relevance. The absence of a statistically significant response to predicted bilateral aid implies no evidence that donor-specific aid exposure affects migration decisions. This conclusion holds across education groups, including those for whom information frictions are expected to be most salient.

C.2 Opportunity-cost/development channel

To the extent that aid-funded projects operate by expanding local employment opportunities, the development channel is likely to dominate. This mechanism implies that development aid may deter migration precisely when it succeeds in strengthening local labor markets, providing a natural motivation to test whether the migration response to aid operates through improvements in local employment conditions.

A growing literature documents that development interventions, particularly those targeting economic infrastructure, productive capacity, and human capital, generate sizable and persistent effects on local employment. Investments in education and basic infrastructure raise labor productivity and expand local labor demand, both directly through construction and service provision and indirectly through firm entry, market integration, and occupational upgrading (Dinkelman, 2011, Donaldson and Hornbeck,

Table C.3: Migration responses to bilateral development aid

	All	By education level		
	All	High skilled	Medium skilled	Low skilled
	(1)	(2)	(3)	(4)
Panel A. OLS				
Bilateral aid project (0/1)	-0.037 (0.072)	-0.033 (0.078)	-0.100 (0.079)	-0.061 (0.080)
Panel B. PPML				
Bilateral aid project (0/1)	0.000 (0.043)	0.012 (0.035)	-0.053 (0.055)	0.014 (0.048)
Panel C. IV estimates				
Bilateral aid project (0/1)	-0.133 (0.100)	-0.113 (0.107)	-0.106 (0.107)	-0.053 (0.109)
Origin $\times$ Year FE	✓	✓	✓	✓
Origin $\times$ Destination FE	✓	✓	✓	✓
Destination $\times$ Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
K. Paap F-Stat	27.42	27.42	27.42	27.42
N	1,038	1,038	1,038	1,038

Notes: This table reports estimates of the effect of bilateral development project aid on migration flows. Panel A reports ordinary least squares (OLS) estimates, Panel B Poisson pseudo-maximum likelihood (PPML) estimates, and Panel C instrumental-variables (IV) estimates. The dependent variable is the logarithm (OLS and IV) or level (PPML) of five-year migration flows from origin to destination, disaggregated by education level. In the IV specifications, bilateral project aid is instrumented using a leave-one-out strategy following [Shrestha et al. \(2026\)](#) and [Bryan and Morten \(2019\)](#). The instrument is constructed as the predicted probability that district o receives aid from donor country d , computed using aid allocations from donor d to all other Senegalese districts, excluding district o . All specifications include origin $\times$ year, destination $\times$ year, and origin $\times$ destination fixed effects, as well as origin-destination-time controls capturing migrant networks. Robust standard errors clustered at the origin-destination level are reported in parentheses. The Kleibergen-Paap F-statistic reports first-stage instrument strength. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

2016, Duflo, 2001, Faber, 2014). Evidence from multi-component development programs further suggests that relaxing productive constraints can induce durable shifts in labor supply and earnings, leading to sustained engagement in local labor markets (Bandiera et al., 2017).

Econometric specification. To test whether this employment channel is at play in Senegal, I combine historical Senegalese census data (1988, 2002, 2013, 2023) with GODAD aid projects level data based on locations and project-level disbursements and check whether employment probabilities respond to them. The estimating equation relates individual employment status to lagged district-level aid exposure across four census rounds.⁴ For individual i residing in district o at census year $\tau \in \{1988, 2002, 2013, 2023\}$:

$$y_{i\tau} = \alpha_o + \alpha_\tau + \beta \ln(\text{Aidpc}_{o,\tau-5:\tau-9}) + X_i' \gamma + V_{o,\tau}' \psi + \varepsilon_{i\tau}, \quad (18)$$

where $y_{i\tau}$ is an indicator equal to one if individual i is employed at the time of the census. District fixed effects α_o absorb all time-invariant differences across districts, including geographic endowments, historical economic structure, that could jointly determine the spatial allocation of aid projects and employment levels. Census-round fixed effects, α_τ , absorb aggregate shocks common to all districts in a given round, including national business cycle conditions and economy-wide policy changes. The vector X_i includes individual-level controls for sex, age group, and education levels. The vector $V_{o\tau}$ includes district-level controls such as population size, population density, and the share of working age population. Aid exposure $\text{Aidpc}_{o,\tau-5:\tau-9}$ is total disbursements per capita received by district o in the five years immediately preceding the migration windows of each census round, ensuring that aid predates the measured employment outcome. Standard errors are clustered at the district-census round level.

