

Working, but in Slums: Climate Migration and Integration in Urban Africa

Mélanie Gittard* and Eleanor Wiseman†

September 2026

[[Link to Latest Version](#)]

Abstract

Does climate-induced migration contribute to the growth of low-quality housing in African cities? We construct bilateral yearly migration flows between districts of eight African countries from harmonized census microdata (censuses 1990-2019, migration windows covering 1981-2018) and instrument migration with drought intensity at origin, selected by post-IV LASSO among candidate climate variables. Drought-induced migration raises the inflow into low-quality housing: of every hundred migrants that drought pushes into a city, about twenty-eight settle in low-quality housing. These migrants are also more housing-vulnerable than others: the marginal drought-induced migrant is about 11 percentage points more likely to live in low-quality housing than the average arrival on the same corridor – 51% against 41%. Yet they are not economically excluded, but absorbed into employment everywhere. The cost of climate migration is thus residential rather than occupational, and it concentrates in secondary and small, rapidly urbanizing cities: the effect is of the same magnitude in towns below 100k and in cities of 100k–1M, and fades significantly in metropolises above one million. It also has an educational gradient: drought does not displace the least-educated adults, who remain trapped at origin, and among those it displaces, it is the less educated who are channeled into low-quality housing. Aggregating over our corridors, drought placed roughly 110,000 additional people in low-quality housing. These results shift the policy focus toward urban housing and toward the fast-urbanizing secondary cities where the pressure concentrates, through both greater absorption capacity and the targeting of the most vulnerable migrants.

Keywords: Climate Change, Migration, Low-Quality Housing, Urbanization, Labor Markets

JEL codes: Q54, R23, O15, O18, J61

*Stanford University, Center on Food Security and the Environment (FSE). mgittard@stanford.edu

†University of Wisconsin-Madison. wiseman3@wisc.edu

We are grateful to David Lobell, Marshall Burke, Jennifer Burney, Rosamond Naylor, Pascal Polonik, Emmett Reynier, Sam Heft-Neal, Andrew Wilson, Aria Grossman, Lucie Boilet, and our many colleagues at Stanford University. We thank Marlon Seror, Valerie Mueller, Anja Benshaul-Tolonen, Victoire Girard, and Denis Cogneau for valuable feedback, as well as audiences at NOVAFRICA, the EEA, and the ICDE.

1 Introduction

African cities stand at the intersection of two major transitions: rapid urbanization and intensifying climate change and variability. By 2050, an estimated 68% of the continent’s population will reside in urban areas, up from 41% in 2020 (Un-Habitat, 2020). This urban transition comes with major deficits in both housing and employment: 62.7% of Sub-Saharan Africa’s urban population currently lives in slums (Un-Habitat, 2022), and a substantial share of urban residents remain employed in agriculture rather than in manufacturing or services (Henderson and Kriticos, 2018). Africa appears to be urbanizing without full structural transformation. At the same time, climate variability increasingly affects rural livelihoods: droughts deteriorate agricultural productivity, deplete household incomes, and push rural populations to migrate as a coping strategy (Henderson et al., 2017; Mueller et al., 2020). Yet where these climate-driven migrants settle, under what conditions, and whether they find work at destination remain poorly understood. The question matters because low-quality housing compounds residents’ exposure to climate hazards such as flooding and extreme heat (Marx et al., 2019): populations fleeing climate-vulnerable rural areas may end up exposed anew in the city. Whether climate migration contributes to the growth of low-quality housing in African cities, and whether climate migrants are more likely to settle in such housing than other migrants, is therefore a question of first-order policy relevance.

This paper asks how drought-induced migrants are absorbed into African cities, in both housing and employment. For each outcome, we distinguish two channels. The *scale* channel asks whether drought-induced migration raises the inflow into low-quality housing and into employment by increasing the number of arrivals. The *propensity* channel asks whether the migrants that drought sends are more likely to end up in low-quality housing, and differ in their employment, from the migrants who would have moved anyway. Separating the two matters for policy. The scale channel points to a congestion effect from the number of additional migrants, calling for greater urban absorption capacity. The propensity channel asks whether climate migrants are a distinct, more vulnerable population, calling instead for targeted support.

We identify the causal effect of migration on low-quality housing by instrumenting bilateral migration flows with drought at origin, in a yearly panel built from census microdata across eight African countries. The first challenge is measurement: studying migrants’ outcomes at destination requires knowing how many migrants move from each origin to each destination in each year - bilateral flows at annual frequency, which existing data for Africa seldom provide. We overcome this by building bilateral migration flows from IPUMS census microdata. We use retrospective questions recording each migrant’s administrative unit of origin and year of arrival, letting us reconstruct yearly cohorts of migrants for every origin-destination pair at the admin-2 level. This yields a cohort-bilateral panel of migration flows between districts, built from censuses 1990-2019 with migration windows covering 1981-2018, matched at the household level with a multi-criteria measure of housing quality based on the UN-Habitat slum definition, as well as each migrant’s employment status. The second challenge is endogeneity: corridors with more migration differ systematically in ways correlated with destination housing. We instrument bilateral migration with drought intensity at origin, selected by a post-IV LASSO procedure among candidate climate variables to re-

move the discretionary choice of instrument. Drought strongly predicts migration, and the identifying assumption is that it affects destination housing only through migration along the corridor, conditional on corridor and arrival-year fixed effects.

A first result of the paper is the first stage, showing that drought moves people to cities. Where rural incomes depend on rain-fed agriculture, a failed growing season is an income shock, and moving to the city is one of the main margins of adaptation. Comparing the same origin-destination corridor across years, we find that in years when the origin's growing season is unusually dry, more people move to urban areas than in normal years. We show that a one-standard-deviation increase in drought at origin raises the corridor's urban migration rate by about 7% of its mean. The response is specific to the dry tail of the rainfall distribution - symmetric wet shocks do not move people - consistent with an agricultural push rather than generic weather variation. This first stage is the variation on which all subsequent results rest.

The main finding of the paper concerns where these additional migrants end up: the inflow of climate migrants contributes to the growth of low-quality housing in the cities that receive them. A 1 percentage-point (pp) increase in the drought-induced migration rate raises the inflow into low-quality housing at destination by 0.28 pp: of every hundred migrants whom drought pushes into a city - people who, absent the dry shock, would have stayed at origin - 28 settle in low-quality housing. In the counterfactual without drought, neither these migrants nor their housing demand would be in the city at all. We also find evidence of selection: drought-induced migrants are more likely to end up in low-quality housing than the migrants who would have moved anyway. Holding the size of the inflow fixed, the marginal drought-induced migrant is about 11 pp more likely to live in low-quality housing than the average migrant arriving on the same corridor - 51% do, against 41% of baseline arrivals.

Decomposing the total effect into these two margins, the scale channel - more migrants arriving, sorting into housing roughly like those already flowing along the corridor - accounts for about 79% of the additional inflow into low-quality housing, and the propensity channel for the remaining 21%. A back-of-the-envelope aggregation puts magnitudes on these estimates: over 1981-2018, drought generated roughly 390,000 of the 11.6 million urban moves recorded in our corridors (3.3%), and thereby placed about 110,000 additional people in low-quality housing.

The second main finding is that climate migration strains cities residentially rather than economically: drought migrants are absorbed into work, but into low-quality housing. Using the same design on employment outcomes, of every hundred marginal migrants about 43 arrive into employment, 7 into unemployment, and 35 into inactivity (the remainder having undetermined status). Drought-induced migrants are also positively selected into work: the marginal migrant is about 12 pp more likely to be employed than the average arrival - 63% against 51% - and about 15 pp less likely to be inactive, with unemployment unchanged. Together with the housing results, this characterizes the drought-induced migrant precisely: employed more often than the typical arrival, yet housed worse. The margin of exclusion is residential, not occupational.

The effect operates where the agricultural-push mechanism predicts, and its geography is distinctive. At origin, it is strongest along rural-to-city corridors: drought displaces rural households whose livelihoods depend most directly on rainfall, and these are the migrants

who most raise low-quality housing inflows. Instrumenting with a cumulative three-season drought instead of the contemporaneous shock leaves the estimate essentially unchanged (0.28 against 0.28), so the effect does not appear to intensify with the persistence of the shock. At destination, the two margins diverge: urban labor markets absorb drought migrants everywhere, but the sorting into low-quality housing concentrates in secondary and small, rapidly urbanizing cities. Classifying destinations by agglomeration size with the Africapolis database, the effect is of the same magnitude in small towns below 100k and in secondary cities of 100k–1M (0.30 and 0.33, statistically indistinguishable from each other) and falls to 0.10 in metropolises above one million, where it is no longer distinguishable from zero - an attenuation significant at the 5% level. Splitting destinations instead by their census urban share gives the same declining direction but no significant contrast, and how close one comes depends on where the rural/urban cutoff is placed; we therefore rest the inference on the size-based classification, which requires no such choice. The working poor that drought sends to the cities - migrants who find work but not decent housing - are thus a phenomenon of the secondary and small urbanizing tier rather than of the primate metropolises.

The same selectivity runs across people, with an educational gradient that cuts in opposite directions at the two stages. Drought does not displace the least-educated adults: their migration response is indistinguishable from zero - a trapped-population pattern. Among the adults it does displace, education shields from low-quality housing but not from displacement: the housing effect is twice as strong for the primary-educated as for the secondary-educated, while their absorption into employment is identical. The least educated stay behind, the secondary-educated escape bad housing, and the lower-educated among the displaced bear the residential burden.

We subject these results to an extensive robustness analysis. The main threat to identification is that drought at origin could affect destination housing through channels other than migration along the corridor; corridor and arrival-year fixed effects absorb the time-invariant and common components, the wet-side placebo rules out channels shared by all rainfall variation, and we discuss the remaining ones - notably food prices at destination and conflict-induced displacement - in Section 6. The estimates are stable across fixed-effects structures and across the alternative deprivation margins of the UN-Habitat definition, including a continuous count of deprived dimensions. Section 6 further reports checks on the instrument itself - alternative drought measures, out-of-sample validation of the LASSO selection - as well as alternative population denominators, an extension to fifteen countries using five-year migration brackets, and a leave-one-country-out exercise.

This paper makes four contributions. First, it provides the first causal evidence linking climate-induced migration to the growth of low-quality housing in Sub-Saharan Africa. A large literature establishes that climate shocks drive migration, from rural-rural moves (Gittard, 2024; McGuirk and Nunn, 2021; Gray and Mueller, 2012) to rural-urban (Henderson et al., 2017; Cattaneo and Peri, 2016; Mueller et al., 2020) and cross-border migration (Beine and Parsons, 2016; Missirian and Schlenker, 2017), while a separate literature debates whether informal settlements are transitional or persistent poverty traps and documents the vulnerability of their residents (Marx et al., 2013, 2019). This paper connects the two: it shifts the focus from out-migration at origin to the housing absorption of climate migrants at destination, and decomposes the effect into a scale and a propensity channel. Second, it characterizes the labor-market integration of these migrants. Whereas existing work em-

phasizes barriers that keep rural-urban migrants in low-productivity activities (Gollin et al., 2014; Bryan and Morten, 2019; Alvarez, 2020; Hamory et al., 2021) and negative labor-market spillovers of climate migration (Kleemans and Magruder, 2018; Gittard, 2024), we find that drought-induced migrants are absorbed into employment: the margin of exclusion is residential, not occupational, a distinction that matters for targeting policy. The individual heterogeneity sharpens it: education shields migrants from low-quality housing but not from displacement, and the least-educated adults are not displaced at all - the trapped-population margin documented by Cattaneo and Peri (2016). Third, it contributes to the literature on African urbanization and structural transformation (Henderson and Kriticos, 2018; McMillan et al., 2014; Jedwab et al., 2017; Gollin et al., 2016) by showing that the housing pressure of climate migration concentrates in urbanizing intermediate cities, shifting attention from primate cities to the secondary tier where urban growth is fastest and infrastructure thinnest. Fourth, methodologically, it illustrates a disciplined treatment of instrument choice in climate econometrics: a post-IV LASSO selection among 25 candidate climate variables (Belloni et al., 2014), validated by a wet-side placebo, replaces the discretionary choice of climate instrument.

The remainder of the paper is organized as follows. Section 2 describes the data and the construction of the migration, weather, and housing-quality variables. Section 3 presents the identification strategy, including the post-IV LASSO selection of the instrument. Section 4 reports the main results on housing and employment, the scale-propensity decomposition, and the aggregate magnitude. Section 5 examines heterogeneity by destination city size and by migrant profile. Section 6 presents robustness checks, Section 7 develops a two-margin model of post-migration integration that organizes the findings, and Section 8 concludes.

2 Data

This section describes the construction of the weather, migration, urbanization, and housing-quality variables. Appendix Table 7 reports descriptive statistics for all variables used in the analysis, and Appendix Figures 3–5 map the climate exposure of origins and the migration, urbanization, and migrant outcomes of destinations. Appendix Figure 6 maps the observed urban in-migration corridors across the eight countries, illustrating the bilateral structure of the data: each origin–destination pair contributes a yearly flow, and identification rests on within-corridor variation in the timing of droughts at origin.

2.1 Construction of Weather Variables

Overview - The empirical model selects among a large set of candidate weather variables constructed from raw climate data, using a post-IV LASSO procedure (Belloni et al., 2014). The choice of candidate variables is grounded in recent climate-econ damage function literature (Hultgren et al., 2025; Auffhammer, 2018), which identifies the most influential weather terms for agricultural production. All weather variables are computed on a common historical reference period of 1983–2023. We distinguish temperature and precipitation terms. Weather variables are defined over the growing season (superscript g) for the main analysis, with yearly versions (superscript y) used in robustness checks. Growing seasons are

defined at the national level using the main rainy season identified from historical monthly precipitation distributions for each country.

Raw data - Precipitation variables are derived from CHIRPS daily data (Funk et al., 2015). Temperature variables and evapotranspiration are derived from ERA5-Land reanalysis data (Muñoz-Sabater et al., 2021).

Pixel-level versus admin-unit-level transformations - We construct strongly nonlinear transformations (squared precipitation, squared temperature, GDD) at the pixel level before spatial aggregation, following Hsiang (2016), to preserve local nonlinear responses. Threshold-based drought variables ($\text{Drought}_{p,J}^g$ and drought_int_z) are constructed using admin-unit-level historical precipitation distributions, consistent with the spatial scale of our migration outcomes.

Precipitation variables - We construct several precipitation candidate variables for the post-LASSO IV. The cumulative precipitation Precip_{cum}^g is the spatial mean of pixel-level monthly cumulative precipitation over the growing season. Its squared version $(\text{Precip}_{cum}^g)^2$ is computed at the pixel level before spatial averaging, to capture non-linear responses. The number of dry days N_{dry}^g counts days with daily precipitation below 1 mm at the pixel level and is spatially averaged at the admin-unit level. Two threshold-based drought dummies Drought_{p10}^g and Drought_{p15}^g equal 1 if Precip_{cum}^g falls below the 10th and 15th percentile, respectively, of its historical distribution at the admin-unit level. Finally, the standardized drought intensity $\text{drought_int_z}_{o,t} = \max(0, -\tilde{z}_{o,t})$, where $\tilde{z}_{o,t} = (\text{Precip}_{o,t}^g - \overline{\text{Precip}_o^g}) / \sigma_{\text{Precip},o}^g$ is the seasonal precipitation anomaly, captures drought intensity asymmetrically: it activates only for below-average rainfall and measures the number of standard deviations below the historical mean.

For an extreme-precipitation analysis conducted separately, we construct two symmetric extreme-wet variables: $\text{ExtremeWet}_{p,J}^g$ is a dummy equal to 1 if Precip_{cum}^g exceeds the 85th or 90th percentile of the historical distribution, and $N_{extreme}^g$ counts days with daily precipitation above 20 mm at the pixel level, spatially averaged at the admin-unit level.

Temperature variables - We construct three temperature candidate variables. The seasonal averages of daily minimum and mean 2m air temperatures, T_{min}^g and T_{mean}^g , are first computed at the pixel level and then spatially averaged. Their squared versions $(T_{min}^g)^2$ and $(T_{mean}^g)^2$ are computed at the pixel level (after conversion to Celsius) before spatial averaging, to capture non-linear thermal responses. Growing Degree Days GDD^g are computed daily at the pixel level as $\max(0, \frac{1}{2}(\tilde{T}_{min,d} + \tilde{T}_{max,d}) - 8)$, where $\tilde{T}_{min,d}$ and $\tilde{T}_{max,d}$ are daily minimum and maximum 2m temperatures capped at 30°C. The daily GDD are then summed over the growing season and spatially averaged at the admin-unit level. Extreme Degree Days EDD^g are in preparation for future versions of the candidate set.

Moisture variables - Moisture-related variables include potential evapotranspiration PET^g and actual evapotranspiration ET^g , both extracted from ERA5-Land at the daily level and aggregated over the growing season. PET captures atmospheric moisture demand (irrigation

requirements), while ET captures realized evaporation. Vapor pressure deficit (VPD) is work in progress for future versions of the candidate set.

LASSO selection - The candidate set used in the post-IV LASSO selection comprises thirteen contemporaneous weather variables, all measured at the admin-unit level over the growing season - the drought dummies Drought_{p10}^g and Drought_{p15}^g , the standardized drought intensity drought_int_z , N_{dry}^g , the precipitation and temperature levels and polynomials, and the degree-day and evapotranspiration measures described above - augmented with the lags and three-season means of the drought-family variables, for 25 candidates in total (Section 3.3). The procedure selects drought_int_z as the unique relevant instrument for bilateral migration — a result we view as informative in itself, isolating an interpretable and asymmetric drought-intensity measure as the dominant climatic driver. We use this single-variable selection for the main analysis to maintain clarity in interpretation.

Robustness of the selection - The selection is probed in several directions (Section 3.3 and Appendix Figure 9). It is stable across LASSO penalty values, from the theoretical benchmark of Belloni et al. (2014) to more permissive thresholds; among lagged candidates, the three-season mean of drought intensity dominates, consistent with the cumulative instrument used in robustness checks; and a wet-side placebo - re-running the exercise on symmetric excess-rainfall candidates - finds essentially no predictive power, confirming that the instrument’s strength is specific to the dry tail of the rainfall distribution. Extensions of the candidate set to the additional non-linear and moisture terms described above, a comparison with Double Machine Learning (Chernozhukov et al., 2018), and an out-of-sample validation across countries following Hultgren et al. (2025) are in progress.

2.2 Migration data

Sources - Both migration and housing quality data come from IPUMS census microdata (Minnesota Population Center, 2019) for fifteen African countries over the period 1987–2019. The sample includes exhaustive population censuses from Benin (2002, 2013), Burkina Faso (1996, 2006), Cameroon (1987, 2005), Egypt (1986, 1996, 2006), Ghana (2000, 2010), Kenya (1989, 1999, 2009, 2019), Malawi (1998, 2008), Mali (1998, 2009), Mauritius (1990, 2000, 2011), Mozambique (1997, 2007, 2017), Senegal (1988, 2002, 2013), Sierra Leone (2015), South Africa (1996, 2001, 2011), Togo (2010), Uganda (2002, 2014), and Zambia (1990, 2000, 2010). These censuses provide harmonized individual-level information on demographic characteristics, migration history (place of birth, year of arrival, previous residence), and housing conditions (dwelling type, access to basic services, urban/rural classification at the enumeration area level). One contribution of this paper is the construction of *bilateral yearly migration flows* at the IPUMS administrative level for 8 countries with available year-of-arrival information, and *five-year migration brackets* for the remaining 7 countries.¹

¹In a robustness exercise, we compare our IPUMS-derived migration measures for Kenya with the exhaustive population censuses for 1989, 1999, and 2009 obtained directly from the Kenya National Bureau of Statistics. Results are consistent across both sources.

Migrant identification - We define an individual i observed at destination d in census year t_c as a migrant from origin o in year t if (i) the IPUMS variable `migyrs1` reports a years-since-arrival value τ such that $t = t_c - \tau$, and (ii) the IPUMS variable for previous administrative unit of residence (or, when unavailable, place of birth) identifies origin o . The destination d is the administrative unit where the individual resides at the census date. For all censuses, we trace migration histories back twelve years prior to each census for the yearly migration analysis, capping the window in multi-census countries so that it does not overlap the preceding census. The individual migration indicator is:

$$\text{InMigrants}_{o,d,t} = \sum_i \mathbf{1}\{i \text{ at } d \text{ in } t_c, \text{ arrived from } o \text{ in } t\} \cdot \text{perwt}_i,$$

where perwt_i is the IPUMS person-weight and the sum runs over all individuals observed at destination d in census t_c .

We exclude from the main analysis origin–destination pairs that record zero migrants in every observed year. These always-zero corridors contribute no within-pair variation to identification under our origin–destination fixed effects specification, and their inclusion would mechanically inflate the sample size without adding identifying information.

Urban Migrant - We use the IPUMS variable `urban` to identify whether an individual resides in an urban or rural enumeration area within their destination administrative unit. This classification is finer than the admin unit level: a single administrative unit may contain both urban and rural enumeration areas. Urban migrants are individuals classified as residing in urban enumeration areas at destination (`urban` = 2 in IPUMS).

Population denominators - We construct two destination population measures. The *census-year population* $\text{Pop}_{d,T_c}^{\text{census}}$ is the sum of person-weights observed at destination d in census year T_c . The *urban census-year population* $\text{Pop}_{d,T_c}^{\text{census,urban}}$ restricts to individuals classified as urban residents (`urban` = 2). For robustness, we also reconstruct annual destination populations $\text{Pop}_{d,t}$ backward from the census year via the accounting identity:

$$\text{Pop}_{d,t} = \text{Pop}_{d,T_c} - \text{InMig}_{d,[t+1,T_c]} - \text{Births}_{d,[t+1,T_c]} + \text{OutMig}_{d,[t+1,T_c]} + \text{Deaths}_{d,[t+1,T_c]},$$

where births are estimated from age or birth-year information in the census, and deaths are derived from country-specific mortality rates in the corresponding period. Our main results use the census-year population, which is more stable; the annual reconstruction is used in robustness checks.

Migration rates. - The bilateral migration rate from origin o to destination d in year t is:

$$\text{MigRate}_{o,d,t} = \frac{\text{InMigrants}_{o,d,t}}{\text{Pop}_{d,T_c}^{\text{census}}} \times 100,$$

expressed in percentage points. The urban migration rate restricts the numerator to urban-resident migrants and uses the urban census-year population as denominator:

$$\text{MigRate}_{o,d,t}^{\text{urban}} = \frac{\text{InMigrants}_{o,d,t}^{\text{urban}}}{\text{Pop}_{d,T_c}^{\text{census,urban}}} \times 100.$$

For countries with only five-year migration brackets, we construct the equivalent five-year migration rate as the sum of migrants arriving over $[t - 4, t]$, normalized by the same census-year denominator. In robustness checks, we re-estimate the main results using the annual reconstructed destination population $\text{Pop}_{d,t}$ as denominator in place of the census-year population, to verify that our results are not driven by the choice of a fixed denominator.

2.3 Urbanization data

Admin-unit urbanization. - For our heterogeneity and robustness analyses, we classify IPUMS administrative units by their urban population share at the most recent census:

$$\text{UrbanShare}_a = \frac{\sum_{i \in a, T_c} \text{perwt}_i \cdot \mathbf{1}\{\text{urban}_i = 2\}}{\sum_{i \in a, T_c} \text{perwt}_i}.$$

This admin-unit classification is distinct from the individual-level urban variable: a rural admin unit may contain some urban residents, and a city admin unit some rural enumeration areas. For the cities and rural sample restrictions, we label an admin unit as a *city* if $\text{UrbanShare}_a \geq 0.55$ and as *rural* otherwise.² For the urbanization heterogeneity analysis, we use three groups: rural-like ($\text{UrbanShare}_a < 0.15$), urbanizing cities ($[0.15, 0.55)$), and main cities (≥ 0.55).

We validate this classification against the Africapolis database (OECD/SWAC, 2024), which delineates all urban agglomerations above 10,000 inhabitants in Africa. Matching each destination to its largest agglomeration confirms that the classification orders city sizes correctly: the median largest agglomeration rises from about 47,000 inhabitants in rural-like destinations to 86,000 in urbanizing cities and 880,000 in main cities, and the share of destinations hosting a city of at least 100,000 inhabitants rises from 28% to 43% to 70%. The main-cities group contains the principal cities of each country (Nairobi, Kampala, Yaoundé, Bamako), the urbanizing-cities group classic secondary cities (Nakuru, Garoua, Fort Portal), and the rural-like group mostly small towns. Appendix 9.1.3 reports the full correspondence (Figure 8 and Table 9).

