Trapped at Home: Weather Shocks, Income, and International Migration*

Céline Azémar[a], Rodolphe Desbordes[b],
Markus Eberhardt[c], and Eric Neumayer[d]

[a]Rennes School of Business
[b]SKEMA Business School, Paris and Université Côte d'Azur, Nice
[c]University of Nottingham and Centre for Economic Policy Research
[d]London School of Economics and Political Science

This version: June 15, 2026

Abstract: By mid-century, climate change is expected to push millions of migrants across borders, yet direct empirical evidence is limited. Using recent monthly bilateral flow data for 127 developing-country origins and 180 destinations, we estimate migration responses to weather shocks. High-dimensional fixed-effects absorb all origin-, destination-, and migration corridor-by-year confounders, seasonality, and global shocks; identification exploits residual within-corridor weather anomalies. We find no evidence that climate stress raises emigration: higher origin temperatures reduce outflows, while precipitation has no effect. Complementary analysis pinpoints depressed origin income as an underlying mechanism: climate stress tightens household budgets, decreasing the financial viability of international movement.

Keywords: climate change, international migration, gravity, immobility trap, budget constraint

JEL codes: F22, J61, Q54, O15, Q15

*Correspondence: Rodolphe Desbordes, SKEMA Business School Paris, 5 Quai Marcel Dassault, 92150 Suresnes, France; email: rodolphe.desbordes@skema.edu

1 Introduction

The notion of the climate refugee has become a fixture of public debate and a focus of international policy. For instance, the Institute for Economics and Peace (2020) argues that ecological threats associated with climate change put 1.2 billion people at risk of displacement by 2050. Even if only a small fraction of those affected were to move across borders,[1] this would likely be sufficient to generate large destabilising economic and political effects in the receiving countries. However, in a chapter devoted to climate-driven migration, the IPCC (2022) concludes that the evidence concerning the interactions between climate change and migration is unclear.[2] Moreover, papers directly addressing the climate change-international migration nexus are relatively scarce, especially those with a global coverage, and are constrained by a lack of detailed migration data for the recent time period. For instance, Beine & Parsons (2015) and Cattaneo & Peri (2016) use changes (net migration flows) in the decadal bilateral migrant stocks computed by Özden et al. (2011) over the 1960–2000 period.[3] These data limitations preclude a fixed-effect structure rich enough to absorb the many confounders potentially correlated with climate change, a difficulty that echoes long-standing concerns in the climate-income literature (Dell et al. 2014, Burke et al. 2015, Kalkuhl & Wenz 2020, Desbordes & Eberhardt 2024).

We address these constraints by combining granular data on migration and weather, as well as (in extensions) origin income and domestic conflict, with a battery of fixed effects tight enough to identify the effects and channels of interest under minimal omitted-variable concerns. Our weather-migration analysis draws on the bilateral flow data of Chi et al. (2025a), released through the Humanitarian Data Exchange, which aggregate Meta (Facebook) cross-border movement counts into a monthly dyadic panel covering 127 developing-country origins and 180 destinations over 2019–2022. The monthly frequency permits identification under a four-way fixed-effect structure: origin-destination-year, origin-by-calendar-month, destination-by-calendar-month, and year-by-calendar-month. This is richer than the origin-year, destination-year, and origin-destination design standard in the gravity literature. The origin-destination-year effect absorbs not only persistent corridor characteristics (such as distance, shared language, bilateral labour agreements) but, being finer than an origin-year or a destination-

[1] According to Abel & Cohen (2019), the existing average yearly global migration flow is between 6 and 20 million, about 0.5-1% of the predicted people at risk.

[2] The IPCC writes "A general theme across studies from all regions is that climate-related migration outcomes are diverse (high confidence) and may be manifest as decreases or increases in migration flows, and may lead to changes in the timing or duration of migration and to changes in migration source locations and destinations" (p. 1079). Recent meta-analyses (Cattaneo et al. 2019, Beine & Jeusette 2021, Hoffmann et al. 2020) confirm this reading of the literature.

[3] Beine & Parsons (2015) and Gröschl & Steinwachs (2017) report no robust direct effect of climate change, whereas Cattaneo & Peri (2016) document a non-monotonic response: warmer years reduce emigration from poor origins and raise it from middle-income ones. Marchiori et al. (2012) find that weather anomalies increase international migration from Sub-Saharan African countries, while Missirian & Schlenker (2017) show that higher source-country temperature raises asylum applications to the European Union. However, both studies employ a monadic source-level specification (net international migration rate or log annual applications on source-country weather with source and year fixed effects).

year effect, it also absorbs (i) every origin-year shock (e.g., origin GDP growth, exchange-rate movements, emigration or passport policy, elections and political crises); (ii) every destination-year shock (e.g., destination labour demand and business cycle, immigration and visa-quota reforms, prevailing wages); and indeed (iii) any shock specific to the country-pair itself in a given year (e.g., a new bilateral labour agreement, a corridor-specific visa change, or a spell of bilateral tension). The two calendar-month effects absorb origin- and destination-specific seasonality (e.g., planting and harvest calendars, the academic year, seasonal labour demand at the destination), and the year-by-calendar-month effect absorbs global month-on-month shocks common to all dyads, including the COVID mobility collapse. This rich fixed-effect structure is precisely what insulates the estimates from the COVID shock during our sample period: the common monthly path of the pandemic and any migration corridor-year COVID shock are absorbed by the fixed effects. Any remaining COVID influences are captured by additional origin- and destination-level mobility controls.[4] Identification of the weather elasticity rests on the residual within-corridor, within-calendar-month variation. We capture climate exposure in both origin- and destination-countries as the population-weighted country-month aggregates of temperature and precipitation, taken from the Climatic Research Unit Time Series (CRU TS) v4.09 (Harris et al. 2020), using alternative weighting schemes for robustness. To capture dynamics we employ distributed lags of various lengths of these variables.

We find that warmer conditions in origin-countries are associated with *less* recorded bilateral movement, not more, whereas precipitation exerts no significant effect. Specifically, the cumulative twelve-month effect of a one-degree origin-temperature anomaly is a reduction in bilateral migration flow by 5.7 percent. Climate stress in the form of abnormally high temperatures keeps households at home rather than driving them abroad. On the destination side, both climate variables matter and both enter negatively: over the same twelve-month horizon, higher destination temperature reduces inflows by 5 percent and higher destination precipitation (per 30 mm/month) reduces inflows by 8.8 percent. These estimates are stable across alternative fixed-effect structures, sample restrictions, different climate-aggregation weights, COVID mobility parameterisations, spatial-autocorrelation controls, a leave-one-destination-out test, and a placebo permutation test. We further corroborate our findings using annual bilateral arrival flows built from three decades of IPUMS microdata for 177 origins. This external-validity exercise is imperfect, with coarser identification and only 25 destinations, but in contrast to alternative data sources it preserves the annual frequency the mechanism demands.[5]

The bilateral migration response is heterogeneous along one dimension and homogeneous along another in ways that bear directly on its interpretation. Splitting the developing origins at the median of 2018 per capita GDP shows that the response does not vary detectably

[4]The bilateral migration variation we exploit is itself standard by historical comparison, not unique to this time window (Appendix A.4).

[5]The five-yearly sampling of standard bilateral flow series, such as Abel & Cohen, 2019, averages away the transitory weather anomalies that identify the effect.

with origin income: the pooled elasticity is representative across the developing country-origin income distribution. Splitting at the median of baseline temperature (each origin's 1990–2018 average), by contrast, shows that the negative response is concentrated in already-hot origins, and the two halves of the distribution differ significantly at the 5 percent level, consistent with the concave temperature–income damage function the macro-climate literature documents. We further exploit the bilateral structure of the panel to test how the migration response varies across corridors. Interacting origin temperature with the 2015 bilateral migrant stock in isolation yields a negative interaction, the opposite of the network-driven amplification documented by the literature on United States-bound corridors (Mahajan & Yang 2020, Ibáñez et al. 2026). However, once destination GDP per capita enters the same horse race alongside standard gravity controls, the migrant stock interaction loses significance and the destination-income interaction is significantly negative: the migration response is stronger in corridors to richer destinations. This is the direct prediction of a budget-constraint mechanism, since reaching a high-income destination is more expensive (visas, travel, the financial reserves that richer-country immigration regimes increasingly require), so a climate-driven income loss binds harder on those routes. The negative stock interaction effect in isolation was simply picking up this destination-income gradient.

Our results support a resource-constrained 'immobility-trap' hypothesis (Black et al. 2013, Benveniste et al. 2022), which we present formally in a stylised model (Section 2): a negative income shock reduces the financial ability of households to cover the cost of a cross-border move.[6] To probe this channel, we apply the same identification strategy used for the weather-migration relationship to the income response itself. A sub-national 0.5° grid of per capita GDP at annual frequency (Rossi-Hansberg & Zhang 2025) carries cell and country-year fixed effects that absorb persistent within-country averages and unobserved time-varying country-level confounders. A country-month nighttime-lights panel from NASA's Black Marble product (Román et al. 2018), estimated under the bilateral regression's grouped distributed-lag specification, carries country and year-month fixed effects that absorb persistent country differences and the global month-on-month path. Both panels return a precisely estimated negative temperature elasticity of origin income, while the precipitation channel is null at the origin in both the cell-level grid and the bilateral panel. The gridded estimates imply that a one-degree annual temperature anomaly reduces origin income by about 2 percent, a magnitude in line with the existing macro literature (Dell et al. 2012, Burke et al. 2015, Kalkuhl & Wenz 2020, Desbordes & Eberhardt 2024).

Lastly, we show that other destructive shocks to origin livelihoods also exert a negative effect on out-migration. We enter a monthly conflict-risk indicator from Aizenman et al. (2026) into the bilateral regression under the same grouped distributed-lag structure as the weather variables. Origin internal conflict risk carries a negative cumulative effect on bilateral migration flows, and weather effects are essentially unchanged once we include temperature

[6] The migration-development literature supports the same interpretation: emigration rises with income in poor origins before falling at higher levels of development (Clemens 2014).

and precipitation (again in grouped lags). A non-climate origin shock that erodes income and the capacity to finance a move thus also exerts a negative effect on bilateral migration: the constraint mechanism is general, and the weather elasticity is not proxying for latent conflict.

The paper makes three contributions, unified by a single, theory-based methodological strategy and constituting data at three levels of resolution, each paired with a fixed-effect structure dense enough to hold the leading confounders fixed. *First*, on a monthly panel of 127 developing-country origins and 180 destinations, the bilateral weather-migration elasticity is identified under a four-way fixed-effect structure, tighter than the annual cross-country gravity literature can run on its data. The finding provides reduced-form bilateral panel evidence for the climate-immobility hypothesis that the climate refugee discourse (Myers 1995, Brown 2008) has commonly overlooked. *Second*, the income channel is identified at two resolutions finer than the existing macro climate-income evidence: a 0.5° sub-national GDP grid (Rossi-Hansberg & Zhang 2025) at annual frequency with cell and country-year fixed effects, finer than the 1,500-region panel of Kalkuhl & Wenz (2020); and a country-month nighttime-lights panel (Román et al. 2018) with country and year-month fixed effects, which adds temporal resolution beyond the annual cross-country standard (Dell et al. 2012, Burke et al. 2015, Desbordes & Eberhardt 2024). *Third*, we show that the income-constraint mechanism is not specific to climate: adding a monthly domestic conflict-risk indicator to the model reveals that origin internal conflict risk affects bilateral migration flow with the same negative sign while leaving the weather coefficients intact. This points to a budget-constraint mechanism that any destructive origin income shock can trigger, rather than a channel specific to climate change.

The remainder of this paper is structured as follows. Section 2 sets out the theoretical framework. Section 3 describes the bilateral panel and the empirical strategy. Section 4 reports the pooled elasticity, documents its robustness and heterogeneity, and tests how it varies across corridors. Section 5 introduces the two income panels, tests the income channel on both, recovers the implied migration-income elasticity, and shows that a non-climate origin shock (internal conflict) similarly exerts a negative effect on outward migration, independent of the negative effect arising from higher temperatures. Section 6 concludes.

2 A model of constrained migration

A unified household at origin i decides whether to migrate to destination j at time t. Migration is a binary choice $m \in \{0, 1\}$ carrying a fixed bilateral cost c_{ij}: for example, visa fees, recruitment fees, travel, and the financial reserves that destination immigration regimes frequently demand. Household income is $y_{i,t}$; subsistence requires $\underline{c}$. The value of staying combines consumption utility over disposable income with a location amenity $A_{i,t}$ that captures the physical and other viability and attractiveness of the origin (housing, productive assets,

basic infrastructure, social network, but also personal safety):

$$V_{i,t}^{\text{stay}} \;=\; u\big(y_{i,t} - \underline{c}\big) \;+\; A_{i,t}, \tag{1}$$

with $u' > 0$, $u'' < 0$. The household solves

$$\max_{m \in \{0,1\}} \; (1-m)\, V_{i,t}^{\text{stay}} \;+\; m\,\big[V_{ij,t}^{\text{move}} - c_{ij}\big], \tag{2}$$

subject to a *financing constraint* requiring that disposable income cover the migration cost up front:

$$y_{i,t} - \underline{c} \;\geq\; c_{ij}. \tag{3}$$

Without (3) the household migrates whenever $V^{\text{move}} - c_{ij} > V^{\text{stay}}$; with it, migration additionally requires the financial slack to finance the move. The origin weather shock $\delta_{i,t}$ introduced below enters the value of staying $V_{i,t}^{\text{stay}}$ only; the move value $V_{ij,t}^{\text{move}}$ is determined by conditions at the destination and is taken to be independent of the origin shock, so all comparative statics below operate through $V_{i,t}^{\text{stay}}$. The object we take to data, the bilateral flow num_{ijym}, is the count of origin-i households executing $m = 1$ towards destination j. With households heterogeneous in income $y_{i,t}^0$ and in the cost c_{ij} they face, this count is a smooth function of a common shock even though each household's choice is discrete; the comparative statics below describe the response of that corridor-level rate, traced out by the mass of households near the two margins the model identifies.

A weather shock $\delta_{i,t} \geq 0$ enters through two channels. We write the model for a single generic stressor, but in the empirics temperature and precipitation enter as separate regressors, and the logic below applies to whichever anomaly erodes origin income or amenity. First, the *income channel* is a reduced-form semi-elasticity $\varepsilon_{y,T} < 0$ (the proportional change in income per temperature degree or per precipitation increment),

$$y_{i,t}(\delta_{i,t}) \;=\; y_{i,t}^0\,\big(1 + \varepsilon_{y,T}\,\delta_{i,t}\big), \tag{4}$$

whose magnitude reflects the share of income tied to climate-sensitive activities. It is a local object: equation (4) linearises a relationship the macro-climate literature finds concave in the temperature level, so $\varepsilon_{y,T}$ varies across origins and steepens in hotter locations, consistent with the baseline-temperature split we explore in Section 4.3.2. The semi-elasticity is also a local linearisation around each origin's seasonal norm (mean): while the model is written in stressor terms with $\delta_{i,t} \geq 0$, our empirical specification (Section 3.4) enters temperature and precipitation in levels and identifies $\varepsilon_{y,T}$ from within-origin deviations in either direction, with the origin-by-month fixed effects absorbing each origin's local norm. The linear-symmetric response is appropriate for a developing-country sample at the small deviations our identification exploits: both above-norm temperatures and above-norm precipitation are documented stressors on agricultural output and household productivity in

already-warm, semi-tropical and tropical regimes, where droughts are also harmful, so that the income response need not be sign-symmetric in the population. The linear-symmetric specification is decomposed into mutually exclusive compound climate-event categories (heat-only, cold-only, drought-only, flood-only, hot+dry, hot+wet, cold+dry, cold+wet) in Section 4.3.1, which traces the baseline weather elasticity to heat-only and drought-only events operating through the income channel and identifies compound hot+dry and hot+wet events as the corner where amenity-dominated distress migration emerges.