⁴The employment channel analysis is conducted on the subsample of stayers, individuals residing in the same district at the census date as five years prior. This restriction is motivated by the logic of the opportunity-cost channel itself. Aid projects raise the returns to remaining by expanding local labor demand; the relevant population for testing this mechanism is therefore those who are actually exposed to the local labor market, not those who have already left. Including migrants in the estimation would contaminate the employment measure in two ways. First, migrants are no longer exposed to local employment conditions created by aid projects at the time of the census, introducing measurement error in the treatment. Second, migration and non-employment are jointly determined, individuals who migrate do so partly because local employment opportunities are insufficient, creating a mechanical negative correlation between local employment rates and the migrant share of the sample that is unrelated to aid. Restricting to stayers resolves both problems and delivers a clean estimand: the effect of local aid exposure on the employment probability of individuals who remained in their district of origin, which is precisely the margin that governs the migration decision under the opportunity-cost hypothesis.

Threats to Identification. OLS estimates of β face two threats. First, aid is endogenously allocated: donors and implementing agencies likely target projects toward locations with weaker economic conditions, lower employment rates, and greater development needs, themselves correlated with local labor market outcomes. Moreover, conditional on receiving aid, the timing and intensity of disbursements may respond to contemporaneous local shocks (Kraay, 2012, 2014). Both channels imply that OLS estimates of the aid-employment nexus are likely biased, with the direction of the bias depending on whether aid targets locations with worse or better baseline conditions. If aid projects are disproportionately directed toward districts with structurally weaker labor markets, OLS will bias the true employment effect downward. A credible source of exogenous variation is needed to recover β .

C.3 Identification strategy

The two instruments used in the main text: (i) the vote share-predicted placement instrument and (ii) the predetermined donor-sector shift-share, both require historical district-level information that is unavailable for the period underlying this mechanism analysis. I observe neither historical presidential vote shares nor predetermined donor-sector aid shares at the district level, so neither strategy can be implemented here. I therefore rely, in this section only, on a synthetic disbursement instrument in the spirit of Kraay (2012, 2014), which isolates the component of local aid driven by the mechanical timing of disbursements rather than by contemporaneous local conditions.

The disbursement profile. I first estimate a typical disbursement profile $\phi_{s,a}$ giving the share of a project’s committed amount that is paid out at each project age a , where age is measured in years since the project’s start year. For a project i in sector s , let $A_{i,t}$ denote disbursements in calendar year t , let start_i be its start year so that $a = t - \text{start}_i$, and let com_i be its commitment fixed at approval. Define the project-age disbursement rate

$$r_{i,a} = \frac{A_{i,\text{start}_i+a}}{\text{com}_i}. \quad (19)$$

I discard disbursements occurring in the project’s start year ($a = 0$) and use only ages $a \in \{1, \dots, \bar{a}\}$, with $\bar{a} = 10$. The raw profile is the average rate within each sector-age cell,

$$\tilde{\phi}_{s,a} = \frac{1}{|\mathcal{I}_s|} \sum_{i \in \mathcal{I}_s} r_{i,a}, \quad a = 1, \dots, \bar{a}, \quad (20)$$

where $\mathcal{I}_s$ is the set of projects in sector s . I normalize the profile to sum to one within each sector,

$$\phi_{s,a} = \frac{\tilde{\phi}_{s,a}}{\sum_{a'=1}^{\bar{a}} \tilde{\phi}_{s,a'}}. \quad (21)$$

The profile is therefore a function of project age and sector alone: it summarizes the average disbursement life-cycle common to projects of a given type and carries no information specific to any individual district.

The instrument. I combine each project-district's commitment $\text{com}_{i,o}$, fixed at approval, with the externally estimated profile to obtain a predicted local disbursement for project i in district o at age a ,

$$\widehat{\text{aid}}_{i,o,t} = \text{com}_{i,o} \cdot \phi_{s(i),a(i,t)}, \quad a(i,t) = t - \text{start}_i. \quad (22)$$

Aggregating across all projects active in district o and summing over the years t falling in exposure window τ yields the district-window instrument

$$Z_{o,\tau} = \sum_{t \in \tau} \sum_{i \in \mathcal{P}_{o,\tau}} \text{com}_{i,o} \cdot \phi_{s(i),a(i,t)}, \quad (23)$$

where $\mathcal{P}_{o,\tau}$ is the set of projects in district o that contribute to window τ . The variation in $Z_{d,\tau}$ thus comes from two predetermined sources only: the cross-sectional allocation of commitments across districts, and the mechanical schedule $\phi_{s,a}$ by which committed funds are disbursed over project age.

Exclusion restriction. The instrument enters the first stage of the local exposure to aid project measure $\text{aid}_{o,\tau}$,

$$\text{aid}_{o,\tau} = \pi Z_{o,\tau} + \lambda_o + \lambda_\tau + u_{o,\tau}, \quad (24)$$

where λ_o and λ_τ are district and census-wave fixed effects, included throughout. Identification therefore rests on within-district variation across waves: time-invariant district characteristics (geography, distance to the coast or to borders, persistent institutional or ethnic composition) and any common shock affecting all districts in a given wave are absorbed by λ_o and λ_τ and cannot drive the estimate. The exclusion restriction requires that $Z_{o,\tau}$ affect migration only through realized local aid $\text{aid}_{o,\tau}$, conditional on these fixed effects.