2.4 Housing Quality

Individual definition of housing quality. - The household-level low-quality housing classification follows the UN-Habitat multi-criteria slum definition (UN-Habitat, 2003).³ The classification relies on IPUMS variables covering five housing dimensions grouped in two categories. Structural quality of the dwelling captures the materials of the floor, walls, and

²The 55% threshold is calibrated on the median urban share across our sample of African administrative units; the rolling-window estimates in Figure S4 trace the estimates continuously along the urban-share distribution, showing that no single threshold drives the results.

³The census-based classification captures the material dimension of housing precarity - structural quality and access to basic services - but not the legal or tenure dimension that defines informal settlements in a stricter sense (unauthorized land occupation or unregulated construction). Throughout the paper, low-quality housing refers to the UN-Habitat statistical slum definition rather than to the legal sense of informal settlement.

roof. Access to basic services and living conditions captures sanitation and household crowding (persons per room).⁴ In the main definition, metal-sheet roofing is classified as durable, consistent with the UN-Habitat durability criterion that durable materials protect occupants from climatic conditions; the robustness definition instead counts metal-sheet roofing as deficient, reflecting its use as a low-cost material in the Sub-Saharan African urban context.⁵

A household is classified as residing in low-quality housing if at least one of these dimensions is deficient, following the UN-Habitat operational slum definition, under which a household lacking any one of the deprivation criteria is classified as slum (a logical OR across dimensions) (UN-Habitat, 2003). As a conservative robustness definition, we also use a stricter threshold requiring at least two deficient dimensions. The coding of each dimension - which materials, services, and crowding levels are classified as deficient, in harmonized IPUMS categories identical across countries, with dimension availability varying by census - is reported in Appendix Table 8.

From the household-level classification, two individual-level housing-quality measures are constructed: the binary indicator $\text{deprived}_i \in \{0, 1\}$ equal to 1 if individual i resides in a household classified as low-quality, and the continuous count $\text{deprived_count}_i \in \{0, 1, \dots, 5\}$ equal to the number of deficient housing dimensions for individual i 's household.

Aggregate housing quality outcomes. - The main analysis focuses on origin-destination-year cohorts of migrants. From the individual-level housing-quality classification above, two aggregate outcomes are constructed at the corridor level. The first outcome, used in the main results to preserve the same estimation sample between first and second stages, is the *scale* outcome $S_{o,d,t}$, the low-quality housing inflow rate at destination:

$$S_{o,d,t} = \frac{\sum_i \text{perwt}_i \cdot \mathbf{1}\{i \text{ at } d \text{ in } t_c, \text{ arrived from } o \text{ in } t, \text{ deprived}_i = 1\}}{\text{Pop}_{d,T_c}^{\text{census}}} \times 100,$$

the share of destination population accounted for by migrants in low-quality housing arriving from origin o in year t . The second outcome, designed to isolate the propensity channel, is the *propensity* outcome $P_{o,d,t}$, the share of migrants residing in low-quality housing in each corridor:

$$P_{o,d,t} = \frac{\sum_i \text{perwt}_i \cdot \mathbf{1}\{i \text{ at } d \text{ in } t_c, \text{ arrived from } o \text{ in } t, \text{ deprived}_i = 1\}}{\sum_i \text{perwt}_i \cdot \mathbf{1}\{i \text{ at } d \text{ in } t_c, \text{ arrived from } o \text{ in } t\}}.$$

This second outcome is defined only on cells with at least one migrant, which yields a smaller estimation sample than $S_{o,d,t}$. Analogous urban-restricted versions, $S_{o,d,t}^{\text{urb}}$ and $P_{o,d,t}^{\text{urb}}$, restrict

⁴Water access is not included. The UN-Habitat criterion counts boreholes and protected wells as improved sources, but the harmonized IPUMS water-supply variable in these censuses distinguishes types of piped supply rather than improved from unimproved sources, so it cannot reliably separate deficient from adequate access in this setting. Including it would have classified households with unknown water supply as deprived while leaving genuinely unimproved sources uncounted.

⁵The UN-Habitat definition comprises five deprivations (durable housing, sufficient living area, improved water, improved sanitation, and secure tenure); we adapt it to the dimensions the harmonized census micro-data can measure reliably, and cannot capture tenure security, water access, or, in some countries, sanitation. Our measure therefore reflects the material and service dimensions of housing precarity rather than the full UN-Habitat set.

the numerator to migrants residing in urban enumeration areas at destination and use the urban census-year population as denominator.

Because $S_{o,d,t}^{\text{urb}}$ is a rate per urban resident, destinations with a very small urban population produce rates that are arithmetically valid but economically meaningless – a handful of arrivals divided by a few dozen residents. We therefore restrict the sample to destinations with at least 2,000 urban residents, the conventional minimum size of an urban locality and the threshold underlying the Africapolis classification we use below. The restriction is stated *ex ante*, applies uniformly to every country, and binds on eight destinations (3,752 corridor-years, 1.5% of the panel). Five of those eight are census-to-census reclassification artefacts: their recorded urban population collapses by 77–99% between consecutive censuses while their total population grows, which is not a change in settlement but a change in what the census codes as urban. No outcome is winsorized: winsorizing the numerator of a rate whose denominator is untouched breaks the accounting identity $S = P \times M$ used in the decomposition below, and the size restriction removes the problematic observations at source instead. The estimation panel contains 241,540 corridor-years.

Decomposing scale and propensity. - The scale outcome embeds two distinct mechanisms by construction. The accounting identity

$$S_{o,d,t} = \frac{\text{MigRate}_{o,d,t}}{100} \times P_{o,d,t}$$

makes explicit that a causal effect of drought-induced migration on the low-quality housing inflow can operate through two channels. The *scale channel* reflects the mechanical contribution of drought-induced migrants arriving at the baseline propensity to reside in low-quality housing. The *propensity channel* reflects a shift in that propensity among migrants themselves: drought-induced migrants may be more or less likely to settle in low-quality housing than baseline migrants on the same corridor. Supplementary Appendix Figure S3 illustrates the distinction with a stylized example.

A first-order decomposition of the 2SLS coefficient on the scale outcome yields:

$$\hat{\beta}_S^{2SLS} \approx \underbrace{\bar{P}}_{\beta^{\text{scale}}} + \underbrace{\frac{\overline{\text{MigRate}}}{100} \times \hat{\beta}_P^{2SLS}}_{\beta^{\text{propensity}}},$$

where β^{scale} is the effect that would obtain if the propensity were fixed at its sample mean, and $\beta^{\text{propensity}}$ captures the additional contribution from a shift in migrants' propensity.

Estimating the propensity channel directly requires a separate 2SLS specification with $P_{o,d,t}$ as outcome. This approach faces three identification challenges. First, *ratio bias*: the endogenous migration rate enters the denominator of the propensity outcome, so variation in $P_{o,d,t}$ partly reflects mechanical denominator shifts rather than changes in migrant composition. Second, *measurement noise*: $P_{o,d,t}$ is noisily measured on cells with few migrants; this is addressed by weighting regressions by the time-mean of migrant count per corridor (Section 4). Third, *interpretive ambiguity*: a positive coefficient on $P_{o,d,t}$ cannot be cleanly separated between sorting of migrants into low-quality housing and the mechanical denominator effect just described.

2.5 Other Variables

Aridity index for heterogeneity analyses - For heterogeneity analyses by climate baseline, we use the CGIAR-CSI Global Aridity Index (Trabucco and Zomer, 2019), computed as the long-run mean ratio of precipitation to potential evapotranspiration. This index is extracted from the CGIAR-CSI global 30-arcsec dataset and spatially averaged at the IPUMS administrative level. Lower values indicate more arid conditions; we use within-country medians to classify origins as more arid or less arid, avoiding conflation between aridity and country fixed effects.

2.6 Descriptive Overview

Migration patterns - Appendix Figure 2 displays the average urban in-migration rate over time in the eight countries of the estimation sample, constructed from the yearly bilateral panel. Rates differ by an order of magnitude across countries, and within each census window they rise toward the census year - the mechanical recall and onward-migration attrition of earlier cohorts that the distance-to-census control absorbs in estimation (Section 3). The corresponding five-year-bracket rates for the full set of fifteen IPUMS countries are shown in Supplementary Appendix Figure S1.⁶

Spatial variation in climate and migration - Beyond temporal patterns, the data reveal substantial spatial heterogeneity within and across countries. Appendix Figure 3 maps mean growing-season precipitation and mean drought intensity by administrative unit, illustrating the diversity of agricultural climates spanned by the sample - from the arid regions of northern Mali and Kenya (below 500 mm per season) to the humid coastal zones of southern Cameroon (exceeding 2,500 mm). Appendix Figure 4 maps urban in-migration rates at the administrative unit level, revealing considerable within-country variation in destination patterns, with concentrated flows toward major coastal and capital cities. This spatial granularity allows us to construct origin-destination migration matrices at a fine geographic scale and to match each origin district with its specific climate exposure.

Low-quality housing prevalence - Low-quality housing is widespread among the urban in-migrants we study: across the estimation panel, 41 percent of urban in-migrants in the average corridor reside in housing classified as low-quality under the main UN-Habitat definition, and 36 percent once corridors are weighted by their migrant flow (Appendix Table 7), a magnitude consistent with the prevalence UN-Habitat reports for Sub-Saharan Africa’s urban population (Un-Habitat, 2022). Appendix Figure 5 maps this share by destination, showing pronounced variation both across and within countries. Appendix Figure 7 reports the index dimension by dimension, for urban residents and for in-migrants separately, and Supplementary Appendix Figures S6–S9 show which material and facility categories fall on the deprived side of each rule. Cross-country levels partly reflect census vintage - Malawi and Mali, the highest in the sample, are measured in 2008 and 2009, while Zambia’s are

⁶A potential concern with retrospective migration data is recall bias: individuals may report their year of arrival with diminishing accuracy as the gap between migration year t and census year t_c widens. All specifications include a linear control for the survey-to-arrival gap $|t - t_c|$, absorbing systematic differences in measurement quality across migration cohorts.

measured in its 1990 and 2000 censuses, the earliest we observe - and are absorbed by the corridor fixed effects in estimation; as a further check, Appendix Figure 16 re-estimates the main result leaving out one country at a time: the effect stays positive throughout and every subsample estimate falls inside the full-sample confidence interval, so no single country drives its magnitude.

3 Identification Strategy

3.1 Empirical Model

This paper investigates whether weather-shock-induced migration increases migrant inflows into low-quality housing in African cities and identifies the channels through which this effect operates. Two related questions structure the analysis. First, does drought-induced migration increase the number of migrants arriving in deprived housing, with more migrants arriving and sorting into housing at the usual rate (the *scale* channel)? Second, are drought-induced migrants more likely than other migrants to settle in low-quality housing at destination (the *propensity* channel)? We then use the same design to characterize the labor-market integration of these marginal migrants (Section 4).

To address these questions, we implement an instrumental-variable strategy using the cohort-bilateral panel of migration flows (o, d, t) . The first stage isolates variation in bilateral migration generated by exogenous weather shocks at origin:

$$\text{MigRate}_{o,d,t} = \alpha_1 Z_{o,t} + \gamma_1 |t - t_c| + \delta_{o,d} + \lambda_t + \varepsilon_{o,d,t}, \quad (1)$$

where $Z_{o,t}$ denotes the weather shock at origin o in year t , selected from a set of candidate climatic variables using post-IV LASSO (Section 3.3). The selected instrument is standardized drought intensity,

$$Z_{o,t} = \max(0, -\tilde{z}_{o,t}),$$

as defined in Section 2.1. Origin-destination pair fixed effects, $\delta_{o,d}$, absorb all time-invariant corridor characteristics, so identification relies on within-corridor variation in the timing of droughts. Arrival-year fixed effects, λ_t , absorb shocks common to all corridors, while $|t - t_c|$ controls linearly for the distance between migration year t and the reference census year t_c . Standard errors are clustered at the origin administrative-unit level, where the instrument varies.

The second stage estimates the causal effect of bilateral migration on migrant inflows into low-quality housing at destination:

$$S_{o,d,t} = \beta_1 \widehat{\text{MigRate}}_{o,d,t} + \gamma_2 |t - t_c| + \mu_{o,d} + \lambda_t + \nu_{o,d,t}, \quad (2)$$

where $\widehat{\text{MigRate}}_{o,d,t}$ is the fitted migration rate from Equation 1, and the fixed effects and controls are defined as above. The scale outcome $S_{o,d,t}$, the inflow of migrants into low-quality housing per 100 destination residents, is defined as

$$S_{o,d,t} = \frac{\sum_{i \in \text{Migrants}(o,d,t)} \text{perwt}_i \mathbf{1}\{\text{deprived}_i = 1\}}{\text{Pop}_{d,t_c}^{\text{census}}} \times 100.$$

Because both the outcome and the migration rate are expressed per 100 residents at destination, β_1 measures the share of the marginal drought-induced migrant inflow that settles in low-quality housing. This specification identifies the *scale* channel.

We isolate the *propensity* channel using an analogous specification in which the outcome is the propensity outcome $P_{o,d,t}$, the share of arriving migrants living in deprived housing. A positive coefficient indicates that migrants induced to move by drought are more likely to live in low-quality housing than the average migrant (Section 4). We also estimate urban-restricted versions of both specifications using the urban migration rate, $\text{MigRate}_{o,d,t}^{\text{urb}}$, and the urban scale outcome, $S_{o,d,t}^{\text{urb}}$.

The distance-to-census control, $|t - t_c|$, addresses three sources of measurement error in retrospective migration data. First, recall error may increase with the time elapsed between migration and the census. Second, recent arrival cohorts may be more fully captured than earlier cohorts. Third, the census records migrants at their place of residence at enumeration, so members of earlier cohorts who subsequently moved to another destination or returned to their origin are no longer observed in the original corridor. Migrant counts are corrected for mortality attrition using crude death rates when constructing the panel (Section 2), while $|t - t_c|$ controls for remaining attrition that varies systematically with the distance to the census. Because these sources of measurement error affect both the migration rate and the outcome, the same control enters both stages.

3.2 Discussion

Interpretation of the IV estimate - The 2SLS estimate β_1 in equation (2) identifies a Local Average Treatment Effect: the causal effect of bilateral migration on the low-quality housing inflow for the subpopulation of drought-induced migrants - the marginal migrants whose decision to move from origin o in year t responds to drought at origin, the compliers. The estimate does not identify the average treatment effect across all migrants, but the effect on the climate-responsive margin.

Discussion of the compliers - In a standard application, these compliers are a technical by-product of the instrument, and the LATE is often viewed as a limitation relative to an average effect over all migrants. Here the interpretation is more direct: the population the instrument selects - households induced to migrate by drought - *is* the population of interest, since the paper asks how climate shocks reshape urban housing through migration. The LATE is therefore the estimand of interest rather than a nuisance parameter to be reconciled with a broader average. This alignment also gives the propensity results their meaning: the same drought variation that identifies the causal effect of migration on housing also characterizes the housing outcomes of climate migrants, because these are one and the same group.

Validity of the instrument and exclusion restriction - Identification rests on two key assumptions. The relevance condition requires that drought intensity predicts bilateral mi-

gration; the strong first-stage relationships reported in Section 3.4 ($F > 20$) confirm this. The exclusion restriction requires that drought at origin o affects the low-quality housing inflow at destination d only through its effect on migration from o to d , conditional on origin-destination pair and arrival-year fixed effects. This assumption is plausible in our setting: drought is a local agricultural shock at origin, and the origin-destination fixed effects absorb any time-invariant correlation between origin climate and destination housing conditions. Potential threats to this assumption - including direct effects of origin drought on destination outcomes via food prices or return migration - are discussed in Section 3.4.

3.3 Instrument Selection: Post-IV LASSO

The choice of the instrument $Z_{o,t}$ is validated by a data-driven selection procedure based on the post-IV LASSO of Belloni et al. (2014). The procedure runs two LASSO regressions on a candidate set of climate variables at origin - one for the endogenous migration rate and one for the outcome, both on Frisch-Waugh-Lovell residuals that partial out the origin-destination pair and arrival-year fixed effects and the distance-to-census control - and retains the union of the selected sets. This removes the discretionary choice of instrument among ex-ante plausible candidates and provides valid post-selection inference under the conditions of Belloni et al. (2014).

The candidate set comprises thirteen contemporaneous climate variables described in Section 2.1 - drought intensity, drought dummies at two thresholds, dry days, precipitation and temperature polynomials, degree-day and evapotranspiration measures - augmented with the lags and 3-season means of the drought-family variables (25 candidates in total). We implement the selection under three values of the penalty multiplier c , from the theoretical benchmark of Belloni et al. (2014) ($c = 1.1$) to more permissive thresholds ($c = 0.5, 0.25$) that let more instruments enter.

The selection confirms the paper’s instrument. Drought intensity is the strongest single candidate on both samples ($F = 23.9$ on the scale sample, $F = 13.0$ on the propensity sample) and is selected at the default penalty in both; among lagged candidates its 3-season mean dominates ($F = 27.4$), consistent with the cumulative instrument used throughout. We use the single strongest candidate, keeping the model just-identified so that Anderson-Rubin inference remains exact. As a placebo, we repeat the exercise on wet-side mirror candidates: excess-rainfall intensity has essentially no predictive power ($F = 4.0$ and $F = 0.03$), and the flood dummies are similarly weak. The instrument’s strength is thus specific to the dry tail of the rainfall distribution. Appendix 9.2.1 reports the full selection results (Figure 9).

3.4 First Stage

Table 1 reports the first stage under alternative sets of fixed effects and controls, for the sample of urban migrations to all destinations. The table shows the strength and the stability of the instrument for in-migration towards urban areas. Column (1) is the paper’s main specification: origin-destination pair and arrival-year fixed effects with the distance-to-census control. Comparing columns (1) and (2) shows that this control is not driving the first stage: the raw specification in column (2), without any control, already delivers a positive and highly significant coefficient, and adding the distance to the census mainly improves

Table 1: First Stage: Drought Intensity and Urban Migration

	<i>DV: Urban migration rate</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Main spec.					
Drought intensity	0.00362*** (0.00066)	0.00244*** (0.00068)	0.00396*** (0.00076)	0.00389*** (0.00071)	0.00400*** (0.00077)	0.00159** (0.00073)
First-stage F	29.65	12.72	26.77	29.67	26.91	4.74
Effect of 1-SD drought (% mean)	+6.8%	+4.6%	+7.4%	+7.3%	+7.5%	+3.0%
Origin FE	No	No	No	No	No	No
Destination FE	No	No	No	No	No	No
Origin-Destination FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country $\times$ Year FE	No	No	No	No	No	No
Destination $\times$ Year FE	No	No	No	No	No	Yes
Dist. to census control	Yes	No	Yes	Yes	Yes	No ^a
Destination controls	No	No	No	Yes	Yes	No
Origin controls	No	No	Yes	No	Yes	No
R ²	0.486	0.481	0.486	0.486	0.486	0.512
Within R ²	0.010	0.000	0.010	0.010	0.010	0.000
Mean migration rate	0.026820	0.026820	0.026820	0.026820	0.026820	0.026820
Observations	241,540	241,540	241,540	241,540	241,540	241,540
OD pairs	16,846	16,846	16,846	16,846	16,846	16,846

Notes: Sample: urban migrants, all destinations; 8 countries, 12-year window. Instrument: contemporaneous drought intensity at origin, $\max(0, -\tilde{z})$. Column (1) is the paper's main specification; fixed effects and controls as indicated (^a`dist_census` mechanically absorbed). First-stage F : cluster-robust Wald statistic on the instrument. Effect of 1-SD drought: coefficient $\times$ SD(instrument) $\div$ mean migration rate $\times$ 100, on each column's own sample. Standard errors, clustered at `admin_origin`, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

precision. Column (3) confirms that the instrument captures rainfall deficits rather than heat, and column (4) adds temperature at destination to absorb a potential climatic pull factor; column (5) includes all controls jointly. In all three cases the coefficient is essentially unchanged and the first stage remains strong. When accounting for destination-by-year fixed effects (column (6)), the coefficient remains positive and significant at the 5% level, but the first-stage F drops below conventional thresholds, as these fixed effects absorb a large share of the identifying variation, rainfall shocks being spatially correlated across origins and destinations within a country.

Quantitatively, the main specification in column (1) implies that a one standard deviation increase in the dryness of the growing season at origin raises in-migration to urban areas by 6.8% of the mean migration rate. Let's consider a destination of 10,000 urban inhabitants. At the mean migration rate, this corridor receives 2.68 migrants in a normal year. A one standard deviation drier growing season raises the migration rate by $0.00362 \times \text{SD}(\text{drought})$, which scaled by the destination population adds about 0.18 migrants — bringing the inflow from 2.68 to 2.86, a 6.8% increase.

This magnitude is in line with existing estimates of drought-induced migration in sub-Saharan Africa (Henderson et al., 2017; Mueller et al., 2020; Gittard, 2024; Hoffmann et al., 2020). Relative to this literature, we focus on in-migration and identify the first stage within origin-destination corridors at annual frequency: a corridor receives more migrants precisely in the years its origin is hit by drought, net of fixed corridor characteristics and common year shocks.

Table 10 in Appendix 9.3 assesses the robustness of the first stage along two dimensions: the sample and the drought exposure window, reporting the paper's main specification throughout. Columns (3)-(4) retain only destinations classified as main cities.⁷ Within this subsample, the first stage remains positive and statistically significant, although weaker: a one-standard-deviation increase in drought intensity raises in-migration by 6.4% of the sample mean, with a first-stage F -statistic of 10.5. Columns (5)-(6) further restrict the sample to migrants originating from rural areas and moving to these main cities, thereby isolating the rural-to-urban migration margin most relevant to our mechanism. In this subsample, the first stage is strongest among the restricted samples: a one-standard-deviation increase raises in-migration by 7.7% of the sample mean ($F = 21.7$). The even-numbered columns re-estimate each first stage with cumulative drought intensity, averaged over the migration year and the two preceding seasons. The cumulative measure is a stronger predictor in the full sample - a one-standard-deviation increase raises in-migration by 10% of the sample mean, against 6.9% for the contemporaneous measure - consistent with an asset-erosion mechanism, whereby households migrate after several consecutive poor seasons rather than in response to a single shock. We nonetheless retain the contemporaneous measure as our preferred instrument, for two reasons : the cumulative first stage collapses in the restricted samples (columns (4) and (6)), while the contemporaneous one remains strong; and cumulative exposure raises exclusion-restriction concerns - a multi-year accumulation is more likely to affect destination outcomes through channels other than the timing of migration - that are

⁷Main cities are defined as destinations for which `share_urban_known_dest` ≥ 0.55 , measured at the country's most recent census wave; rural origins are the complement, with an urban share strictly below the threshold or unmeasured.

less salient for a single-season shock. Across both sample restrictions, the contemporaneous instrument remains a robust predictor of urban in-migration.

3.5 Reduced Form

Table 2 and Table 11 report the reduced-form relationship between drought at origin and the scale and propensity outcomes, respectively, estimated as

$$Y_{odt} = \beta D_{ot} + \gamma_1 |t - t_c| + \gamma_{od} + \delta_t + \varepsilon_{odt}, \quad (3)$$

where Y_{odt} is either the low-quality housing inflow rate S_{odt} (the scale outcome) or the share of migrants in low-quality housing P_{odt} (the propensity outcome) in corridor (o, d) at time t ; D_{ot} is drought intensity at origin; γ_{od} and δ_t are origin-destination pair and year fixed effects; and $|t - t_c|$ is the distance-to-census control. The reduced form regresses the outcome directly on the instrument, without instrumenting migration; it captures the total effect of a drought shock at origin on the outcome and rests only on the exogeneity of rainfall variation at origin.

Residential status - Scale - Columns (1)-(4) use the contemporaneous drought instrument and columns (5)-(8) its cumulative (three-season mean) counterpart on the lag-restricted sample; within each group, the first column is our main specification and the remaining columns progressively saturate the fixed effects, up to destination-by-year effects that absorb all destination-year shocks. The table shows the stability of the estimate: under the contemporaneous instrument the coefficient is essentially identical across all four configurations, indicating that the relationship is driven by within-corridor variation in drought rather than by fixed differences across origins or destinations or by destination-level shocks. The main specification in column (1) implies that a one standard deviation increase in drought intensity at origin raises the low-quality housing inflow rate by 3.3% of its mean; under the cumulative instrument (column (5)), the same shock raises it by 6.7%, consistent with several consecutive poor seasons displacing more households than a single one.