Second, the *amenity channel* is an elasticity $\varepsilon_{A,T}$ on which we impose only a sign restriction: it is non-positive but otherwise unrestricted, admitting both the zero case (climate stress leaves amenity unchanged) and any strictly negative magnitude (climate stress erodes amenity, the more so the larger $|\varepsilon_{A,T}|$),

$$A_{i,t}(\delta_{i,t}) \;=\; A^0_{i,t} + \varepsilon_{A,T}\,\delta_{i,t}, \qquad \varepsilon_{A,T} \le 0. \tag{5}$$

The migration response to a weather shock operates through two margins, depending on whether the financing constraint (3) is slack or binds for a given household. When the constraint is *slack* ($y^0_{i,t} - \underline{c}$ comfortably exceeds c_{ij}), it holds with room to spare, so the move-stay decision turns purely on desirability. Differentiating the indifference condition $V^{\text{move}}_{ij,t} - c_{ij} = V^{\text{stay}}_{i,t}$ and using (4)–(5) gives the response of the flow at this *desirability margin*,

$$\left.\frac{\partial m}{\partial \delta_{i,t}}\right|_{\text{slack}} \propto \; -\underbrace{u'\big(y^0_{i,t} - \underline{c}\big)\cdot \varepsilon_{y,T}\, y^0_{i,t}}_{\text{income channel}} - \underbrace{\varepsilon_{A,T}}_{\text{amenity channel}} \;>\; 0. \tag{6}$$

With $\varepsilon_{y,T} < 0$ and $\varepsilon_{A,T} \le 0$, both terms in (6) are non-negative, the income channel strictly so, leaving the right-hand side positive: V^{stay} falls, households at the desirability margin flip into migrating, and the bilateral flow rises. This is the income-push channel of the standard climate-migration narrative. Instead, when the constraint *binds* ($y^0_{i,t} - \underline{c} \approx c_{ij}$), a second margin is active. The same income shock pulls disposable income below c_{ij} for households sitting at the financing threshold; (3) fails, and a household that was willing and able to move is mechanically forced from $m = 1$ to $m = 0$. Its contribution to the flow is governed by

$$\left.\frac{\partial m}{\partial \delta_{i,t}}\right|_{\text{binding}} \propto \; \underbrace{y^0_{i,t}\,\varepsilon_{y,T}}_{\text{financing margin}} \;<\; 0, \tag{7}$$

the (resource-constrained) immobility trap: climate-driven income loss reduces realised migration not because the value of moving falls but because the migration option is mechanically removed. The two margins coexist within any migration corridor, and the aggregate climate response of the corridor flow is their density-weighted sum,

$$\frac{\partial m}{\partial \delta_{i,t}} \;=\; \kappa^{\text{des}}\Big[-u'\big(y^0_{i,t} - \underline{c}\big)\,\varepsilon_{y,T}\,y^0_{i,t} - \varepsilon_{A,T}\Big] \;+\; \kappa^{\text{fin}}\, y^0_{i,t}\,\varepsilon_{y,T}, \tag{8}$$

with $\kappa^{\text{des}}, \kappa^{\text{fin}} \geq 0$ the masses of households at the desirability and financing margins. The first bracket is non-negative and the second term non-positive, so the sign of the aggregate response is an empirical question, settled by which margin carries the greater mass. Slack and binding describe a household's own constraint, not that of an origin and, therefore, not of a corridor: every origin and every corridor mixes households at both margins, and the sign records which margin dominates the mix. Acute amenity shocks act only through the first bracket: a large negative $\varepsilon_{A,T}$ raises it, so even where the financing margin is wide the net response can turn positive, as feasible households who would not otherwise have moved are driven out by an origin that has become physically unviable.[7]

The sign of the migration response is not a model prediction but an empirical question. By equation (8) it depends on the relative mass at the two margins (κ^{des} versus κ^{fin}) and on the relative magnitudes of the income and amenity elasticities, and the two signs are not equally revealing. A *negative* response can arise only when the financing margin dominates with the income channel in control: the desirability bracket in (8) is non-negative, so only its second (financing) term can pull the aggregate below zero. A *positive* response is less revealing, being consistent either with desirability-margin dominance (κ^{des} large) or, even where the financing margin carries substantial mass, with an acute amenity shock ($|\varepsilon_{A,T}|$ large) that lifts the desirability bracket above the financing term. The sign is therefore informative but not, on its own, decisive about the mechanism. Nevertheless, the financing-margin reading that a negative elasticity points to carries two auxiliary implications that an unconstrained income-push model does not. First, the same weather shocks that depress outmigration must also depress origin income, since that is the causal mechanism ($\varepsilon_{y,T} < 0$ on the same origins). Second, any negative non-climate origin income shock must also exert a negative effect on bilateral migration.

The empirical strategy speaking to these theoretical considerations thus has three steps:

1. **Sign and dominant margin.** We estimate the pooled bilateral elasticity of migration to weather shocks in the dyadic migration panel (Section 4).

2. **Income channel.** A binding-constraint reading requires $\varepsilon_{y,T} < 0$ on the same origin countries: whichever weather shock (temperature, precipitation, or both) contracts the bilateral flow must also depress origin-country income. We test this using a sub-national grid of GDP per capita (Section 5.1) and a country-month panel of nighttime-lights (Section 5.2).

3. **Mechanism generality.** If the constraint is general and not climate-specific, a negative non-climate income shock in the origin country should also exert a negative effect on

[7] The positive source-temperature effect on asylum applications reported by Missirian & Schlenker (2017) can be read as such an amenity-dominated case. Distress migration in which *constrained* households relax the financing constraint itself (informal borrowing, distress asset sales, irregular routes) lies outside this stylised model, which reads the amenity case through the desirability margin only. Capturing it requires replacing the hard constraint (3) with a costly-financing technology in which a shortfall $\max\{0,\, c_{ij} - (y_{i,t} - \underline{c})\}$ is met at a convex penalty; that extension, derived in Appendix A.1, nests the present model and opens a second route by which the amenity channel can overturn the sign for constrained households.

bilateral migration. We test this by adopting a calibrated monthly conflict-risk indicator in specifications on its own and alongside weather shocks (Section 5.4).

3 Climate and migration: data and empirical strategy

3.1 Bilateral migration flows

Our dependent variable is the monthly count of bilateral migration arrivals from origin i to destination j in the year-month ym, num_{ijym}. Flows are drawn from the Humanitarian Data Exchange release of Chi et al. (2025a), which converts the location histories of roughly three billion monthly active Meta (Facebook) users into a bilateral migration panel covering 181 countries at monthly frequency. The data construction (see Chi et al. 2025b, for details) implements the UN definition of migration through a segment-based algorithm: for each user, the algorithm identifies segments of consecutive days in which the predicted home country is the same, allowing gaps of up to 60 days. A user is counted as migrating from i to j when two adjacent segments belong to different countries, each segment lasts at least 12 months, and the user is present in the segment country on at least 50% of its days. Transit stays shorter than 12 months are not counted. The twelve-month residence requirement on each side excludes short-stay cross-border travel; the count is closer to a migration decision than to a flow of border crossings, but rare permanent moves with intermediate-length transits are missed. The same forward-residence requirement right-censors the final months of the panel: a household arriving shortly before the sample endpoint has not yet accumulated the twelve months of destination residence the definition requires, so the latest months understate the true flow. Because this censoring is common to all corridors in a given calendar month, it is absorbed by the year-by-calendar-month fixed effects of our specification (Section 3.4) and leaves the within-origin-month weather variation that identifies the weather shock effect unaffected. The bilateral flow variation over our sample is, moreover, not atypical by historical standards: once each country's overall migration level is netted out, the corridor-level flow shocks over 2019–2022 are in line with the 1990–2015 record (Abel & Cohen 2019, from), allowing for the latter's coarser five-yearly sampling (Appendix A.4), hence our monthly panel captures normal migration dynamics rather than an unusual episode.

A population-level estimate is obtained from the Meta data by applying a *selection-rate* weighting whose denominator combines the country's Facebook penetration rate with an income-scaled adjustment. The income-scaled term receives more weight in low-income countries, where selection into Facebook usage and selection into migration both load on the same wealthier sub-population. Unfilled origin-destination-month cells are imputed as zeros, allowing extensive-margin adjustment. Our sample retains those origins classified as developing countries by the 2018 World Bank income taxonomy ($n = 127$), is unrestricted on destinations ($n = 180$), and covers 2019–2022 (48 monthly flow observations; climate data are extended back to January 2018 to capture the longest twelve-month distributed lag).[8]

[8] Twelve developing economies have no Meta-recorded outflow and do not enter the panel as origins; the four populous cases (China, Iran, North Korea, and Cuba) are countries where Meta platforms are blocked or severely restricted, so their absence is geopolitical rather than evidence of zero emigration; the remainder are small island micro-states, disputed territories, and Somalia, where recorded Facebook coverage is too sparse to register outflows. The estimation sample is therefore confined to developing origins where Facebook

Figure 1 visualises the panel's migration corridor structure as a flow diagram aggregated to the seven World Bank regions. South Asia, Latin America & Caribbean, and East Asia & Pacific are the largest sending regions, and the Middle East & North Africa, Europe & Central Asia, and North America the largest recipients. Our four-way fixed-effect structure (Section 3.4) should absorb any persistent corridor-level discrepancy between the population-representative estimate and the unobserved true flow.[9]

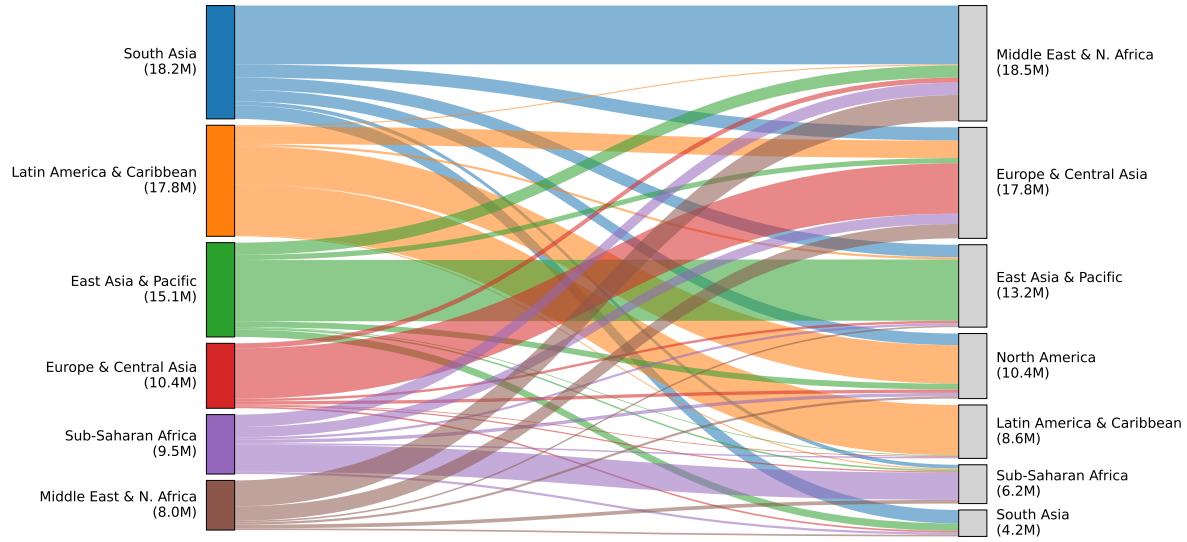

Figure 1: Migration flows by World Bank region, 2019-2022. Sending region on the left, receiving region on the right; ribbon width is proportional to total bilateral migration count over the panel window in millions of migrants. Origins restricted to the developing-country estimation sample; destinations unrestricted. Region-by-region corridors with total flow below 100,000 migrants omitted.

3.2 Climate exposures

Monthly country-level temperature (TMP) and precipitation (PRE) are aggregated from the CRU TS v4.09 0.5° gridded series as

$$\bar{x}_{i,t,m} = \frac{\sum_{c \in i} w_c \, A_{c,i} \, x_{c,t,m}}{\sum_{c \in i} w_c \, A_{c,i}}, \tag{9}$$

where c indexes cells inside country i, $A_{c,i}$ is the fractional area of cell c inside country i computed from the Database of Global Administrative Areas (GADM), and w_c is a cell weight. Precipitation is divided by 30 mm before estimation so that one unit corresponds to roughly the same share of within-origin-month variability as one degree Celsius of temperature,

coverage is sufficient to observe migration.

[9]For further corroboration, we compare our data against Eurostat's annual bilateral immigration counts for EU-27 destinations over the same 2019–2022 window. The two series correlate at 0.89 in logs across 10,257 corridor-year observations.

putting the temperature and precipitation coefficients on comparable magnitudes.[10]

We construct five weighting schemes, each fixed at the 2015 vintage to keep the weight exogenous to the 2019–2022 climate variation captured in the regressions. The baseline uses population weights: migration is initiated by people, and the relevant exposure is the climate the average resident faces. The *population* weight uses Gridded Population of the World v4 rev11 (Center for International Earth Science Information Network (CIESIN), Columbia University 2018) cell counts at 0.5°, adjusted to UN World Population Prospects 2015 country totals. The *area* weight uses $\cos(\phi_c)$ at cell latitude ϕ_c. The *cropland* weight uses Global Agro-Ecological Zones (GAEZ+) 2015 monthly cropland data (Grogan et al. 2021). The *GDP* weight uses cell-level gridded GDP from the Becker Friedman Institute (BFI) Data Studio sub-national release (Rossi-Hansberg & Zhang 2025). The *low-income* weight uses $\text{pop}_c^2/\text{GDP}_c$, equivalent to population weighted by inverse per capita GDP; both numerator and denominator are from BFI. The five weights span the within-country distribution of climate exposure: population (the baseline) weights country by the average resident, area weights as a unit of land, GDP the average dollar of output, low-income the population weighted by inverse GDP per capita, and cropland the average farmer. Section 4.2 exploits the differences. The representativeness of the 2019–2022 weather variation that identifies our estimates is documented in Appendix A.3 and is discussed where it bears on identification in Section 3.4.

Table 1: Definitions of corridor-quality flags used in the sample-restriction robustness.

Flag	Definition (corridor $i \to j$; panel totals over 2019–2022)	Share of observations
Gulf	Destination is a Gulf Cooperation Council state (Saudi Arabia, United Arab Emirates, Qatar, Kuwait, Bahrain, Oman).	3.4%
Transit	Destination is one of thirteen documented transit hubs: Belarus, Colombia, Ecuador, Greece, Guatemala, Honduras, Libya, Morocco, Mexico, Nicaragua, Panama, Serbia, Türkiye.	7.2%
Circular	Two-way corridor: the smaller directional flow exceeds half the larger ($\min / \max > 0.5$) and the combined two-way flow falls in the top decile across all corridors.	4.2%
Any	Any of the three above; the higher-quality sample is its complement (86% of observations).	14.0%

3.3 Controls and migration corridor classifiers

We control for COVID mobility restrictions using the Oxford COVID-19 Government Response Tracker (OxCGRT). $\text{mob}_{i,ym}^{\text{COVID}}$ and $\text{mob}_{j,ym}^{\text{COVID}}$ are the first principal component of

[10]The rescaling only changes the absolute scale of the reported precipitation elasticity (multiplying any reported precipitation coefficient by 30 recovers the mm-per-month interpretation) and does not affect statistical significance.

four OxCGRT sub-indicators (public transport closure, stay-at-home requirements, internal-movement restrictions, international-travel controls) at origin and destination, respectively, entered contemporaneously and with a one-month lag. The PCA score is normalised to mean zero and unit variance over 2020–2022 and set to zero in 2018–2019.

Three corridor-quality 'flags' (Gulf labour, circular, transit) enter the sample-restriction robustness check of Section 4.2.1, intended to gauge the robustness of our estimates against anomalous/exceptional migration patterns; their definitions are provided in Table 1, which reports the rule used to construct each flag and the share of observations it marks. Table 2 reports summary statistics for the estimation panel.

Table 2: Descriptive statistics, monthly bilateral migration panel, 2019–2022.

	Mean	SD	p10	Median	p90	Min	Max
Bilateral migration count (thousands)	0.08	1.03	0.00	0.00	0.05	0.00	163.94
Origin temperature (°C)	21.87	7.46	11.58	24.02	28.37	-21.93	36.94
Origin precipitation (mm/day)	3.49	3.64	0.15	2.23	8.52	0.00	34.90
Destination temperature (°C)	19.88	8.72	6.22	22.58	28.23	-21.93	39.41
Destination precipitation (mm/day)	3.21	3.34	0.15	2.16	7.82	0.00	34.90
Origin COVID mobility index	0.04	1.73	-1.39	-0.93	2.87	-1.39	5.00
Destination COVID mobility index	-0.04	1.66	-1.39	-1.01	2.70	-1.39	5.00

Notes: Monthly bilateral migration panel of Chi et al. (2025a), 2019–2022, restricted to developing-country origins (WB 2018 classification). $n = 919{,}532$ origin-destination-month observations across 127 origin countries and 180 destinations. Approximately one-third of dyad-months have a zero bilateral migration count. Origin and destination climate variables are population-weighted country-month aggregates of CRU TS v4.09 temperature and precipitation. The COVID mobility index is the first principal component of four OxCGRT sub-indicators (public transport closure, stay-at-home requirements, internal-movement restrictions, international-travel controls), normalised to mean zero and unit variance over 2020–2022 and set to zero in 2018–2019.

3.4 Specification

The estimation equation for bilateral flows is

$$\begin{aligned}\mathsf{E}[\mathsf{num}_{ijym} \mid \cdot] \;=\; \exp\Big\{ & \sum_{\ell=0}^{L}\big(\beta_\ell^o \mathrm{TMP}_{i,ym-\ell} + \gamma_\ell^o \mathrm{PRE}_{i,ym-\ell} + \beta_\ell^d \mathrm{TMP}_{j,ym-\ell} + \gamma_\ell^d \mathrm{PRE}_{j,ym-\ell}\big) \\ & + \sum_{\ell=0}^{1}\big(\eta_\ell^o\, \mathsf{mob}^{\mathsf{COVID}}_{i,ym-\ell} + \eta_\ell^d\, \mathsf{mob}^{\mathsf{COVID}}_{j,ym-\ell}\big) \\ & + \;\phi\, \ln(\mathsf{num}_{ij,ym-1}+1) \;+\; \psi\, \mathbf{1}[\mathsf{num}_{ij,ym-1}=0] \\ & + \;\alpha_{ijt} + \alpha_{im} + \alpha_{jm} + \alpha_{ym} \Big\}, \end{aligned} \tag{10}$$

where TMP and PRE are the origin's (superscript o) and destination's (d) population-weighted temperature and precipitation in year-month ym lagged by up to $L = 12$ months; the COVID controls, lagged dependent variable, and extensive-margin indicator in rows 2 and 3 are detailed above (COVID) and below; and $\alpha_{ijt}, \alpha_{im}, \alpha_{jm}, \alpha_{ym}$ are origin-destination-year, origin-calendar-

month, destination-calendar-month, and year-calendar-month fixed effects. We estimate this model by Poisson Pseudo-Maximum-Likelihood (PPML) with high-dimensional fixed effects (Correia et al. 2020), standard for gravity equations with positive and zero flows (Santos Silva & Tenreyro 2006).