Two features of the construction support this restriction. First, because $\phi_{s,a}$ is estimated from sector-by-age cells, district-specific shocks that accelerate or delay actual disbursement, local implementation capacity, political interference, conflict, weather, do not enter $Z_{o,\tau}$; the instrument uses only the typical life-cycle, purged of district-specific

realization, and the start-year payment most exposed to launch-time conditions is excluded altogether. Second, because the instrument is anchored to commitments fixed at project approval, predicted exposure is detached from contemporaneous local determinants of migration. The identifying assumption is therefore that, conditional on district and wave fixed effects, a district’s predicted disbursement $Z_{o,\tau}$ is uncorrelated with unobserved within-district drivers of migration, except through the aid flows it mechanically predicts.

Limits. This assumption is stronger than for the core instruments. The cross-sectional commitments that anchor $Z_{o,\tau}$ are not randomly assigned, and the district fixed effects λ_o absorb only their time-invariant component: if the *timing* of large commitments across waves is itself correlated with evolving local migration pressures, the exclusion restriction would be threatened. The instrument addresses the timing of disbursement, not the selection of where aid is committed. I therefore treat these estimates as evidence on mechanisms rather than as primary causal estimates, and read them alongside, the core results.

First-Stage Results. Panel B of Table C.4 reports the first-stage estimates. It confirms that the synthetic disbursement is a strong predictor of local aid intensity. The magnitude of the first-stage coefficient implies that a 10 percent increase in the predicted aid is associated with a 2.52 percent increase in realized local aid per capita. The Kleibergen–Paap F-statistic is 15.72, comfortably exceeds the conventional weak-instrument threshold of 10.

IV estimates. The IV estimates reported in Panel C of Table C.4 confirm and sharpen the OLS results reported in Panel A, revealing a positive and statistically significant effect of local aid exposure on overall employment probability. A 10 percent increase in local project aid per capita raises the probability of being employed by approximately 0.13 percentage points. Dis-aggregating by employment type reveals that this aggregate effect masks meaningful heterogeneity across employment forms. Aid exposure significantly increases self-employment, while simultaneously reducing unpaid employment. Specifically, a 10 percent increase in project aid per capita leads to 0.52 percentage point increase in self employment and 0.31 percentage point decrease in unpaid employment. Taken together, these results suggest that aid-funded projects expand local labor demand primarily through self-employment and own-account activities rather than through formal wage jobs, a pattern consistent with the broader evidence that development interventions in developing settings generate employment gains largely

Table C.4: Effect of Local Aid Exposure on Employment

	Any employment	By type of employment		
	Any	Wage	Self	Unpaid
	(1)	(2)	(3)	(4)
<i>Panel A. OLS</i>				
Log Project Aid p.c.	0.008** (0.004)	-0.008** (0.004)	0.021*** (0.006)	-0.005 (0.004)
<i>Panel B. First stage</i> (Dep. var: Log Project Aid p.c.)				
Synthetic disbursement	0.252*** (0.064)			
<i>Panel C. IV</i>				
Log Project Aid p.c.	0.013* (0.007)	-0.005 (0.010)	0.052*** (0.019)	-0.031** (0.015)
Observations	930,259	930,259	930,259	930,259
K–Papp F-stat	15.72	15.72	15.72	15.72

Notes: This table reports the effect of local development aid exposure on individual employment outcomes across four Senegalese census rounds (1988, 2002, 2013, and 2023). The dependent variables are indicators for current employment (Employed), wage employment (Wage), self-employment (Self), and unpaid work (Unpaid). Panel A reports OLS estimates. Panel B reports the first-stage regression of log aid per capita on the shift–share instrument. Panel C reports IV estimates in which log aid per capita is instrumented by the shift-share log IV constructed from donor-sector disbursement shocks to non-WAEMU Sub-Saharan African countries, interacted with predetermined district-level portfolio shares. All specifications include district, and census-round fixed effects, control for individual characteristics such as gender, age group, and education and for time varying district level factors such as population size, population density, and the share of working age population. Standard errors clustered at the district $\times$ census-round level are reported in parentheses. The Kleibergen–Paap F -statistic tests for weak instruments.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

through micro-enterprise creation and agricultural productivity improvements rather than through formal sector job creation.

Figure E.15 dis-aggregates the estimates by five-year age group. The positive effect of local aid exposure on any employment (panel a) is concentrated among prime-age workers, peaking around ages 20–29 and declining steadily thereafter to reach zero by the 60–64 group. Decomposing by employment type isolates the source of these gains: the self-employment response (panel c) mirrors the any-employment profile, positive and significant across most prime-age groups, whereas wage employment (panel b) responds negatively—significantly so for the 30–54 groups—and unpaid employment (panel d) is negative or null throughout. The aggregate employment gain is therefore driven entirely by own-account activity, and indeed understates the shift toward self-employment, since rising self-employment is partly offset by contracting wage work. This pattern is consistent with aid stimulating employment through household and informal enterprise rather than formal job creation.