Residential status - Propensity - Beyond raising the number of arrivals, drought may shift *who* arrives toward more housing-vulnerable migrants. Table 11 in Appendix 9.3.2 reports the reduced form for this propensity channel: the direct effect of drought at origin on the composition of the migrant inflow. Columns (1)-(2) use the contemporaneous drought instrument and columns (3)-(4) the cumulative (three-season mean) instrument on the lag-restricted sample. Odd columns report the extensive margin, based on the paper's main dependent variable: the share of migrants in corridor (o, d, t) living in low-quality housing,

$$P_{o,d,t} = \frac{\sum_{i \in \text{migrants}(o,d,t)} \text{perwt}_i \cdot \mathbf{1}\{\text{slum}_i = 1\}}{\sum_{i \in \text{migrants}(o,d,t)} \text{perwt}_i},$$

where slum_i equals one if migrant i 's household lacks at least one housing dimension under the main definition. Even columns report the depth of deprivation: the mean number of deficient housing dimensions among the same migrants. With the contemporaneous instrument, a one standard deviation drought raises the share of arriving migrants in low-quality housing by

Table 2: Reduced Form: Drought Intensity and Low-Quality Housing Inflows, Both Instruments

	<i>Contemporaneous drought</i>				<i>3-season mean drought (lagged)</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Main spec.				Main spec.			
Drought intensity	0.00103*** (0.000350)	0.00106*** (0.000354)	0.000546 (0.000431)	0.000546 (0.000437)	0.00266*** (0.000735)	0.00278*** (0.000751)	0.00151 (0.000999)	0.00150 (0.00101)
Effect of 1-SD drought (% mean)	+4.3%	+4.5%	+2.3%	+2.3%	+6.6%	+6.9%	+3.8%	+3.7%
Origin FE	No	Yes	No	No	No	Yes	No	No
Destination FE	No	Yes	No	No	No	Yes	No	No
Origin-Destination FE	Yes	No	Yes	Yes	Yes	No	Yes	Yes
Year FE	Yes	Yes	No	Yes	Yes	Yes	No	Yes
Country $\times$ Year FE	No	No	Yes	No	No	No	Yes	No
Destination $\times$ Year FE	No	No	No	Yes	No	No	No	Yes
Dist. to census control	Yes	Yes	No ^a	No ^a	Yes	Yes	No ^a	No ^a
Destination controls	No	No	No	No	No	No	No	No
Origin controls	No	No	No	No	No	No	No	No
R ²	0.370	0.088	0.372	0.393	0.383	0.091	0.386	0.407
Within R ²	0.003	0.002	0.000	0.000	0.003	0.002	0.000	0.000
Mean outcome	0.01196	0.01196	0.01196	0.01196	0.01189	0.01189	0.01189	0.01189
Observations	241,540	241,540	241,540	241,540	235,920	235,920	235,920	235,920
OD pairs	16,846	16,846	16,846	16,846	16,846	16,846	16,846	16,846

Notes: Reduced form: the low-quality-housing inflow rate (main UN-Habitat definition, per 100 urban residents) regressed directly on drought intensity at origin. Columns (1)–(4) use the contemporaneous instrument; columns (5)–(8) its cumulative (three-season mean) counterpart on the lag-restricted sample. Within each group, the first column is the paper’s main specification and the remaining columns saturate the fixed effects; the year-FE-only specification is omitted (its coefficient is several times larger and is reported in the replication package). ^aThe distance-to-census control is mechanically absorbed under country $\times$ year or destination $\times$ year fixed effects. Effect of 1-SD drought: coefficient $\times$ SD(instrument) $\div$ mean outcome $\times$ 100, on each column’s own sample. Sample: urban migrants, all destinations; 8 countries, 12-year window; SEs clustered at *admin.origin*, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

0.34 percentage points, off a 46% baseline, while the depth of deprivation moves little. The cumulative instrument delivers a stronger and broader pattern: the extensive-margin shift roughly doubles and the depth of deprivation also rises. Because the reduced form does not rely on the first stage, these lagged estimates remain interpretable despite the weak first stage of the cumulative instrument. These results indicate that persistent drought sends not only more migrants but also sends migrants that are more likely to live in low-quality housing.

Figure 11 in Appendix 9.3.3 plots binned migration and low-quality housing inflows against drought intensity at origin, after residualizing on the main specification’s fixed effects and controls. Both the first-stage and reduced-form relationships are close to linear and hold across the full distribution of drought shocks, rather than being driven by extreme outliers.

Labor Statuts - Tables 12 and 13 report the reduced forms for the employment results, regressing the scale and propensity outcomes directly on origin drought for both instruments. On the scale margin, all six inflows rise significantly at the 1% level: a one-standard-deviation drought raises the employed, unemployed, and inactive inflows by 6.8%, 10.7%, and 4.9% of their means with the contemporaneous instrument, and by 8.2%, 21.3%, and 8.2% with the cumulative one. On the propensity margin, the same shock shifts the composition of arrivals by about half a percentage point from inactivity toward employment under both instruments (+0.43 and -0.51 pp contemporaneous; +0.42 and -0.56 pp lagged), while the unemployed share barely moves. Since the reduced form does not rely on the first stage, the lagged columns provide valid inference where the lagged first stage of Table 17 is weak, and confirm the same reallocation.

4 Main Results

4.1 Residential Integration

4.1.1 Scale

Table 3 reports the main estimates of the causal effect of drought-induced migration on the urban low-quality housing inflow (the scale outcome) at destination. Columns (1)-(3) report the OLS, first-stage and 2SLS estimates using the contemporaneous drought instrument - column (3) is the paper’s main estimate - and columns (4)-(6) report the same estimates using the cumulative instrument, on the sample restricted to complete lag coverage. The sample is urban migrants to all destinations, with origin–destination pair and year fixed effects and the distance-to-census control throughout.

The main scale result shows that drought-induced migration causally raises the inflow into low-quality housing at destination. The main scale result is the 2SLS coefficient in column (3): a 1 pp increase in the drought-induced urban migration rate raises the low-quality housing inflow rate at destination by 0.285 pp. Since both rates are expressed per hundred urban residents, the coefficient reads directly as a ratio of flows: for every additional hundred migrants that drought pushes into a destination, about 28 settle in low-quality housing (see Table 14 for more details on the calculation).

Table 3: Drought-Induced Migration and Low-Quality Housing Inflows: Main Results

	<i>Contemporaneous drought</i>			<i>3-season mean drought (lagged)</i>		
	OLS (1)	FS (2)	2SLS (3)	OLS (4)	FS (5)	2SLS (6)
Urban migration rate	0.46865*** (0.02498)		0.28455*** (0.06323)	0.46443*** (0.02536)		0.28082*** (0.04813)
Drought intensity		0.00362*** (0.00066)			0.00947*** (0.00164)	
First-stage F			29.65			33.42
AR 95% CI			[0.132, 0.397]			[0.171, 0.370]
Effect of 1-SD drought (% mean)	+7.1%	+6.8%	+4.3%	+11.0%	+10.5%	+6.6%
Origin-Destination FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dist. to census control	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.697	0.486	0.647	0.702	0.499	0.652
Within R ²	0.521	0.010	0.441	0.518	0.011	0.437
Mean outcome		0.011964			0.011892	
Mean migration rate		0.026820			0.026804	
SD(drought intensity)		0.502			0.297	
Observations		241,540			235,920	
OD pairs		16,846			16,846	

Notes: Columns (1)–(3): contemporaneous instrument; columns (4)–(6): 3-season mean instrument on the lag-restricted sample, reported as robustness – column (3) is the paper’s main estimate. OLS and 2SLS columns report the effect of the urban migration rate on the low-quality housing inflow rate (main UN-Habitat definition: at least one deprived dimension among floor, wall, roof, sanitation and crowding, metal roofs durable). Water is excluded from the index: the harmonized IPUMS water-supply codes in these censuses cannot reliably separate deficient from adequate access. FS columns report the first stage of the migration rate on drought intensity, reproduced from Table 1 and its lagged counterpart. First-stage F : cluster-robust Wald on the instrument. AR 95% CI: cluster-robust Anderson–Rubin confidence interval. Effect of 1-SD drought: for OLS/2SLS columns, coefficient $\times$ first-stage coefficient $\times$ SD(instrument) $\div$ mean outcome $\times$ 100; for FS columns, coefficient $\times$ SD(instrument) $\div$ mean migration rate $\times$ 100. Sample: urban migrants, all destinations; 8 countries (Egypt excluded), 12-year window, Mali 1998 included. SEs clustered at *admin_origin*, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

To interpret the magnitude of the 2SLS estimate, consider a representative origin-destination corridor with 10,000 urban residents at destination. At the baseline urban migration rate of 0.0268 migrants per 100 residents (the sample mean reported in Table 3), this corridor receives $0.0268 \times 100 = 2.68$ migrants in a year without drought, of whom $0.0120 \times 100 = 1.20$ settle in low-quality housing, the mean inflow rate reported in the same table. Now let drought intensity at origin rise by one standard deviation ($SD = 0.502$, as reported in the table). The first stage translates this shock into migration: the migration rate rises by $\beta_{FS} \times SD = 0.00362 \times 0.502 = 0.0018$, or 6.8% of its mean, where β_{FS} is the first-stage coefficient from column (2) of Table 3 - 0.18 additional migrants for the corridor, bringing the inflow to 2.86. Combining the two stages translates the same shock into housing: the low-quality housing inflow rate rises by $\beta_{2SLS} \times \beta_{FS} \times SD = 0.285 \times 0.00362 \times 0.502 = 0.00052$ - equivalently, the reduced-form coefficient of Table 2 times SD , since with a single instrument the reduced form equals the product of the two stages - or 0.052 additional residents in low-quality housing. The ratio of the two changes is, by construction, the 2SLS coefficient itself: because the low-quality housing inflow and the migration rate share the same denominator - destination urban population - $\beta_{2SLS} = 0.052/0.18 = 0.28$ reads directly as the share of the marginal inflow that settles in low-quality housing. *Of every hundred migrants that drought pushes into a city - migrants who would not have moved otherwise - roughly twenty-eight end up in low-quality housing.* This 28% is the central quantitative result of the paper; Appendix Table 14 reports the calculation step by step.

Column (6) reports the 2SLS estimate under the cumulative instrument: 0.281, significant at the 1% level, with a strong first stage ($F = 33.42$) and an Anderson-Rubin interval of $[0.17, 0.37]$. The interpretation is the same ratio reading as for the main estimate, applied to a different set of marginal migrants: of every hundred migrants displaced by *persistent* drought - three consecutive poor seasons rather than a single shock - about 28 settle in low-quality housing, the same proportion as for those displaced by a single-season shock. The two estimates are statistically indistinguishable and each lies well inside the other's Anderson-Rubin interval. The effect therefore does not appear to intensify with the persistence of the shock: households pushed out by repeated poor seasons are absorbed into low-quality housing at the same rate as those pushed out by a single one.

The OLS estimates (columns (1) and (4)) give the raw within-corridor correlation between migration and low-quality housing inflows, without the causal interpretation of the 2SLS estimates. They measure the effect of an additional percentage point of migration within the origin-destination corridor on the low-quality housing inflow rate at destination. Because corridors with more migration may differ systematically in ways correlated with housing quality, this reflects a correlation rather than the causal effect of drought-induced migration. Under the contemporaneous instrument, the OLS estimate in column (1), 0.229, is nearly identical to the 2SLS estimate: a 1 pp increase in migration flows within corridors implies that about 23 of every 100 additional migrants settle in low-quality housing. Under the cumulative instrument, the 2SLS estimate in column (6) is larger than its OLS counterpart in column (4) (0.315 against 0.219), consistent with persistent drought selecting more housing-vulnerable migrants, though the difference is imprecisely estimated and we do not push this comparison. Neither comparison directly answers whether the migrants that drought adds are *individually* more likely to end up in low-quality housing than those who would have moved anyway - a question about the composition of the inflow, which the propensity

analysis addresses next.

Table 15 in Appendix 9.4 restricts the sample first to destinations that are main cities, and then to rural-to-main-city corridors. The effect is robust across both restrictions and is strongest along the rural-to-urban margin, consistent with the agricultural-push mechanism at the core of the paper: drought displaces rural households whose livelihoods depend most directly on rainfall, and these are the migrants who most raise urban low-quality housing inflows.

4.1.2 Propensity

Table 4: Propensity for Low-Quality Housing: Main Results

	<i>Contemporaneous drought</i>				<i>3-season mean drought (lagged)</i>			
	Deprived		Number of deprivations		Deprived		Number of deprivations	
	Dummy - Main def.		Continuous		Dummy - Main def.		Continuous	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)	OLS (7)	2SLS (8)
Urban migration rate	0.02983** (0.01509)	1.21591** (0.59790)	0.01818 (0.03971)	0.94933 (1.06501)	0.03382** (0.01534)	1.98110** (0.92526)	0.03086 (0.03767)	3.19449* (1.63623)
First-stage F	15.50		15.50		8.24		8.24	
AR 95% CI	[0.17, 2.94]		[-1.27, 3.52]		[0.64, 6.47]		[0.60, 10.48]	
Effect of 1-SD drought	+0.01 pp	+0.37 pp	+0.0001	+0.0029	+0.01 pp	+0.58 pp	+0.0001	+0.0093
R ²	0.392	0.295	0.450	0.436	0.396	0.141	0.455	0.303
Within R ²	0.007	-0.151	0.002	-0.023	0.007	-0.411	0.002	-0.278
Mean outcome (average migrant)	0.406	0.406	0.662	0.662	0.404	0.404	0.656	0.656
Mean migration rate	0.089	0.089	0.089	0.089	0.089	0.089	0.089	0.089
Implied marginal migrant		0.515		0.747		0.580		0.940
Observations	72,723	72,723	72,723	72,723	71,207	71,207	71,207	71,207
OD pairs	10,537	10,537	10,537	10,537	10,510	10,510	10,510	10,510

Notes: Shares among urban in-migrants under the main UN-Habitat definition – at least one deprived dimension among floor, wall, roof, sanitation and crowding, metal roofs durable – and the mean number of deprived dimensions. Water is excluded: the harmonized IPUMS water-supply codes in these censuses cannot reliably separate deficient from adequate access. 2SLS with migration instrumented by contemporaneous (cols (1)–(4)) or 3-season mean (cols (5)–(8), lag-restricted sample) drought. Weak lagged first stage ($F \approx 7$): lagged point magnitudes not interpretable – AR intervals only. Main specification (od-pair and year FE, `dist_census` control); SEs clustered at `admin_origin`. Effect of 1-SD drought in pp of the share / dimension units. ‘Mean outcome (average migrant)’ and ‘Mean migration rate’ are means of each column’s own estimation sample. ‘Implied marginal migrant’ is $\bar{P} + \beta \bar{M}$, the propensity of the marginal migrant himself: since $S = P \times M$, the 2SLS coefficient equals $(p_{\text{marg}} - \bar{P})/\bar{M}$. Both means are taken on that same sample – never on the full panel, where the mean migration rate is three times smaller and would shrink the implied gap by the same factor. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4 examines the propensity channel: whether the migrants that drought adds at the margin – the compliers – are more likely to settle in low-quality housing than the average migrant already on the corridor. A positive coefficient indicates that the drought-induced compliers are more housing-vulnerable than the infra-marginal migrants who would have moved regardless. We estimate the same specification as for the scale outcome, replacing

the inflow rate with the composition of the inflow: the share of in-migrants living in low-quality housing under the paper’s main definition, $P_{o,d,t}$ (the extensive margin, columns (1)-(2) and (5)-(6)), and the mean number of deprived housing dimensions among the same migrants (the depth of deprivation, columns (3)-(4) and (7)-(8)). Columns (1)-(4) use the contemporaneous drought instrument and columns (5)-(8) the cumulative instrument on the lag-restricted sample. Robustness across alternative deprivation definitions is reported in Appendix Table 21.

The main result of this section is that drought shifts the composition of arrivals toward low-quality housing, and not only their number. The main propensity result is the 2SLS estimate in column (2): a 1 pp increase in drought-induced migration raises the share of migrants living in low-quality housing by 1.216 pp (significant at the 5% level, with a first-stage F of 15.50 and an Anderson-Rubin interval of $[0.17, 2.94]$ excluding zero). Two readings follow. For a realistic shock, a one-standard-deviation drought raises the share of arriving migrants living in low-quality housing by 0.37 pp, from 40.6% to 41.0%. Scaling instead by the mean migration rate gives the propensity of the marginal migrant himself: since $S = P \times M$, the coefficient equals $(p_{\text{marg}} - \bar{P})/\bar{M}$, so the marginal drought-induced migrant settles in low-quality housing with probability $\bar{P} + \beta\bar{M} = 0.406 + 1.216 \times 0.089 = 0.515$. He is thus about 11 percentage points more likely to live in low-quality housing than the average migrant arriving on the same corridor – 51.5% against 40.6%.⁸ The depth of deprivation, by contrast, does not respond to a single-season shock (column (4)): contemporaneous drought shifts *whether* migrants end up in low-quality housing, not *how* deprived that housing is.

Under the cumulative instrument, the effect appears on both margins: persistent drought raises both the share of migrants in low-quality housing (column (6)) and the depth of their deprivation (column (8)), the latter now significant unlike under the single-season shock. Because the first stage is weak on this subsample ($F = 8.24$), inference rests on the Anderson-Rubin intervals, which remain valid under weak instruments and exclude zero on both margins ($[0.64, 6.47]$ and $[0.60, 10.48]$); the point magnitudes below are therefore indicative. The estimate in column (6) implies that a 1 pp increase in persistent-drought-induced migration raises the share of migrants living in low-quality housing by 1.98 pp; for a realistic shock, a one-standard-deviation cumulative drought raises the share of arriving migrants in low-quality housing by 0.58 pp, off a baseline of 40.4% – about half again the contemporaneous shift. The depth of deprivation now responds as well (column (8)): a one-standard-deviation cumulative drought raises the mean number of deprived housing dimensions among arriving migrants by 0.009, off a baseline of 0.66. Migrants displaced by persistent drought are thus not only more often in low-quality housing, but in worse housing.

Reassuringly, these patterns are consistent with the reduced-form estimates (Table 11), which do not depend on first-stage strength, though they cannot net out the migration channel. Persistent drought thus displaces a more housing-vulnerable population across a broader range of margins, consistent with repeated poor seasons eroding household resources before departure.

⁸Both means are taken on the estimation sample of the propensity regression itself – the corridor-years with a positive urban flow, where the share is defined and where $\bar{M} = 0.089$. Using the full-panel mean (0.027) would be inconsistent, since the coefficient is not identified on that sample; it would shrink the implied gap by a factor of three. The same rule is applied to every marginal-versus-average reading in the paper, each on its own sample.

The OLS coefficient in column (1) is small (0.040): across all migration, larger inflows barely change the deprived-housing share of arrivals - ordinary migration comes at roughly the average housing mix. The 2SLS estimate is more than twenty times larger. Unlike at the scale margin, where OLS and 2SLS coincide, drought-induced migration shifts the *composition* of the inflow far more than migration does on average. This shows the selection result: drought does not add migrants at the average housing mix, but migrants who are disproportionately housing-vulnerable.

For the restricted samples, the propensity point estimates are consistent with the full-sample result but too imprecise to establish it independently (Appendix Table 16).

4.2 Labor-Market Integration

The housing results show that drought-induced migration raises the inflow into low-quality housing. In this section, we address a complement: whether and how drought-induced migrants - the compliers of our IV estimate - are absorbed into the urban labor market, in terms of employment status: employed, unemployed, and inactive. For all three, we use the same identification strategy, counting the number of migrants within the corridor (o, d, t) that are employed, unemployed, and inactive per 100 urban residents, and we measure both the scale and the propensity channel.

4.2.1 Scale

Table 5: Drought-Induced Migration and the Labor-Force Status of the Inflow: Scale

	<i>Employed</i>		<i>Unemployed</i>		<i>Inactive</i>	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
Urban migration rate	0.28216*** (0.01362)	0.42760*** (0.04330)	0.03573*** (0.00416)	0.07477*** (0.01391)	0.48244*** (0.01048)	0.35477*** (0.04038)
First-stage F		29.65		29.65		29.65
AR 95% CI		[0.349, 0.531]		[0.051, 0.110]		[0.254, 0.425]
Effect of 1-SD drought (% mean)	+4.8%	+7.3%	+4.8%	+10.1%	+7.4%	+5.5%
R^2	0.725	0.658	0.372	0.284	0.819	0.788
Within R^2	0.480	0.354	0.106	-0.019	0.711	0.661
Mean outcome	0.0106	0.0106	0.00134	0.00134	0.0118	0.0118
Observations	241,540	241,540	241,540	241,540	241,540	241,540

Notes: Outcomes: employed / unemployed / inactive in-migrants per 100 urban residents, unwinsorized. 2SLS: migration instrumented by contemporaneous drought. Main specification (od-pair and year FE, `dist_census` control); SEs clustered at `admin_origin`. Undetermined labor-force status excluded from the known-status denominator. Sample: urban migrants, all destinations; 8 countries, 12-year window. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5 addresses the scale margin, replacing the outcome of the main specification with the number of employed, unemployed, and inactive in-migrants per 100 urban residents;

migrants whose labor-force status is undetermined in the census are excluded from the counts. The specification, instrument, and sample are otherwise identical to Table 3.

All three inflows increase significantly with migration, both in OLS and 2SLS. The 2SLS coefficient for the employed inflow is 0.428, significant at the 1% level: an increase in the migration rate of one in-migrant per 100 urban residents raises the number of employed in-migrants by 0.428 per 100 urban residents. The corresponding coefficients are 0.075 for the unemployed inflow and 0.355 for the inactive inflow, both also significant at the 1% level. The first stage is unchanged from the main table ($F = 29.65$), and the Anderson-Rubin intervals are bounded away from zero for all three outcomes. Because each outcome shares the denominator of the migration rate, these coefficients read directly as the labor-force status of the marginal migrants themselves: of every 100 additional drought-induced migrants, about 43 arrive into employment, 7 into unemployment, and 35 into inactivity, the remainder having undetermined status.

In relative terms, a one-standard-deviation drought raises the employed, unemployed, and inactive inflows by 7.3%, 10.1%, and 5.5% of their respective means. The larger relative rise in unemployment comes off a very small base - unemployed migrants are a negligible share of the inflow in levels - so in absolute terms employment dominates: the urban labor market absorbs a substantial part of the drought-induced inflow into work. These migrants find employment far more often than they join the ranks of the unemployed, even as they disproportionately settle in low-quality housing.

In comparison, the OLS estimates are smaller (0.097 for the employed, 0.009 for the unemployed, and 0.160 for the inactive inflow), and their ordering differs: in OLS the inactive inflow is the largest component of migration, whereas in 2SLS it is employment. The OLS coefficients pool all variation in migration - including flows unrelated to drought and measurement error in census-recalled moves, which attenuates the estimates - while the 2SLS isolates the drought-driven component. The comparison therefore suggests that the drought-induced inflow is more work-oriented than the typical migration flow, though the two estimates answer different questions and the gap cannot be read as a measure of bias. Whether drought-induced migrants indeed differ from the average arrival in their labor-force status is precisely what the propensity margin tests directly, holding the size of the inflow fixed; we turn to it next.

4.2.2 Propensity

Table 6 turns to the propensity margin: the outcomes are now the shares of employed, unemployed, and inactive migrants among in-migrants with known labor-force status. As in the housing propensity analysis, the denominator is the number of migrants rather than the urban population, so the estimates hold the size of the inflow fixed and isolate changes in its composition. Because the three shares sum to one, the 2SLS coefficients sum to zero by construction: drought can only reallocate migrants across statuses. The specification and instrument are identical to the main table; on this smaller sample of corridors with observed labor-force status, the first stage remains strong ($F = 15.51$).

Drought tilts the inflow toward employment. The 2SLS coefficient on the employed share is 1.395, significant at the 5% level, and its mirror image is a decline in the inactive share of -1.685 , also significant at the 5% level; the unemployed share is unchanged (0.291,

Table 6: Drought-Induced Migration and the Labor-Force Status of the Inflow: Propensity

	<i>Employed</i>		<i>Unemployed</i>		<i>Inactive</i>	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
Urban migration rate	-0.12996*** (0.01513)	1.39452** (0.61688)	-0.00449 (0.00405)	0.29061 (0.24176)	0.13445*** (0.01591)	-1.68513** (0.65293)
First-stage F		15.51		15.51		15.51
AR 95% CI		[0.397, 3.317]		[-0.157, 0.947]		[-3.807, -0.673]
Effect of 1-SD drought	-0.04 pp	+0.42 pp	-0.00 pp	+0.09 pp	+0.04 pp	-0.51 pp
R^2	0.331	0.149	0.243	0.205	0.326	0.061
Within R^2	0.005	-0.264	0.000	-0.049	0.007	-0.384
Mean outcome (average migrant)	0.5055	0.5055	0.0578	0.0578	0.4367	0.4367
Mean migration rate	0.089	0.089	0.089	0.089	0.089	0.089
Implied marginal migrant		0.630		0.084		0.286
Observations	71,503	71,503	71,503	71,503	71,503	71,503

Notes: Outcomes: employed / unemployed / inactive shares among known-status in-migrants; the three shares sum to one, so the 2SLS coefficients sum to zero by construction. ‘Mean outcome (average migrant)’ and ‘Mean migration rate’ are means of this table’s own estimation sample – the corridor-years with known labor-force status, not the full panel. ‘Implied marginal migrant’ is $\bar{P} + \beta \bar{M}$, the marginal migrant’s own share, both means taken on that sample. 2SLS: migration instrumented by contemporaneous drought. Main specification (od-pair and year FE, `dist_census` control); SEs clustered at `admin_origin`. Undetermined labor-force status excluded from the known-status denominator. Sample: urban migrants, all destinations; 8 countries, 12-year window. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

not significant, with an Anderson-Rubin interval that includes zero). For a realistic shock, a one-standard-deviation drought shifts the composition of arrivals by +0.42 percentage points toward employment and -0.51 percentage points away from inactivity. Applying the same conversion as on the housing margin, and again with both means taken on this table's own estimation sample ($\bar{M} = 0.089$ among corridor-years with known labor-force status), the marginal migrant is employed with probability $0.506 + 1.395 \times 0.089 = 0.630$: about 12 percentage points more likely to be employed than the average arrival - 63.0% against 50.6% - and about 15 percentage points less likely to be inactive, 28.6% against 43.7%. The unemployment margin is unchanged.