The set of four fixed-effects absorb the following. The origin-destination-year effect α_{ijt} removes everything constant within a corridor in a given year. Because it is finer than a simple origin-year or destination-year effect, it absorbs not only persistent dyadic migration intensity (e.g., diaspora networks, distance, shared language, colonial ties, bilateral labour agreements) but also any *origin-year* shock (e.g., origin GDP growth, exchange-rate movements, national emigration or passport policy, elections and political crises) and any *destination-year* shock (e.g., destination labour demand and business cycle, immigration and visa-quota reforms, prevailing wages, asylum-regime changes). The origin-calendar-month effect α_{im} removes origin-specific seasonal cycles (e.g., planting and harvest calendars, the academic year, religious-calendar timing such as Ramadan and the Hajj). The destination-calendar-month effect α_{jm} removes destination-specific seasonal labour demand (e.g., farming, construction, and tourism seasons). The year-calendar-month effect α_{ym} removes global month-on-month shocks common to all dyads, including the global COVID mobility collapse, world commodity prices, and the global business cycle. The variation remaining in weather shocks is the deviation of country i's temperature in year-month ym from its origin-by-month seasonal norm, purged of the year-by-month global path. Identification of β^o_ℓ rests on this within-origin-month residual variation, distributed over the 12-month lag. Similarly for γ^o_ℓ and the equivalent destination terms.

Although the panel spans only four years, the identifying variation is not particular to the window: the distribution of weather shocks over 2019–2022 is statistically indistinguishable from a 1990–2018 baseline, so the elasticity we recover is not an artifact of an anomalous weather pattern in the estimation window (see Appendix A.3). The maintained identifying assumption is that, conditional on the four-way fixed effects, no within-origin-month determinant of bilateral flow co-moves with the origin weather anomaly; the temporal-permutation test of Section 4.2.5 probes, but cannot by itself prove, this assumption. Standard errors are clustered at the origin-destination dyad. The cumulative effect at horizon h reported in our results tables is the sum of the grouped-lag coefficients over $\ell = 0, \ldots, h$, with contemporaneous ($\ell = 0$), short-lag ($\ell = 1, \ldots, 4$), medium-lag ($\ell = 5, \ldots, 8$), and long-lag ($\ell = 9, \ldots, 12$) groups. Grouping is less susceptible to inter-lag multicollinearity than the individually-unconstrained alternative. Within each multi-month window, the lagged regressor is the mean of the monthly values over that window (the contemporaneous group is the current month), so each grouped coefficient is the response to a one-degree change in the average temperature over its window. The twelve-month cumulative origin-temperature effect is accordingly the sum of the four window coefficients, $\beta^o_0 + \beta^o_{1-4} + \beta^o_{5-8} + \beta^o_{9-12}$ (the cumulative impulse response). Similarly for precipitation and, in our extension in Section 5.4, conflict risk. We also construct the cumulative effects from 12 unconstrained lags (Figure 2).

Lagged-flow controls. The bilateral flow $\text{num}_{ij,ym}$ combines an intensive margin (how many migrants move through an already-active migration corridor) and an extensive margin (whether the corridor is active at all). Roughly one-third of dyad-months in our panel have $\text{num}_{ij,ym} = 0$ (Table 2). The four-way fixed-effect structure pins down a corridor-year average level of activity and the calendar-month and global timing baselines, but does not absorb within-corridor month-to-month dynamics in either margin. The lagged dependent variable $\ln(\text{num}_{ij,ym-1} + 1)$ enters with coefficient ϕ and identifies intensive-margin persistence on active dyads. The extensive-margin indicator $\mathbf{1}[\text{num}_{ij,ym-1} = 0]$ enters with coefficient ψ and captures the discrete shift for dyads that were dormant in $ym-1$; without it, ϕ would absorb both the smooth persistence on active corridors and the kink at $\text{num}_{ij,ym-1} = 0$. Together the two controls prevent the weather elasticity from absorbing persistence that is itself auto-correlated with monthly weather. The Nickell-type bias on ϕ is of order $1/T$ in our 48-month panel, and any contamination of the weather coefficient through ϕ is bounded by ϕ times the small weather-LDV covariance. Because the lagged dependent variable enters the specification, the cumulative weather effects reported below are conditional on lagged flow and are best read as short-run, conservative estimates of the total dynamic response; we do not rescale them to a long-run elasticity by a $1/(1-\phi)$ multiplier, since the origin-destination-year fixed effects absorb each corridor's annual mean flow and a long-run steady-state level is not identified.

4 Climate and migration: the pooled bilateral elasticity

4.1 The pooled elasticity

We estimate the grouped-lag specification (10) of bilateral migration on origin and destination temperature and precipitation, pooled over the developing country-origin sample.

The upper panel of Figure 2 plots the cumulative effect of origin temperature and precipitation on bilateral flow as the distributed-lag window widens from the contemporaneous month out to twelve months, each lag entered without constraint. The origin-temperature profile descends into negative territory over the first months, settles at its mature value by about the eighth month, and is flat thereafter. The origin-precipitation profile stays close to zero throughout.

Table 3 reports the cumulative grouped-lag coefficients at the 8- and 12-month horizons. The coefficient on origin temperature is negative and statistically significant at both horizons. Read as an elasticity,[11] the cumulative effect of a one-degree origin-temperature anomaly is a reduction in bilateral flow of about 6.3 percent at the eight-month horizon and 5.7 percent at twelve months. The coefficient on origin precipitation is small and not statistically

[11] We use 'elasticity' throughout as shorthand for the response of log bilateral flow to a one-unit shift in a weather variable (one degree Celsius for temperature, 30 mm/month for precipitation). Because the weather variables enter in levels, this is strictly a semi-elasticity. The implied migration-income object of Section 5.3, $\varepsilon_{m,y}$, is by contrast a true (log-log) elasticity.

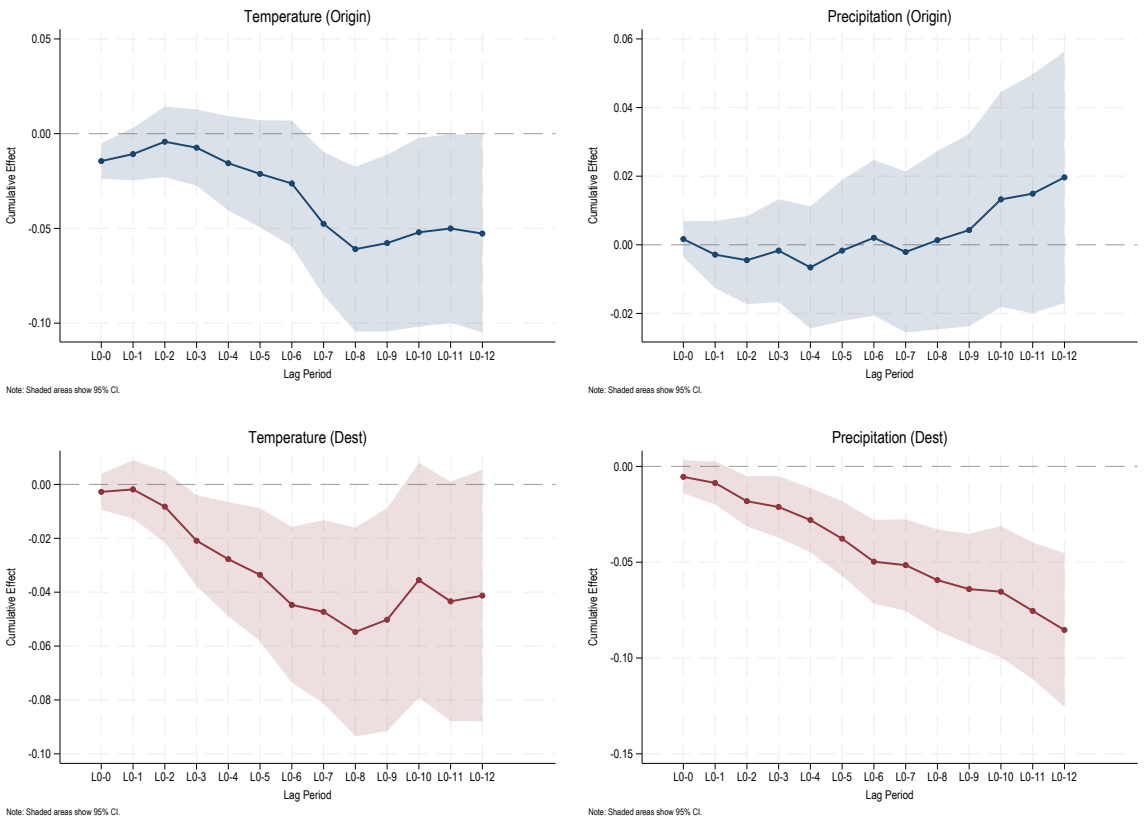

Figure 2: Cumulative individual-lag effects of origin temperature and precipitation on bilateral migration flows. Lags $\ell = 0, \ldots, 12$ are individually unconstrained. Shaded area: 95% CI.

distinguishable from zero at either horizon.

The equivalent plots for destinations in the lower panel of Figure 2 are also downward-sloping. In Table 3, the coefficient on destination temperature is negative and statistically significant, of similar magnitude to the origin one, with a one-degree anomaly lowering inflows by about 6.4 percent after eight months. The coefficient on destination precipitation is negative and statistically significant, implying that a 30 mm/month anomaly lowers inflows by about 6.2 percent at eight months and 8.8 percent at twelve. The precipitation response is thus asymmetric, operating at the receiving end of the migration corridor but not at the sending end; the destination coefficients are consistent with a labour-demand channel in the sectors that hire the marginal migrant, and the eight- to twelve-month accumulation rules out a purely mechanical short-run transport disruption. Our specification identifies the reduced-form destination effect but cannot separate the candidate channels behind it (labour demand, amenity, or administrative processing); doing so would require destination-side labour-market and other data that the present panel lacks.

In the remainder of this paper we focus mainly on origin weather shocks, and especially on origin-temperature anomalies, the variation that carries the income channel at the heart of the constraint mechanism; the origin-precipitation and destination responses are reported throughout but are not our central focus.

Our baseline origin-temperature coefficient is stable under alternative fixed-effect

Table 3: Pooled bilateral elasticity of migration to climate: cumulative grouped-lag effects of origin and destination temperature and precipitation.

Climate grouped-lag	8-month	12-month
Origin temperature	-0.063^{***}	-0.057^{**}
	(0.022)	(0.027)
Origin precipitation	$+0.002$	$+0.019$
	(0.013)	(0.018)
Destination temperature	-0.064^{***}	-0.051^{**}
	(0.020)	(0.023)
Destination precipitation	-0.062^{***}	-0.088^{***}
	(0.013)	(0.020)
Observations	919,532	

Notes: PPML with four-way fixed effects on the developing-origin bilateral panel. Each climate variable enters under the grouped distributed-lag windows $\{0, 1\text{-}4, 5\text{-}8, 9\text{-}12\}$; cells are cumulative effects at the 8- and 12-month horizons (temperature per degree Celsius, precipitation per 30 mm/month). These comprise origin-destination-year, origin-month, destination-month, and year-month effects; standard errors clustered at the origin-destination dyad in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

structures, sample restrictions, climate-aggregation weights, COVID parameterisations, a leave-one-destination-out test, and a temporal-permutation placebo (Section 4.2). Our analysis in Section 4.3 looks beyond the average elasticity, decomposing it by event type and documenting its heterogeneity across origins, and Section 4.4 extends the analysis to pre-existing migrant networks.

4.2 Robustness

4.2.1 Alternative fixed effects and corridor restrictions

Table 4 reports the cumulative 8- and 12-month coefficients on the four weather variables under six alternative specifications. Replacing the additive origin-by-month and destination-by-month fixed effects with a single corridor-by-month fixed effect α_{ijm} in column (2) leaves the origin-temperature 8-month coefficient at -0.066, indistinguishable from the baseline; the result does not merely reflect corridor-specific seasonal patterns that the additive structure fails to absorb. Interacting every climate grouped-lag with the binary flag combining the Gulf, circular, and transit indicators in (3a) yields a clean-corridor (86 percent of all observations) elasticity of -0.075 and a *differential* interaction on the flagged 14 percent of $+0.022$, reported in (3b), that is not statistically significant (the implied flagged-corridor elasticity is -0.053). Dropping the ten dyads with the largest total bilateral flow (Gulf labour pipelines, MEX-USA,

VEN-COL, UKR-POL, among others) in (4) yields an 8-month origin-temperature elasticity of -0.054: modestly attenuated, but still significant. The weather-flow elasticity is therefore not a property of Gulf or circular corridors and is not driven by a small number of large dyads.

Columns (5) and (6) address spatial autocorrelation in weather shocks, which could inflate the precision of the dyad-clustered baseline. Weather systems are spatially extended, so residuals across observations whose origins share weather systems may be correlated even after the four-way fixed-effect structure. Column (5) re-clusters the baseline regression three-ways: on the dyad and on the M49 sub-region $\times$ year-month at both ends of the corridor.[12] Column (6) augments the baseline regression with row-standardised inverse-distance neighbour-weighted climate (within 2,000 km) as additional regressor (a spatial-lag-of-X, or SLX), so the reported coefficient identifies the own-country effect holding the neighbour-spillover constant. At eight months the origin- and destination-temperature elasticities survive both checks at the 5% level, the destination-precipitation channel at the 1% level, and origin precipitation is insignificant throughout. More generally, where the temperature channels weaken across the robustness specifications, it is at the twelve-month horizon; this is expected, since the twelve-month cumulative effect is essentially identical to the eight-month one (the impulse response is flat beyond eight months, Figure 2) but less precisely estimated, the 9–12 month window adding variance without signal.

4.2.2 Alternative weighting schemes

Figure 3 compares the 8- and 12-month cumulative weather shock coefficients under the five weighting schemes of Section 3.2. These weights load on different subsets of the country's spatial distribution: population on the average resident, area weights on the average square kilometre, GDP on the average dollar, low-income on the population weighted by inverse GDP per capita, cropland on the average farmer. The estimates are stable across all five schemes because, despite within-country cell-level population, GDP, and cropland correlating only at around 0.4, weather shocks tend to span the full country territory, so the weighting choice barely moves the country-month aggregate. If the weights matter at all, it is only in the geographically largest developing countries (Brazil, Mexico, Indonesia, and India), where the population, GDP, and cropland weights pick out different regions of a vast territory, across which an origin temperature anomaly need not be uniform.

4.2.3 COVID specification

The 2020–2021 portion of the panel coincides with the steepest peacetime mobility restrictions on record. We test for confounding with the COVID mobility collapse in two ways. First, we

[12] M49 is the United Nations Statistics Division geographic classification (*Standard Country or Area Codes for Statistical Use*, Series M, No. 49), which partitions the world's countries into 22 sub-regions (e.g. Western Africa, South-Eastern Asia, Northern Europe). Re-clustering at the M49 sub-region $\times$ year-month allows residuals to be correlated across all corridors whose origin (or destination) sits in the same sub-region in the same month, the relevant unit of common weather exposure for a synoptic-scale shock.