The heterogeneity by gender and education (Figure E.16) shows the response is broadly shared across subgroups. The age profile is similar for men and women (panel a), with the strongest effects for both among workers under 30. By skill (panel b), the gains are concentrated among the low-skilled, who form the large majority of the stayer population, while high- and medium-skilled estimates are smaller and largely indistinguishable from zero. Because low-skilled prime-age workers also exhibit the highest baseline propensity to migrate irregularly, the employment response is largest precisely for the cohorts most exposed to that margin, corroborating the opportunity-cost interpretation advanced in the main analysis.

D Tables

Table D.5: Robustness to alternative instruments

	OLS	PPML	PPML-IV	
			IV-1	IV-2
	(1)	(2)	(3)	(4)
(5)				
Log Aid per capita	-0.023** (0.010)	-0.026** (0.010)	-0.143*** (0.039)	-0.066*** (0.014)
Residuals			0.125*** (0.040)	0.058*** (0.015)
Origin $\times$ Destination FE	✓	✓	✓	✓
Destination $\times$ Time FE	✓	✓	✓	✓
Origin-Time controls	✓	✓	✓	✓
Kleibergen–Paap F-stat	.	.	11.10	16.56
Observations	10,644	10,644	10,644	9,708

Notes: The table reports full sample estimates of η from equation 2 under alternative instrumental variables, as indicated by the column headers. Columns (1) and (2) report the OLS and PPML estimates respectively. Columns (3) and (4) report the PPML-IV points estimates when per capita disbursement is instrumented using placement adjusted synthetic disbursement (IV-1), and a canonical Shift share instrument (IV-2). All specifications include origin district $\times$ destination fixed effects, destination $\times$ time fixed effects, and time-varying origin district controls, including climate shocks, nighttime light intensity, conflict exposure, population size, and urbanization, as well as origin $\times$ destination $\times$ time controls capturing migrant network. Standard errors, reported in parentheses, are clustered at the origin-time level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.6: Extensive margins: Robustness to various cutoffs

	Benchmark	Cutoffs (minimum number of migrants)				
	Full sample	> 0	> 1	> 2	> 8	≥ 10
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. OLS						
Log Aid per capita	-0.023** (0.010)	-0.033*** (0.011)	-0.036*** (0.012)	-0.034*** (0.012)	-0.026* (0.014)	-0.027* (0.014)
Panel B. PPML						
Log Aid per capita	-0.026** (0.010)	-0.026** (0.010)	-0.026** (0.010)	-0.025** (0.010)	-0.020* (0.011)	-0.020* (0.011)
Panel C. PPML-IV						
Log Aid per capita	-0.143*** (0.039)	-0.146*** (0.039)	-0.148*** (0.038)	-0.143*** (0.038)	-0.143*** (0.039)	-0.146*** (0.039)
Residuals	0.125*** (0.040)	0.127*** (0.040)	0.129*** (0.040)	0.125*** (0.040)	0.130*** (0.041)	0.134*** (0.041)
Origin $\times$ Destination FE	✓	✓	✓	✓	✓	✓
Destination $\times$ Time FE	✓	✓	✓	✓	✓	✓
Origin-Time controls	✓	✓	✓	✓	✓	✓
Mean dep. (baseline)	161	222	255	279	365	375
K. Paap F-Stat	11.10	11.34	11.31	11.56	11.28	11.27
N	10,644	7,680	6,686	6,104	4,606	4,468

Notes: The table reports estimates of η from equation 2 under alternative sample restrictions, as indicated by the column headers. Column (1) presents estimates for the full sample. Columns (2)–(6) progressively restrict the sample to origin–destination corridors with a minimum number of migrants observed across all time windows, as specified in the column titles. Panel A reports ordinary least squares estimates using the inverse hyperbolic sine (arsinh) transformation of migration flows as the dependent variable. Panel B reports instrumental-variables estimates using the Bartik shift–share instrument described in Section 5.2. All specifications include origin district $\times$ destination fixed effects, destination $\times$ time fixed effects, and time-varying origin district controls, including climate shocks, conflict exposure, population size, and urbanization, as well as origin district $\times$ destination $\times$ time controls capturing migrant network effects. Standard errors, reported in parentheses, are clustered at the origin–time level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.7: Project aid across various dimensions

	Commitments and disbursements (million USD)							
	2004–2008		2008–2012		2014–2018		2018–2022	
	C	D	C	D	C	D	C	D
Overall local project aid								
Overall	811	791	1,819	722	1,738	1,001	1,885	1,363
Share disbursed	-	97.5%	-	39.7%	-	57.6%	-	72.3%
By donor type								
Multilateral	716	671	1,515	519	1,140	650	1,668	1,027
Share disbursed	-	93.7%	-	34.3%	-	57.0%	-	61.6%
Bilateral	94	119	305	204	598	351	217	337
Share disbursed	-	126.6%	-	66.9%	-	58.7%	-	155.3%
By Sector								
Economic infra. and serv.	50	62	53	24	313	215	74	141
Share disbursed	-	124.0%	-	45.3%	-	68.7%	-	190.5%
Social infra. and serv.	34	49	107	104	203	63	55	159
Share disbursed	-	144.1%	-	97.2%	-	31.0%	-	289.1%
Production sector	6	4	2	2	4	2	1	2
Share disbursed	-	66.7%	-	100.0%	-	50.0%	-	200.0%
Other sectors	5	4	143	74	77	- 71	- 87	36
Share disbursed	-	80.0%	-	51.7%	-	92.2%	-	41.4%