The OLS estimates again differ from the 2SLS in a telling way: they are small and of the opposite sign (-0.097 for the employed share, +0.095 for the inactive share, both significant at the 1% level), so typical variation in migration is associated with slightly less employed inflows, while the drought-driven component is more employed. This is the direct confirmation of the pattern suggested by the scale margin: the marginal, drought-induced migrant differs from the average arrival, and differs in the direction of work. Together with the housing results, this points to a precise characterization of drought-induced migrants: they integrate the urban labor market more than the typical arrival, yet disproportionately settle in low-quality housing.

Table 17 in Appendix 9.4.3 re-estimates both margins with the cumulative drought instrument: the scale results are unchanged, and the shift of arrivals from inactivity toward employment persists.

4.3 Interpreting the Effect

4.3.1 Decomposing the Effect: Scale versus Propensity

Residential Integration - To assess the relative importance of the two channels, the additional low-quality housing inflow from a one-standard-deviation drought decomposes as

$$\Delta S \approx \underbrace{\Delta \text{MigRate} \times \bar{P}}_{\text{scale}} + \underbrace{\overline{\text{MigRate}} \times \Delta P}_{\text{propensity}},$$

where the scale term captures additional migrants arriving with the average propensity for low-quality housing, and the propensity term the baseline flow of migrants arriving with a higher propensity. The identity is an accounting split only if all four of its inputs come from one sample. The propensity is defined only on corridor-years with a positive urban migrant flow, so we evaluate the decomposition there and re-estimate the scale coefficient on that same sub-sample ($N = 72,723$); the headline scale estimate of Table 3 remains the full-panel one, since restricting to positive-flow corridors would condition on the endogenous outcome. On the common sample, the scale component is $0.00302 \times 0.406 = 0.00123$ additional migrants in low-quality housing per 100 urban residents, and the propensity component $0.08908 \times 0.00368 = 0.00033$: about 79% of the additional inflow operates through the number of migrants and 21% through their composition. The reason is transparent in the two sets of estimates - drought raises the number of migrants by 6.8% while raising their propensity for low-quality housing by about four tenths of a pp (from 40.6 to 41.0 percent of arrivals). Drought therefore fills low-quality housing primarily by sending more migrants,

and secondarily by sending more vulnerable ones. Appendix Table 18 details the calculation. One caveat remains: the identity evaluates the scale channel at the average migrant’s propensity rather than the marginal migrant’s, so it measures the relative weight of the two channels rather than the level of the effect.

This decomposition also explains why OLS and 2SLS coincide at the scale margin (Table 3) while differing sharply at the propensity margin. A gap between OLS and 2SLS on the scale outcome would require a large propensity channel - drought-induced migrants entering low-quality housing at a markedly different rate than ordinary migrants. But since this channel accounts for only about a fifth of the total effect, the selection it reflects is too small to move the scale estimate detectably, and OLS and 2SLS are nearly identical there. The selection is not absent - it is directly visible on the propensity outcome, where the 2SLS estimate is more than forty times the OLS one - but it is dominated by the scale channel. The two findings are consistent: drought-induced migrants are more housing-vulnerable than average, yet this composition shift is second-order relative to the increase in migrant numbers.

Labor Integration - The employment results decompose along the same lines as the housing results. Applying the identity of Section 4.3.1 to the employed inflow, and again evaluating both channels on their common sample - the corridor-years where labor-force status is observed ($N = 71,503$) - the scale component is $0.00302 \times 0.506 = 0.00153$ additional employed migrants per 100 urban residents and the propensity component $0.08923 \times 0.00422 = 0.00038$: about 80% of the additional employed inflow comes from the larger number of migrants and 20% from their higher propensity to be employed, with the same caveat as before - the identity is evaluated at the average arrival’s employment share rather than the marginal arrival’s. Drought therefore expands urban employment among migrants mainly by scale, exactly as it fills low-quality housing: the two outcomes share not only the same shock but the same anatomy.

4.3.2 The Aggregate Picture

Residential Integration - The corridor-level estimates can be aggregated to gauge the size of the phenomenon. Over 1981-2018, the censuses record about 11.6 million urban in-migrants in the corridors we observe across the eight countries. Drought accounts for roughly 390,000 of these moves - about 3.3% of observed urban in-migration.⁹ Applying the 2SLS estimate, under which 28 of every 100 marginal migrants settle in low-quality housing (Table 3), drought placed roughly 110,000 additional people in low-quality housing over the period. Drought degraded urban housing conditions mostly by adding people to the cities, and secondarily by re-sorting a given flow of migrants into worse housing - the scale and propensity weights of Section 4.3.1. We do not split the 110,000 between the two channels: the decomposition is identified on the corridor-years where the propensity is defined, whereas this

⁹The number of drought-induced migrants is obtained by multiplying each corridor-year’s realized drought intensity by the first-stage coefficient and the destination’s urban population, and summing across corridors and years, which yields $\approx 387,700$ migrants, or 3.3% of the 11.6 million observed. Applying the 2SLS estimate to this flow gives the additional residents of low-quality housing: $387,700 \times 0.285 \approx 110,300$. Both coefficients are taken from Table 3.

aggregate is computed on the full panel, and combining the two would mix samples.

These magnitudes are conservative. They cover only the corridors observed in the censuses, within the twelve-year windows preceding each census; they capture permanent, inter-regional moves and miss temporary or short-distance migration; and they hold destination conditions fixed, abstracting from any equilibrium response of housing supply.

Labor Integration - The same aggregation applies to the labor-market results: of the roughly 390,000 drought-induced migrants, about 166,000 arrived into employment, 29,000 into un-employment, and 138,000 into inactivity, the remainder having undetermined status (Table 5). Drought thus added workers to the cities on a similar scale as it added residents of low-quality housing.

Working Poor - Table 19 in Appendix 9.4.4 cross-tabulates the two dimensions on the same individuals, splitting the inflow into four employment $\times$ housing cells. All four inflows rise with drought-induced migration, but unevenly: the two employed cells and the not-employed cell outside low-quality housing increase significantly at the 1% level, while the most deprived cell - not employed and in low-quality housing - shows the weakest and least precise response, with an Anderson-Rubin interval that includes zero. The key ratio concerns the composition of the low-quality housing inflow: about 63% of the drought-induced migrants entering low-quality housing arrive employed (0.145/0.230), against 44% of the baseline inflow (0.004566/0.010421, from the mean outcomes in columns (1) and (3)). The marginal residents of low-quality housing are thus predominantly working poor rather than idle poor - drought-induced migrants integrate the labor market yet still settle in deprived housing, now verified on the same individuals rather than inferred from the separate housing and employment results.


5 Heterogeneity

5.1 Heterogeneity by Destination City Size

We examine whether the effect varies with the type of destination, classifying destinations by the population of the largest Africapoli agglomeration located in or within 10 km of the destination unit (OECD/SWAC, 2024): small towns ($< 100k$), secondary cities (100k–1M) and metropolises ($\geq 1M$). We use this size-based classification as our main cut for three reasons: it is external to the census data that produce our outcomes, it does not require choosing a threshold on a continuous and smoothly distributed urban share, and the agglomeration boundaries are those of the urban-settlement literature rather than of administrative units. Section S1.3 reports the same exercise on the census urban share as a complementary check. We estimate a single specification in which the migration rate is fully interacted with the destination-type indicators:

$$S_{o,d,t} = \sum_{k=1}^3 \beta_k (\text{MigRate}_{o,d,t} \times \mathbf{1}\{d \in k\}) + \sum_{k=1}^3 \delta_k (\text{dist}_{o,d,t} \times \mathbf{1}\{d \in k\}) + \alpha_{od} + \mu_{t \times k} + \varepsilon_{o,d,t}, \quad (4)$$

Figure 1: Heterogeneity by Destination City Size (Africapolis): Low-Quality Housing and Employment



Notes: Heterogeneity by destination city size: destinations are classified by the population of the largest Africapolis (OECD/SWAC, 2024) agglomeration located in or within 10 km of the destination unit (2020 population; zones: small towns < 100k, secondary cities 100k–1M, metropolises ≥ 1 M). The classification is independent of the census urban share and recovers Nairobi and 15 further Kenya units via the harmonized boundary file used for the climate extraction ($N = 240,340$ corridor-years). Panels (a)–(c): 3-bin saturated interacted-IV estimates (migration $\times$ bin instrumented by drought $\times$ bin, origin–destination and arrival-year $\times$ bin fixed effects, bin-specific distance-to-census control, SEs clustered at the origin district), plotted at zone midpoints; whiskers: 95% Wald intervals, shaded bars: Anderson–Rubin 95% intervals. Pairwise contrasts, housing: secondary cities – metropolises $p = 0.035$, small towns – metropolises $p = 0.084$, small towns – secondary cities $p = 0.809$; employment: secondary cities – metropolises $p = 0.138$, small towns – metropolises $p = 0.624$, small towns – secondary cities $p = 0.051$. Metropolises versus the other two zones pooled: $p = 0.032$ (housing), $p = 0.847$ (employment). Panels (d)–(f): rolling-window local IV along the logarithm of destination city size (main specification on corridors into the ~ 60 destinations nearest each grid point; pointwise 95% Wald bands; indicative only). The shaded zone marks where the window no longer slides; the binned estimates carry the inference there. Hollow points: weak first stage (Sanderson–Windmeijer conditional $F < 10$). Housing and employment panels share a common coefficient axis; Anderson–Rubin intervals and rolling bands extending beyond it are truncated. Bottom row: distribution of destination city sizes.

Source: Authors’ computations using IPUMS census, CHIRPS climate, and Africapolis 2024 data.

where each interacted migration term is instrumented by the corresponding interaction of drought intensity at origin, $Z_{o,t} \times \mathbf{1}\{d \in k\}$ (three endogenous variables, three instruments: just-identified), α_{od} are corridor fixed effects, and $\mu_{t \times k}$ arrival-year $\times$ bin fixed effects. No uninteracted migration term is included, so each β_k measures the absolute effect within destinations of type k rather than a difference from a reference category. The system is numerically equivalent to estimating the main specification separately by destination type, while the pooling delivers the covariance across the β_k used to test for differences.¹⁰ Figure 1 reports the binned estimates with the first stage by group (panels (a)–(c)), alongside a continuous rolling-window counterpart that re-estimates the main specification along the logarithm of destination city size (panels (d)–(f), indicative only) and the distribution of destinations.

The housing effect concentrates in secondary and small, rapidly urbanizing cities, and fades in the largest metropolises (panel (b)). Of every 100 marginal migrants, about 33 enter low-quality housing in a secondary city of 100k–1M (0.326, Anderson–Rubin interval [0.121, 0.523]) and about 30 in a small town below 100k (0.297, [0.045, 0.454]); the two are statistically indistinguishable from each other ($p = 0.809$), so the effect is of the same magnitude across the whole non-primate tier. In metropolises above one million the coefficient falls to 0.100, with an Anderson–Rubin interval that includes zero ([−0.147, 0.225]). The attenuation at the top is the one difference the data establish: the metropolis coefficient lies significantly below the other two zones pooled ($p = 0.032$) and below secondary cities pairwise ($p = 0.035$). The rolling estimates in panel (e) trace the same profile continuously along city size. Employment absorption does not follow this pattern (panels (c) and (f)): drought migrants are absorbed into work in destinations of every size - 0.393 in small towns, 0.533 in secondary cities, 0.425 in metropolises, each with an Anderson–Rubin interval well away from zero - and metropolises do not differ from the rest ($p = 0.847$). The only employment contrast approaching conventional significance runs between small towns and secondary cities ($p = 0.051$), with absorption lower in the smallest destinations. The two margins therefore separate: urban labor markets absorb drought migrants everywhere, while the housing burden falls on the secondary and small cities rather than on the primate ones.

Classifying destinations instead by their census urban share yields the same direction - the coefficient declines as destinations urbanize - but no contrast along that gradient reaches significance, and how close it comes depends on where the rural/urban cutoff is placed; Section S1.3 reports that exercise and a sweep over the cutoff. The geography of the inflows themselves is consistent with the size-based result: Appendix Figure 15 maps the share of each destination’s in-migration that the first stage attributes to drought, and that share declines with destination size, averaging 7.4% among destinations receiving fewer than 50,000 migrants over the panel against 3.3% among those receiving more than 200,000. It reaches 10 to 14% in secondary destinations such as Bushenyi, Kisii and Sikasso, against 1.2% in Mombasa. The map illustrates where the estimated inflows land; the identified result remains the bin-level analysis above.

A causal forest on the first stage confirms that this concentration is not selection. Drought’s

¹⁰Saturation by bin is required for identification: with pooled year fixed effects or a continuous interaction, the interacted instruments do not separately identify the interacted endogenous variables (Sanderson–Windmeijer conditional F near zero).

effect on migration does vary across corridors, but with corridor distance and origin rurality only: no destination characteristic – urban share, baseline low-quality housing prevalence, baseline employment, density or aridity – predicts where the migration response is stronger (Supplementary Appendix Table S7). Since drought does not tilt migration toward a particular type of destination in the first place, the geography of the housing effect is second-stage sorting – what happens to migrants once they arrive – rather than a first-stage choice of where drought sends them.

Supplementary Appendix Figure S5 shows a consistent (though noisier) pattern for the propensity margin: the sorting of climate migrants into low-quality housing is detectable only in urbanizing cities, the one bin where the weak-instrument-robust interval excludes zero; the propensity first stages are weak throughout, so this margin carries little heterogeneity information.

5.2 Individual Heterogeneity: Who Are the Migrants?

Appendix Table 20 profiles the drought-induced migrants themselves, comparing arrivals in high- versus no-drought years with the complier characterization of Abadie (2003). Both readings agree: drought-induced migrants are observationally similar to other migrants – the same gender balance and comparable or slightly higher education – but younger, with an overrepresentation of children under fifteen (42% versus 29%), consistent with whole-household displacement rather than the selective migration of young workers.

The effect itself is equally unselective on gender and age. Appendix Figures 13 and 14 re-estimate the main specification with subgroup-specific inflows and migration rates: drought displaces women and men, children and adults, proportionally alike (panel (a): relative first stages of +6 to +7% of each group’s mean migration rate), and the housing and employment responses are correspondingly uniform across both dimensions.

Education is the exception, and it cuts in opposite directions on the two stages. Drought does not displace the least-educated adults: their migration response is indistinguishable from zero and significantly below that of the secondary-educated (–8.6 percentage points of relative response, $p = 0.001$) – a trapped-population pattern. Among the adults it does displace, the migration response tilts if anything toward the more educated (+8.8% against +5.6%, $p = 0.064$), yet the housing effect runs the other way – twice as strong for the primary-educated ($p = 0.036$, both groups well identified with conditional first stages above 14) – while employment absorption is education-blind ($p = 0.787$). Education thus eases displacement but shields from low-quality housing: the least educated stay behind, the secondary-educated escape bad housing, and the lower-educated among the displaced bear the residential burden. We verify that this selection operates along a single, additive dimension: a fully saturated model crossing gender, age at arrival, and education into twelve profile cells rejects equality of the first-stage responses (joint $p = 0.024$) but finds no interaction beyond the additive margins (joint $p = 0.59$) – drought selects migrants along a monotone educational gradient, from near-zero mobility among the least educated to the strongest response among the secondary-educated, without any distinct combined profile. The separate margins of Figure 13 are therefore a sufficient statistic for who drought moves (Supplementary Appendix Table S6). Attainment is measured at census and may post-date arrival, so the two displacement-side gradients remain suggestive. Full estimates are in Supplementary

Appendix Tables S2–S5. Together with the profile evidence, the pattern points to barriers rather than composition: similar people, absorbed uniformly into work, sorted selectively into housing.

6 Robustness

This section presents a series of robustness checks and alternative specifications designed to assess the stability of the main results and the validity of the identification strategy. Several of these checks are reported in the appendix - the fixed-effects saturation in Table 2, the definition sweep, and the wet-side placebo of Figure 9 - and the remaining estimates are in progress.

Alternative specifications. - We consider four alternative specifications. First, we estimate a manually specified two-instrument design combining $\text{Drought}^g * p10$ and $N^g * \text{dry}$ to examine whether the results are robust to multi-instrument identification despite a weaker joint first stage. Second, we replace drought intensity with extreme precipitation measures based on high-percentile rainfall shocks to test whether migration responses are symmetric across the rainfall distribution or primarily driven by adverse drought events. Third, following Hultgren et al. (2025), we perform an out-of-sample validation to evaluate the stability of the post-LASSO instrument selection across held-out countries. Fourth, we replace the census-based destination population denominator with the annually reconstructed population series, $\text{Pop}_{d,t}$, to assess the sensitivity of the estimated migration rates to the choice of normalization.

Alternative samples. - To assess the external validity of the main findings, we extend the analysis to all fifteen countries in the dataset by exploiting five-year migration intervals rather than restricting attention to the eight countries with annual migration information. This exercise evaluates whether the estimated scale effect generalizes beyond the baseline sample.

Sensitivity analyses. - We conduct two additional sensitivity analyses. First, we replace the binary low-quality housing indicator with a continuous measure of housing deprivation defined as the number of deficient UN-Habitat dimensions (ranging from 0 to 5), allowing us to test whether drought-induced migration affects the entire distribution of housing quality rather than only the extensive margin of low-quality housing residence. Second, we perform a leave-one-country-out exercise to verify that the estimated effects are not driven by any single country in the sample.

Leave-one-country-out. - Because the eight countries differ in size, census vintage and housing levels, we re-estimate the main specification eight times, each time dropping one country (Appendix Figure 16). The estimate stays positive in every subsample and ranges from 0.197 to 0.340 against 0.285 on the full sample, with each subsample estimate inside the full-sample confidence interval. Precision falls where the first stage weakens: dropping

Benin or Kenya lowers the first-stage F to about 13 or 15, and in those two subsamples the Anderson–Rubin interval includes zero, a loss of power rather than a change in the estimate.

Exclusion restriction. - The identifying assumption may be violated if drought at origin affects destination outcomes through channels other than migration, such as food price transmission, remittance flows, or government infrastructure responses. Origin–destination pair fixed effects absorb time-invariant correlations between origin characteristics and destination outcomes, while arrival-year fixed effects capture aggregate shocks common to all migration corridors. To further assess the plausibility of the exclusion restriction, we are investigating whether origin-specific drought shocks predict destination food prices after conditioning on the baseline specification. Evidence from this exercise will be reported in future versions of the paper.

Crowding-out. - The scale outcome counts every migrant of the arrival cohort living in low-quality housing, not only those whom drought displaced. A positive coefficient is therefore consistent with two mechanisms: drought-induced migrants settling in low-quality housing themselves, or drought-induced migrants crowding other arrivals out of decent housing into it. The aggregate margins cannot separate them, because on the sample where housing is observed the decent-housing response is the exact arithmetic complement of the low-quality one.

Two checks bound the concern. First, drought-induced migration does not degrade the housing of migrants who arrived on the same corridor in *earlier* years, who cannot be compliers of the current shock: the weak-instrument-robust interval on that outcome excludes any effect above 0.138 per marginal migrant, and a placebo on their total number is flat (Appendix Table 22). Crowding-out, if it operates, does not spill onto previously settled migrants. Second, the estimated effect (0.285) lies *below* the benchmark that no crowding-out and average sorting would imply — the baseline low-quality-housing share among migrants, 0.359 — whereas crowding-out would push it above.

Neither check identifies crowding-out *within* the arrival cohort, which is the channel that would bias the estimate. A direct test on the one migrant group whose own inflow does not respond to drought — adults with less than primary education, and therefore almost entirely always-takers — is inconclusive: their low-quality-housing propensity rises with drought-induced migration by 1.25, of the same order as the aggregate propensity effect, but the Anderson–Rubin interval, $[-0.21, +3.53]$, contains both zero and substantial crowding-out (Appendix Table 23). We therefore read the coefficient as a total effect on the corridor — each hundred additional drought-induced migrants is associated with twenty-eight additional migrants living in low-quality housing — rather than as a claim about which individuals occupy which dwellings.

Comparison with other migration drivers. - Two additional exercises will compare drought-induced migration with migration generated by alternative shocks. The first instruments migration using conflict events at origin (ACLED data); the second uses exogenous variation in mine openings at the admin-unit level. These comparisons sharpen the interpretation of

the LATE estimated in the main analysis by contrasting push driven (drought, conflict) and pull-driven (mining) migration mechanisms and their respective urban absorption patterns.

7 A Model of Post-Migration Housing Integration

7.1 Motivation

The empirical analysis establishes three main results. First, drought at origin causally increases rural-to-urban migration. Second, this drought-induced migration raises urban low-quality housing inflows: of every hundred additional climate migrants, about twenty-eight settle in low-quality housing, mainly through the scale of arrivals but also through selection - the marginal drought-induced migrant is about 11 percentage points more likely to end up in low-quality housing than the average migrant arriving on the same corridor. Third, these same migrants are absorbed into urban employment: the housing and employment margins decouple, and the integration gap of climate migrants is residential rather than occupational.

While these reduced-form results document a systematic residential integration gap, they do not quantify the extent to which this gap reflects binding integration barriers - deposits and guarantor requirements in the formal rental market, documentation, lack of local networks - as opposed to selection on income or preferences, nor do they allow for counterfactual policy analysis. The framework below addresses the propensity margin only. The scale margin - how many migrants arrive - is taken as given, since it is governed by the migration decision the empirical strategy instruments rather than by post-arrival sorting within the city. To make progress on these questions, we sketch a simple conceptual framework of post-migration housing integration within cities.

The framework is static and partial equilibrium. Migrants have already arrived in the city, and it abstracts from migration decisions across locations, focusing on the housing integration choice after arrival. The decoupling documented above disciplines its key ingredient: because climate migrants find work at least as often as other arrivals, the barrier they face is specific to the housing margin.

7.2 A Two-State Example

Consider migrants who have already arrived in a destination city, and two housing integration states: a low-integration state L (low-quality housing) and a high-integration state H (formal, adequate housing). Each migrant draws a pair of state-specific capacities (s_{iL}, s_{iH}) , indexing how well the migrant fares in each state. Draws are margin-specific: (s_{iL}^E, s_{iH}^E) for employment and (s_{iL}^R, s_{iH}^R) for housing. Formal housing is valued for the bundle of amenities and services it delivers rather than for a differential return to labor, which keeps the two margins separable and consistent with the decoupling documented above. Let w_H and w_L denote the value per unit of capacity in each state - on the housing margin, inclusive of the amenity bundle. Absent barriers, migrants choose the state that maximizes value: those with $w_H s_{iH} > w_L s_{iL}$ sort into formal housing. Graphically (Appendix Figure 17), these migrants lie above the ray OC with slope w_L/w_H .

Residential integration barriers - deposits for formal rentals, guarantor and documentation requirements, lack of local networks - reduce the effective value of the high-integration state. If climate migrants face a proportional integration cost τ , their effective return in formal housing becomes $(1 - \tau)w_H$, and only those above the steeper ray OE , with slope $w_L/[(1 - \tau)w_H]$, choose it. The rays partition migrants into four groups. In Region I, above OE , migrants integrate into formal housing even in the presence of barriers. Regions II and III, between OE and OC , contain the migrants the barrier misallocates: they would choose formal housing absent barriers - whether their comparative advantage lies in the high-integration state (Region II, above the 45° line) or not (Region III) - yet remain in low-quality housing. Region IV, below OC , comprises migrants who optimally remain in low-quality housing regardless of barriers.