Table 4: Headline cumulative effects

	(1) Baseline	(2) Corridor-month FE	(3a) Clean main effect	(3b) Differential flagged	(4) Drop top-10 dyads	(5) Three-way cluster (M49)	(6) SLX 2000km dyad cluster
Origin temperature							
8m	-0.063^{***}	-0.066^{***}	-0.075^{***}	$+0.022$	-0.054^{**}	-0.063^{**}	-0.073^{**}
	(0.022)	(0.022)	(0.026)	(0.033)	(0.022)	(0.030)	(0.035)
12m	-0.057^{**}	-0.064^{**}	-0.081^{**}	$+0.045$	-0.046^{*}	-0.057	-0.085^{*}
	(0.027)	(0.026)	(0.032)	(0.047)	(0.026)	(0.036)	(0.045)
Origin precipitation (per 30mm)							
8m	$+0.002$	$+0.002$	$+0.008$	-0.014	-0.016	$+0.002$	$+0.012$
	(0.013)	(0.013)	(0.016)	(0.022)	(0.013)	(0.016)	(0.013)
12m	$+0.019$	$+0.018$	$+0.028$	-0.016	-0.005	$+0.019$	$+0.030$
	(0.018)	(0.018)	(0.023)	(0.033)	(0.019)	(0.021)	(0.018)
Destination temperature							
8m	-0.064^{***}	-0.057^{***}	-0.079^{***}	$+0.031$	-0.042^{**}	-0.064^{**}	-0.063^{***}
	(0.020)	(0.020)	(0.021)	(0.030)	(0.019)	(0.032)	(0.018)
12m	-0.051^{**}	-0.042^{*}	-0.065^{**}	$+0.029$	-0.028	-0.051	-0.052^{**}
	(0.023)	(0.024)	(0.027)	(0.042)	(0.023)	(0.040)	(0.022)
Destination precipitation (per 30mm)							
8m	-0.062^{***}	-0.061^{***}	-0.050^{***}	-0.024	-0.064^{***}	-0.062^{***}	-0.058^{***}
	(0.013)	(0.013)	(0.019)	(0.026)	(0.015)	(0.017)	(0.013)
12m	-0.088^{***}	-0.089^{***}	-0.069^{**}	-0.035	-0.083^{***}	-0.088^{***}	-0.085^{***}
	(0.020)	(0.020)	(0.030)	(0.037)	(0.021)	(0.024)	(0.020)
Observations	$919,532$	$830,990$	$919,532$		$919,062$	$919,532$	$919,532$

Notes: PPML with four-way fixed effects on the developing-origin bilateral panel. Cumulative effects are sums of grouped-lag coefficients. Standard errors in parentheses. Column (1) baseline (dyad cluster); Column (2) replaces $\alpha_{im} + \alpha_{jm}$ with α_{ijm}; Column (3) interacts every climate grouped-lag with the binary corridor-quality flag (Gulf, circular, or transit), column (3a) the main effect (clean corridors), column (3b) the differential on the flagged sub-sample (flagged-corridor elasticity = (3a)+(3b)); Column (4) drops the ten dyads with the largest total bilateral flow. Column (5) re-clusters the baseline three-way at the dyad and at the M49 sub-region × year-month at both ends of the corridor. Column (6) augments the baseline regression with row-standardised inverse-distance neighbour-weighted climate (2000 km cutoff) constructed inline from CEPII bilateral distances; reported coefficients are the own-country effects holding the neighbour spillover constant, dyad-clustered. Origin and destination COVID mobility controls, lagged $\ln$(flow), and lagged zero-flow indicator included throughout. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

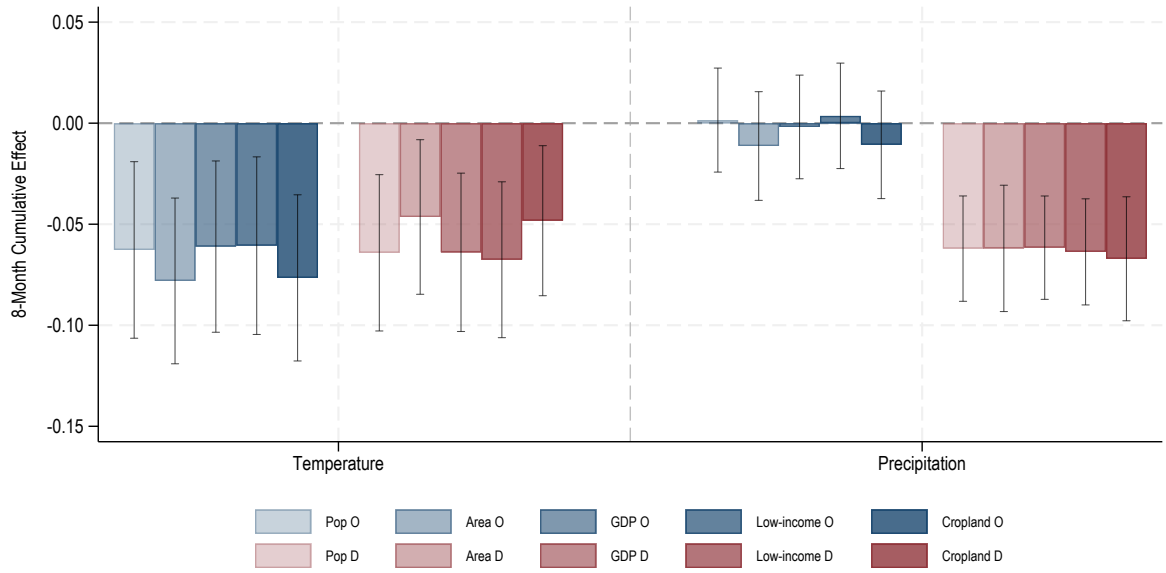

(a) 8-month cumulative effect

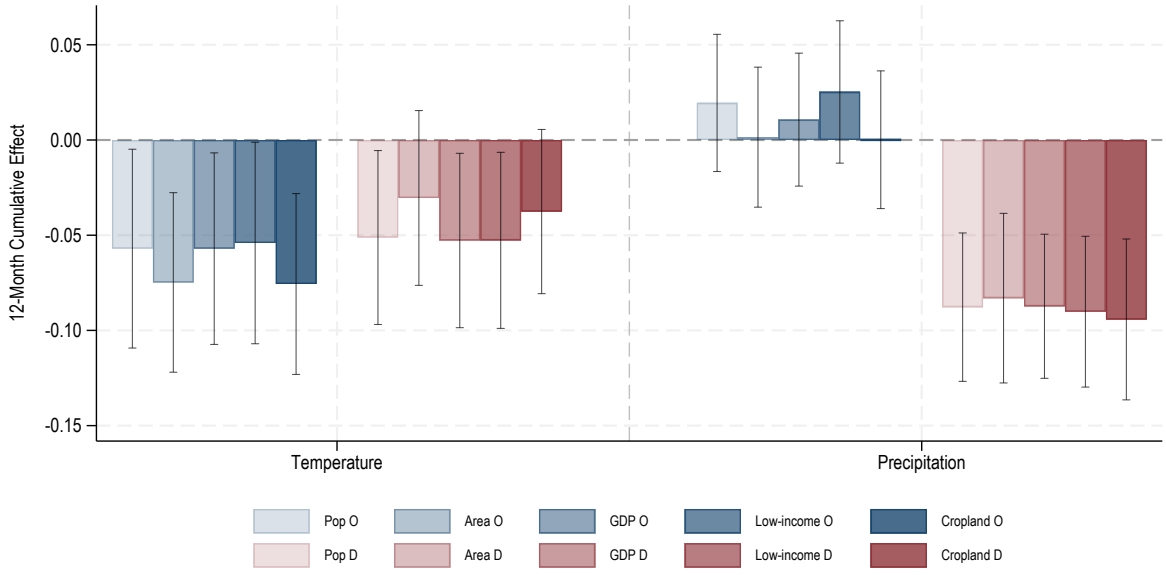

(b) 12-month cumulative effect

Figure 3: Climate effects under five country-level weighting schemes (area, population, GDP, low-income, cropland) at the 8-month (top) and 12-month (bottom) cumulative horizon. The low-income weight is pop^2/GDP from BFI 2015, equivalent to population weighted by inverse GDP per capita. The GDP and low-income weights share a common sample (BFI 2015 coverage); the cropland-weighted regression drops country-months with zero cropland coverage (very high-latitude or hyper-arid months), which slightly reduces its sample relative to the area- and population-weighted regressions. Navy bars for origin, maroon for destination estimates. Error bars show 95 percent confidence intervals.

vary the functional form of the COVID controls: the baseline contemporaneous-and-lagged linear form is replaced by year-varying linear interactions and by a quadratic in the mobility index. Second, we drop the COVID controls entirely and rely on the four-way fixed-effect structure alone. Figure 4 reports the 8- and 12-month cumulative weather coefficients across all four specifications. The estimates are visually indistinguishable: the baseline negative origin-temperature elasticity, the destination-temperature effect, and the destination-precipitation effect retain their sign, magnitude, and significance in every parameterisation.

A stricter test interacts every weather grouped lag with a binary indicator for the COVID-peak window of March 2020 through December 2021. Table 5 reports the cumulative 8- and 12-month effects in three columns: the non-peak weather–migration elasticity (the base coefficients), the differential peak effect (the interaction coefficient, whose standard error gives the formal difference test), and the implied peak elasticity (their sum). The non-peak elasticity is, in every case, at least as large in absolute value as the baseline pooled estimate. The interactions are systematically opposite-signed to the base coefficients at the eight-month horizon and statistically significant: the COVID peak attenuates the origin-temperature mobility trap effect, the destination-precipitation deterrent, and the origin-precipitation channel by roughly thirty to forty percent of the non-peak coefficient. At the twelve-month horizon the temperature interactions are essentially zero, indicating that the long-horizon weather response is identical in peak and non-peak periods.

4.2.4 Leave-one-destination-out specifications

We re-estimate the baseline twenty times, each time dropping one of the twenty destinations with the largest cumulative inflows from developing origins. The left panel of Figure 5 reports the 8-month origin-temperature elasticity in each leave-one-out specification. The estimates cluster tightly around the full-sample point, ranging over $[-0.053, -0.069]$, with every 95 percent confidence interval comfortably excluding zero. Three exceptions (USA, ARE, POL) arise at the 12-month horizon, where the cumulative effect is less precisely estimated, the 9–12 month window adding variance without signal (as noted in Section 4.2.1).

4.2.5 Temporal permutation: a placebo test of causality

We implement a permutation exercise as a placebo test of causality: it asks whether the estimated weather-migration relationship requires the actual timing of weather shocks, or whether a similar coefficient would arise from randomly reshuffled weather that preserves only each origin's marginal climate distribution. We randomly reassign each origin's observed monthly temperature and precipitation across months, independently for each origin, preserving each origin's own climate distribution but scrambling the calendar timing, which destroys the alignment between a weather anomaly and the bilateral flow in the same month. The reshuffling severs the hypothesised causal channel while leaving each origin's marginal climate distribution intact. If the estimated effect reflects the causal effect of the timing of

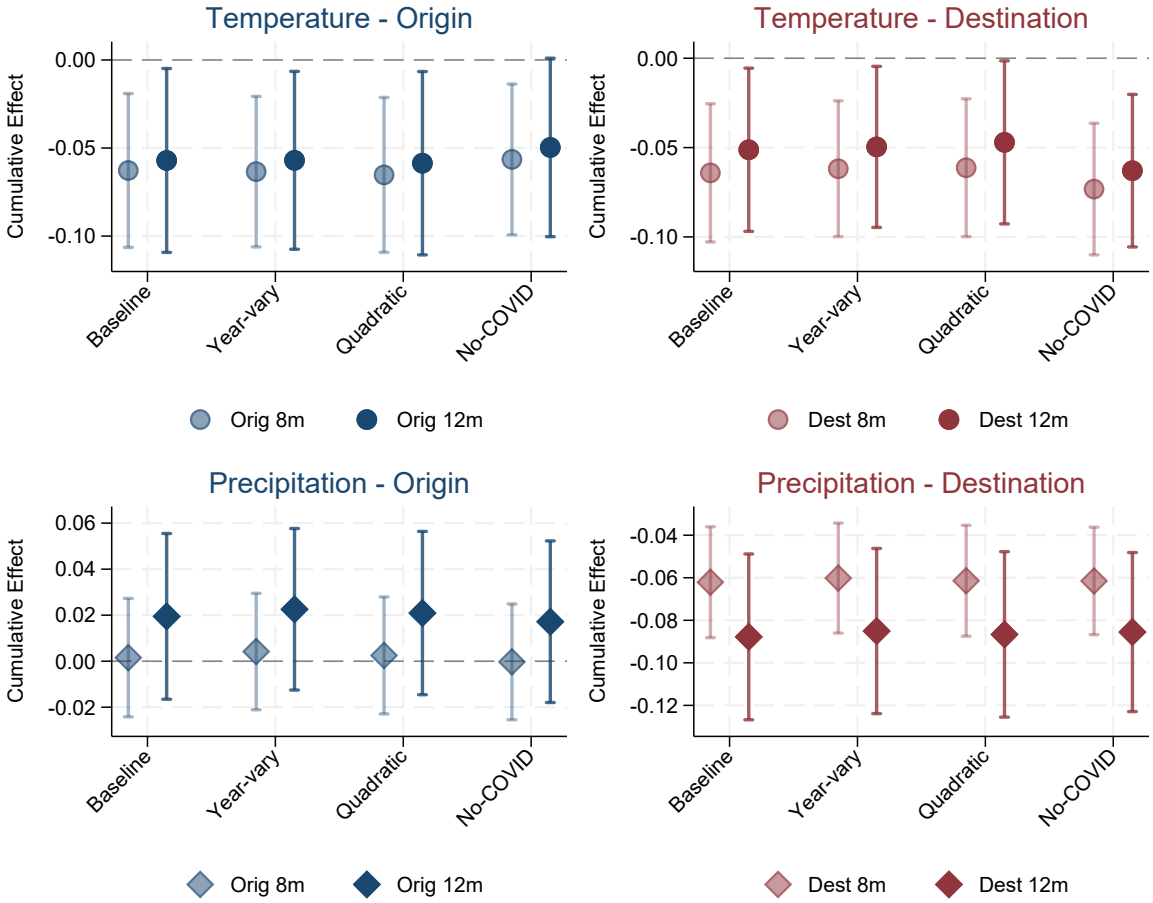

Figure 4: Robustness to COVID specification: 8-month and 12-month cumulative weather effects under alternative COVID parameterisations. Navy bars for origin, maroon for destination estimates. Light markers – 8m, dark markers – 12m estimates. Error bars show 95 percent confidence intervals.

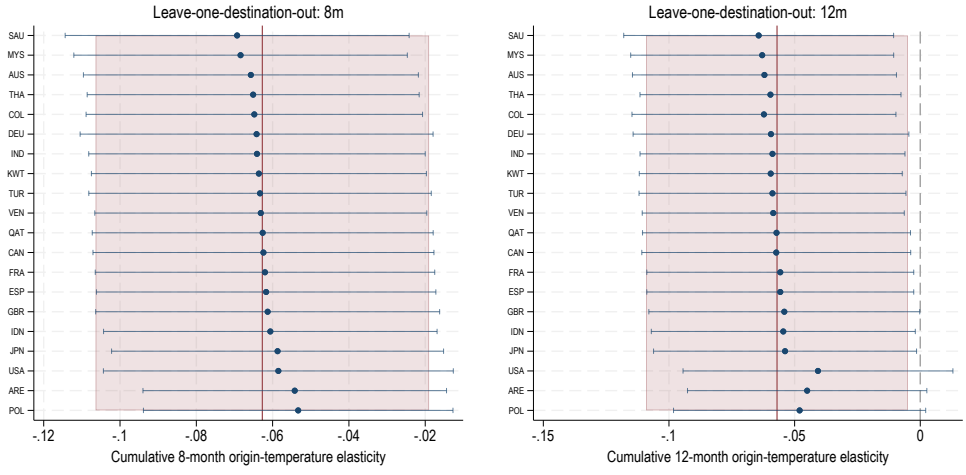

Figure 5: Leave-one-destination-out robustness check, 8-month (left plot) and 12-month (right plot) cumulative origin-temperature elasticity. Each row drops one of the top-20 destinations by total flows and re-estimates equation (10) on the remainder. Error bars show 95 percent confidence intervals. Vertical reference line: full-sample point estimates with respective confidence interval shaded red.

Table 5: Weather–migration elasticity in the COVID-peak window: interaction specification.

	Non-peak base coef.	Interaction peak − non-peak	Peak base + interaction
Temperature (°C)			
Origin, 8m	-0.081^{***}	$+0.024^{***}$	-0.057^{***}
	(0.022)	(0.006)	(0.020)
Origin, 12m	-0.080^{***}	-0.000	-0.080^{***}
	(0.025)	(0.003)	(0.025)
Dest., 8m	-0.056^{***}	-0.012^{***}	-0.068^{***}
	(0.020)	(0.004)	(0.020)
Dest., 12m	-0.047^{**}	$+0.001$	-0.047^{**}
	(0.023)	(0.003)	(0.024)
Precipitation (per 30 mm/month)			
Origin, 8m	$+0.016$	-0.030^{***}	-0.014
	(0.013)	(0.008)	(0.014)
Origin, 12m	$+0.043^{**}$	-0.035^{***}	$+0.008$
	(0.019)	(0.011)	(0.018)
Dest., 8m	-0.075^{***}	$+0.030^{***}$	-0.045^{***}
	(0.014)	(0.010)	(0.014)
Dest., 12m	-0.103^{***}	$+0.037^{***}$	-0.066^{***}
	(0.021)	(0.012)	(0.020)
Observations		$919{,}532$	

Notes: PPML with four-way fixed effects on the developing-origin bilateral panel. Every weather grouped lag is interacted with a binary indicator equal to one for March 2020 through December 2021. The Non-peak column reports the cumulative 8- and 12-month effects on the base coefficients (the elasticity that prevails outside the COVID-peak window); the Interaction column reports the differential effect during the peak (the cumulative sum of the climate-by-peak interaction terms), with the standard error in parentheses providing the formal difference test; the Peak column is the sum of base and interaction, i.e. the implied elasticity during the peak window. Specification follows the headline (COVID mobility controls, lagged dependent variable, and extensive-margin indicator), with standard errors clustered at the origin-destination dyad. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