Notes: This table reports total commitments (C) and disbursements (D) of local development aid projects in Senegal, expressed in million U.S. dollars, aggregated over four five-year windows between 2004 and 2022. Commitments correspond to the total value of approved projects, while disbursements capture actual payments made during each window, regardless of the year in which projects were approved. Aid is disaggregated by donor type (multilateral versus bilateral) and by broad sector of intervention. “Share disbursed” is defined as the ratio of disbursements to commitments within a given window. Disbursement rates above 100 percent reflect a larger disbursement relative to committed amounts. It is worth highlighting that disbursement rates below 100 percent do not necessarily indicate delayed or withheld funding, as the analytical windows span five years while the average duration of completed projects is 7.6 years (Figure E.9, Panel E.9a), implying that a share of committed funds is disbursed outside the window in which commitments are recorded. All amounts are reported in constant U.S. dollars.

Table D.8: Regular corridors only $\min(M_{od}) \geq 10$

	Full sample	Internal	International
	(1)	(2)	(3)
PPML-IV			
Log Aid per capita	-0.146*** (0.039)	-0.151*** (0.044)	0.005 (0.047)
Residuals	0.134*** (0.041)	0.140*** (0.044)	-0.033 (0.051)
Origin $\times$ Destination FE	✓	✓	✓
Destination $\times$ Time FE	✓	✓	✓
Origin-Destination-Time controls	✓	✓	✓
Origin-Time controls	✓	✓	✓
Mean dep. (baseline)	375	467	200
K. Paap F-Stat	11.27	11.06	11.33
N	4,468	2,928	1,540

Notes: The dependent variable is the five-year bilateral migration flow from origin district o to destination d over period $\tau:\tau-5$. The sample is restricted to “regular” origin–destination corridors with at least ten migrants observed in each five-year window. Column (1) pools all destinations, while columns (2) and (3) separately estimate internal (within-country) and international migration flows. All specifications include origin district $\times$ destination fixed effects and destination $\times$ time fixed effects, as well as time-varying origin district controls. Standard errors, reported in parentheses, are clustered at the origin–time level. The Kleibergen–Paap F-statistic reports first-stage instrument strength. The mean dependent variable refers to the baseline five-year migration flow (2009–2013). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.9: Exclude districts belonging to the capita region

	Full sample	Internal	International
	(1)	(2)	(3)
PPML-IV			
Log Aid per capita	-0.130*** (0.047)	-0.136** (0.058)	0.005 (0.046)
Residuals	0.104** (0.048)	0.113** (0.057)	-0.032 (0.050)
Origin $\times$ Destination FE	✓	✓	✓
Destination $\times$ Time FE	✓	✓	✓
Origin-Destination-Time controls	✓	✓	✓
Origin-Time controls	✓	✓	✓
Mean dep. (baseline)	116	242	48
K. Paap F-Stat	11.12	10.80	11.58
N	10,292	3,592	6,700

Notes: The dependent variable is the five-year bilateral migration flow from origin district o to destination d over period $\tau : \tau - 5$. Destination districts belonging to the capital region of Senegal—Dakar, Pikine, Rufisque, and Guédiawaye—are excluded from the sample. Column (1) pools all remaining destinations, while columns (2) and (3) separately estimate internal (within-country) and international migration flows. All specifications include origin district $\times$ destination fixed effects and destination $\times$ time fixed effects, as well as time-varying origin district controls. Standard errors, reported in parentheses, are clustered at the origin–time level. The Kleibergen–Paap F-statistic reports first-stage instrument strength. The mean dependent variable refers to the baseline five-year migration flow (2009–2013). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.10: Robustness to alternative instruments (irregular migration)

	OLS	PPML	PPML-IV	
			IV-1	IV-2
	(1)	(2)	(3)	(4)
Log Aid per capita	-0.028*	-0.085***	-0.147***	-0.126**
	(0.014)	(0.031)	(0.037)	(0.057)
Residuals			0.135*	0.116
			(0.073)	(0.143)
Origin region $\times$ Destination group FE	✓	✓	✓	✓
Destination country $\times$ FE	✓	✓	✓	✓
Origin-Destination controls	✓	✓	✓	✓
Origin controls	✓	✓	✓	✓
Kleibergen–Paap F-stat	.	.	17.9	16.6
Observations	3,348	3,082	3,082	2,820