Two qualifications on how the rays are read. First, baseline migrants face their own barriers $\tau^0 \geq 0$, so OC is their cutoff rather than a frictionless benchmark: the propensity gap identifies the differential barrier $\tau - \tau^0$, and recovering its level - and hence the welfare gain from a given reduction in τ - requires an external restriction on τ^0 . Second, the trapped fan is defined in the space of housing draws, so labelling it the working poor presumes that migrants in that fan also clear the employment margin. This holds if capacities are positively correlated across margins, which the cross-tabulation supports but does not impose.

This diagram organizes the interpretation of our propensity result: a higher τ for drought-induced migrants shifts part of their distribution from Region I into the trapped Regions II–III, raising their share in low-quality housing relative to baseline migrants on the same corridor - precisely the three-percentage-point gap we estimate - without affecting their employment. The same logic applies margin by margin, and this is where the framework connects to the paper’s distinctive finding. Consider two integration margins - employment and housing - each with its own two-state choice and its own proportional barrier, τ_E and τ_R (Appendix Figure 17). Our employment results discipline the first wedge: drought-induced migrants are absorbed into work at least as often as baseline arrivals, so $\tau_E \approx 0$ and the employment margin has no trapped region (panel (a)). The housing results discipline the second: a residential gap alongside absorption into work requires $\tau_R > 0$, opening the trapped fan of panel (b). Jointly, a residential-only barrier predicts that climate migrants concentrate in the employed-but-badly-housed cell - the working poor - which is precisely what the cross-tabulation of housing and employment outcomes shows (Appendix Table 19).

The diagram is an interpretive device rather than an identification result. Two alternatives generate the same propensity gap without any barrier. Migrants may sort on preferences or prices, choosing low-quality housing to economise on rent. More importantly, the gap may reflect the composition of draws rather than a rotation of the cutoff: drought-induced migrants originate from agricultural areas and may draw a lower s_H/s_L ratio, shifting the distribution across a fixed OC . The diagram does not separate the two. A proportional barrier rotates the cutoff while a fixed cost - a rental deposit - shifts it, so the two are distinguishable along the income distribution with price and income data. Doing so, estimating the size of the trapped region for climate migrants relative to other migrants, and quantifying the welfare gains from policies that reduce τ (housing vouchers, formal rental access programs), is the object of the structural extension of this paper.

8 Conclusion

This paper provides causal evidence on the contribution of drought-induced migration to the growth of low-quality housing in Sub-Saharan Africa. Using harmonized census microdata for eight countries, censuses 1990–2019 with migration windows covering 1981–2018, it combines a cohort-bilateral panel of migration flows with high-resolution rainfall data to construct exogenous variation in drought intensity at origin. A post-IV LASSO procedure selects drought intensity as the instrument for bilateral migration, and the 2SLS estimates identify a Local Average Treatment Effect on drought-induced migrants, precisely the population climate policy concerns.

Drought-induced migration raises urban low-quality housing inflows at destination: of every hundred migrants that drought pushes into a city, about twenty-eight settle in low-quality housing, and over 1981–2018 drought placed roughly 110,000 additional people in low-quality housing across our corridors. This effect operates mainly through the scale of arrivals - about four fifths of it - but also through selection: the marginal drought-induced migrant is about 11 percentage points more likely to settle in low-quality housing than the migrants who would have moved anyway - 51% against 41%. Yet these migrants are absorbed into urban employment across destinations; the margin of exclusion is residential, not occupational. The housing effect is concentrated along rural-origin corridors, consistent with an agricultural-push mechanism, and it concentrates in secondary and small, rapidly urbanizing cities: classifying destinations by agglomeration size using Africapolis data, the effect is of the same magnitude in towns below 100k and in cities of 100k–1M, and falls to a third of that level in metropolises above one million, a gap significant at the 5% level.

These findings have direct implications for urban and climate policy in Sub-Saharan Africa. Because the effect runs mainly through the scale of arrivals, the immediate priority is expanding urban absorption capacity; the selection channel calls, in addition, for targeting the most vulnerable climate migrants. That the effect concentrates in secondary and small urbanizing cities, and fades in the largest metropolises, suggests that housing policies focused on primate metropolitan areas may overlook precisely the destinations where climate-induced migration weighs most heavily on housing. If low-quality housing itself exposes residents to climate hazards (Marx et al., 2019), climate-induced growth in low-quality housing may amplify rather than mitigate the costs of climate variability.

Several avenues remain open for future research. A structural model of migrant integration, jointly modeling housing and labor-market choices, would quantify the welfare cost borne by climate migrants and allow the evaluation of targeted policies such as housing support in small and intermediate cities. Subnational crop data, such as HarvestStat Africa, would allow a direct test of the agricultural mechanism - drought reducing harvests, and harvests driving migration. As climate change intensifies and urbanization accelerates, understanding the interaction between climate shocks, migration, and urban integration will be increasingly important for policies promoting both adaptation and inclusive urban development.

References

- Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics*, 113(2):231–263.
- Alvarez, J. A. (2020). The agricultural wage gap: Evidence from brazilian micro-data. *American Economic Journal: Macroeconomics*, 12(1):153–173.
- Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives*, 32(4):33–52.
- Beine, M. and Parsons, C. R. (2016). Climatic factors as determinants of international migration: Redux. CREA Discussion Paper Series 16-11, Center for Research in Economic Analysis, University of Luxembourg.
- Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650.
- Bryan, G. and Morten, M. (2019). The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy*, 127(5):2229–2268.
- Cattaneo, C. and Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 122(Supplement C):127 – 146.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.
- Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., and Michaelsen, J. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Scientific Data*, 2(1):1–21.
- Gittard, M. (2024). Droughts, migration and population in kenya.
- Gollin, D., Jedwab, R., and Vollrath, D. (2016). Urbanization with and without industrialization. *Journal of Economic Growth*, 21(1):35–70.
- Gollin, D., Lagakos, D., and Waugh, M. E. (2014). The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993.
- Gray, C. L. and Mueller, V. (2012). Natural disasters and population mobility in bangladesh. *TProceedings of the National Academy of Sciences*, 109(16):6000–6005.
- Hamory, J., Kleemans, M., Li, N. Y., and Miguel, E. (2021). Reevaluating agricultural productivity gaps with longitudinal microdata. *Journal of the European Economic Association*, 19(3):1522–1555.

- Henderson, J. V. and Kriticos, S. (2018). The development of the african system of cities. *Annual Review of Economics*, 10(1):287–314.
- Henderson, J. V., Storeygard, A., and Deichmann, U. (2017). Has climate change driven urbanization in africa? *Journal of development economics*, 124:60–82.
- Hoffmann, R., Dimitrova, A., Muttarak, R., Crespo Cuaresma, J., and Peisker, J. (2020). A meta-analysis of country-level studies on environmental change and migration. *Nature Climate Change*, 10(10):904–912.
- Hsiang, S. (2016). Climate econometrics. *Annual Review of Resource Economics*, 8:43–75.
- Hultgren, A., Carleton, T., Delgado, M., Gergel, D. R., Greenstone, M., Houser, T., Hsiang, S., Jina, A., Kopp, R. E., Malevich, S. B., McCusker, K. E., Mayer, T., Nath, I., Rising, J., Rode, A., and Yuan, J. (2025). Estimating global impacts to agriculture from climate change accounting for adaptation. *Nature*, 641:644–652.
- Jedwab, R., Christiaensen, L., and Gindelsky, M. (2017). Demography, urbanization and development: Rural push, urban pull and... urban push? *Journal of Urban Economics*, 98:6–16.
- Kleemans, M. and Magruder, J. (2018). Labour market responses to immigration: Evidence from internal migration driven by weather shocks. *The Economic Journal*, 128(613):2032–2065.
- Marx, B., Stoker, T., and Suri, T. (2013). The economics of slums in the developing world. *Journal of Economic perspectives*, 27(4):187–210.
- Marx, B., Stoker, T. M., and Suri, T. (2019). There is no free house: Ethnic patronage in a kenyan slum. *American Economic Journal: Applied Economics*, 11(4):36–70.
- McGuirk, E. and Nunn, N. (2021). Transhumant pastoralism, climate change, and conflict in africa. *Working Paper*.
- McMillan, M., Rodrik, D., and Verduzco-Gallo, Í. (2014). Globalization, structural change, and productivity growth, with an update on africa. *World development*, 63:11–32.
- Minnesota Population Center (2019). Integrated public use microdata series, international: Version 7.2 [dataset]. Minneapolis, MN: IPUMS.
- Missirian, A. and Schlenker, W. (2017). Asylum applications respond to temperature fluctuations. *Science*, 358(6370).
- Mueller, V., Gray, C., and Hopping, D. (2020). Climate-induced migration and unemployment in middle-income africa. *Global Environmental Change*, 65:102183.
- Muñoz-Sabater, J., Dutra, E., Agustí-Panareda, A., Albergel, C., Arduini, G., Balsamo, G., Boussetta, S., Choulga, M., Harrigan, S., Hersbach, H., Martens, B., Miralles, D. G., Piles, M., Rodríguez-Fernández, N. J., Zsoter, E., Buontempo, C., and Thépaut, J.-N. (2021).

- Era5-land: A state-of-the-art global reanalysis dataset for land applications. *Earth System Science Data*, 13(9):4349–4383.
- OECD/SWAC (2024). Africapolis (database). <https://africapolis.org>.
- Trabucco, A. and Zomer, R. (2019). Global aridity index and potential evapotranspiration (ET0) climate database v2. CGIAR Consortium for Spatial Information (CGIAR-CSI).
- UN-Habitat (2003). *The Challenge of Slums: Global Report on Human Settlements 2003*. Earthscan, London.
- Un-Habitat (2020). Strengthen un-habitat regional representation in sub-saharan africa. *Un-Habitat report*.
- Un-Habitat (2022). Envisaging the future of cities. *World Cities Report*.

9 Appendix

9.1 Descriptive Statistics


Table 7 reports descriptive statistics for all variables used in the analysis: the corridor-year estimation panel (Panel A), the migrant propensity outcomes (Panel B), and destination characteristics (Panel C).

Table 7: Descriptive Statistics

Variable	Mean	SD	P25	Median	P75	<i>N</i>
<i>Panel A: Corridor-year panel (main estimation sample)</i>						
Urban migration rate (per 100 urban residents)	0.0268	0.1003	0.0000	0.0000	0.0099	241,540
Low-quality housing inflow rate (per 100)	0.0120	0.0589	0.0000	0.0000	0.0000	241,540
Employed inflow rate (per 100)	0.0106	0.0405	0.0000	0.0000	0.0000	241,540
Unemployed inflow rate (per 100)	0.0013	0.0095	0.0000	0.0000	0.0000	241,540
Inactive inflow rate (per 100)	0.0118	0.0522	0.0000	0.0000	0.0000	241,540
Drought intensity at origin	0.350	0.502	0.000	0.000	0.625	241,540
Cumulative 3-season drought at origin	0.366	0.297	0.106	0.347	0.551	235,920
Distance to census $ t - t_c $ (years)	6.08	3.29	3.00	6.00	9.00	241,540
Urban population at destination ('000)	123.8	315.6	16.0	41.4	95.9	241,540
<i>Panel B: Migrant propensity outcomes (corridors with positive denominators)</i>						
Share of migrants in low-quality housing	0.406	0.388	0.000	0.333	0.750	72,723
Mean number of deprived dimensions (0–5)	0.662	0.818	0.000	0.429	1.000	72,723
Share of migrants employed	0.505	0.362	0.200	0.500	0.842	71,503
Share of migrants unemployed	0.058	0.155	0.000	0.000	0.036	71,503
Share of migrants inactive	0.437	0.358	0.000	0.429	0.667	71,503
<i>Panel C: Destination characteristics (one row per destination, most recent census)</i>						
Urban population share of the destination	0.256	0.241	0.085	0.188	0.351	351

Notes: Descriptive statistics for the variables used in the analysis. Panel A: the main corridor-year estimation panel (8 countries, Egypt excluded; Mali 1998 included; 12-year capped windows with the census year excluded). The panel is the full origin–destination rectangle and retains pairs that never record an urban migrant: 37% of the 16,846 origin–destination pairs register no urban migrant at all over the window, and 70% of corridor-years record zero urban migrant arrivals, which accounts for the zero lower quartiles of the rate variables. Zero-flow corridor-years are kept because dropping them would condition the sample on the endogenous regressor. Rates are counts per 100 urban residents at destination; no outcome is winsorized, and the sample is restricted to destinations with at least 2,000 urban residents (the conventional minimum size of an urban locality). Drought intensity is the standardized negative seasonal rainfall anomaly at origin, $\max(0, -\tilde{z})$, equal to zero by construction in wetter-than-average seasons; the cumulative version averages the current and two preceding seasons (defined where lags exist). The inactive inflow counts migrants out of the labor force with determined status. Panel B: shares among migrants (housing: denominator = migrants; employment: denominator = migrants with known labor-force status); defined only for corridor-years with a positive denominator. Panel C: census urban population share, one observation per destination at its most recent census.

Figure 2: Average Urban In-Migration Rates Over Time, by Country




Notes: Mean urban in-migration rate (in-migrant arrivals per 100 urban residents at destination) across corridor-year observations of the yearly estimation panel, by calendar year and country. Dashed vertical lines mark census years; each census contributes a capped retrospective window of yearly migration cohorts. Within each window, rates rise mechanically toward the census year, reflecting recall and onward-migration attrition of earlier cohorts; the distance-to-census control absorbs this gradient in estimation.

Source: Authors' computations from IPUMS census microdata.

9.1.1 Descriptive Maps


Figure 3: Descriptive Maps: Climate at Origin



Notes: The eight countries of the estimation sample are outlined in black, with their administrative units (second level; first level for Malawi and Uganda); other African countries covered by IPUMS are outlined in grey. Panel (a): mean cumulative precipitation over the growing season (mm), by origin unit – the $Precip_{cum}^g$ measure underlying the instrument. Panel (b): mean of the standardized drought intensity $\max(0, -\tilde{z})$ over the panel years, by origin unit.

Source: Authors' computations from CHIRPS data and IPUMS harmonized shapefiles.


Figure 4: Descriptive Maps: Migration and Urbanization at Destination



Notes: Panel (a): mean yearly urban in-migrant arrivals per 100 urban residents, by destination unit, aggregated over all origins. Panel (b): urban population share of the destination (urban over total census population, most recent census). Grey units: not observed in the estimation panel.

Source: Authors' computations from IPUMS census microdata.

Figure 5: Descriptive Maps: Migrant Outcomes at Destination




Notes: The propensity outcomes of the paper, aggregated by destination over the panel: among urban in-migrants, the share residing in low-quality housing (panel (a), main UN-Habitat definition; denominator: migrants) and the share employed (panel (b); denominator: migrants with known labor-force status). Grey units: no in-migrants observed.

Source: Authors' computations from IPUMS census microdata.

9.1.2 Housing-Quality Definition

Figure 7 shows how the index behaves in the data. For each country it reports the share of households counted deprived on each dimension, separately for all urban residents and for the urban in-migrants on whom the paper's outcome is built. Three features are worth noting. First, in-migrants are less deprived than the resident urban population in every country and on almost every dimension, by 6 to 19 percentage points on the summary measure: recent arrivals are not, in the raw data, the worst-housed group in the city. This is not in tension with the paper's estimates, which compare the marginal drought-induced migrant with the average migrant arriving on the same corridor, holding the corridor fixed, and never compare migrants with residents. If anything it works against a mechanical reading of the results, since the unconditional correlation between migration and housing deprivation runs the other way. Second, crowding is the one dimension on which the two populations coincide: migrants occupy dwellings of comparable or better materials but at similar densities, which locates the pressure of in-migration on living space rather than on the built structure. Third, the dimensions contribute very unequally. Floors, walls and crowding do most of the work; roofs fire rarely once sheet metal is treated as adequate; and sanitation fires only where latrines are absent rather than unimproved. Water is absent from the index altogether, for the reason given in Section 2: these censuses distinguish only types of piped supply, while UN-Habitat counts boreholes and protected dug wells as improved sources, so the variable cannot separate deficient from adequate access. The index should therefore be read as a measure of housing durability and living space, not as the full four-pillar UN-Habitat definition, and its levels are not comparable to published slum headcounts. The category composition behind each dimension is reported in Supplementary Appendix Figures S6–S9.

Figure 6: Observed Urban In-Migration Corridors, by Country



Notes: Observed bilateral migration corridors, one panel per country (241,540 corridor-years). Units are shaded by mean drought intensity at origin, on the scale of Figure 3(b); grey units are never observed as an origin. The faint web shows all 10,399 corridors with a positive urban flow; highlighted arcs are the 189 largest (top 25 per country), 44% of urban in-migration. Arcs run from origin (open dot) to destination (arrowhead), with width proportional to the square root of the mean ANNUAL flow – totals are not comparable across countries, since Kenya, Mali and Zambia contribute two censuses. White circles mark main-city destinations (census urban share ≥ 0.55). Flows count urban in-migrants, the numerator of the endogenous variable. Centroids are interior points of administrative polygons, not city locations.


Source: Authors' computations from IPUMS census microdata and IPUMS harmonized shapefiles.

Table 8: Coding of the Housing-Quality Dimensions

Dimension	Classified as deficient
<i>Panel A: Structural quality of the dwelling</i>	
Floor	Unfinished natural floor – earth, sand, or dung – or salvaged wood (codes 100–120, 235)
Walls	No walls, or rudimentary, organic, or salvaged materials – cane, palm, logs, grass, scrap (codes 100–299, 405, 407) – and non-permanent earth materials: mud, adobe, clay, unburnt brick, pole and mud (codes 513, 515, 527, 533, 535, 539, 540, 541). Cement-stabilised earth, burnt or stabilised brick, and wood are treated as permanent
Roof	Wood, thatch, or palm; mud or earth; cardboard, plastic, or salvaged materials; other (codes 40–53, 60–61, 70–73, 90)
<i>Panel B: Access to basic services and living conditions</i>	
Sanitation	No toilet facility (code 10)
Crowding	More than three persons per habitable room

Notes: Coding of the five housing-quality dimensions, in harmonized IPUMS categories (codes in parentheses). The coding is identical across countries; the availability of each dimension varies by census (Egypt records no floor, wall, or roof variable and is excluded from the housing analysis; sanitation is unavailable in Kenya 2009, and crowding in Cameroon 2005 and Uganda 2014, whose censuses record only rooms used for sleeping). Water access is not part of the index: the UN-Habitat criterion counts boreholes and protected wells as improved sources, but the harmonized IPUMS water-supply variable in these censuses distinguishes types of piped supply rather than improved from unimproved sources, so it cannot reliably separate deficient from adequate access in this setting. Sanitation follows the UN-Habitat standard, under which a pit latrine with a slab counts as improved, so only the complete absence of a facility is coded deficient. The crowding denominator is the habitable room; for Zambia, where the census records living rooms and bedrooms separately, it is their sum. Sheet-metal roofing (codes 30–39) is classified as durable in the main definition and as deficient in the robustness definition. A household resides in low-quality housing if at least one available dimension is deficient (main definition); the robustness definition requires at least two. The continuous deprivation count sums the deficient dimensions (0–5).

Figure 7: Housing Deprivation among Urban Residents and Urban In-Migrants



Notes: Share of households counted deprived on each dimension of the housing index, by country. Grey points are all urban residents, blue points urban in-migrants, and the segment joins the two for the same country and dimension. An in-migrant is a person whose unit of origin differs from the unit of enumeration and whose arrival year falls in that census's retrospective window; both populations are restricted to urban areas and person-weighted. The first panel applies the paper's threshold, at least one deprived dimension. Water deprivation is virtually absent in all samples and is therefore omitted from the figure. Crowding cannot be constructed for Cameroon 2005 and Uganda 2014, whose censuses record only rooms used for sleeping, so those two countries are absent from the crowding panel and mechanically lower in the summary panel. Sheet-metal roofs count as adequate.

Source: Authors' computations from IPUMS census microdata.

9.1.3 Urbanization Data

Table 9 and Figure 8 validate the census-based urbanization classification against an external measure of city size, the population of the largest Africapolis agglomeration at destination. The two measures line up where it matters: across the three bins, the median agglomeration size rises from 47,000 to 86,000 to 880,000, the share of destinations hosting a city of at least 100,000 inhabitants rises from 28 to 43 to 70 percent, and the main-cities bin contains the region’s actual metropolises (Nairobi, Kampala, Yaoundé, Bamako). The destination-level rank correlation is positive but moderate (Spearman 0.33), as expected: the census share captures the urban composition of the destination’s population, while Africapolis measures the physical size of the nearest agglomeration, and coarser administrative units in Malawi and Uganda mechanically dilute urban shares (panel (b)). Appendix Figure S2 reports the country-specific cumulative distributions of both measures.

Table 9: Validation of the Urbanization Classification: Census Urban Share versus Africapolis City Size

Bin (urban share)	Destinations	Corridor-years	Largest agglomeration		Example cities
			Median ('000)	$\geq 100k$ (%)	
Rural-like (< 0.15)	142	91,019	36	23%	Naivasha, Arua, Molo, Kitui
Urbanizing cities ($[0.15, 0.55]$)	169	122,821	90	43%	Nakuru, Lilongwe, Kisumu, Fort Portal
Main cities (≥ 0.55)	36	27,700	652	78%	Nairobi, Kampala, Yaounde, Bamako

Notes: Correspondence between the census urbanization bins and Africapolis city sizes. Bins: census urban population share of the destination. Agglomeration populations and names: largest Africapolis (2024) agglomeration in or within 10 km of the destination unit (2020 population; recovered boundary file, Nairobi and 15 further Kenya units included). Example cities are assigned to the bin of their own unit (highest urban share among the units they map to). Four rows carrying conurbation-cluster totals (Kisumu, Embu, Mbale, Kikima) are corrected using each agglomeration’s pre-merge decadal growth (e.g., Kisumu 15.5M $\rightarrow$ 901k). All eight countries enter; corridor-years from the estimation sample ($N = 241,540$).


9.2 Identification Strategy

9.2.1 Post-IV Lasso Selection

Figure 9 reports a data-driven validation of the instrument choice. We build a pool of candidate instruments - drought intensity, drought dummies at alternative thresholds, precipitation and temperature polynomials, degree-day and evapotranspiration measures, and their lags and 3-season means - and apply the post-IV-LASSO selection of Belloni et al. (2014), alongside the single-candidate first-stage F . We split the exercise between drought shocks (panels (a)-(b)) and their extreme-wet counterparts (panels (c)-(d)).

The drought-intensity instrument is the strongest single candidate on both samples ($F = 23.9$ on the scale sample, $F = 13.0$ on the propensity sample) and is selected by the LASSO in both; adding lagged candidates, its 3-season mean dominates ($F = 27.4$), consistent with the cumulative instrument used throughout. At the default penalty the LASSO retains a small set of dry-side candidates rather than a single winner (drought intensity, dry days, and extreme degree days on the scale sample); we nevertheless use the single strongest

Figure 8: Census Urban Share versus Africapolis City Size at Destination



Notes: Panel (a) plots one observation per destination, colored by country. The horizontal axis reports the destination's urban population share in the most recent census. The vertical axis reports the 2020 population of the largest agglomeration in the Africapolis 2024 database located within, or within 10 km of, the destination unit, on a logarithmic scale. Panel (b) shows the distribution of census urban shares across destinations, separately by country. Dashed vertical lines indicate the classification thresholds of 0.15 and 0.55. Destinations without a matched Africapolis agglomeration of at least 10,000 inhabitants are excluded from panel (a). Four Africapolis agglomeration records reporting conurbation-cluster totals are corrected as described in the notes to Table 9. The Spearman rank correlation between the two measures is 0.33. Appendix Figure S2 reports their country-specific cumulative distributions.

Source: Authors' computations using IPUMS census data and the Africapolis 2024 database.

candidate, both for transparency and to keep the model just-identified, so that Anderson-Rubin inference remains exact and the complier population is not mixed across instruments. The robustness section reports the main results under the LASSO-selected instrument sets, showing that our findings are not tied to this choice.

By contrast, Figure 10 shows the results for the excess-rainfall intensity - the exact positive mirror of the drought instrument - has essentially no predictive power ($F = 4.0$ and $F = 0.03$), and the flood dummies are similarly weak. The instrument is thus data-selected, and its power comes specifically from the dry tail of the rainfall distribution.


9.3 First Stage and Reduced Form

9.3.1 First Stage

Table 10 assesses the robustness of the first stage to sample restrictions and to the drought exposure window, reporting the paper's main specification for each sample and instrument; specification robustness itself is established on the full sample in Table 1. Columns (1)-(2) cover the main sample of urban migrations to all destinations, columns (3)-(4) restrict destinations to main cities,

footnoteMain cities are defined as destinations with an urban population share weakly above the threshold, $\text{share}_d^{\text{urb}} \geq 0.55$, measured at the country's most recent census wave and held fixed across waves. Rural origins are defined as the complement: locations with $\text{share}_o^{\text{urb}} < 0.55$ or with no measured urban share. Robustness checks vary this threshold and use


Figure 9: Post-IV-LASSO Instrument Selection: Drought Candidates



Notes: Each point reports the first-stage F -statistic of a single candidate instrument for the urban migration rate under the main specification, including origin–destination and year fixed effects and the distance-to-census control. Standard errors are clustered at the origin-district level. The panels compare candidate sets without and with lagged instruments for the scale and propensity samples, respectively. Filled markers denote candidates selected by post-IV-LASSO.