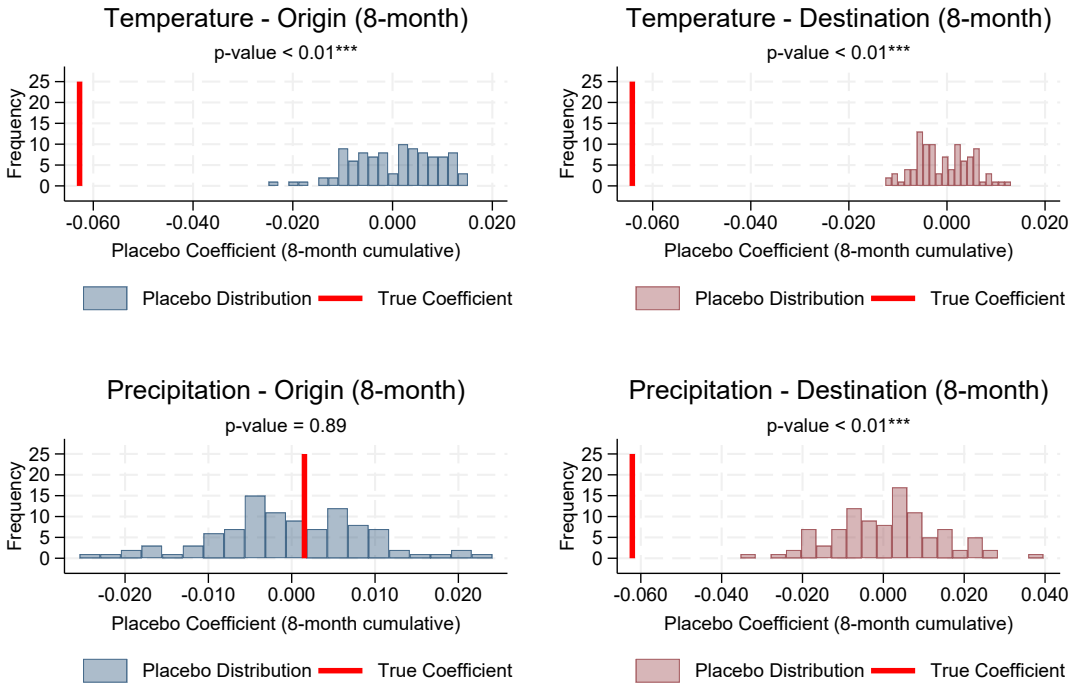

(a) 8-month cumulative effect

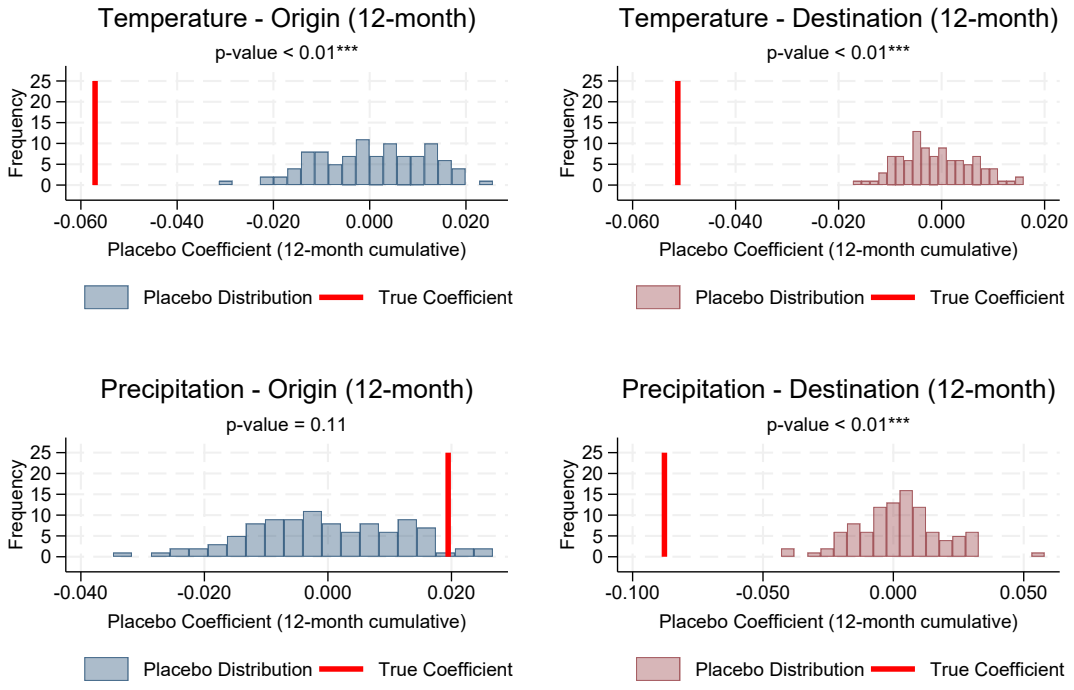

(b) 12-month cumulative effect

Figure 6: Temporal permutation distribution of 8-month (top) and 12-month (bottom) cumulative origin and destination temperature and precipitation effects. Vertical reference lines: observed point estimates. Distributions from 100 temporally shuffled placebo panels. Navy bars for origin, maroon for destination placebo estimates.

weather shocks on migration, scrambling that timing should drive the coefficient toward zero; if it were instead a spurious feature of the panel's structure, it would survive the reshuffling.

We draw 100 placebo panels under this reshuffle, the months of each origin permuted independently, and repeat the exercise separately at the destination side; on each panel we re-estimate (10) and store the cumulative 8- and 12-month coefficients.

Figure 6 plots the placebo distributions. The observed origin- and destination-temperature coefficient estimates (marked by red vertical lines) at both horizons, and destination precipitation, fall well outside their placebo distributions: the effects are present only when the true climate–flow timing is preserved and vanish once it is scrambled, which is the pattern expected of a causal response, and not of a timing-independent artifact such as seasonality or the panel's finite length. Origin precipitation lies inside its placebo distribution, mirroring the baseline origin-precipitation null finding. Spatial autocorrelation is addressed separately, and conservatively, by the sub-region re-clustering and the inverse-distance spatial-lag control of Section 4.2.1. Procedural details are provided in Appendix A.2.

4.2.6 Census-microdata corroboration

We use IPUMS International census extracts of almost 900 million person records to build annual bilateral arrival flows for 1990–2019 (Appendix A.6 details the construction). A bilateral arrival is a person born in origin o, enumerated in destination d at census c, who reports having immigrated in calendar year y. The panel covers 177 origin and 25 destination countries (the immigration societies of the Americas plus Spain, Greece, Switzerland, South Africa, Kenya and Malaysia), drawing on 18 distinct census years. Annual frequency is what makes the data usable: it matches the calendar timing of arrivals to the calendar year of the weather signal, whereas at the five-year or decadal frequency of widely used bilateral flow series (e.g., Abel & Cohen 2019), transitory temperature anomalies become hard to separate from slower-moving unobserved trends.

The exercise is an external-validity check on our bilateral results from the monthly data. Spanning thirty years and 177 origins, with none of the Meta panel's ingredients, it shows whether our findings reflect a genuine annual migration response or merely a short-run (quasi-monthly) feature of the 2019–2022 sample window. Nevertheless, two limitations should be borne in mind. First, identification is coarser than the monthly four-way design of our main analysis, since origin-year confounders are absorbed only up to regional-by-year fixed effects along with the origin's neighbours' weather. Second, destination coverage is partial: the 25 destinations received 29 percent of developing-origin migration flows over 2000–2020 in the Abel & Cohen (2019) reconstruction, with the Gulf states and many large European destinations not in the panel since they do not field the year-of-arrival question.

Table 6 reports estimates for developing origins, based on PPML regressions (with standard errors clustered at the dyad-level), controlling for log origin population, log destination population, and log destination per capita GDP. A one-degree origin-temperature anomaly in the arrival year reduces arrivals by about 10 percent (column 1); the cumulative

effect over the arrival year and the year before is about 20 percent. Destination temperature is also strongly negative and significant, -0.19. Origin precipitation has no effect; destination precipitation has the same negative sign as in our main analysis but is insignificant in this sample of only 25 destinations. The patterns described survive World Bank region×year fixed effects (column 2) and the inclusion of origin-neighbour weather as a spatial-lag control (column 3), with both neighbour terms zero. Adding log origin per capita GDP as a mediator (column 4) leaves the temperature coefficient virtually unchanged.

Table 6: Census-based annual bilateral arrival flows and origin weather, 1990–2019.

	(1) Baseline	(2) WB region ×year	(3) SLX	(4) Origin income
Origin temperature	-0.103^{***}	-0.077^{**}	-0.118^{***}	-0.096^{***}
	(0.025)	(0.035)	(0.046)	(0.024)
Origin precipitation	-0.024	-0.016	-0.022	-0.022
	(0.022)	(0.021)	(0.023)	(0.022)
Destination temperature	-0.191^{**}	-0.223^{***}	-0.192^{**}	-0.192^{**}
	(0.082)	(0.056)	(0.082)	(0.079)
Destination precipitation	-0.083	-0.027	-0.082	-0.082
	(0.062)	(0.058)	(0.064)	(0.062)
Neighbour temperature			0.025	
			(0.057)	
Neighbour precipitation			-0.007	
			(0.027)	
ln origin GDP per capita				-0.177
				(0.108)
Origin temperature, cumulative $(y,\ y-1)$	-0.205^{***}	-0.182^{**}	-0.187^{**}	-0.195^{***}
	(0.050)	(0.072)	(0.090)	(0.049)
Destination temperature, cumulative $(y,\ y-1)$	-0.112	-0.156^{***}	-0.102	-0.117
	(0.085)	(0.054)	(0.085)	(0.079)
Time fixed effects	Year	WB reg.×yr	Year	Year
Observations	11,473	11,473	11,473	11,473

Notes: PPML on the zero-filled census arrival panel, developing origins. The dependent variable is the weighted count of persons born in origin o, enumerated in destination d at census c, who immigrated in year y (within ten years of the census, $y \geq 1990$); corridor-census-year cells with no sampled arrival are coded zero. All columns include origin×destination×census fixed effects, the time effects indicated in the panel, and control for log origin population, log destination population, and log destination GDP per capita (PWT 11.0, annual); the sample is held fixed across columns. Weather is the population-weighted CRU TS country mean in the arrival year (precipitation per 30 mm/month); cumulative rows enter y and $y-1$ jointly and report their sum. Column (3) adds row-standardised inverse-distance neighbour weather of the origin within 2,000 km; column (4) adds log origin GDP per capita, which absorbs part of the income channel through which weather operates. Standard errors clustered at the origin-destination dyad. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The cumulative destination-temperature effect is smaller than the contemporaneous one because the lagged destination term enters with the opposite sign (about $+0.07$ across the four columns, against -0.19 for the contemporaneous term). One reading would be intertemporal substitution at the destination: a hot destination year may delay rather than cancel completed arrivals, since visa processing, job searches and housing can displace planned moves into the following year. The origin side shows no such reversal, with effects in year y and $y-1$ both negative at about -0.10 each.

4.3 Decomposition and heterogeneity

Two exercises look beyond the average elasticity: a decomposition of the response by weather event type, and a simple split of the origin countries by income and by baseline temperature.

4.3.1 Decomposition into compound extreme weather events

The linear temperature and precipitation specification rolls together the response to qualitatively different weather phenomena. We decompose the weather–migration elasticity into eight mutually exclusive compound climate-event categories: four single-event isolates (heat-only, cold-only, drought-only, flood-only) and four compound events (hot+dry, hot+wet, cold+dry, cold+wet), entered at origin and destination on the same grouped distributed-lag structure as the linear baseline. Event categories are constructed from the standardised climate indices of the World Meteorological Organization (WMO) at the base layer, with detailed definitions, thresholds, and aggregation reported in Appendix A.7. The unit of each event is the share of country population exposed in the country-month, rescaled to $[0,1]$ so that each reported coefficient is the implied log-flow response to a counterfactual full-country exposure, cumulated over the indicated horizon. Appendix Figure A3 shows how often each event occurs and how much of the population it exposes when it does: heat-only events dominate both margins, whereas cold-related events are rare.

Tables 7 (base layer, $|\text{index}| \geq 1.0$) and 8 (severe layer, $|\text{index}| \geq 1.5$, conditions outside the 7th/93rd percentile of the long-run climatological distribution) report the cumulative 8- and 12-month effects of the eight compound categories.[13] We focus on the coefficients that retain statistical significance at conventional levels in both tables, on the rationale that a substantive climate channel should appear at both the base *and* the rarer severe layer.

At origins, four patterns recur. *Heat-only* exposure is the cleanest version of the migration trap result: the coefficient is negative and significant at both horizons in both tables. The average base-layer heat-only event exposes about 39 percent of the population, implying a

[13] The event-decomposition sample is slightly smaller than the baseline one ($n = 916{,}210$ rather than $919{,}532$, a loss of about 0.4 percent). The Standardised Precipitation-Evapotranspiration Index (SPEI) used to flag drought, flood, and the wet/dry compounds requires monthly potential evapotranspiration, which CRU TS does not produce for a handful of remote micro-state cells (Bermuda, Maldives, Marshall Islands, Nauru, Tuvalu, and seven other small island territories), even though it does produce temperature and precipitation there. The linear baseline of Section 4.1 only uses temperature and precipitation, so it is unaffected; the event decomposition drops corridors involving these countries as destinations.

migration flow reduction of 7 to 10 percent; the average severe heat-only event exposes about 35 percent and implies a 5 to 7 percent reduction. The base-layer magnitude matches the linear baseline, now isolated to the one event type where the income channel operates without amenity confounds. *Compound hot+dry and hot+wet* exposure carry the opposite sign at the twelve-month horizon, significantly raising origin outflows in both tables. The average base-layer hot+dry event exposes about 23 percent of the population and implies an increase in out-migration of 4.5 percent; the average severe hot+dry event exposes only 13 percent, but the per-100%-exposure coefficient is roughly three times larger, so the implied rise in migration climbs to 7.8 percent. Hot+wet is similar in magnitude. These patterns can be interpreted as the 'corner' result formalised in Appendix A.1: when a stressor erodes location amenity faster than it tightens the financing constraint, distress migration rises despite the income loss; the effect sharpens at the severe threshold, where amenity damage compounds. *Cold+wet* exposure at origins carries a robustly negative twelve-month coefficient, extending the migration trap interpretation to the cold tail. The average base-layer cold+wet event exposes about 11 percent of the population, implying a flow reduction of about 8 percent; at the severe layer, exposure falls to 8 percent but the coefficient is over twice as large, implying a 15 percent reduction at the rare severe cold+wet event.

At destinations, only the positive cold-only coefficient survives both thresholds, and we read it descriptively rather than causally: the destinations cold enough to trigger the indicator are predominantly high-income Northern countries where the cold-only flag is unlikely to capture the relevant climate-economic stress. The remaining destination compound coefficients are statistically significant at the base layer but not at the severe layer, suggesting they capture moderate-event labour-market frictions too fragile to support a causal interpretation; the linear destination-precipitation channel of Section 4.1 subsumes the base-layer flood signal with greater precision.

4.3.2 Origin heterogeneity: income and baseline-temperature splits

We probe the heterogeneity of the origin-temperature elasticity along two origin characteristics. In each case the developing-country origins are split at the median of the characteristic, origin temperature is interacted with an above-median dummy in the bilateral regression of Section 4.1, and destination climate variables enter in the same way as before. Figures 7 and 8 plot the cumulative 8- and 12-month origin-temperature elasticity for each half, with 95 percent confidence intervals.

Figure 7 splits origins at the median of per capita GDP in 2018 (expenditure-side chained PPPs in 2021 value US dollars, PWT 11.0; Feenstra et al. 2015). The coefficient on origin temperature is negative in both halves and similar in magnitude (above-median -0.082 at 8m and -0.067 at 12m; below-median -0.067 and -0.082); the formal test of the interaction term does not reject equality across the two halves at either horizon. The elasticity therefore does not vary detectably with origin income, and the pooled estimate is representative across the developing-origin income distribution.

Table 7: Decomposition of the bilateral weather–migration elasticity into compound extreme events: base layer threshold.

	Origin		Destination	
	8m	12m	8m	12m
Single-event categories				
Heat-only	-0.189^{***}	-0.252^{***}	-0.104	-0.040
	(0.061)	(0.079)	(0.086)	(0.107)
Cold-only	$+0.082$	$+0.129$	$+0.446^{***}$	$+0.638^{**}$
	(0.198)	(0.235)	(0.165)	(0.262)
Drought-only	$+0.024$	-0.293^{**}	-0.157	-0.135
	(0.108)	(0.148)	(0.099)	(0.144)
Flood-only	$+0.102$	$+0.107$	-0.212^{**}	-0.291^{**}
	(0.098)	(0.140)	(0.094)	(0.130)
Compound events				
Hot + Dry	$+0.109$	$+0.194^{*}$	$+0.191^{***}$	$+0.256^{***}$
	(0.095)	(0.111)	(0.073)	(0.087)
Hot + Wet	$+0.138$	$+0.291^{**}$	-0.304^{**}	-0.292^{*}
	(0.115)	(0.148)	(0.119)	(0.152)
Cold + Dry	-0.550	-0.113	-0.854^{*}	-0.750
	(0.462)	(0.514)	(0.479)	(0.573)
Cold + Wet	-0.324	-0.698^{**}	-0.037	-0.233
	(0.264)	(0.298)	(0.236)	(0.292)
Observations	$916{,}210$			

Notes: PPML with four-way fixed effects on the developing-origin bilateral panel, 2019–2022. Each row reports the cumulative 8- and 12-month coefficients on the grouped distributed lags $\{0, 1\text{-}4, 5\text{-}8, 9\text{-}12\}$ of a compound climate-event category at origin (columns 2–3) and destination (columns 4–5). Events are defined at the base layer of the WMO threshold convention (conditions beyond the 16th/84th percentile of the long-run standardised index, SPEI or STI; see Appendix A.7). The eight event categories are mutually exclusive: heat-only isolates a heatwave that does not coincide with concurrent drought or flood; cold-only is the analogous cold-wave isolate; drought-only and flood-only isolate water-stress shocks that do not coincide with concurrent heat or cold; and hot+dry, hot+wet, cold+dry, and cold+wet are the compound counterparts. The unit of each event is the share of country population exposed in the country-month, scaled to $[0, 1]$, so each reported coefficient is the implied response of log bilateral flow to a full-country exposure to the event in question, cumulated over the indicated horizon. Specification follows the headline (COVID mobility controls, lagged dependent variable, and extensive-margin indicator), with standard errors clustered at the origin-destination dyad. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Decomposition of the bilateral weather–migration elasticity into compound extreme events: severe-layer threshold.