Notes: This table reports robustness of the irregular-migration results to alternative instrumental variables, as indicated by the column headers. Column (1) reports OLS and column (2) PPML estimates. Columns (3) and (4) report PPML-IV estimates implemented via a control-function approach, where per capita disbursement is instrumented using the placement-adjusted synthetic disbursement instrument (IV-1) and a canonical shift-share instrument (IV-2), respectively; Residuals denotes the first-stage residual term. The dependent variable is the five-year irregular migration flow from origin district o to destination country d . All specifications include origin region $\times$ destination group and destination country fixed effects, as well as origin and origin–destination controls. The Kleibergen–Paap F-statistic reports first-stage instrument strength. Standard errors clustered at the origin district level are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.11: Robustness to regular corridors (irregular migration)

	Full sample	Regular	Irregular
	(1)	(2)	(3)
Panel A. PPML			
Log Aid per capita	-0.013 (0.021)	-0.017 (0.028)	-0.088*** (0.033)
Panel B. PPML-IV			
Log Aid per capita	-0.054* (0.031)	-0.028 (0.043)	-0.163*** (0.039)
Residuals	0.099 (0.063)	0.023 (0.072)	0.158** (0.077)
Origin region $\times$ Destination group FE	✓	✓	✓
Destination FE	✓	✓	✓
Origin controls	✓	✓	✓
Mean dep. (baseline)	401	160	44
K. Paap F-Stat	17.33	17.33	17.33
N	2,231	767	767

Notes: This table reports estimates of the effect of local exposure to development projects on irregular international migration, restricting the sample to “regular” origin–destination corridors with at least ten migrants observed in each five-year window. Column (1) pools all international flows, while columns (2) and (3) separately estimate regular and irregular migration flows. Panel A reports Poisson Pseudo–Maximum Likelihood (PPML) estimates, and Panel B instrumental-variables (PPML-IV) estimates implemented via a control-function approach, where Residuals denotes the first-stage residual term. The dependent variable is the five-year migration flow from origin district o to destination country d . All specifications include origin region $\times$ destination group and destination country fixed effects, as well as origin and dyadic controls. The Kleibergen–Paap F-statistic reports first-stage instrument strength. The mean dependent variable refers to the baseline five-year migration flow (2009–2013). Standard errors clustered at the origin district level are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Table D.12: Heterogeneous response of irregular migration by destination groups

	Full sample	By destination				
	(1)	(2)	(3)	(4)	(5)	(6)
		Neighbors	Europe	Other Africa	North America	RoW
Panel A. PPML						
Log Aid per capita	-0.085*** (0.031)	0.089*** (0.029)	-0.132*** (0.044)	-0.090** (0.038)	-0.062 (0.058)	-0.052 (0.050)
Panel B. PPML-IV						
Log Aid per capita	-0.147*** (0.037)	0.076 (0.048)	-0.215*** (0.056)	-0.120* (0.065)	-0.007 (0.091)	-0.055 (0.083)
Residuals	0.135* (0.073)	0.029 (0.077)	0.194* (0.111)	0.065 (0.093)	-0.131 (0.147)	0.007 (0.146)
Origin region $\times$ Destination group FE	✓	✓	✓	✓	✓	✓
Destination FE	✓	✓	✓	✓	✓	✓
Origin controls	✓	✓	✓	✓	✓	✓
Mean dep. (baseline)	11	28	25	6	6	1
K. Paap F-Stat	17.97					
N	3,082	223	809	1,158	83	809

Notes: This table reports estimates of the effect of exposure to development project on irregular international emigration by destination group. Panel A reports Pseudo Maximum Likelihood (PPML) estimates, and Panel B instrumental-variables (PPML-IV) estimates. The dependent variable is the five-year irregular migration flows from origin district o to destination country d . All specifications include higher origin region $\times$ destination group, destination country fixed effects, origin and dyadic level controls. Standard errors clustered at the origin level are reported in parentheses. The Kleibergen–Paap F-statistic reports first-stage instrument strength. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

E Figures


Figure E.1: Senegalese migration relative to Sub-Saharan Africa (SSA)



Notes: Figure E.1 assesses the external validity of the Senegalese migration patterns relative to Sub-Saharan Africa (SSA). Panel E.1a compares the relative size of irregular migration routes, expressed as percentages, for Senegal and four SSA sub-regions: West, Central, Eastern, and Southern Africa. Panel E.1b reports average internal migration rates for Senegal and six SSA comparator countries available in IPUMS, weighted by person weights. The dashed vertical line marks the overall weighted cross-country average, equal to 0.059, while Senegal is highlighted in maroon. Panel E.1c plots average internal migration rates by age, measured five years prior to the census year, for Senegal and the same comparator countries. Panel E.1d compares the destination-region distribution of Senegalese migrants with the SSA average across World Bank destination regions, expressed as shares of total migrants. Internal migration estimates are weighted by person weights. Data sources: Integrated Public Use Microdata Series (IPUMS); Abel and Cohen (2022); full Senegalese population censuses.