Source: Authors' computations using IPUMS census and CHIRPS climate data.

Figure 10: Placebo Instrument Selection: Wet-Side Mirror Candidates



Notes: This figure replaces each drought candidate with its wet-side mirror. Excess-rainfall intensity is defined as the positive mirror of drought intensity. Each point reports the single-instrument first-stage F -statistic under the main specification. The weaker first stages indicate that instrument strength is specific to the dry tail of the rainfall distribution.

Source: Authors' computations using IPUMS census and CHIRPS climate data.

alternative definitions and datasets. and columns (5)-(6) further restrict to moves from rural origins to main cities. Within each pair, the first column uses the contemporaneous drought instrument and the second its cumulative (three-season mean) counterpart on the lag-restricted sample.

The contemporaneous instrument is a strong and significant predictor of migration in every sample, including moves from rural origins to major cities. On the full sample, the cumulative instrument delivers an even stronger first stage ($F = 27.4$ against 23.9), but its strength collapses once the sample is restricted to city destinations (F below conventional thresholds in columns (4) and (6)). This asymmetry, together with exclusion-restriction considerations—a multi-year accumulation of drought is more likely to affect destination outcomes through channels other than the timing of migration, such as persistent income losses at origin—leads us to retain the contemporaneous instrument as the paper’s main one. The cumulative instrument is nonetheless reported alongside the main results throughout, and delivers consistent conclusions wherever its first stage is strong.

Table 10: First Stage: Restricted Samples and Lagged Instrument (Main Specification)

	<i>All destinations</i>		<i>Main cities</i>		<i>Rural → main cities</i>	
	Contemp. (1)	Lagged (2)	Contemp. (3)	Lagged (4)	Contemp. (5)	Lagged (6)
Drought intensity	0.00362*** (0.00066)	0.00947*** (0.00164)	0.00312*** (0.00092)	0.00386 (0.00317)	0.00254*** (0.00065)	0.00375* (0.00195)
First-stage F	29.65	33.42	11.61	1.48	15.16	3.70
Effect of 1-SD drought (% mean)	+6.8%	+10.5%	+4.8%	+3.4%	+5.2%	+4.4%
R^2	0.486	0.499	0.770	0.773	0.757	0.758
Within R^2	0.010	0.011	0.025	0.025	0.042	0.043
Mean migration rate	0.026820	0.026804	0.033169	0.033500	0.024660	0.024872
Observations	241,540	235,920	27,700	26,764	24,280	23,488
OD pairs	16,846	16,846	1,948	1,948	1,704	1,704

Notes: First stage (DV: urban migration rate) at the paper’s main specification (origin–destination pair and arrival-year fixed effects, distance-to-census control) in each sample and for each instrument – specification robustness is established on the full sample in Table 1. ‘Contemp.’: $\max(0, -\tilde{z})$ at origin; ‘Lagged’: its 3-season mean over $\{t, t-1, t-2\}$ (lag-restricted sample). Main cities: destinations with *share_urban_known_dest* ≥ 0.55 ; rural origins additionally restrict to origins below the threshold. The lagged instrument has essentially no first stage in the restricted samples – the reason the city-level analyses use the contemporaneous instrument. First-stage F : cluster-robust Wald on the instrument. Effect of 1-SD drought: coefficient $\times$ SD(instrument) $\div$ mean migration rate $\times 100$. 8 countries, 12-year window; SEs clustered at *admin_origin*, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

9.3.2 Reduced Form

Residential Status - This section reports the reduced-form counterparts of the paper’s main estimates: each outcome is regressed directly on the drought instrument at origin, under the main specification, so that no result depends on the strength of the first stage. Table 11

reports the reduced form for the propensity margin. Columns (1)-(2) use the contemporaneous instrument and columns (3)-(4) the cumulative instrument on the lag-restricted sample; within each pair, the first column reports the share of in-migrants living in low-quality housing under the main definition (dummy) and the second the mean number of deprived housing dimensions among the same migrants (continuous). Its significance pattern mirrors the Anderson-Rubin intervals of Table 4, as the two tests coincide.

Table 11: Reduced Form: Drought Intensity and the Propensity for Low-Quality Housing

	<i>Contemporaneous drought</i>		<i>3-season mean (lagged)</i>	
	Deprived Dummy - Main def. (1)	Number of deprivations Continuous (2)	Deprived Dummy - Main def. (3)	Number of deprivations Continuous (4)
Drought intensity	0.00731** (0.00325)	0.00571 (0.00637)	0.01995*** (0.00687)	0.03217** (0.01360)
Effect of 1-SD drought	+0.37 pp	+0.0029	+0.58 pp	+0.0093
R ²	0.392	0.450	0.396	0.455
Within R ²	0.007	0.002	0.007	0.002
Mean outcome	0.406	0.662	0.404	0.656
Observations	72,723	72,723	71,207	71,207
OD pairs	10,537	10,537	10,510	10,510

Notes: Reduced form: propensity outcomes (main UN-Habitat share – floor, wall, roof, sanitation, crowding, water excluded as unmeasurable in these censuses; mean nb. of deprived dimensions) regressed directly on origin drought – no first-stage reliance; RF t -test = AR test of zero in the 2SLS table. Cols (3)–(4): lagged instrument, lag-restricted sample, directly interpretable. Main specification; SEs clustered at `admin.origin`. 1-SD effects in pp / dimension units. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Labor-Force Status - Tables 12 and 13 report the reduced forms for the employment results, regressing the scale and propensity outcomes directly on origin drought for both instruments. On the scale margin, all six inflows rise significantly at the 1% level. On the propensity margin, drought shifts the composition of arrivals from inactivity toward employment under both instruments; since the reduced form does not rely on the first stage, the lagged columns provide valid inference where the lagged first stage of Table 17 is weak, and confirm the same reallocation.

9.3.3 Identifying Variation

Figures 11 and 12 provide a graphical representation of the relationships underlying the identification strategy. Each panel plots mean outcomes within forty equal-sized bins of drought intensity at origin, after residualizing all variables on origin-destination pair and arrival-year fixed effects and the distance-to-census control; the fitted line is estimated on the underlying corridor-year data, so its slope equals the corresponding regression coefficient exactly, and the shaded area is the 95% confidence band based on standard errors clustered at the origin level. All relationships are positive, close to linear, and supported across the

Table 12: Reduced Form: Drought Intensity and the Labor-Force Status of the Inflow: Scale

	<i>Contemporaneous drought</i>			<i>3-season mean (lagged)</i>		
	Employed (1)	Unemployed (2)	Inactive (3)	Employed (4)	Unemployed (5)	Inactive (6)
Drought intensity	0.00155*** (0.00027)	0.00027*** (0.00005)	0.00128*** (0.00033)	0.00315*** (0.00063)	0.00089*** (0.00013)	0.00343*** (0.00078)
Effect of 1-SD drought	+7.3%	+10.1%	+5.5%	+8.8%	+19.5%	+8.7%
R ²	0.474	0.298	0.377	0.486	0.303	0.387
Within R ²	0.006	0.001	0.007	0.006	0.002	0.007
Mean outcome	0.0106	0.00134	0.0118	0.0106	0.00136	0.0118
Observations	241,540	241,540	241,540	235,920	235,920	235,920
OD pairs	16,846	16,846	16,846	16,846	16,846	16,846


Notes: Reduced form: employed / unemployed / inactive in-migrants per 100 urban residents (unwinsorized) regressed directly on origin drought – no first-stage reliance; the RF t -test = AR test of zero in the 2SLS table. Cols (4)–(6): lagged instrument, lag-restricted sample. Effect of 1-SD drought in % of the mean outcome. Main specification (od-pair and year FE, **dist_census** control); SEs clustered at **admin_origin**. Undetermined labor-force status excluded. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 13: Reduced Form: Drought Intensity and the Labor-Force Status of the Inflow: Propensity

	<i>Contemporaneous drought</i>			<i>3-season mean (lagged)</i>		
	Employed (1)	Unemployed (2)	Inactive (3)	Employed (4)	Unemployed (5)	Inactive (6)
Drought intensity	0.00839*** (0.00303)	0.00175 (0.00137)	-0.01014*** (0.00293)	0.01403** (0.00599)	0.00500* (0.00270)	-0.01904*** (0.00576)
Effect of 1-SD drought	+0.42 pp	+0.09 pp	−0.51 pp	+0.41 pp	+0.14 pp	−0.55 pp
R ²	0.329	0.243	0.324	0.335	0.245	0.329
Within R ²	0.003	0.000	0.005	0.003	0.000	0.005
Mean outcome	0.505	0.0578	0.437	0.507	0.0586	0.435
Observations	71,503	71,503	71,503	70,036	70,036	70,036
OD pairs	10,452	10,452	10,452	10,420	10,420	10,420

Notes: Reduced form: employed / unemployed / inactive shares among known-status in-migrants regressed directly on origin drought – no first-stage reliance; the RF t -test = AR test of zero in the 2SLS table, the honest inference for the weak lagged first stage there. The three shares sum to one, so coefficients sum to zero by construction within each instrument. Cols (4)–(6): lagged instrument, lag-restricted sample. Effect of 1-SD drought in percentage points. Main specification (od-pair and year FE, **dist_census** control); SEs clustered at **admin_origin**. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 11: First Stage and Reduced Form: Binned Scatterplots under the Main Specification



Notes: Binned scatterplots of the first-stage (panel a) and reduced-form (panel b) relationships under the main specification. All variables are residualized on origin–destination pair and arrival-year fixed effects and the distance-to-census control (Frisch–Waugh); each dot is the mean within one of forty equal-sized bins of residualized drought intensity at origin. Panel (a) plots the urban migration rate; panel (b) the low-quality-housing inflow rate (main slum definition, unwinsorized). The fitted line is estimated on the underlying corridor-year data and passes through the origin by construction; its slope equals the corresponding regression coefficient exactly (0.00362 in panel a, 0.00103 in panel b). The shaded area is the 95% confidence band based on standard errors clustered at the origin administrative-unit level. Sample: urban migrants, all destinations; 8 countries; 241,540 corridor-years.

Source: Authors' computations from IPUMS census and CHIRPS climate data.

entire distribution of drought shocks rather than driven by a few extreme bins. Corridors experiencing unusually dry growing seasons send more migrants to urban areas, more migrants into low-quality housing, and more employed migrants (Figure 11), the visual counterparts of the first-stage and reduced-form estimates of Tables 1 and 2. The same holds for the propensity outcomes (Figure 12): among migrants, the shares in low-quality housing and in employment both rise with origin drought, the reduced-form counterpart of the sorting results.

9.4 Main Results

9.4.1 Residential Integration

Scale- Table 14 details the calculation behind the interpretation of the main estimate: for a representative corridor with 10,000 urban residents, it translates a one-standard-deviation drought at origin into additional migrants and additional residents in low-quality housing, whose ratio equals, by construction, the 2SLS coefficient of Table 3. Table 15 replicates the main table under the two sample restrictions used in the paper - destinations limited to main cities, and rural origins moving to main cities - with the contemporaneous instrument, whose first stage remains strong in both samples.

9.4.2 Propensity

Propensity - Table 16 re-estimates the propensity outcomes - the share of in-migrants in low-quality housing under the main definition and the mean number of deprived dimensions - in the two restricted samples, with the contemporaneous instrument. The propensity

Table 14: Interpreting the Main Estimate: A Representative Corridor

Panel A. Step-by-step calculation				
	Formula			Value
<i>Baseline (no drought)</i>				
Migrants	$\overline{\text{MigRate}}/100 \times 10,000$			$0.02682 \times 100 = 2.68$
Migrants in low-quality housing	$\overline{S}/100 \times 10,000$			$0.01196 \times 100 = 1.20$
<i>Effect of a 1-SD drought shock</i>				
Δ Migration rate	$\beta_{FS} \times SD(\text{drought})$			$0.00362 \times 0.502 = 0.00182$
Δ Migrants	$\Delta \text{MigRate}/100 \times 10,000$			+0.18
Δ Low-quality housing inflow (S)	$\beta_{2SLS} \times \beta_{FS} \times SD(\text{drought})$			$0.285 \times 0.00362 \times 0.502 = 0.00052$
Δ Migrants in low-quality housing	$\Delta S/100 \times 10,000$			+0.052
<i>Share of new migrants in low-quality housing</i>				
Δ Low-quality / Δ Migrants	β_{2SLS}			$0.052/0.18 = 28\%$
Panel B. Summary				
	Without drought	With 1-SD drought	Change	Share of new migrants in low-quality housing
Number of migrants	2.68	2.86	+0.18 (+6.8%)	–
Migrants in low-quality housing	1.20	1.25	+0.052 (+4.3%)	28%


Notes: Illustration for a single origin–destination corridor with 10,000 urban residents at destination (all-destinations sample, main definition of low-quality housing). Notation as in the text: MigRate is the urban migration rate and S the low-quality housing inflow rate, both per 100 urban residents; bars denote sample means. Because S and MigRate share the same denominator (destination urban population), β_{2SLS} reads directly as the share of marginal, drought-induced migrants who settle in low-quality housing – about 28 of every 100. $\beta_{FS} = 0.00362$, $\beta_{2SLS} = 0.28455$, $SD(\text{drought}) = 0.502$; all values from Table 3.

Table 15: Drought-Induced Migration and Low-Quality Housing Inflows: Main Cities and Rural-Origin Samples

	<i>Main cities</i>			<i>Rural origins → main cities</i>		
	OLS (1)	FS (2)	2SLS (3)	OLS (4)	FS (5)	2SLS (6)
Urban migration rate	0.34155*** (0.02545)		0.13052* (0.07793)	0.32717*** (0.02831)		0.13922* (0.07379)
Drought intensity		0.00312*** (0.00092)			0.00254*** (0.00065)	
First-stage F			11.61			15.16
AR 95% CI			[-0.123, 0.267]			[-0.066, 0.273]
Effect of 1-SD drought (% mean)	+4.6%	+4.8%	+1.8%	+4.7%	+5.2%	+2.0%
Origin-Destination FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dist. to census control	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.879	0.770	0.814	0.856	0.757	0.804
Within R ²	0.589	0.025	0.365	0.528	0.042	0.355
Mean outcome		0.011634			0.008890	
Mean migration rate		0.033169			0.024660	
SD(drought intensity)		0.507			0.505	
Observations		27,700			24,280	
OD pairs		1,948			1,704	

Notes: Contemporaneous instrument throughout; the cumulative instrument is omitted, as its first stage collapses in these samples (see the first-stage appendix table). Columns (1)–(3): destinations restricted to main cities (*share_urban_known_dest* $\geq$ 0.55 at the country’s most recent census). Columns (4)–(6): additionally, rural origins only. OLS and 2SLS columns report the effect of the urban migration rate on the low-quality housing inflow rate (main UN-Habitat definition,); FS columns report the sample-specific first stage of the migration rate on drought intensity. First-stage F : cluster-robust Wald on the instrument. AR 95% CI: cluster-robust Anderson–Rubin confidence interval. Effect of 1-SD drought: for OLS/2SLS columns, coefficient $\times$ first-stage coefficient $\times$ SD(instrument) $\div$ mean outcome $\times$ 100; for FS columns, coefficient $\times$ SD(instrument) $\div$ mean migration rate $\times$ 100. 8 countries (Egypt excluded), 12-year window, Mali 1998 included. SEs clustered at *admin_origin*, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 12: First Stage and Reduced Form for the Propensity Outcomes: Binned Scatterplots



Notes: Binned scatterplots of the first-stage (panel a) and propensity reduced-form (panels b–c) relationships under the main specification. All variables are residualized on origin–destination pair and arrival-year fixed effects and the distance-to-census control (Frisch–Waugh); each dot is the mean within one of forty equal-sized bins of residualized drought intensity at origin. Panel (a) plots the urban migration rate; panel (b) the share of migrants in low-quality housing (main definition, denominator: migrants); panel (c) the share of migrants employed (denominator: known-status migrants). Shares are unwinsorized, as are all outcomes in the paper. The fitted line is estimated on the underlying corridor-year data and passes through the origin by construction; its slope equals the corresponding regression coefficient exactly (0.00628 in panel a, 0.00679 in panel b, 0.00850 in panel c). The shaded area is the 95% confidence band based on standard errors clustered at the origin administrative-unit level. Sample: urban migrants, all destinations; 8 countries; 72,821 corridor-years in panels (a)–(b) (corridors with at least one migrant) and 71,587 in panel (c) (corridors with at least one known-status migrant).

Source: Authors’ computations from IPUMS census and CHIRPS climate data.

subsamples are considerably smaller than for the scale outcome, since they require corridors with observed housing characteristics, and both first stages are moderate ($F = 6.1$ in the main-cities sample, $F = 10.4$ in the rural-to-cities sample), so the Anderson-Rubin intervals carry the inference. The 2SLS point estimates on the extensive margin (2.22 and 2.41) are larger than the full-sample estimate of 1.15, but both Anderson-Rubin intervals are very wide and include zero. These subsamples therefore neither confirm nor contradict the main propensity result; they are too imprecise to be informative on their own.

9.4.3 Labor Market Integration

Table 17 re-estimates the employment results of Tables 5 and 6 with the 3-season cumulative drought instrument. On the scale margin the first stage is strong ($F = 27.4$) and the employed, unemployed, and inactive inflows all rise significantly. On the propensity margin the first stage is weak ($F = 7.1$), so inference relies on the Anderson-Rubin intervals, which confirm the contemporaneous pattern: the employed share of arrivals rises, the inactive share falls, and the unemployed share is unchanged.

9.4.4 Interpreting the effect

Decomposing the Effect for housing - Table 18 details the decomposition of the additional low-quality housing inflow from a one-standard-deviation drought into its scale and propensity components, reporting the formula, inputs, and implied contribution of each channel - roughly 90% through the number of migrants and 10% through their propensity to live in low-quality housing.

Table 16: Propensity for Low-Quality Housing: Main Cities and Rural-Origin Samples

	<i>Main cities</i>				<i>Rural origins → main cities</i>			
	Deprived		Number of deprivations		Deprived		Number of deprivations	
	Dummy - Main def.		Continuous		Dummy - Main def.		Continuous	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)	OLS (7)	2SLS (8)
Urban migration rate	0.07472*** (0.02551)	2.21915 (1.80309)	0.10170** (0.04359)	0.66641 (2.61755)	0.14962** (0.05970)	2.40963 (2.23782)	0.22845** (0.10708)	-0.75675 (3.49205)
First-stage F		6.07		6.07		10.44		10.44
AR 95% CI		[-0.85, 12.82]		[-6.93, 9.90]		[-1.83, 9.44]		[-9.88, 7.20]
Effect of 1-SD drought	+0.01 pp	+0.36 pp	+0.0002	+0.0011	+0.02 pp	+0.33 pp	+0.0003	-0.0010
R ²	0.361	0.229	0.384	0.381	0.364	0.305	0.383	0.379
Within R ²	0.011	-0.193	0.004	-0.001	0.011	-0.080	0.004	-0.002
Mean outcome	0.345	0.345	0.472	0.472	0.346	0.346	0.470	0.470
Observations	17,273	17,273	17,273	17,273	14,720	14,720	14,720	14,720
OD pairs	1,633	1,633	1,633	1,633	1,414	1,414	1,414	1,414

Notes: Propensity (composition) outcomes in the restricted samples, contemporaneous instrument only (the cumulative instrument has no first stage in these samples; see the first-stage appendix table). Columns (1)–(4): destinations restricted to main cities ($share_urban_known_dest \geq 0.55$); columns (5)–(8): additionally, rural origins only. Within each sample, the first pair of columns reports the share of in-migrants living in low-quality housing under the main definition (dummy) and the second pair the mean number of deprived housing dimensions (continuous); denominator: migrants, unweighted, not winsorized. Main specification (origin–destination pair and arrival-year FE, distance-to-census control); SEs clustered at *admin_origin*, in parentheses. First-stage F : cluster-robust Wald on the instrument, on each column's estimation sample. AR 95% CI: cluster-robust Anderson–Rubin confidence interval. Effect of 1-SD drought: coefficient $\times$ first-stage coefficient $\times$ SD(instrument), in pp of the share (dummy) or dimension units (continuous). 8 countries, 12-year window. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 17: Drought-Induced Migration and the Labor-Force Status of the Inflow: Lagged Instrument

	<i>Scale</i>			<i>Propensity</i>		
	Inflows per 100 urban residents			Shares of known-status in-migrants		
	Employed	Unemployed	Inactive	Employed	Unemployed	Inactive
	2SLS (1)	2SLS (2)	2SLS (3)	2SLS (4)	2SLS (5)	2SLS (6)
Urban migration rate	0.33300*** (0.03297)	0.09442*** (0.01398)	0.36242*** (0.03467)	1.38997* (0.74028)	0.49531 (0.31496)	-1.88528** (0.83418)
First-stage F	33.42	33.42	33.42	8.20	8.20	8.20
AR 95% CI	[0.265, 0.402]	[0.071, 0.131]	[0.280, 0.424]	[0.243, 4.758]	[-0.024, 1.855]	[-6.126, -0.731]
Effect of 1-SD drought	+8.8%	+19.5%	+8.7%	+0.41 pp	+0.14 pp	-0.55 pp
R^2	0.725	0.195	0.794	0.159	0.142	0.010
Within R^2	0.468	-0.152	0.667	-0.259	-0.136	-0.469
Mean outcome	0.0106	0.00136	0.0118	0.507	0.0586	0.435
Observations	235,920	235,920	235,920	70,036	70,036	70,036
OD pairs	16,846	16,846	16,846	10,420	10,420	10,420

Notes: 2SLS estimates with the 3-season cumulative drought instrument on the lag-restricted sample. Columns (1)–(3): employed / unemployed / inactive in-migrants per 100 urban residents, unwinsorized; the effect of a 1-SD drought is in % of the mean. Columns (4)–(6): shares among known-status in-migrants (summing to one, so coefficients sum to zero by construction); the effect of a 1-SD drought is in percentage points. The propensity-side first stage is weak ($F \approx 8$): the lagged point magnitudes in columns (4)–(6) are not interpretable and inference should rely on the Anderson-Rubin intervals. Main specification (od-pair and year FE, `dist_census` control); SEs clustered at `admin_origin`. Undetermined labor-force status excluded. Sample: urban migrants, all destinations; 8 countries, 12-year window. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 18: Decomposition of the Drought Effect on Low-Quality Housing Inflows: Scale versus Propensity

	Low-quality housing inflow rate (scale) (1)	Share of migrants in low-quality housing (propensity) (2)
<i>Panel A: Both channels on the common (positive-flow) sample</i>		
Urban migration rate	0.32022*** (0.10891)	1.21591** (0.59790)
First-stage F	15.50	15.50
AR 95% CI	[0.032, 0.525]	[0.170, 2.940]
Mean outcome	0.03974	0.4064
Observations	72,723	72,723
<i>Panel B: Decomposition of a one-standard-deviation drought</i>		
	Calculation	Share
Scale: $\Delta \text{MigRate} \times \bar{P}$ (more migrants, same propensity)	$0.00302 \times 0.4064 = 0.00123$	79%
Propensity: $\bar{\text{MigRate}} \times \Delta P$ (same migrants, higher propensity)	$0.08908 \times 0.00368 = 0.00033$	21%
Total	0.00156	100%

Notes: Decomposition of the additional low-quality housing inflow caused by a one-standard-deviation drought, $\Delta S \approx \Delta \text{MigRate} \times \bar{P} + \bar{\text{MigRate}} \times \Delta P$, per 100 urban residents at destination. The identity is an accounting split only if all four inputs come from the same sample, so both channels are estimated on the sub-sample where the propensity is defined: corridor-years with a positive urban migrant flow ($N = 72,723$). Panel A reports those two 2SLS estimates; column (1) is therefore *not* the paper’s headline scale estimate, which is 0.28455 on the full panel ($N = 241,540$, Table 3). Restricting to positive-flow corridors conditions on the endogenous outcome and is not a defensible main specification; it is used here only so the two channels share a denominator. Panel B inputs, all from that sub-sample: $\Delta \text{MigRate} = \beta_{FS} \times SD(\text{drought}) = 0.00601 \times 0.503 = 0.00302$; $\bar{P} = 40.6\%$, the baseline share of in-migrants in low-quality housing; $\bar{\text{MigRate}} = 0.08908$; and $\Delta P = +0.37$ percentage points, the effect of the same shock on that share. The implied total (0.00156 per 100 residents) exceeds the direct estimate of 0.00097 on the same sample because the identity evaluates the scale channel at the *average* migrant’s propensity (40.6%) rather than the *marginal* migrant’s, which is the scale coefficient itself (32.0%). The decomposition therefore measures the relative weight of the two channels, not the level of the effect. Main specification throughout: origin–destination pair and arrival-year fixed effects, distance-to-census control, migration instrumented by drought intensity at origin, standard errors clustered at the origin district. AR: cluster-robust Anderson–Rubin confidence interval. Housing measure: definition F (five UN-Habitat dimensions, water excluded), unwinsorized; sample restricted to destinations with at least 2,000 urban residents. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Working Poor - Table 19 splits the drought-induced inflow into four employment $\times$ housing cells, measured on the same individuals. About 63% of the migrants entering low-quality housing arrive employed (0.145/0.230), against 44% of the baseline inflow (mean outcomes, columns (1) and (3)): the marginal residents of low-quality housing are predominantly working poor rather than idle poor.