	Origin		Destination	
	8m	12m	8m	12m
Single-event categories				
Heat-only	-0.151^{**}	-0.201^{**}	-0.057	$+0.001$
	(0.063)	(0.088)	(0.078)	(0.098)
Cold-only	$+0.172$	$+0.391$	$+0.951^{***}$	$+1.134^{***}$
	(0.393)	(0.467)	(0.273)	(0.387)
Drought-only	-0.031	-0.256	$+0.132$	$+0.280$
	(0.132)	(0.174)	(0.135)	(0.189)
Flood-only	$+0.204^{*}$	$+0.270^{*}$	-0.045	-0.093
	(0.117)	(0.159)	(0.118)	(0.165)
Compound events				
Hot + Dry	$+0.342^{**}$	$+0.599^{***}$	$+0.196$	$+0.091$
	(0.150)	(0.192)	(0.141)	(0.191)
Hot + Wet	$+0.119$	$+0.626^{*}$	$+0.196$	$+0.381$
	(0.258)	(0.343)	(0.306)	(0.417)
Cold + Dry	$+1.638$	-0.113	-4.065^{**}	-1.562
	(2.003)	(3.028)	(1.845)	(2.753)
Cold + Wet	-0.564	-1.864^{***}	-0.528	-0.593
	(0.404)	(0.683)	(0.593)	(0.760)
Observations	$916{,}210$			

Notes: Same specification and sample as Table 7; the only difference is the event-construction threshold, which is here raised from the WMO base (events outside the 16th/84th percentile of the long-run climatological distribution, $|\text{index}| \geq 1.0$) to the WMO severe layer (7th/93rd percentile, $|\text{index}| \geq 1.5$). Each reported coefficient is the implied response of log bilateral flow to a counterfactual full-country exposure to the event in question, cumulated over the indicated horizon. Severe-layer events are rarer in the developing-origin sample (in particular, severe cold-dry and severe cold-wet exposure occur in fewer than 2% of origin-months), so the cold-compound coefficients in particular are estimated with much wider confidence intervals than their base-layer counterparts and should be interpreted with caution. Specification follows the headline (COVID mobility controls, lagged dependent variable, and extensive-margin indicator), with standard errors clustered at the origin-destination dyad. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 8 instead splits origins at the median of baseline temperature, defined as each origin's average temperature over 1990–2018, separating hot origins from cold(er) ones using a fixed window prior to the sample period.[14] The coefficient on origin temperature is small and statistically indistinguishable from zero in the cooler half (-0.025 at 8m, -0.011 at 12m), and negative, large, and statistically significant in the hotter half (-0.150 at 8m, -0.164 at 12m); the formal test rejects the null at the 5 percent level for both horizons. The temperature split displays an asymmetry, and the negative origin-temperature response is concentrated in already-hot origins. The pattern is consistent with the concave temperature-income damage function in the macro-climate literature, in which $\varepsilon_{y,T}$ steepens at higher baseline temperatures (Burke et al. 2015, Desbordes & Eberhardt 2024).

We apply the same two exercises to origin precipitation, which displays no comparable heterogeneity: in both splits it is close to zero in each half, and the two halves never differ significantly (income differential $+0.02$, $p = 0.58$; baseline-temperature differential -0.05, $p = 0.21$, at twelve months). The heterogeneity is thus specific to the temperature channel that carries the income response.

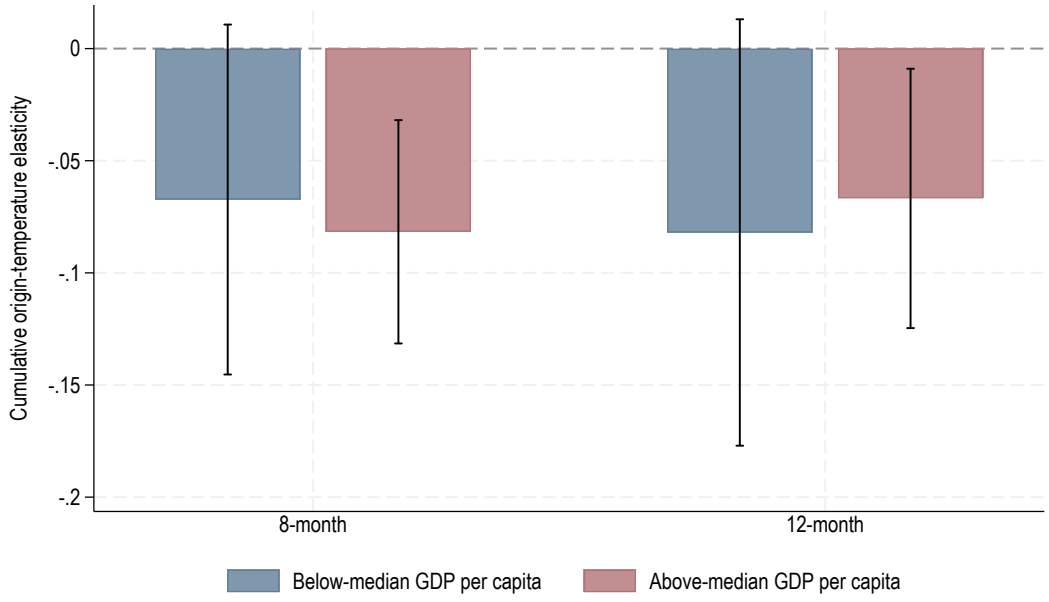

Figure 7: Origin-temperature elasticity of bilateral migration by origin income. Developing origins split at the median of 2018 GDP per capita; cumulative 8- and 12-month origin-temperature elasticity for each half, with 95 percent confidence intervals.

4.4 Extension: corridor heterogeneity and the budget constraint

The bilateral structure of the panel allows us to test how the climate response varies along migration corridor characteristics. Two of these carry competing predictions for the budget-

[14] The split is hence on a predetermined climatology rather than the in-sample mean, but the two are interchangeable here: across the developing origins the 1990–2018 and 2019–2022 average temperatures correlate at 0.999, and only three of the origins change side of the median between them.

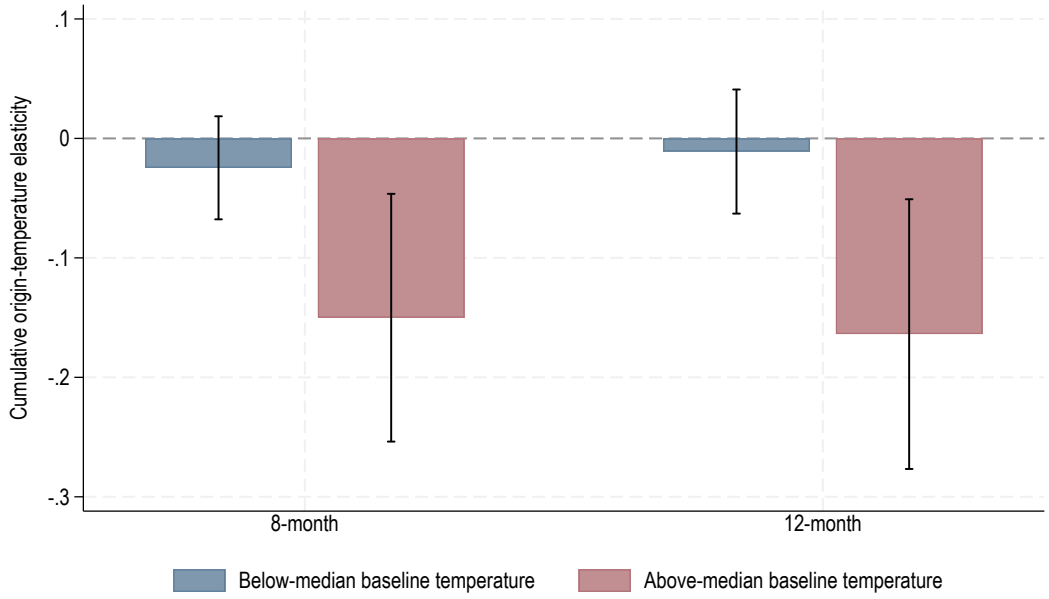

Figure 8: Origin-temperature elasticity of bilateral migration by baseline temperature. Developing origins split at the median of baseline origin temperature, each origin's average temperature over 1990–2018 (a window predetermined relative to the 2019–2022 estimation period); cumulative 8- and 12-month origin-temperature elasticity for each half, with 95 percent confidence intervals.

constraint mechanism. First, a larger pre-existing migrant network lowers the cost of moving and, on specific United States-bound corridors, can turn a weather shock into a positive emigration push (Mahajan & Yang 2020, Ibáñez et al. 2026). Second, a wealthier destination, by contrast, raises the cost of reaching it: visas, travel, and the financial reserves richer-country immigration regimes increasingly require, all weigh more heavily on a household whose income has just been cut by a weather shock. We test both predictions by interacting the origin-temperature grouped lags with the 2015 bilateral migrant stock (in logs) from United Nations Department of Economic and Social Affairs, Population Division (2020) and with the 2018 destination per capita GDP (in logs) from PWT 11.0 (Feenstra et al. 2015), alongside standard gravity controls from CEPII. The stock variable is constructed as $\ln(\text{stock} + 1)$ so that corridors with zero recorded 2015 stock are not dropped; this matters, because about half of the corridors with strictly positive 2019–2022 migration flow have no recorded 2015 bilateral stock in the UN DESA data. Each moderator is time-invariant within a corridor and absorbed by the origin-destination-year fixed effect, so only its interaction with the time-varying weather shock is identified. The migrant stock is centred at the median over corridors with strictly positive 2015 stock (about 500 bilateral migrants), and the other moderators at the median over corridors with strictly positive 2019–2022 flows. The main origin-temperature coefficient therefore reads as the climate response at a typical active corridor with median diaspora, distance, and other characteristics, directly comparable to the baseline elasticity of Section 4.1.

For a more reliable migration stock specification we drop corridors whose 2015 UN DESA bilateral figure is least informative. Two grounds apply. The first is a reporting revision: the corridor's 2015 reported stock fell to zero from a positive 2010 value, a pattern occurring in eight destinations (Egypt, Greece, Iceland, Latvia, Slovenia, Sierra Leone, Estonia, and Bulgaria) that appear to have stopped reporting bilateral detail in the 2015 wave. The second is a conflict-affected origin, one of eleven countries (Syria, Somalia, South Sudan, Sudan, the Central African Republic, DR Congo, Eritrea, Yemen, Afghanistan, Myanmar, and Iraq) where bilateral migrant-stock figures are known to misclassify refugee populations.

Table 9 reports the result. In isolation, the stock interaction is negative and marginally significant in the full sample in column (1), and more precisely estimated in the narrower sample in column (3): the climate response is stronger in high-network corridors than in thinly-connected ones, the opposite of network-driven amplification reported in the above-cited literature on the US as destination. The pattern flips once the destination-income interaction enters along with standard gravity controls. The stock interaction loses significance in columns (2) and (4), while the destination-income interaction is negative and statistically significant at both horizons in both samples. The climate response is stronger in corridors to higher-income destinations, the direct prediction of the budget-constraint mechanism: reaching a wealthier destination costs more, so a climate-driven income loss binds harder on those routes. The stock interaction in columns (1) and (3) was picking up this destination-income gradient rather than a direct network amplification, the 2015 bilateral stock and destination GDP per capita being highly correlated across developing-country corridors (richer destinations have attracted larger diasporas).

Among the other five gravity moderators, only common official language carries a statistically significant interaction (positive, of the opposite sign to the destination-income one). The pattern is consistent with shared language modestly easing arrival at the desirability margin, partly offsetting the budget constraint where the latter binds. Distance, contiguity, common religion, and colonial ties show no significant interaction effects.

Table 9: Horse race: migrant-stock interaction versus standard gravity interactions.

	(1)	(2)	(3)	(4)
Sample	Full developing-origin		Excl. unreliable stock	
Gravity interactions	No	Yes	No	Yes
Cumulative 8-month effect				
Origin temperature (main)	-0.0505^{***}	-0.0487^{*}	-0.0656^{***}	-0.0646^{**}
	(0.0191)	(0.0251)	(0.0200)	(0.0256)
Interactions with origin-destination characteristics:				
ln(migrant stock 2015)	-0.0038^{*}	-0.0015	-0.0051^{**}	-0.0023
	(0.0023)	(0.0030)	(0.0024)	(0.0031)
ln(distance)		0.0073		0.0049
		(0.0171)		(0.0175)
common official language		0.0565^{*}		0.0709^{**}
		(0.0312)		(0.0326)
contiguity		0.0227		0.0179
		(0.0565)		(0.0606)
common religion		-0.0179		-0.0318
		(0.0585)		(0.0636)
colonial dependence ever		-0.0202		-0.0421
		(0.0512)		(0.0448)
ln(destination GDP per capita)		-0.0443^{**}		-0.0497^{**}
		(0.0187)		(0.0193)
Cumulative 12-month effect				
Origin temperature (main)	-0.0397^{*}	-0.0425	-0.0559^{**}	-0.0618^{*}
	(0.0217)	(0.0318)	(0.0228)	(0.0327)
Interactions with origin-destination characteristics:				
ln(migrant stock 2015)	-0.0055	-0.0025	-0.0072^{**}	-0.0037
	(0.0034)	(0.0043)	(0.0035)	(0.0044)
ln(distance)		0.0003		-0.0034
		(0.0243)		(0.0250)
common official language		0.0804^{*}		0.1038^{**}
		(0.0445)		(0.0468)
contiguity		0.0204		0.0155
		(0.0770)		(0.0824)
common religion		-0.0233		-0.0424
		(0.0804)		(0.0877)
colonial dependence ever		-0.0615		-0.0951
		(0.0708)		(0.0635)
ln(destination GDP per capita)		-0.0586^{**}		-0.0654^{**}
		(0.0262)		(0.0271)
Observations	$919,532$	$813,666$	$896,300$	$794,241$

Notes: Each cell reports the cumulative interaction effect of origin temperature on log bilateral flow at the indicated horizon, multiplied by the indicated origin-destination characteristic. The migrant-stock interaction is centred at the median of corridors with strictly positive 2015 stock; the other characteristics are centred at the median over corridors with strictly positive bilateral flow. The main origin-temperature coefficient therefore reads as the response at the representative active corridor. Columns (1) and (3) enter the migrant stock interaction only; columns (2) and (4) add interactions with the other corridor characteristics and drop corridors for which CEPII reports no common-religion value or PWT reports no destination GDP per capita. Columns (1) and (2) use the full developing-origin sample; columns (3) and (4) exclude corridors whose 2015 UN DESA bilateral stock is least reliable. Standard errors are clustered at the origin-destination dyad. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5 The constraint mechanism: income and a non-climate shock

Section 4 delivered Step 1 of our empirical strategy: we demonstrated that the pooled bilateral elasticity is negative, so the financing margin dominates in the typical developing-origin corridor, with the income channel in control. We now execute Steps 2 and 3. Step 2 tests the income channel on two income panel datasets for the same origin countries at different temporal and spatial resolutions: an annual cell-level GDP per capita panel from Rossi-Hansberg & Zhang (2025) (Section 5.1) and a country-month nighttime-lights panel we build from the NASA Black Marble VNP46A2 daily product (Section 5.2). The two panels measure the same income response at the annual within-cell and the monthly within-country-year margins. Section 5.3 combines the cell-level elasticity with the bilateral elasticity through equation (13) to recover an implied migration-income elasticity. Step 3, in Section 5.4, shows that another, non-climate, negative origin income shock (internal conflict) similarly exerts a negative effect on bilateral migration.

5.1 The income response at the grid level

The constraint mechanism rests on a premise the bilateral panel cannot test directly: temperature anomalies at the origin reduce household income. We now test this on a sub-national grid of per capita GDP built for the same developing-country origins.

The grid data come from the BFI Data Studio sub-national income panel (Rossi-Hansberg & Zhang 2025), which assigns annual per capita GDP to roughly 73,000 0.5° land cells through a random-forest predictor trained on county- and state-level GDP for the subset of countries that report such data.[15] Cell-level estimates are produced as shares of country GDP and converted to constant 2021 PPP US$. No climate variable enters the predictor, so the cell-level GDP series is exogenous to the within-cell temperature variation that identifies our regression. The series covers 2012–2022. We use the full BFI window rather than restricting to the 2019–2022 bilateral-flow sample because the cell and country-year fixed effects identify the temperature-migration elasticity, β in equation (11) below, from within-cell year-to-year temperature variation; truncating to four years would sharply reduce that variation and the precision of β without changing what is identified. We pair the cells with CRU TS v4.09 (Harris et al. 2020) monthly temperature and precipitation aggregated to annual means on the same 0.5° grid. Restricting to cells in developing countries and dropping cells with the synthetic floor-coded value the BFI allocator uses for non-positive interpolations yields a panel of roughly 460,000 cell-years across 133 developing countries (the gridded panel is defined on all developing countries, a slightly broader set than the 127 bilateral-migration origins).

[15] The predictor inputs are unsaturated nighttime-light radiance from the VIIRS Black Marble product (building on Henderson et al. 2012), LandScan gridded population, MODIS land-use shares, MODIS net primary productivity, EDGAR sectoral CO_2 emissions, terrain ruggedness, and national per capita GDP.