Figure E.2: Validating Abel and Cohen (2022) against Senegalese census (2013, 2023)

(a) Abel and Cohen vs. Census



Note: Weighted means by destination region. Data sources: Abel and Cohen (2022); Full Senegalese population censuses (2013, 2023).

(b) Destinations by wave



Note: Weighted means by destination region. Data source: Abel and Cohen (2022)


Notes: Figure E.2 compares the distribution of Senegalese migrants across World Bank destination regions with the corresponding distribution observed in the full Senegalese population censuses, expressed as shares of total migrants. Panel E.2a compares the bilateral migration estimates from Abel and Cohen (2022) with census-based shares for two migration periods: the left-hand chart matches the 2008–2013 census period with the 2005–2010 estimates from Abel and Cohen (2022), while the right-hand chart matches the 2018–2023 census period with their 2015–2020 estimates. Panel E.2b compares Senegal's destination distribution with the average distribution for Sub-Saharan Africa (SSA) over the same two waves. Data sources: Abel and Cohen (2022); full Senegalese population censuses, 2013 and 2023.

(a) Census-based vs. Asylum applications
(b) Census-based vs. Asylum seekers




26

Figure E.4: Robustness to leave destination group out




Notes: This figure reports leave-one-destination-group-out robustness checks for the estimated effect of local exposure to aid projects on irregular international migration. Each point corresponds to the coefficient on log aid per capita obtained after excluding all migration flows to the indicated destination group from the estimation sample. Blue circles report PPML estimates and red diamonds report PPML-IV estimates; vertical bars denote 95 percent confidence intervals. Standard errors are clustered at the district of origin level.

Figure E.5: Heterogeneous response of irregular migration by age, gender and education




Notes: The figure plots heterogeneous effects of development aid on irregular international migration, estimated separately by subgroup. Panel E.5a splits the sample by age group at migration, panel E.5b by age group interacted with gender, and panel E.5c by skill group. In each panel, blue circles report PPML point estimates and red diamonds report PPML-IV (control-function) estimates, with the instrument as in the baseline specification. In panel E.5b, filled markers denote women and hollow markers denote men. Vertical bars are 95% confidence intervals, with standard errors clustered by origin district level. The “All” category reproduces the pooled baseline estimate for reference. All specifications include the same fixed effects and controls as the corresponding baseline regression.

Figure E.6: Project locations (coordinates) across Senegalese districts




Notes: This figure maps the geographic locations of development aid projects across Senegalese districts over two periods. Panels [E.6a](#) and [E.6b](#) display project coordinates for projects initiated in 2004–2008 and 2014–2018, respectively. Each point corresponds to the reported geographic centroid of a project location, as recorded in administrative project-level data. The figure illustrates changes in the spatial density and dispersion of project implementation across periods but does not scale points by project size or financial volume; all projects are weighted equally. Differences across panels therefore reflect changes in the number and location of projects rather than changes in project intensity.

Figure E.7: Spatial distribution of migration across districts




Notes: This figure maps district-level five-year migration rates constructed from the 2013 and 2023 population censuses for individuals aged 15-64 five years preceding the census year. Panels E.7a and E.7b report internal migration rates for the periods 2009–2013 and 2018–2023, respectively, defined as the share of individuals whose district of usual residence at the census differs from that five years earlier. Panels E.7c and E.7d report international migration rates over the same periods, defined as the share of individuals residing abroad at the census year who left Senegal during the preceding five years, as reported by the household. Migration rates are computed at the district of origin and expressed in percent. All maps use a common set of district boundaries, and color intensities correspond to fixed bins shown in the legends, allowing for direct visual comparison across periods within each migration type. Differences across panels therefore reflect changes in migration intensities.

Figure E.8: Dynamic and sectoral composition of aid projects in Senegal




Notes: This figure plots the dynamic of aggregate project aid disbursement in Senegal as well as its sectoral composition. Panel E.8a plots the annual volume of development aid disbursements in Senegal, measured in constant million U.S. dollars and dated by the year in which payments occurred. The series aggregates disbursements across all projects and donors. Panel E.8b reports the sectoral composition of total disbursements over the full sample period, grouped into broad sectors: economic infrastructure and services, social infrastructure and services, production sectors, and other sectors. Bars show the share of cumulative disbursements accounted for by each sector, expressed in percent. Sector classifications follow the OECD Creditor Reporting System. All amounts are deflated to constant U.S. dollars.

Figure E.9: Development aid project duration




Notes: This figure describes the duration of development aid projects in Senegal. Panel E.9a plots the distribution of project duration for closed projects, measured as the number of years between project approval and project closure. The dashed vertical line indicates the mean project duration, equal to 7.6 years. Panel E.9b reports the average duration of closed projects by approval-year window, with windows defined in five-year intervals. Project durations are computed using administrative project-level data and include all donors and sectors. The relative stability of average durations across windows indicates limited time variation in project length over the sample period.