Table 19: Drought-Induced Migration across Employment and Housing Cells: The Working Poor

	<i>Employed</i>		<i>Not employed</i>	
	Low-quality housing 2SLS (1)	Not low-quality housing 2SLS (2)	Low-quality housing 2SLS (3)	Not low-quality housing 2SLS (4)
Urban migration rate	0.14531*** (0.02710)	0.28099*** (0.04028)	0.08435** (0.04179)	0.34414*** (0.04341)
First-stage F	29.65	29.65	29.65	29.65
AR 95% CI	[0.090, 0.203]	[0.208, 0.377]	[-0.022, 0.155]	[0.262, 0.443]
Mean outcome	0.004566	0.006037	0.005855	0.007303
Observations	241,540	241,540	241,540	241,540

Notes: Outcomes: in-migrants per 100 urban residents in the four employment $\times$ low-quality housing cells (main definition), measured on the same individuals – migrants with both labor-force status and housing observed (94.6% of urban migrants). No outcome is winsorized; the sample is restricted to destinations with at least 2,000 urban residents. Column (1) is the “working poor” cell. 2SLS: migration instrumented by contemporaneous drought. Main specification (od-pair and year FE, `dist_census` control); SEs clustered at `admin_origin`. The estimates describe the corridor-level composition of the marginal inflow, not an individual-level causal link between employment and housing. Column (3) is imprecise: the Anderson–Rubin interval includes zero. Sample: urban migrants, all destinations; 8 countries, 12-year window. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

9.5 Individual Heterogeneity: Migrant Profiles

This section asks two distinct questions about the migrants behind the main result: who the drought-induced migrants are (Table 20), and whether the effect varies with their profile - for low-quality housing (Figure 13) and for employment (Figure 14).

Who Are the Drought-Induced Migrants? - Table 20 profiles the drought-induced migrants on gender, age at arrival, and education, comparing the naive read - migrants arriving in high- versus no-drought years - with the complier characterization of Abadie (2003), which describes the marginal migrants identified by the instrument, the same population as the main result. The two readings agree: drought-induced migrants are observationally similar to other migrants - the same gender balance and comparable or slightly higher education - but younger, with an over-representation of children under fifteen, consistent with whole-household displacement rather than the selective migration of young workers. Their higher

propensity for low-quality housing is therefore not accounted for by observably lower human capital.

Table 20: Who Are the Drought-Induced Migrants? Selection on Observables

	<i>Arrival-year drought (descriptive)</i>			<i>Complier characterization (2SLS)</i>		
	No drought (1)	High drought (2)	Diff. (3)	Drought-induced (4)	All migrants (5)	Diff. (6)
Share female (%)	50.5	50.9	+0.4	51.8	50.4	+1.4
Mean age at arrival (years)	21.0	20.3	-0.7***	16.4	21.1	-4.8***
Share below 15 at arrival (%)	29.1	30.3	+1.2	41.6	28.5	+13.1***
Share 15–29 at arrival (%)	52.5	52.2	-0.3	41.9	53.1	-11.2**
Share below primary, 25+ (%)	15.7	14.7	-1.0	11.2	15.2	-4.0
Share secondary or above, 25+ (%)	43.9	45.1	+1.2	47.0	44.7	+2.3
Migrants represented (millions)	5.72	1.29		11.60		

Notes: Characteristics of urban in-migrants, canonical estimation panel (8 countries, Egypt excluded; $N = 241,540$ corridor-years; first-stage $F = 29.65$). Columns (1)–(3), descriptive: migrant-weighted means for arrivals in no-drought ($z = 0$) versus high-drought (top quartile of positive z) corridor-years; the difference is a migrant-weighted WLS contrast. Columns (4)–(6), complier characterization (Abadie, 2003): 2SLS of each characteristic-specific inflow rate (count per 100 urban residents, unwinsorized as in the main tables) on the migration rate instrumented by drought; column (4) is the ratio of coefficients and describes the IV’s marginal drought-induced migrants – the same LATE population as the main result; the difference (6) is tested by 2SLS on the excess outcome. Age at arrival = census age minus years since migration. Education is measured at census and may post-date arrival; adults 25+ with known attainment (~ 40 – 55% of migrants). SEs clustered at `admin_origin`. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Does the Effect Vary by Profile? - Figure 13 estimates the effect separately by subgroup, on both margins, with subgroup-specific inflows and migration rates under the main specification. Effects and propensity shifts are broadly uniform across gender and age - every identified estimate sits close to the aggregate anchor, and no within-dimension contrast is significant. The only gradient is educational: on the scale margin, the effect is twice as strong for primary-educated adults as for those with secondary education or above ($p = 0.036$), suggestive since education is measured at census. Full estimates and contrasts are in Supplementary Appendix Tables S2 and S3.

Figure 14 repeats the exercise for the employment margins. The mirror image completes the picture: employment absorption is uniform where housing is not - there is no educational gradient in the employed inflow ($p = 0.787$), while the housing effect is twice as strong for the primary-educated. Full estimates are in Supplementary Appendix Tables S4 and S5.

9.6 Geography of the Drought-Induced Inflows


Figure 15 maps the destination side of the first stage: for each destination, the share of its observed urban in-migration that the first stage attributes to drought. It illustrates the destination heterogeneity of Section 5.1 geographically – the share is highest in intermediate

Figure 13: Heterogeneous Effects by Migrant Profile



Notes: Panels (a)–(b): relative first stage - the effect of a 1-SD drought on the group's migration rate, as a % of the group's own mean (OLS, so valid regardless of instrument strength; dashed line: the aggregate +6.9%). Panels (c)–(f): subgroup-specific 2SLS under the main specification - the subgroup's low-quality-housing inflow per 100 urban residents (panel (c), unwinsorized; coefficients read per 100 induced migrants of the group) and the share of its migrants in low-quality housing (panel (e), positive-flow corridor-years), with the subgroup's migration rate instrumented by drought at origin. Whiskers: 95% Wald intervals; shaded bars: Anderson–Rubin 95%; dashed lines: the aggregate anchors (0.285, 1.216); shades of blue index the dimension. In the 2SLS panels, subgroups with conditional first-stage $F < 8$ appear as empty positions (n.i.) and are reported AR-only in Supplementary Appendix Tables S2–S3; hollow diamonds: $F < 10$; propensity first stages are weak throughout (F 8–13), so inference rests on the AR intervals. Panels (b), (d), (f): within-dimension contrasts from the saturated stacked interaction. Education is measured at census and may post-date arrival. *Source:* Authors' computations using IPUMS census and CHIRPS climate data.

Figure 14: Heterogeneous Effects by Migrant Profile: Employment




Notes: As Figure 13, for the employment margins (the first stage, common to both figures, is in Figure 13(a)): the subgroup's employed inflow per 100 urban residents (panel (a)) and the employed share among its status-known migrants (panel (c)); dashed lines: the aggregate anchors (0.428, 1.395); shades of orange index the dimension; same display rule (empty n.i. positions: $F < 8$; hollow: F between 8 and 10; weak propensity first stages, F 9–13, AR inference). The age contrast in panel (b) is mechanical - children are rarely employed - and the under-15 column of panel (c) is child employment, conceptually thin. Full estimates: Supplementary Appendix Tables S4–S5.

Source: Authors' computations using IPUMS census and CHIRPS climate data.

destinations and lowest in the largest metropolises – and is an illustration of where the estimated inflows land, not an identified result.

Figure 15: Drought-Induced Share of Urban In-Migration by Destination, by Country



Notes: Circles mark destinations receiving at least 1,000 urban in-migrants (299 of 352); area is proportional to total in-migration, shade to the drought-induced share of it – the first-stage-implied inflow $\sum_{o,t} \hat{\beta}_{k(o,d)} \text{drought}_{o,t} \text{pop}_d^{urb}/100$ divided by the observed one. The share is bounded at destination level (p5 1.3%, median 4.5%, p95 18.7%) and capped at 22% on the scale; the corridor-level analogue is not usable. Grey lines show all 10,399 corridors with a positive urban flow, the darker arcs the 189 largest (top 25 per country). The first stage is estimated by within-country distance tercile rather than pooled, since one coefficient would give short and long corridors the same per-capita response: it falls from 0.00706 on the nearest third of corridors to 0.00323 and 0.00076. The mapped pattern is robust to this (rank correlation 0.935 with the pooled version), and the corrected allocation totals 430,494 drought-induced migrants against 400,139 pooled. The figure illustrates the geography of estimated inflows; the identified result is the bin-level 2SLS analysis of Figure S4. Centroids are interior points of administrative polygons, not city locations.

Source: Authors' computations from IPUMS census microdata, CHIRPS data and IPUMS harmonized shapefiles.

9.7 Robustness


9.7.1 Leave-One-Country-Out

9.7.2 Housing Quality Gradient

9.8 Model

Figure 17 illustrates the two-state housing integration framework of Section 7.

Figure 16: Leave-One-Country-Out Robustness of the Main Effect



Notes: Each row re-estimates the main specification on the estimation sample with one country removed; the diamond is the full eight-country estimate and the shaded band its 95% confidence interval. Thin whiskers are cluster-robust Wald intervals, thick shaded bars the weak-instrument-robust Anderson–Rubin 95% intervals, and the annotation gives the first-stage F of each subsample. The outcome is the low-quality-housing inflow per 100 urban residents (main UN-Habitat definition, unwinsorized); migration is instrumented by drought intensity at origin under the main specification (origin–destination-pair and arrival-year fixed effects, distance-to-census control, standard errors clustered at `admin_origin`). The estimate ranges from 0.197 to 0.340 against 0.285 on the full sample; it stays positive in every subsample and every one of the eight falls inside the full-sample confidence interval, so no single country moves the magnitude. Precision, by contrast, depends on the countries that carry the first stage: dropping Benin or Kenya lowers the first-stage F to about 11 and 12 respectively, and in Benin, Kenya and Mali the Anderson–Rubin interval then includes zero. That reflects lost power rather than a change in the estimate, since the point estimates in those subsamples (0.094, 0.092, 0.134) sit inside the full-sample confidence interval.

Source: Authors’ computations from IPUMS census microdata and CHIRPS climate data.

Table 21: Robustness: Scale and Propensity across Definitions of Low-Quality Housing

	<i>Metal roofs deprived</i>		<i>Metal roofs durable</i>		Nb. of deprived dimensions
	≥ 1 dim.	≥ 2 dim. (conservative)	≥ 1 dim. (main)	≥ 2 dim.	
	(1)	(2)	(3)	(4)	(5)
<i>Panel A. Scale: definition-specific inflow per 100 urban residents</i>					
Urban migration rate	0.94102*** (0.04162)	0.30075*** (0.05247)	0.28455*** (0.06323)	0.13222*** (0.05065)	-0.00382 (0.00469)
First-stage F			29.65		
AR 95% CI	[0.86, 1.04]	[0.18, 0.40]	[0.13, 0.40]	[0.01, 0.23]	[-0.02, 0.00]
Effect of 1-SD drought (% mean)	+7.5%	+5.3%	+4.3%	+4.2%	-1.1%
Mean outcome	0.02275	0.01034	0.01196	0.00566	0.00065
Observations			241,540		
<i>Panel B. Propensity: outcome measured among in-migrants</i>					
Urban migration rate	0.41202 (0.36910)	1.05043* (0.56732)	1.21591** (0.59790)	0.04763 (0.40178)	0.94933 (1.06501)
First-stage F			15.50		
AR 95% CI	[-0.32, 1.34]	[0.06, 2.69]	[0.17, 2.94]	[-0.87, 0.92]	[-1.27, 3.52]
Effect of 1-SD drought	+0.12 pp	+0.32 pp	+0.37 pp	+0.01 pp	+0.0029
Mean outcome	0.8508	0.3475	0.4064	0.1637	0.6620
Observations			72,723		
Origin-Destination FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Dist. to census control	Yes	Yes	Yes	Yes	Yes

Notes: Each column re-estimates both channels under one definition of low-quality housing (contemporaneous instrument, main specification; SEs clustered at `admin_origin`). Column (3) is the paper's main definition, the UN-Habitat operational criterion (≥ 1 deprivation, metal roofing durable). Panel A: inflow per 100 urban residents, unwinsorized. Panel B: share among in-migrants, unweighted. Column (5) uses the mean number of deprived dimensions instead of a share, so its coefficients and 1-SD effects are in dimension units. The first stage is common to all columns within a panel. Water is excluded from every definition: the harmonized IPUMS water-supply codes in these censuses cannot reliably separate deficient from adequate access. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 22: Crowding-Out: Does Drought-Induced Migration Degrade the Housing of Earlier Cohorts?

	Concurrent cohort		Earlier cohorts	
	Full sample (1)	Restricted (2)	Low-quality housing (3)	Total inflow (4)
Urban migration rate	0.28455*** (0.06323)	0.19956 (0.12759)	-1.01066 (0.82838)	0.24737 (1.38003)
First-stage F	29.65	10.80	10.80	10.80
AR 95% CI	[0.130, 0.398]	[-0.287, 0.396]	[-4.499, 0.138]	[-4.944, 2.403]
Mean outcome	0.01196	0.01252	0.04644	0.10137
Observations	241,540	219,589	219,589	219,589
OD pairs	16,846	16,846	16,846	16,846


Notes: Test of whether drought-induced migration crowds previously settled migrants out of decent housing. All columns are the paper’s main specification – the urban migration rate instrumented by drought intensity at origin, origin–destination pair and arrival-year fixed effects, distance-to-census control, standard errors clustered at the origin district. Column (1) is the headline scale estimate (Table 3). Column (2) repeats it on the corridors that have at least one earlier arrival cohort, the sample used in columns (3)–(4). Column (3) replaces the outcome with the low-quality-housing count of migrants who arrived on the same corridor in *earlier* years, per 100 urban residents; those migrants arrived before the shock and cannot be its compliers, so a positive coefficient would be crowding-out. Column (4) is a placebo on their total number, which no mechanism should move. The Anderson–Rubin interval in column (3) excludes any effect above 0.138 per marginal migrant. Note that this channel does not enter the headline outcome, which counts only the concurrent cohort: a null here shows that crowding-out does not spill onto settled migrants, not that it is absent within the arrival cohort (Table 23). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 23: Crowding-Out within the Arrival Cohort: Evidence from Migrants Whose Inflow Does Not Respond to Drought

	Less than primary (1)	Primary (2)	Secondary or more (3)	All migrants (4)
<i>Panel A: Does drought change how many of them arrive?</i>				
Drought intensity	0.000142 (0.000102)	0.000532*** (0.000107)	0.000597*** (0.000119)	0.003618*** (0.000664)
<i>p</i> -value	0.163	0.000	0.000	0.000
<i>Panel B: Does their low-quality-housing propensity rise with drought-induced migration?</i>				
Urban migration rate	1.25066 (0.79559)	1.52444* (0.82700)	1.72953* (1.01358)	1.21591** (0.59790)
First-stage <i>F</i>	13.78	11.57	6.15	15.50
AR 95% CI	[-0.210, 3.526]	[0.192, 4.448]	[0.324, 9.009]	[0.178, 2.942]
Baseline share	0.531	0.395	0.247	0.406
Observations	28,734	37,688	34,071	72,723

Notes: Test of crowding-out *within* the arrival cohort – the channel that would bias the headline estimate. Education is measured at census among adults aged 25 and over. Panel A regresses each group’s own arrival rate per 100 urban residents on drought intensity at origin: the inflow of the least-educated group does not respond to drought (column (1), $p = 0.163$), whereas the other two do, so arrivals of that group in a drought year are almost entirely always-takers. Panel B regresses each group’s share in low-quality housing on the *total* urban migration rate, instrumented by drought, under the main specification. If compliers crowded always-takers out of decent housing, column (1) would be positive. It is positive and of the same order as the aggregate propensity effect, but its Anderson–Rubin interval contains both zero and substantial crowding-out: the test is *inconclusive*, and we do not read it as a null. Column (4) reproduces the aggregate propensity estimate of Table 4 for reference. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 17: The Two-Margin Integration Model: Employment versus Housing



Notes: Each panel is a two-state Roy diagram for one integration margin; each migrant draws state-specific capacities (S_L, S_H) . The two panels show distinct, margin-specific simulated draws, positively correlated across margins. Absent barriers, migrants choose the high-integration state above the ray OC with slope w_L/w_H (blue dashed); OC is the cutoff of baseline migrants, who face their own barriers τ^0 , so the hatched fan represents the differential misallocation $\tau - \tau^0$. Panel (a), employment margin: the low state is valued at the shadow value of non-market time; the employment results imply no differential barrier for drought-induced migrants ($\tau_E \approx 0$), so the cutoff is unchanged and there is no trapped region. Panel (b), housing margin: a residential barrier $\tau_R > 0$ steepens the cutoff to the ray OE with slope $w_L/[(1 - \tau_R)w_H]$ (red dotted); the hatched fan between OE and OC contains the migrants whom the barrier misallocates - they clear the employment margin but remain in low-quality housing, the barrier-induced working poor of Appendix Table 19. This reading presumes positive correlation of capacities across margins, and the position of Region II above the 45° line presumes $(1 - \tau_R)w_H < w_L < w_H$. The hatched region captures the barrier-induced working poor only: migrants below OC may also be employed and badly housed, but optimally so. The 45° line indicates equal capacities across states.

Source: Authors' illustration.

S1 Supplementary Appendix

S1.1 Descriptive Statistics

S1.1.1 Migration Data: Fifteen-Country Sample

Figure S1 reports average in-migration rates over time for the full set of fifteen IPUMS countries, using five-year migration brackets for cross-country comparability; the yearly rates for the eight-country estimation sample are in Appendix Figure 2.

S1.1.2 Urbanization Data

Figure S2 reports the cumulative distributions of the census urban share and of Africapolis city size at destination, separately by country. The height of a country's curve at the 0.15 and 0.55 thresholds gives the share of its destinations in each urbanization bin; the curves confirm that the lower bins are populated by all countries rather than driven by any single one, with Malawi and Mali most concentrated at low urban shares.

S1.2 Scale versus Propensity: A Stylized Example

Figure S3 illustrates the two channels through which drought-induced migration can raise the low-quality housing inflow, using the language of the identification strategy. Migrants who move regardless of the shock (always-takers) settle in low-quality housing at some baseline rate. A drought shock adds compliers - the drought-induced migrants of our 2SLS estimates. Under the *scale* channel alone (middle row), compliers sort into housing exactly like always-takers: the share of migrants in low-quality housing is unchanged, and only the number of arrivals rises - a congestion problem, calling for housing supply where flows land. Under the *propensity* channel (bottom row), compliers are additionally more likely to end up in low-quality housing, shifting the composition of arrivals - a vulnerability problem, calling for support targeted at the marginal migrants themselves. The numbers in the figure are illustrative; in our estimates, the baseline share is 41 percent, the marginal drought-induced migrant is about 11 percentage points more likely to be in low-quality housing, and the propensity channel accounts for roughly 21 percent of the total effect (Section 4.3.1).

S1.3 Heterogeneity by the Census Urban Share

This section repeats the destination heterogeneity of Section 5.1 on the census urban share instead of Africapolis agglomeration size, as a complementary check. Destinations are split into three bins by their urban population share at the most recent census: rural-like (< 0.15), urbanizing cities ($[0.15, 0.55)$) and main cities (≥ 0.55). The specification is identical to the one in the body.

Figure S4 and Table S1 report the estimates. The direction matches the size-based result: the coefficient declines as destinations urbanize, from 0.324 in rural-like destinations to 0.258 in urbanizing cities and 0.131 in main cities, the last significant only at the 10% level. The magnitudes are also consistent with the body, since the rural-like and urbanizing bins





Figure S1: Average In-Migration Rates Over Time by Country.

Notes: This figure displays mean in-migration rates across administrative units for 15 African countries, disaggregated by destination type. Migration flows are calculated from IPUMS census microdata, measuring the number of in-migrants (weighted by `perwt`) divided by the destination population, aggregated into 5-year brackets. The panel titles show the geographic level (GEOLEV1 = region/province, GEOLEV2 = district/commune), the number of administrative units, and total observations. Note that Y-axis scales vary across countries to facilitate within-country comparisons. Countries with a single census year (Sierra Leone, South Africa) provide limited temporal variation. The figure reveals substantial heterogeneity in migration patterns, with urban destinations consistently attracting higher rates than rural areas, particularly in rapidly urbanizing countries like Ghana, Togo, and Uganda.

Source: Authors' elaboration based on IPUMS census data.

Figure S2: Distributions of Census Urban Share and Africapolis City Size, by Country



Notes: The figure reports empirical cumulative distribution functions across destinations, separately by country. Panel (a) shows the destination's urban population share in the most recent census. At each dashed threshold, 0.15 and 0.55, the height of a country's curve gives the proportion of its destinations below that threshold. Panel (b) shows the 2020 population of the largest agglomeration in the Africapolis 2024 database located within, or within 10 km of, each destination unit, on a logarithmic scale. Destinations without a matched Africapolis agglomeration of at least 10,000 inhabitants are excluded from panel (b). Four Africapolis agglomeration records reporting conurbation-cluster totals are corrected as described in the notes to Table 9.

Source: Authors' computations using IPUMS census data and the Africapolis 2024 database.

Table S1: Heterogeneity by Destination Urbanization: 3-Bin Interacted IV Estimates

	<i>Low-quality housing inflow</i>			<i>Employed inflow</i>		
	Rural-like [†] (1)	Urbanizing cities (2)	Main cities (3)	Rural-like [†] (4)	Urbanizing cities (5)	Main cities (6)
Urban migration rate	0.3239*** (0.0994)	0.2579*** (0.0629)	0.1305* (0.0779)	0.4732*** (0.0714)	0.3651*** (0.0497)	0.4947*** (0.0602)
SW conditional F	14.01	24.17	11.62	14.01	24.17	11.62
AR 95% CI	[0.067, 0.524]	[0.102, 0.372]	[-0.123, 0.267]	[0.358, 0.703]	[0.263, 0.474]	[0.381, 0.674]
Mean outcome	0.0153	0.0096	0.0116	0.0122	0.0086	0.0139
Observations	91,019	122,821	27,700	91,019	122,821	27,700
Urbanizing – rural-like		$p = 0.515$			$p = 0.152$	
Urbanizing – main cities		$p = 0.128$			$p = 0.043$	
Urbanizing – rest		$p = 0.596$			$p = 0.092$	

Notes: 3-bin interacted IV: the urban migration rate interacted with each destination-urbanization bin (share urban at the most recent census; zones: rural-like < 0.15 , urbanizing cities $[0.15, 0.55]$, main cities ≥ 0.55), instrumented by contemporaneous drought intensity interacted with the same indicators; origin–destination and arrival-year $\times$ bin fixed effects, bin-specific `dist_census` control, SEs clustered at `admin_origin`. Outcomes per 100 urban residents (no outcome is winsorized). SW F : Sanderson–Windmeijer conditional first-stage F ; [†] weak first stage ($F < 10$) – the AR interval carries the inference there. The first stage is common to both outcomes (identical sample); per-bin coefficients 4.23 / 2.75 / 3.12 ($\times 10^{-3}$). Contrast p -values from the same system's clustered covariance; "rest" pools the rural-like and main-cities bins. Sample: corridors with non-missing destination urbanization ($N = 193,290$). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure S3: Scale versus Propensity: More Migrants, or Different Migrants?



Notes: Stylized example. Grey dots: always-takers (migrants who move regardless of the drought shock); yellow dots: compliers (drought-induced migrants); green outline: higher-quality housing; orange outline: low-quality housing. Top row: no shock - always-takers settle in low-quality housing at the baseline rate. Middle row: a drought shock adds compliers who sort like always-takers - the scale channel; the share in low-quality housing is unchanged. Bottom row: compliers are additionally more likely to end up in low-quality housing - the propensity channel; the share among all migrants rises. Shares are illustrative; see Section 4.3.1 for the estimated counterparts.