The grid-level income regression is

$$\ln \text{GDPpc}_{i,t} \;=\; \beta\, \text{TMP}_{i,t} \;+\; \gamma\, \text{PRE}_{i,t} \;+\; \alpha_i \;+\; \alpha_{c(i),t} \;+\; \varepsilon_{i,t}, \tag{11}$$

where i indexes 0.5° grid cells, $c(i)$ is the country containing cell i, α_i is a cell fixed effect, and $\alpha_{c(i),t}$ is a country-year fixed effect that absorbs every shock common to all cells in a given country-year (e.g., national macroeconomic conditions, COVID). The coefficient β is identified from within-cell year-to-year temperature variation around the cell's climatological mean, net of country-year common shocks. Standard errors are clustered at the country level. We estimate this equation with three within-cell weighting schemes that are the cell-level counterparts of three of the five country-level weights of Section 3.2: population-weighted using 2015 cell population (the baseline, matching the country-level dyadic regression), area-weighted (the cosine of latitude, proportional to physical cell area), and cropland-weighted using GAEZ+ 2015 cropland intensity.[16] Because each weight is time-invariant (area) or a fixed 2015 cross-sectional snapshot (population, cropland), it does not co-move with the within-cell year-to-year weather variation that identifies β, so it stays exogenous to the shock whether or not 2015 falls inside the estimation window. The population-weighted estimate gives the elasticity for the cells where people live; the area-weighted estimate for the average unit of land; and the cropland-weighted estimate for the cells where farming takes place. The population-weighted estimate enters the decomposition (13) in Section 5.3, so that both sides load on the same underlying population.

Figure 9 reports the cell-level elasticities. The coefficient on temperature is negative in all three weighting schemes, significant at the 5 percent level for the population- and cropland-weighted estimates and at the 10 percent level for the area-weighted estimate. Read as an elasticity, a one-degree annual temperature anomaly lowers cell-level per capita GDP by about 2.2 percent where people live (population-weighted), by 1.3 percent for the average unit of land (area-weighted), and by 1.7 percent where cropland is concentrated (cropland-weighted): the income loss is largest in the populated cells, the economically most relevant weight. The coefficient on precipitation is not statistically distinguishable from zero in any scheme and its sign varies, mirroring the origin-precipitation null effect in the bilateral migration panel. The income channel is therefore present and concentrated where households live; the population-weighted elasticity (-0.022) is the magnitude carried into the decomposition in Section 5.3.

[16] These are cell-level counterparts of the dyadic weights, not literal reproductions of its aggregation formula. The fractional-area factor $A_{c,i}$ that the dyadic construction applies to each cell weight apportions a border cell between countries when cell climate is averaged up to a country-level exposure; it does not arise here, because each $0.5°$ cell is assigned whole to one country and the regression is run at the cell level rather than aggregated. Population and cropland enter as the cell's own counts, which already record how much of each is present and so require no area apportionment.

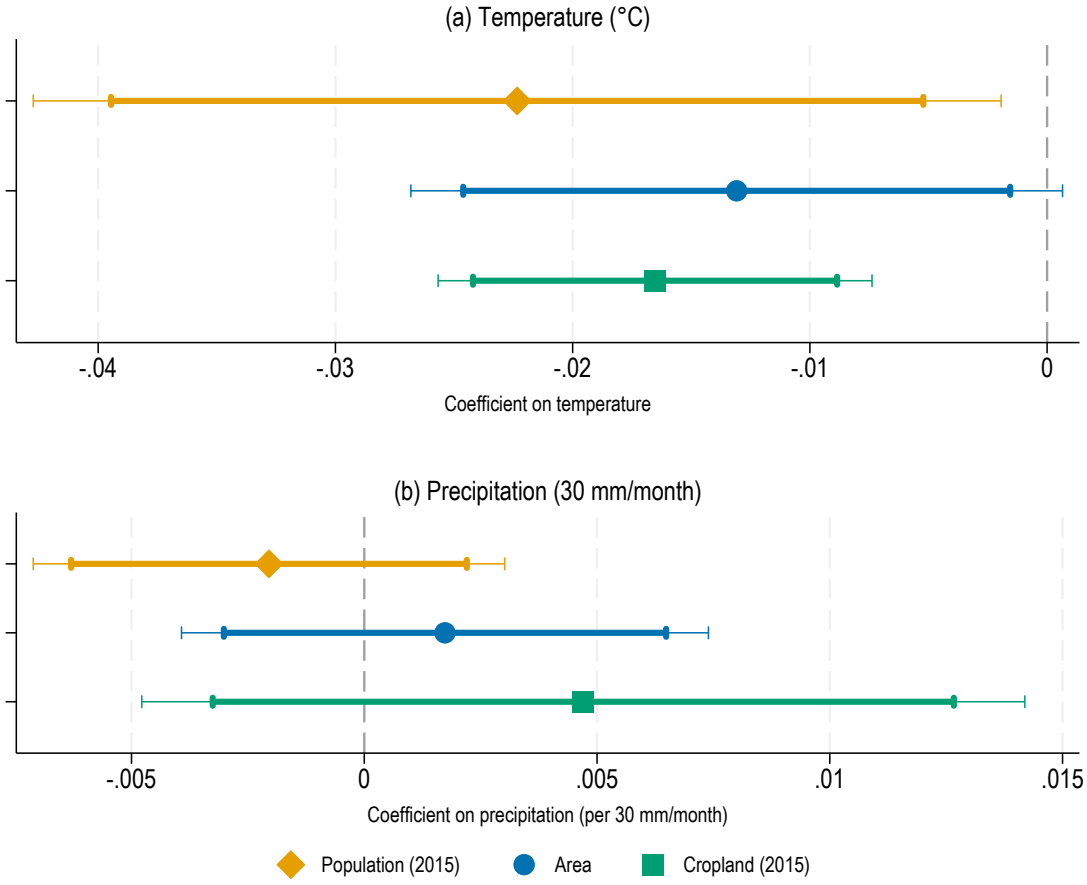

Figure 9: Cell-level climate elasticities of GDP per capita under three weighting schemes, developing-country cells, 2012–2022. Top panel: temperature (per degree Celsius); bottom panel: precipitation (per 30 mm/month of annual mean, the same rescaling as in the bilateral panel). Cell and country-year fixed effects; standard errors clustered at the country level. Thick (thin) error bars show 95 (90) percent confidence intervals.

5.2 Country-month cross-check from satellite nighttime lights

The BFI panel is annual, whereas the bilateral migration panel is monthly. To test whether the activity response operates at the monthly frequency that identifies the bilateral regression, we build a country-month proxy for aggregate economic activity from satellite nighttime luminosity data. The panel covers the same 127 developing-country origins as the bilateral migration sample over April 2012 – December 2024 ($n \approx 18{,}241$ country-months, depending on the dependent variable definition). As with the gridded GDP panel, we keep the full 2012–2024 window rather than restricting to the 2019–2022 flow years: the country-by-year and year-month fixed effects identify the temperature response from within-country-year monthly variation, so truncating to the flow window would cut the panel to roughly a third of its size and reduce precision without changing what is identified. The proxy is the population-weighted sum of country-level radiance from NASA's Black Marble VNP46A2 daily product (Román et al. 2018), after masking for quality, snow, water, the VIIRS noise floor, and detected gas-

flaring sources (see Appendix A.5). We further decompose total brightness (sum radiance) into an *intensive* margin (mean radiance per valid pixel) and an *extensive* margin (lit area in km^2), each estimated on the same weather grouped-lags terms and fixed effects.

The specification mirrors the bilateral migration regression:

$$\ln \mathrm{NTL}_{c,ym} \;=\; \sum_{\ell=0}^{12} \left(\beta_\ell \, \mathrm{TMP}_{c,ym-\ell} + \gamma_\ell \, \mathrm{PRE}_{c,ym-\ell} \right) \;+\; \alpha_{ct} \;+\; \alpha_{ym} \;+\; \varepsilon_{c,ym}, \tag{12}$$

where $\ln \mathrm{NTL}_{c,ym}$ is one of the three log brightness measures (total radiance, mean radiance per valid pixel, or lit area) for country c in year-month ym; the temperature and precipitation lags are grouped into the windows $\{0,\ 1\text{-}4,\ 5\text{-}8,\ 9\text{-}12\}$ of equation (10); α_{ct} is a country-by-year fixed effect that absorbs such potential confounders as national macroeconomic conditions and the COVID shock, and α_{ym} a year-month fixed effect that absorbs the common global month, so β_ℓ is identified from within-country-year monthly weather anomalies. Country-months are weighted by their *coverage share* (the fraction of land pixels that passed every data-quality exclusion and had at least three valid daily retrievals in the month), so observations contribute in proportion to how much of the country the satellite 'saw', and standard errors are clustered by country. Cumulative effects of a sustained $1\,^\circ$C temperature shock or a sustained $+30$ mm/month precipitation shock are obtained by summing the relevant lag-window coefficients.

Table 10: Climate shocks and country-month nighttime lights.

	(1) ln(sum radiance)	(2) ln(mean rad/valid)	(3) ln(lit area)
Temperature (°C)			
8m cumulative	-0.035^{***}	-0.032^{**}	-0.027^{**}
	(0.012)	(0.013)	(0.013)
12m cumulative	-0.053^{***}	-0.040^{***}	-0.042^{***}
	(0.015)	(0.015)	(0.016)
Precipitation (per 30 mm/month)			
8m cumulative	$+0.028^{**}$	$+0.035^{***}$	-0.026^{**}
	(0.012)	(0.012)	(0.010)
12m cumulative	$+0.020$	$+0.034^{*}$	-0.060^{***}
	(0.018)	(0.017)	(0.014)

Notes: OLS on the country-month nighttime-lights panel restricted to the 127 developing-country origins of the bilateral migration sample (April 2012 to December 2024, $n \approx 18{,}241$ depending on the dependent variable), with observations weighted by the country-month coverage share. Column (1) is total country brightness, column (2) the intensive margin (mean radiance per valid pixel), and column (3) the extensive margin (lit area). Each climate variable enters under the grouped distributed-lag windows $\{0,\ 1\text{-}4,\ 5\text{-}8,\ 9\text{-}12\}$; cells are cumulative effects at the 8- and 12-month horizons. Country-by-year and year-month fixed effects; cluster-robust standard errors at the country level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10 reports the cumulative 8- and 12-month effects for the three dependent variables. The coefficient on temperature is negative and statistically significant for total country brightness at both horizons: a $1\,^{\circ}$C anomaly lowers brightness by about 3.5% at 8 months and 5.3% at 12 months, and the intensive and extensive margins are similarly negative and significant. The activity response works through both margins on a timing that matches the bilateral regression. The precipitation response is margin-dependent: a $+30$ mm/month shock reduces lit area (the extensive margin) by about 2.6% at 8 months and 6.0% at 12 months, but raises mean radiance per valid pixel and, at the 8-month horizon, total brightness. Precipitation therefore has no single-signed effect on aggregate activity, whereas the temperature response is uniformly negative; on temperature the NTL panel returns the sign of the cell-year BFI panel at the monthly frequency on which the bilateral regression identifies. The mixed precipitation response is itself informative: with no consistent effect on aggregate origin activity, precipitation offers no sign-coherent income channel through which to move migration. This matches the above panel evidence, where origin precipitation is insignificant in both the cell-level grid (Figure 9) and the bilateral migration panel. The income mechanism thus predicts the origin-precipitation null effect, not only the origin-temperature effect.

5.3 Implied income-migration elasticity

If the weather shock operates mainly through the income channel ($\varepsilon_{A,T} \approx 0$), the total bilateral elasticity factors as

$$\varepsilon_{m,T} \;=\; \frac{d\ln m}{d\delta} \;=\; \underbrace{\frac{\partial \ln m}{\partial \ln y}}_{\varepsilon_{m,y}} \cdot \underbrace{\frac{d\ln y}{d\delta}}_{\varepsilon_{y,T}} \;=\; \varepsilon_{y,T} \times \varepsilon_{m,y}, \tag{13}$$

which identifies $\varepsilon_{m,y}$ as the ratio of the bilateral and income elasticities on the same origins, without instrumenting income with climate. The BFI panel is annual, so the apples-to-apples horizon for the bilateral coefficient is the 12-month cumulative effect. A one-degree temperature anomaly reduces population-weighted cell-level log per capita GDP by 0.022 log-points and reduces bilateral flow by 0.057 log-points at 12 months. The ratio is

$$\varepsilon_{m,y} \;=\; \frac{-0.057}{-0.022} \;\approx\; +2.6.$$

Each one-percent climate-driven decline in origin per capita GDP is associated with roughly a 2.6 percent decline in bilateral migration flow. Using the cropland-weighted income estimate yields $\varepsilon_{m,y} \approx +3.5$; the area-weighted income estimate yields $\varepsilon_{m,y} \approx +4.4$ (all ratios computed from unrounded coefficients). All three are positive and large, consistent with a developing-country origin sample in which the migration-cost constraint binds. The factorisation assumes the shock works through income; if the amenity channel is also operative ($\varepsilon_{A,T} < 0$), it raises outflows and partly offsets the income-driven contraction, so the recovered $\varepsilon_{m,y}$ is a lower

bound on the income-only elasticity.[17]

Table 11: Conflict risk and bilateral migration.

	(1)	(2)	(3)	(4)
	Risk only		Risk + climate	
	8m	12m	8m	12m
Conflict-risk indicators				
Origin internal conflict	-0.444^{**}	-0.828^{***}	-0.343^{*}	-0.700^{***}
	(0.217)	(0.283)	(0.196)	(0.258)
Destination internal conflict	$+0.032$	$+0.408$	$+0.026$	$+0.398$
	(0.234)	(0.300)	(0.236)	(0.303)
Bilateral interstate conflict	-0.202	$+0.114$	-0.351	-0.059
	(0.661)	(1.036)	(0.633)	(1.002)
Climate grouped-lags				
Origin temperature			-0.060^{***}	-0.052^{**}
			(0.021)	(0.024)
Origin precipitation			$+0.000$	$+0.016$
			(0.013)	(0.018)
Destination temperature			-0.058^{***}	-0.047^{**}
			(0.019)	(0.023)
Destination precipitation			-0.062^{***}	-0.085^{***}
			(0.013)	(0.020)
Observations		897,484		

Notes: PPML with four-way fixed effects on the developing-origin bilateral panel. Each conflict-risk indicator (calibrated monthly conflict risk in $[0, 1]$, GDELT-based) enters under the same grouped distributed-lag windows as the climate variables, $\{0,$ 1-4, 5-8, 9-12$\}$; cells are cumulative effects. The risk-only columns include the three conflict indicators; the risk-plus-climate columns add the four climate grouped-lags. A non-climate origin shock reproduces the negative bilateral elasticity, and the headline climate coefficients are unchanged when conflict risk is included. Four-way fixed effects (origin-destination-year, origin-month, destination-month, year-month); standard errors clustered at the origin-destination dyad in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.4 A non-climate origin shock

The constraint mechanism is not specific to climate: any destructive origin shock that erodes income and the capacity to finance a move should contract cross-border migration. We test this with the calibrated monthly conflict-risk indicator of Aizenman et al. (2026), built from GDELT event data and scaled to $[0, 1]$, available at the country-month level (internal conflict

[17] Estimated symmetrically, with a lagged dependent variable on both the bilateral-flow and the cell-income sides, the implied elasticity is $\varepsilon_{m,y} \approx 3.0$, close to the value above: flow and income are comparably persistent, so the short-run versus long-run rescaling largely cancels in the ratio.

risk) and the directed country-pair-month level (interstate risk) over 2015–2026.[18] We enter the indicator in the bilateral regression under the same grouped distributed-lag structure as the weather variables: windows $\{0,\ 1\text{-}4,\ 5\text{-}8,\ 9\text{-}12\}$, the four-way fixed-effect structure, and dyad-clustered standard errors.

Table 11 reports the cumulative grouped-lag effects of the three conflict-risk indicators on bilateral migration, with and without the weather grouped-lags. The coefficient on origin internal conflict risk is negative and statistically significant: a cumulative -0.44 at the 8-month horizon and -0.83 at 12 months in the specifications with only conflict indicators, and -0.34 and -0.70 once the weather grouped-lags are added. Both horizons are statistically significant. Because the indicator runs on $[0, 1]$, the twelve-month coefficient implies that moving an origin from median internal-conflict risk to the top percentile contracts bilateral migration flows by roughly one third. A non-climate origin shock that raises political tensions thus also produces a negative bilateral migration elasticity. The baseline weather coefficients are unchanged by the inclusion of conflict risk (origin temperature -0.060 at eight months, destination temperature -0.058, origin precipitation insignificant, destination precipitation -0.062): the weather response is not proxying for conflict.

A binding-constraint reading of these results carries a further implication: this non-climate shock should erode origin income, in the same way that a temperature anomaly does. We test this mechanism by applying the country-month nighttime-lights income specification of Section 5.2 with the origin conflict-risk index, entered as grouped distributed lags, in place of weather. A rise in origin conflict risk lowers origin brightness: the eight-month cumulative effect is negative and statistically significant, and stays negative when the weather grouped-lags are included (Appendix Table A1). Conflict thus reproduces both halves of the mechanism, depressing origin economic activity as well as contracting bilateral migration flows. Because conflict timing is less plausibly exogenous than weather, we read the conflict effects on migration and income as corroborative rather than as clean causal estimates.

6 Conclusion

Climate change and the accompanying threat to livelihoods for populations endowed with only fragile resilience to economic shocks are widely predicted to dramatically increase the number of 'climate refugees' over the coming decades. In the developing economies that would likely send most of the world's potential climate migrants, we find that the cross-border response to climate stress runs *opposite* to the climate-refugee narrative: a weather shock in the form of a higher-temperature anomaly reduces international migration rather than raising it, while precipitation anomalies have no significant effect.