Figure E.10: Development aid disbursement profile




Notes: This figure shows the average disbursement profile of development aid projects in Senegal. The bars report the fraction of a project's total lifetime disbursement that occurs in each year since project start, computed across all projects with observed disbursement histories. Year 0 corresponds to the approval year. The dashed line plots the same distribution smoothed across years to highlight the typical disbursement pattern. The figure illustrates that disbursements are front-loaded but extend over several years, with a substantial share of total payments occurring well beyond the first five years after project start.

Figure E.11: Robustness to exclusion of important destinations




Notes: This figure assesses the sensitivity of the baseline migration response to the exclusion of individual high-volume destinations. Each position on the horizontal axis reports estimates from a separate regression in which the indicated destination is dropped from the sample; the ten destinations correspond to the largest internal and international corridors (M’baké, Pikine, Rufisque, Dakar, Guédiawaye, Mbour, France, Thiès, Italy, and Spain), with the final position (“All together”) excluding all ten simultaneously. The dependent variable is the five-year bilateral migration flow for any migration. Circles report OLS estimates, triangles PPML estimates, and diamonds PPML-IV estimates; vertical bars denote 95% confidence intervals based on standard errors clustered at the origin–time level. All specifications include origin–destination and destination–time fixed effects, origin–destination–time and time-varying origin-level controls. The stability of the point estimates across exclusions indicates that the estimated deterrent effect of local aid exposure is not driven by any single dominant corridor.

Figure E.12: Migration effects of development aid by gender




Notes: The figure plots point estimates and 95 percent confidence intervals for the effect of local aid per capita on five-year migration flows, disaggregated by education group and gender. Panel (a) reports estimates pooling men and women, while panels (b) and (c) report estimates separately for women and men. Education groups correspond to low-skilled, medium-skilled, and high-skilled individuals. Solid markers report estimates for all migration flows, dashed markers report estimates for internal migration, and dotted markers report estimates for international migration. All specifications include origin–destination fixed effects, destination–time fixed effects, and time-varying origin–destination and origin-level controls. Confidence intervals are based on standard errors clustered at the origin–time level. The vertical dashed line denotes zero.

Figure E.13: Migration effects of development aid by age groups




Notes: The figure plots point estimates and 95 percent confidence intervals for the effect of local aid per capita on five-year migration flows, disaggregated by age group. Solid markers report estimates pooling all migration flows, dashed markers report estimates for international migration, and dotted markers report estimates for internal migration. Age groups correspond to individuals aged 15–24, 25–34, 35–44, 45–54, and 55–64 at the time of migration. All specifications include origin–destination fixed effects, destination–time fixed effects, time-varying origin-level, and origin–destination controls. Confidence intervals are based on standard errors clustered at the origin–time level. The vertical dashed line denotes zero.

Figure E.14: Migration effects of foreign aid by broad destinations




Notes: The figure plots point estimates and 95 percent confidence intervals for the effect of local aid per capita on five-year migration flows, disaggregated by broad destination group. Estimates are obtained from regressions that interact aid per capita with destination-group indicators, and the plotted coefficients represent total effects for each destination group, computed as the sum of the coefficient on the reference category and the corresponding interaction term. The reference category is domestic destinations (i.e., internal migration). All specifications include origin–destination fixed effects, destination–time fixed effects, time-varying origin-level controls, and time-varying origin–destination controls. Confidence intervals are based on standard errors clustered at the origin–time level. The vertical dashed line denotes zero.

Figure E.15: Employment response to aid projects



Notes: This figure plots estimates of the effect of local project aid per capita on individual employment outcomes by age group, constructed from four Senegalese census rounds (1988, 2002, 2013, and 2023). Each point corresponds to a separate regression in which the sample is restricted to individuals in the indicated five-year age group. Panel (E.15a) reports estimates for any employment; Panel (E.15b) for wage employment; Panel (E.15c) for self-employment; and Panel (E.15d) for those employed in unpaid jobs. The dependent variables are unconditional binary indicators defined over the full stayer working age population, equal to one if the individual is in the relevant employment category and zero otherwise. All specifications include district, and census-round fixed effects, and control for both individual characteristics and district level characteristics. Vertical bars represent 95% confidence intervals. Standard errors are clustered at the district $\times$ census-round level.

Figure E.16: Employment response to employment aid projects by gender and education



Notes: This figure plots estimates of the effect of local project aid per capita on individual employment outcomes by age group across gender (E.16a) and education (E.16b) dimensions, constructed from four Senegalese census rounds (1988, 2002, 2013, and 2023). Each point corresponds to a separate regression in which the sample is restricted to individuals in the indicated five-year age group. The dependent variables are unconditional binary indicators defined over the full stayer working age population, equal to one if the individual is employed and zero otherwise. All specifications include district, and census-round fixed effects, and control for both individual characteristics and district level characteristics. Vertical bars represent 95% confidence intervals. Standard errors are clustered at the district $\times$ census-round level.