Source: Authors' illustration.

together cover much of what Africapolis classifies as small towns and secondary cities. What this classification does not deliver is inference: no pairwise contrast along the gradient is significant – urbanizing versus rural-like $p = 0.515$, urbanizing versus main cities $p = 0.128$, rural-like versus main cities $p = 0.119$ – and the urbanizing bin does not differ from the other two pooled ($p = 0.596$). For employment the ordering is different, with absorption lowest in urbanizing cities (0.365 against 0.473 and 0.495), and the urbanizing–main-cities contrast is the one significant difference here ($p = 0.043$).

The reason we do not use this classification in the body is that its inference depends on where the rural/urban cutoff is placed, and the urban-share distribution offers no break to place it at. Across 351 destinations the share rises smoothly to a mode near 0.09–0.10 and declines slowly thereafter; the largest gap between consecutive destination shares below 0.35 is 0.014, against a median gap of 0.0008, so there is no separation anywhere in the range. Moving the lower cutoff therefore moves the contrasts without any data-driven reason to prefer one position:


Lower cutoff	0.05	0.10	0.15	0.20	0.25
Secondary bin (N)	186,843	151,097	122,821	93,884	71,652
Housing: secondary – main cities	+0.167	+0.183	+0.127	+0.146	+0.117
p -value	0.065	0.040	0.128	0.099	0.194
Employment: secondary – main cities	−0.068	−0.077	−0.130	−0.170	−0.139
p -value	0.325	0.220	0.043	0.006	0.026

The housing contrast crosses the 5% level at one cutoff only, and the employment contrast does so at a different and disjoint set of cutoffs. We therefore report the census-urban-share bins at the conventional 0.15 threshold, treat the declining gradient as directional evidence consistent with the body, and rest the inference on the size-based classification, which requires no such choice. The rural-like bin is also weakly identified here (Sanderson–Windmeijer conditional $F = 14.0$ against 24.2 in the urbanizing bin), so its Anderson–Rubin interval carries the inference.

Figure S5 reports the urbanization heterogeneity for the propensity (sorting) margin.

Heterogeneous Effects by Migrant Profile - Tables S2–S5 report the full estimates behind Appendix Figures 13 and 14: the main effect separately by migrant subgroup, on the scale margin (the subgroup’s low-quality-housing inflow) and the propensity margin (the share of the subgroup’s migrants in low-quality housing), each with subgroup-specific migration rates instrumented by drought under the main specification; the reference column reproduces the aggregate anchors. Both margins are broadly uniform across gender and age - no within-dimension contrast is significant. The only gradient is educational: on the scale margin, the effect is twice as strong for primary-educated adults as for those with secondary education or above ($p = 0.036$), suggestive given that attainment is measured at census. The oldest and least-educated adults are AR-only on both margins - drought barely moves their migration, itself informative - and the propensity first stages are weak throughout (F between 8 and 13), so propensity inference rests on the Anderson–Rubin intervals: the sorting shift is detected among men and adults, and imprecise for women and children.

Figure S4: Heterogeneity by Destination Urbanization: Low-Quality Housing and Employment



Notes: Heterogeneity by destination urbanization (urban population share at the most recent census; zones: rural-like < 0.15 , urbanizing cities $[0.15, 0.55]$, main cities ≥ 0.55). Panels (a)-(c): first-stage and 2SLS estimates from a single specification saturated by urbanization bin -the migration rate interacted with each bin indicator, instrumented by drought intensity interacted with the same indicators, with origin-destination and arrival-year $\times$ bin fixed effects and bin-specific distance-to-census controls; standard errors clustered at the origin district. Estimates are plotted at bin midpoints; whiskers are 95% Wald intervals and shaded bars Anderson-Rubin 95% intervals; outcomes in panels (b)-(c) are the low-quality-housing and employed inflows per 100 urban residents. For housing, no pairwise contrast is significant: urbanizing – rural-like $p = 0.515$, urbanizing – main cities $p = 0.128$, rural-like – main cities $p = 0.119$; the urbanizing bin does not differ from the other two pooled ($p = 0.596$). For employment, the urbanizing bin lies significantly below main cities ($p = 0.043$); the other contrasts are not significant. Panels (d)-(f): rolling-window local IV – the main specification re-estimated on corridors into the ~ 60 destinations nearest each urban-share grid point; pointwise 95% Wald bands; indicative only, as windows overlap. The shaded area marks where the window no longer slides (too few destinations above 0.55); the binned estimates in panels (a)–(c) carry the inference there. Hollow points flag a weak first stage (Sanderson–Windmeijer conditional $F < 10$; bin values 14.0 / 24.2 / 11.6). Bottom row: distribution of destination urban shares with the three zones labeled. $N = 241,540$ corridor-years (91,019 / 122,821 / 27,700 across bins).

Table S2: Heterogeneous Effects by Migrant Profile: Low-Quality Housing Inflows (Scale)

	<i>Gender</i>			<i>Age at arrival</i>			<i>Education, 25+</i>		
	All (1)	Women (2)	Men (3)	<15 (4)	15+ (5)	55+ (6)	<Prim. (7)	Primary (8)	Sec.+ (9)
Subgroup urban migration rate	0.3106*** (0.0858)	0.2799*** (0.0645)	0.2862*** (0.0757)	0.3242*** (0.0926)	0.2535*** (0.0574)	AR-only	AR-only	0.3928*** (0.0898)	0.2158*** (0.0606)
First-stage F (SW cond.)	23.86	29.70	25.88	19.96	28.66	0.10	1.95	24.71	25.00
AR 95% CI	[0.08, 0.46]	[0.12, 0.40]	[0.10, 0.42]	[0.08, 0.49]	[0.11, 0.36]	[-16.9, 17.6]	[-1.80, 3.37]	[0.19, 0.57]	[0.09, 0.35]
Contrast p (within dimension)		0.916			0.378			0.092	
Share of migrants (%)	100.0	50.4	49.6	28.5	71.5	1.4	7.8	20.5	22.8
Dist. control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OD-pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean outcome	0.0123	0.0061	0.0059	0.0050	0.0070	0.0003	0.0019	0.0021	0.0010
Observations	245,292	241,540	241,540	241,540	241,540	241,540	241,540	241,540	241,540
OD pairs	16,888	16,846	16,846	16,846	16,846	16,846	16,846	16,846	16,846

Notes: Heterogeneous effect of drought-induced migration on low-quality housing inflows (main definition), scale margin, by migrant profile. Each column is a separate 2SLS: the outcome is the subgroup's low-quality-housing inflow per 100 urban residents (unwinsorized; country), the endogenous variable its migration rate, instrumented by drought at origin – canonical specification (od-pair and year FE, *dist_census* control, SEs clustered at **admin_origin**). Coefficients read as the number per 100 additional drought-induced migrants of the subgroup settling in low-quality housing; column (1) is the aggregate reference and reproduces the published anchor (0.228). Contrast p -values from the fully saturated stacked interaction. Effects are broadly uniform across gender and age – no significant contrast; the only gradient is educational (primary > secondary+, $p = 0.036$), suggestive since education is measured at census and may post-date arrival. Columns with conditional first-stage F below 8 (55+, less than primary) are AR-only – drought barely moves their migration, itself informative. Propensity-margin counterpart: Table S3. Age at arrival = census age minus years since migration. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table S3: Heterogeneous Effects by Migrant Profile: Share in Low-Quality Housing (Propensity)

	<i>Gender</i>			<i>Age at arrival</i>			<i>Education, 25+</i>		
	All (1)	Women (2)	Men (3)	<15 (4)	15+ (5)	55+ (6)	<Prim. (7)	Primary (8)	Sec.+ (9)
Subgroup urban migration rate	1.2159** (0.5979)	0.9462 (1.0923)	3.1375** (1.5943)	0.0117 (1.2202)	3.9751** (1.6423)	AR-only	AR-only	AR-only	AR-only
First-stage F (SW cond.)	15.50	19.48	10.11	12.53	11.27	2.33	0.70	5.16	7.27
AR 95% CI	[0.17, 2.94]	[-1.19, 3.59]	[0.61, 9.29]	[-3.06, 2.64]	[1.6, 10.6]	[-450.4, 568.2]	[-367.9, 430.1]	[1.6, 91.4]	[1.7, 38.9]
Contrast p (within dimension)		0.177			0.039			–	
Share of migrants (%)	100.0	50.4	49.6	28.5	71.5	1.4	7.8	20.5	22.8
Dist. control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OD-pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean outcome	0.406	0.413	0.394	0.424	0.395	0.438	0.531	0.395	0.247
Observations	72,723	57,284	58,921	46,618	64,621	8,305	28,734	37,688	34,071
OD pairs	10,537	9,337	9,573	8,478	9,981	2,948	6,657	7,272	6,712


Notes: Heterogeneous effect of drought-induced migration on the propensity margin, by migrant profile. Each column is a separate 2SLS: the outcome is the share of the subgroup’s migrants in low-quality housing (unweighted, not winsorized), on corridor-years with a positive subgroup flow (per-column N); the endogenous variable is the subgroup’s migration rate, instrumented by drought at origin – canonical specification (od-pair and year FE, *dist.census* control, SEs clustered at `admin.origin`). Column (1) is the aggregate reference and reproduces the published anchor (1.081). Contrast p -values from the fully saturated stacked interaction; the education contrast is not estimable (both columns below the F gate). Propensity shifts are broadly uniform across gender and age – no significant contrast; the shift is detected among men and adults (AR intervals excluding zero) and imprecise for women and children. First stages are weak throughout (F 8–13): inference rests on the cluster-robust Anderson–Rubin intervals. Columns with conditional F below 8 are AR-only – 55+ and less than primary are not identified on this margin either (drought barely moves their migration, itself informative), and all education columns are AR-only here. Education is measured at census and may post-date arrival. Scale-margin counterpart: Table S2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table S4: Heterogeneous Effects by Migrant Profile: Employed Inflows (Scale)

	<i>Gender</i>			<i>Age at arrival</i>			<i>Education, 25+</i>		
	All (1)	Women (2)	Men (3)	<15 (4)	15+ (5)	55+ (6)	<Prim. (7)	Primary (8)	Sec.+ (9)
Subgroup urban migration rate	0.4276*** (0.0433)	0.3455*** (0.0515)	0.5155*** (0.0524)	0.0768*** (0.0269)	0.6775*** (0.0564)	AR-only	AR-only	0.7932*** (0.0805)	0.9139*** (0.0403)
First-stage F (SW cond.)	29.65	29.70	25.88	19.96	28.66	0.10	1.95	24.71	25.00
AR 95% CI	[0.35, 0.53]	[0.25, 0.47]	[0.42, 0.64]	[0.02, 0.14]	[0.57, 0.81]	[-54.6, 52.1]	[-1.84, 3.31]	[0.63, 0.97]	[0.83, 1.00]
Contrast p (within dimension)		0.003			0.000			0.175	
Share of migrants (%)	100.0	50.4	49.6	28.5	71.5	1.4	7.8	20.5	22.8
Dist. control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OD-pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean outcome	0.0106	0.0042	0.0064	0.0009	0.0097	0.0002	0.0017	0.0032	0.0031
Observations	241,540	241,540	241,540	241,540	241,540	241,540	241,540	241,540	241,540
OD pairs	16,846	16,846	16,846	16,846	16,846	16,846	16,846	16,846	16,846

Notes: Heterogeneous effect of drought-induced migration on employed inflows, scale margin, by migrant profile. Each column is a separate 2SLS: the outcome is the subgroup's employed inflow per 100 urban residents (unwinsorized), the endogenous variable its migration rate, instrumented by drought at origin – canonical specification (od-pair and year FE, *dist_census* control, SEs clustered at `admin_origin`; $N = 245,292$ per row). Column (1) reproduces the published anchor (0.339). Contrast p -values from the fully saturated stacked interaction. The age contrast is MECHANICAL – children are rarely employed, so their employed inflow is near zero by construction – and should not be read as economic heterogeneity. Employment absorption is otherwise uniform: no significant gender contrast, and no educational gradient ($p = 0.787$), in contrast with the housing margin (Table S2). Columns with conditional first-stage F below 8 (55+, less than primary) are AR-only – drought barely moves their migration, itself informative. Education is measured at census and may post-date arrival. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure S5: Heterogeneity by Destination Urbanization: Migrant Propensity



Notes: The figure reports estimates from a three-bin, fully interacted IV specification examining heterogeneity in migrant propensity by destination urbanization. Propensity outcomes are defined only for corridor-years with positive migrant flows; the samples contain $N = 55,698$ observations for housing and $N = 55,323$ for employment. Panel (a) reports the first-stage estimates and the corresponding Sanderson–Windmeijer conditional F -statistics. The conditional F -statistic is below 10 in each bin, as indicated by hollow markers; estimates in panels (b) and (c) should therefore be interpreted cautiously. Whiskers report conventional 95% Wald confidence intervals, while shaded bars report weak-instrument-robust 95% Anderson–Rubin confidence intervals. Intervals extending beyond the displayed axis range are truncated. The Anderson–Rubin interval excludes zero only for the share of migrants living in low-quality housing in urbanizing cities, with a confidence interval of $[0.81, 14.2]$. This result is consistent with the scale heterogeneity reported in Figure S4, although none of the pairwise contrasts is statistically significant. The specification otherwise matches that used in Figure S4.

Source: Authors’ computations using IPUMS census and CHIRPS climate data.

Table S5: Heterogeneous Effects by Migrant Profile: Employed Share (Propensity)

	<i>Gender</i>			<i>Age at arrival</i>			<i>Education, 25+</i>		
	All (1)	Women (2)	Men (3)	<15 (4)	15+ (5)	55+ (6)	<Prim. (7)	Primary (8)	Sec.+ (9)
Subgroup urban migration rate	1.3945** (0.6169)	2.4854** (1.0848)	1.2001 (1.3858)	2.4288** (0.9511)	1.0167 (1.0656)	AR-only	AR-only	AR-only	AR-only
First-stage F (SW cond.)	15.51	19.29	10.17	16.33	11.17	2.29	0.67	4.98	7.26
AR 95% CI	[0.40, 3.32]	[0.69, 5.62]	[-1.78, 5.11]	[0.86, 5.27]	[-0.96, 4.31]	[-665.5, 526.8]	[-284.0, 249.7]	[-3.0, 66.6]	[-4.4, 17.4]
Contrast p (within dimension)		0.434			0.314			–	
Share of migrants (%)	100.0	50.4	49.6	28.5	71.5	1.4	7.8	20.5	22.8
Dist. control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OD-pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean outcome	0.505	0.385	0.614	0.143	0.638	0.461	0.592	0.702	0.796
Observations	71,503	56,113	57,537	42,415	64,593	8,300	28,717	37,671	34,057
OD pairs	10,452	9,248	9,459	8,047	9,979	2,945	6,652	7,271	6,710

Notes: Heterogeneous effect of drought-induced migration on the employment propensity margin, by migrant profile. Each column is a separate 2SLS: the outcome is the employed share among the subgroup’s migrants with known labor-force status (unweighted, not winsorized), on corridor-years with a positive status-known subgroup flow (per-column N); the endogenous variable is the subgroup’s migration rate, instrumented by drought at origin – canonical specification. Column (1) reproduces the published anchor (1.319). Contrast p -values from the fully saturated stacked interaction; the education contrast is not estimable (both columns below the F gate). No contrast is significant. The under-15 column is child employment, conceptually thin. First stages are weak throughout (F 9–13): inference rests on the cluster-robust Anderson–Rubin intervals. Columns with conditional F below 8 are AR-only (55+, less than primary, and all education columns on this margin). Education is measured at census and may post-date arrival. Scale counterpart: Table S4.

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

Saturated First Stage by Profile - Table S6 reports the fully saturated first stage across the twelve gender $\times$ age-at-arrival $\times$ education cells - the additivity check behind Section 5.2: selection into migration runs along the educational gradient alone, with all less-than-primary cells near zero across gender and age.

Table S6: Saturated First Stage by Migrant Profile: Who Does Drought Displace?

	Share of migrants (%)	Relative first stage (1-SD drought, % of cell mean)	(SE)
<i>Panel A: Secondary education or above</i>			
Women, arrived <30	13.7	8.38***	(1.99)
Women, arrived 30+	4.6	13.56***	(3.59)
Men, arrived <30	16.9	5.92***	(2.01)
Men, arrived 30+	9.4	8.16***	(2.25)
<i>Panel B: Primary completed</i>			
Women, arrived <30	13.7	5.62***	(1.71)
Women, arrived 30+	6.0	7.30***	(2.69)
Men, arrived <30	12.5	5.06***	(1.70)
Men, arrived 30+	8.0	5.14**	(2.03)
<i>Panel C: Less than primary (the trapped margin)</i>			
Women, arrived <30	4.6	2.77	(2.03)
Women, arrived 30+	4.2	3.82	(2.68)
Men, arrived <30	3.2	1.14	(2.93)
Men, arrived 30+	3.2	0.86	(2.66)
Equality of the 12 responses (Wald)		$p = 0.289$	
Interactions beyond additive margins (Wald, 7 terms)		$p = 0.955$	


Notes: Fully saturated first stage by migrant profile: the corridor-year migration rate of each of the twelve gender $\times$ age-at-arrival $\times$ education cells (urban migrants aged 25+ at census with known attainment), regressed on drought intensity at origin under the canonical specification (od-pair and year FE, *dist_census* control, SEs clustered at *admin_origin*; $N = 245,292$ corridor-years per cell). Because rainfall is exogenous, each cell's OLS first stage is valid regardless of instrument strength. The reported effect is the RELATIVE first stage - a 1-SD drought as a % of the cell's own mean migration rate - the scale comparable across cells (raw coefficients mostly reflect cell size); the aggregate value is +6.9%. Joint tests from the stacked model on the standardized outcome, clusters shared. All four less-than-primary cells are near zero (first-stage $F \leq 0.75$) - the trapped-population result, robust across gender and age; its formal test is the pooled contrast ($p = 0.001$) reported in the main text. Education is measured at census and may post-date arrival; the top cell (women arrived 30+, secondary+) is only 4.6% of migrants. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Does Drought Select on Destination Type? - Table S7 reports the linear projection of a causal forest fitted on the first stage, a data-driven check on whether drought directs migrants toward a particular kind of destination. Because the treatment is rainfall at origin, unconfoundedness holds and the forest recovers how the migration response varies with corridor and destination characteristics; the same exercise would be invalid on the second

stage, where the migration rate is endogenous. The response does vary across corridors – the calibration test rejects homogeneity – but only corridor distance and origin urban share have projection coefficients distinguishable from zero. No destination characteristic does, including the urban share that defines the bins of Figure S4. The forest also recovers, without being told, the distance decay imposed by hand elsewhere in the paper: the predicted response falls from 0.0061 in the nearest within-country distance tercile to 0.0032 and 0.0019.

Composition of the Housing Index - Figures S6 to S9 open up the index dimension by dimension. Figure S6 pools the eight countries and shows which dimensions actually drive the measure; Figures S7 and S8 report, country by country, which material and facility categories fall on the deprived side of each rule; Figure S9 shows the distribution of persons per habitable room against the three-person threshold.

Figure S6: Which Dimensions Drive the Housing Index



Notes: Urban households, person-weighted. An in-migrant is a person whose unit of origin differs from the unit of enumeration and whose arrival year falls in that census's retrospective window. Water is omitted throughout: deprivation is nil in every sample, because these censuses distinguish only types of piped water while UN-Habitat counts boreholes and protected dug wells as improved sources. Share of urban households deprived on each dimension, eight countries pooled. Crowding is the only dimension on which in-migrants and residents coincide; on every material dimension in-migrants are about half as often deprived.


Source: Authors' computations from IPUMS census microdata.

Table S7: Does Drought Select on Destination Type? Linear Projection of the First-Stage Causal Forest

	Projection coefficient (per 1 SD, $\times 10^{-3}$)	(SE)	<i>p</i> -value
<i>Corridor</i>			
Distance, within-country percentile	-2.424**	(1.119)	0.030
Distance (km)	0.017	(0.900)	0.985
<i>Origin</i>			
Urban share	-1.691**	(0.738)	0.022
Growing-season rainfall (mm)	-1.268	(0.968)	0.190
<i>Destination</i>			
Urban share	-0.847	(0.549)	0.123
Baseline low-quality housing share	-1.081*	(0.561)	0.054
Baseline employed share	0.186	(0.582)	0.749
Growing-season rainfall (mm)	-0.061	(0.805)	0.940
Population density (per km ²)	0.055	(0.247)	0.823
Mean $\hat{\tau}(x)$	0.00369	(pooled first stage 0.00362)	
Calibration test, differential forest prediction	1.105	$(p = 1.4 \times 10^{-22})$	
Corridor-years / corridors / origin clusters	241,540 / 16,846 / 354		
Trees	1000		

Notes: Best linear projection of the conditional average partial effect $\hat{\tau}(x)$ estimated by a causal forest (**grf**, 1000 trees) of drought intensity at origin on the corridor's urban migration rate, with standard errors clustered at **admin_origin**. The projection, not the forest's variable-importance ranking, is the object to interpret: importance splits credit arbitrarily between correlated covariates - the two distance measures and the two rainfall measures here - and ranked destination urban share second although its projection coefficient is indistinguishable from zero. The forest has no fixed effects, so the migration rate and drought intensity are first residualized on origin-destination pair and arrival-year fixed effects with the *dist_census* control and the forest is fitted on the residuals ($\hat{Y} = \hat{W} = 0$): the identifying variation stays within corridor, while the covariates, being corridor- and destination-level and time-invariant, carry the heterogeneity across corridors. As an anchor, the mean of $\hat{\tau}(x)$ reproduces the pooled first stage (0.00369 against 0.00362). Coefficients are reported per one standard deviation of the covariate and in units of 10^{-3} , because the raw units are not comparable (kilometres, millimetres, persons per km², and 0–1 shares); this rescaling is monotone and leaves *p*-values unchanged. The calibration test rejects the absence of heterogeneity ($p = 1.4 \times 10^{-22}$), so the response does vary across corridors - but only corridor distance and origin urban share have projection coefficients distinguishable from zero, and no destination characteristic does. The two destination composition covariates are measured among migrants, not residents, since the estimation panel contains no resident population: they are the low-quality-housing and employed shares of urban in-migrants arriving in the destination's earliest census window, computed leave-one-origin-out with a 30-observation floor, so their nulls are partly a measurement limitation rather than established absences of selection. This is a first-stage validation exercise - drought is exogenous, so unconfoundedness holds - and says which destinations drought directs migrants toward; it is not an estimate of the effect of migration on housing. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Figure S7: Durability: Materials on the Deprived Side of the Rule



Notes: Urban households, person-weighted. An in-migrant is a person whose unit of origin differs from the unit of enumeration and whose arrival year falls in that census's retrospective window. Water is omitted throughout: deprivation is nil in every sample, because these censuses distinguish only types of piped water while UN-Habitat counts boreholes and protected dug wells as improved sources. Each bar is the category composition of one dimension for one country and population. Deprived categories are named individually and shown in warm shades when they exceed 5 per cent of a bar; all adequate categories are pooled into a single grey, since the question here is where the deprivation cut falls. Non-permanent wall materials are the unburnt-brick, adobe, mud and plant families; burnt or stabilised brick, cement-stabilised earth, stone, concrete, metal and wood count as permanent. Sheet-metal roofs count as adequate.

Source: Authors' computations from IPUMS census microdata.


Figure S8: Sanitation: Facilities on the Deprived Side of the Rule



Notes: Urban households, person-weighted. An in-migrant is a person whose unit of origin differs from the unit of enumeration and whose arrival year falls in that census's retrospective window. Water is omitted throughout: deprivation is nil in every sample, because these censuses distinguish only types of piped water while UN-Habitat counts boreholes and protected dug wells as improved sources. Only the complete absence of a toilet counts as deprivation. Pit latrines, which account for roughly half of the pooled sample, are not counted: IPUMS records them in a single undifferentiated category, and UN-Habitat treats an improved pit latrine with a slab as adequate sanitation, so the conservative reading is adopted.

Source: Authors' computations from IPUMS census microdata.

Figure S9: Living Space: Persons per Habitable Room



Notes: Urban households, person-weighted. An in-migrant is a person whose unit of origin differs from the unit of enumeration and whose arrival year falls in that census's retrospective window. Water is omitted throughout: deprivation is nil in every sample, because these censuses distinguish only types of piped water while UN-Habitat counts boreholes and protected dug wells as improved sources. A household is deprived when more than three people share a habitable room, the UN-Habitat threshold. Habitable rooms exclude kitchens, corridors, verandas and bathrooms. Cameroon 2005 and Uganda 2014 are absent: their censuses record only the number of rooms used for sleeping, which is not comparable to a room count, so crowding is left unmeasured there rather than approximated.

Source: Authors' computations from IPUMS census microdata.