This pattern is compatible with a resource-constrained immobility trap, and in line with the

[18] The Global Database of Events, Language, and Tone (GDELT) is a machine-coded catalogue of politically relevant events extracted in near real time from worldwide news media. Aizenman et al. (2026) aggregate its event counts into a calibrated monthly index of conflict risk bounded in $[0, 1]$.

migration-cost constraint long documented for income shocks in the migration-development literature. Our results bear directly on how climate-driven migration is projected. Forecasts that multiply a global population at climate risk by a positive emigration elasticity overstate cross-border displacement for precisely the origins where the constraint binds. The populations most exposed to climate damage are typically also those least able to finance an international move, so in the near term climate stress is more likely to dampen orderly cross-border emigration from poor origins than to generate a surge of climate refugees in distant, rich countries. Whether those who stay move within their own country instead is a margin our cross-border data and analysis cannot speak to.

We carefully investigate and confirm that the response we uncover is representative and not tied to our modelling choices, data source, or sample period. The elasticity does not vary across the income distribution of developing origins, and it reappears, with the same negative sign and the same null on precipitation, in independent census microdata spanning three decades. We further establish that the financing constraint is not specific to climate: a non-climate livelihood shock — a rise in conflict risk — reproduces the same contraction, so the binding force is the financing of the move rather than temperature shocks as such.

This negative response is not confined to small or moderate weather shocks. Our panel contains severe and extreme events, and the compound-event decomposition examines these directly. The constraint dominates across the bulk of that variation: even severe heat-only exposure, where the income channel operates without amenity confounds, reduces outflows. The response turns positive only at the rare hot-and-dry and hot-and-wet compound extremes, where the erosion of location amenity outpaces the tightening of the financing constraint and distress migration rises. That rare extreme weather phenomena (single-event *and* compound-event shocks) will become 'less rare' in the future is one of the predictions common across many climate models. For the time being, however, the negative elasticity is the rule across the empirically relevant range. On this evidence the 'climate refugee' is a special case compared to the more common mechanism: hardship tends to immobilise poor populations rather than to move them.

These findings are troubling on both economic and ethical grounds. For many developing origins the principal welfare benefit of emigration accrues not to the migrant but to those who stay behind, through remittances (di Giovanni et al. 2015). Weather-induced immobility is therefore doubly costly: the income shortfall that stops a household financing departure also forecloses the remittances on which the origin economy relies. Cross-border movement is itself an important means of adaptation, so a constraint that tightens as conditions worsen withholds that channel exactly where it is needed most, and from the populations least responsible for the global warming that drives it. Their immobility is imposed by the constraint rather than chosen, the defining feature of a trapped population.

References

Abel, G. J. & Cohen, J. E. (2019), 'Bilateral international migration flow estimates for 200 countries', Scientific Data **6**(1), 82.

Abel, G. J. & Cohen, J. E. (2022), 'Bilateral international migration flow estimates updated and refined by sex', Scientific Data **9**(1), 173.

Aizenman, J., Desbordes, R. & Saadaoui, J. (2026), Bilateral conflict risk and trade: Military wars, trade wars, and diplomatic noise, NBER Working Paper 35077, National Bureau of Economic Research.

Beine, M. & Jeusette, L. (2021), 'A meta-analysis of the literature on climate change and migration', Journal of Demographic Economics **87**(3), 293–344.

Beine, M. & Parsons, C. (2015), 'Climatic factors as determinants of international migration', Scandinavian Journal of Economics **117**(2), 723–767.

Benveniste, H., Oppenheimer, M. & Fleurbaey, M. (2022), 'Climate change increases resource-constrained international immobility', Nature Climate Change **12**(7), 634–641.

Black, R., Arnell, N. W., Adger, W. N., Thomas, D. & Geddes, A. (2013), 'Migration, immobility and displacement outcomes following extreme events', Environmental Science & Policy **27**, S32–S43.

Brown, O. (2008), Migration and climate change, Technical Report 31, International Organization for Migration, Geneva.

Burke, M., Hsiang, S. M. & Miguel, E. (2015), 'Global non-linear effect of temperature on economic production', Nature **527**(7577), 235–239.

Cattaneo, C., Beine, M., Fröhlich, C. J., Kniveton, D., Martinez-Zarzoso, I., Mastrorillo, M., Millock, K., Piguet, E. & Schraven, B. (2019), 'Human migration in the era of climate change', Review of Environmental Economics and Policy **13**(2), 189–206.

Cattaneo, C. & Peri, G. (2016), 'The migration response to increasing temperatures', Journal of Development Economics **122**, 127–146.

Center for International Earth Science Information Network (CIESIN), Columbia University (2018), 'Gridded population of the world, version 4 (GPWv4): Population count, revision 11'.

Chi, G. et al. (2025a), 'A global dataset of monthly bilateral international migration flows', Humanitarian Data Exchange.
URL: *https://data.humdata.org/dataset/international-migration-flows*

Chi, G. et al. (2025b), 'Measuring global migration flows using online data', Proceedings of the National Academy of Sciences **122**(18), e2409418122.

Clemens, M. A. (2014), Does development reduce migration?, in R. E. B. Lucas, ed., 'International Handbook on Migration and Economic Development', Edward Elgar, Cheltenham, UK, chapter 6, pp. 152–185.

Correia, S., Guimarães, P. & Zylkin, T. (2020), 'Fast poisson estimation with high-dimensional fixed effects', Stata Journal **20**(1), 95–115.

Dell, M., Jones, B. F. & Olken, B. A. (2012), 'Temperature shocks and economic growth: Evidence from the last half century', American Economic Journal: Macroeconomics **4**(3), 66–95.

Dell, M., Jones, B. F. & Olken, B. A. (2014), 'What do we learn from the weather? the new climate-economy literature', Journal of Economic Literature **52**(3), 740–798.

Desbordes, R. & Eberhardt, M. (2024), 'Climate change and economic prosperity: Evidence from a flexible damage function', Journal of Environmental Economics and Management **125**, 102974.

di Giovanni, J., Levchenko, A. A. & Ortega, F. (2015), 'A global view of cross-border migration', Journal of the European Economic Association **13**(1), 168–202.

Elvidge, C. D., Zhizhin, M., Hsu, F.-C. & Baugh, K. E. (2013), 'Viirs nightfire: Satellite pyrometry at night', Remote Sensing **5**(9), 4423–4449.

Feenstra, R. C., Inklaar, R. & Timmer, M. P. (2015), 'The next generation of the Penn World Table', American Economic Review **105**(10), 3150–3182. Data: PWT 11.0 (2024 update), https://www.rug.nl/ggdc/productivity/pwt/.

Grogan, D., Frolking, S., Wisser, D., Prusevich, A. & Glidden, S. (2021), 'GAEZ+_2015 monthly cropland'.

Gröschl, J. & Steinwachs, T. (2017), 'Do natural hazards cause international migration?', CESifo Economic Studies **63**(4), 445–480.

Harris, I., Osborn, T. J., Jones, P. & Lister, D. (2020), 'Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset', Scientific Data **7**(109), 1–18.

Henderson, J. V., Storeygard, A. & Weil, D. N. (2012), 'Measuring economic growth from outer space', American Economic Review **102**(2), 994–1028.

Hoffmann, R., Dimitrova, A., Muttarak, R., Crespo Cuaresma, J. & Peisker, J. (2020), 'A meta-analysis of country-level studies on environmental change and migration', Nature Climate Change **10**(10), 904–912.

Hosking, J. R. M. (1990), 'L-moments: Analysis and estimation of distributions using linear combinations of order statistics', Journal of the Royal Statistical Society: Series B (Methodological) **52**(1), 105–124.

Ibáñez, A. M., Quigua, J., Romero, M. J. & Velásquez, A. (2026), 'Responses to extreme temperatures: Migrant networks and international migration from el salvador', American Economic Journal: Economic Policy **18**(2), 212–241.

Institute for Economics and Peace (2020), Ecological threat register 2020: Understanding ecological threats, resilience and peace, Technical report, Institute for Economics and Peace, Sydney.
URL: *http://visionofhumanity.org/reports*

IPCC (2022), Climate change 2022: Impacts, adaptation and vulnerability. contribution of working group ii to the sixth assessment report of the intergovernmental panel on climate change, Technical report, Cambridge University Press.

Kalkuhl, M. & Wenz, L. (2020), 'The impact of climate conditions on economic production: Evidence from a global panel of regions', Journal of Environmental Economics and Management **103**, 102360.

Mahajan, P. & Yang, D. (2020), 'Taken by storm: Hurricanes, migrant networks, and US immigration', American Economic Journal: Applied Economics **12**(2), 250–277.

Marchiori, L., Maystadt, J.-F. & Schumacher, I. (2012), 'The impact of weather anomalies on migration in sub-saharan africa', Journal of Environmental Economics and Management **63**(3), 355–374.

Missirian, A. & Schlenker, W. (2017), 'Asylum applications respond to temperature fluctuations', Science **358**(6370), 1610–1614.

Myers, N. (1995), Environmental exodus: An emergent crisis in the global arena, Climate Institute, Washington, DC.

Özden, C., Parsons, C. R., Schiff, M. & Walmsley, T. L. (2011), 'Where on earth is everybody? the evolution of global bilateral migration 1960–2000', The World Bank Economic Review **25**(1), 12–56.

Pekel, J.-F., Cottam, A., Gorelick, N. & Belward, A. S. (2016), 'High-resolution mapping of global surface water and its long-term changes', Nature **540**(7633), 418–422.

Román, M. O., Wang, Z., Sun, Q. et al. (2018), 'NASA's Black Marble nighttime lights product suite', Remote Sensing of Environment **210**, 113–143.

Rossi-Hansberg, E. & Zhang, J. (2025), Local GDP estimates around the world, BFI Working Paper 2025-17, Becker Friedman Institute for Economics, University of Chicago.

Santos Silva, J. M. C. & Tenreyro, S. (2006), 'The log of gravity', Review of Economics and Statistics **88**(4), 641–658.

United Nations Department of Economic and Social Affairs, Population Division (2020), 'International migrant stock 2020', United Nations database, POP/DB/MIG/Stock/Rev.2020. Bilateral origin–destination migrant-stock matrix.

Vicente-Serrano, S. M., Beguería, S. & López-Moreno, J. I. (2010), 'A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index', Journal of Climate **23**(7), 1696–1718.

World Meteorological Organization (2012), Standardized precipitation index user guide, Technical Report WMO-No. 1090, World Meteorological Organization, Geneva.

Zhizhin, M., Matveev, A., Ghosh, T., Hsu, F.-C., Howells, M. & Elvidge, C. (2021), 'Measuring gas flaring in Russia with multispectral VIIRS nightfire', Remote Sensing **13**(16), 3078.

A Online Appendix (not intended for publication)

A.1 Costly-financing extension

The model of Section 2 imposes the financing requirement (3) as a hard constraint: a household whose disposable income falls short of the bilateral cost cannot move at all. This keeps the two margins clean but rules out, by construction, the case in which a constrained household relaxes the requirement itself through informal borrowing, distress asset sales, or irregular routes. We show here that softening the constraint into a costly-financing technology nests the base model, leaves the empirical content of Section 2 unchanged, and opens a second route by which an acute amenity shock can overturn the sign of the migration response.

Replace the hard constraint with a penalty on the financing shortfall. Define the shortfall

$$s_{ij,t} \;=\; \max\{0,\; c_{ij} - (y_{i,t} - \underline{c})\}, \tag{A1}$$

the amount by which the cost c_{ij} exceeds disposable income. A household with $s_{ij,t} = 0$ can finance the move from its own resources; a household with $s_{ij,t} > 0$ must raise the difference at a cost $\Phi(s_{ij,t})$, with

$$\Phi(0) = 0, \qquad \Phi' > 0, \qquad \Phi'' > 0. \tag{A2}$$

The convexity captures the rising marginal cost of distress financing: small gaps are bridged cheaply (a short informal loan), large gaps only at steep and increasing cost (asset liquidation at fire-sale prices, high-interest debt, the risk premium on an irregular crossing). The household now migrates whenever the surplus from moving, net of both the direct cost and the financing penalty, is positive:

$$S_{ij,t} \;\equiv\; V^{\text{move}}_{ij,t} - c_{ij} - \Phi\big(s_{ij,t}\big) - V^{\text{stay}}_{i,t} \;>\; 0. \tag{A3}$$

When $s_{ij,t} = 0$ the penalty vanishes and (A3) collapses to the desirability comparison $V^{\text{move}} - c_{ij} > V^{\text{stay}}$, so unconstrained households behave exactly as in the base model.

Constrained-household response. Consider a household with $s_{ij,t} > 0$ and differentiate the surplus (A3) with respect to the weather shock $\delta_{i,t}$. A higher δ erodes income through (4), so disposable income falls and the shortfall widens at rate $\partial s_{ij,t}/\partial\delta_{i,t} = -\varepsilon_{y,T}\, y^0_{i,t} > 0$. Using $\partial V^{\text{stay}}/\partial\delta_{i,t} = u'(y_{i,t} - \underline{c})\,\varepsilon_{y,T}\, y^0_{i,t} + \varepsilon_{A,T}$,

$$\frac{\partial S_{ij,t}}{\partial \delta_{i,t}} \;=\; \underbrace{-\varepsilon_{A,T}}_{\text{amenity}} + \big(-\varepsilon_{y,T}\, y^0_{i,t}\big)\Big[\underbrace{u'\big(y_{i,t} - \underline{c}\big)}_{\text{income push}} - \underbrace{\Phi'\big(s_{ij,t}\big)}_{\text{financing}}\Big]. \tag{A4}$$

The income loss $-\varepsilon_{y,T}\, y^0_{i,t} > 0$ now pulls in two opposing directions. As before it lowers the value of staying, raising the surplus from moving (the income push, $u' > 0$). But it also widens the shortfall and, through the convex penalty, raises the marginal cost of financing

the move ($\Phi' > 0$): the financing effect. The amenity term $-\varepsilon_{A,T} \geq 0$ raises the surplus regardless of the financing position. A constrained household therefore responds to the shock by migrating more iff

$$-\varepsilon_{A,T} \; + \; \left(-\varepsilon_{y,T}\, y^0_{i,t}\right)\left[u'(y_{i,t} - \underline{c}) - \Phi'(s_{ij,t})\right] \; > \; 0. \tag{A5}$$

The two limiting cases. The base model is the prohibitive-penalty limit. As $\Phi'(s_{ij,t}) \to \infty$ for any $s_{ij,t} > 0$, the financing term in (A4) dominates and the surplus from moving collapses: a constrained household never moves, whatever the income push, recovering the mechanical immobility of the financing margin (7); unconstrained households ($s_{ij,t} = 0$) are untouched and the desirability margin (6) is unchanged. The hard constraint (3) is thus the $\Phi' \to \infty$ corner of (A2), and a finite but steep Φ simply smooths the knife-edge so that a household near the threshold responds continuously rather than flipping discretely from $m = 1$ to $m = 0$.

The opposite case is distress migration. When the marginal financing cost exceeds the marginal utility of disposable income, $\Phi'(s_{ij,t}) > u'(y_{i,t} - \underline{c})$, the bracket in (A5) is negative, so the income channel on net deters the move: the household is in the trap. Yet a sufficiently large amenity loss can still satisfy (A5): if $-\varepsilon_{A,T}$ exceeds $(-\varepsilon_{y,T}\, y^0_{i,t})[\Phi'(s_{ij,t}) - u'(y_{i,t} - \underline{c})]$, a constrained household migrates despite the widening shortfall, financing the move at the convex penalty. This is the amenity-override route that the hard-constraint model cannot represent: there, an acute amenity shock acts only through unconstrained households at the desirability margin (the first bracket of (8)), whereas here it can also drive out households the base model would have classified as trapped. It offers a distress-financed route to the positive source-temperature effect on asylum applications reported by Missirian & Schlenker (2017): that observation need not arise only from unconstrained households at the desirability margin, the reading offered in the main text, but can equally reflect constrained households financing the move under an acute amenity loss, a distinction the reduced form cannot resolve.

Aggregate response and empirical content. Integrating (A4) over the corridor's households reproduces the density-weighted structure of (8) with the knife-edge financing margin replaced by a smooth band of constrained households indexed by their shortfall $s_{ij,t}$. The sign of the aggregate climate response remains an empirical question, now settled by the relative magnitude of $|\varepsilon_{A,T}|$ against the income-push-net-of-financing term integrated across that band rather than against a single threshold mass. Because the extension changes neither the reduced-form regressors nor the auxiliary implications tested in Steps 2 and 3 of Section 2 (climate must still depress origin income, and a negative non-climate origin income shock must still exert a negative effect on bilateral migration), the empirical strategy and all reported estimates are unaffected. We retain the hard constraint in the main text for transparency and treat the costly-financing technology as the extension that accommodates an amenity-dominated, distress-financed reading where the data call for it.

A.2 Permutation procedure details

For each origin i, we randomly permute the 60 monthly observations of $\mathrm{TMP}_{i,\cdot}$ and $\mathrm{PRE}_{i,\cdot}$ across the months of the panel (2018–2022), breaking the alignment between i's climate at month ym and i's bilateral flow at month ym. Shuffles are independent across origins, so the cross-origin spatial correlation of weather is destroyed in the placebo. Each origin's marginal climate distribution (its mean and variance) is preserved, but the seasonal ordering and the within-origin temporal autocorrelation are not. The seed is fixed at 20260417. We repeat $R = 100$ times; the observed effect is judged causal when it lies outside the resulting placebo distribution, that is, when few or no placebo replications reproduce a coefficient as large in absolute value as the one estimated on the correctly aligned panel.

A.3 Representativeness of the 2019–2022 weather variation

A four-year window invites the concern that the 2019–2022 weather was atypical, so that the estimated elasticity might reflect a peculiar period rather than a general response. It does not. The variation that identifies the elasticity is the within-origin-month anomaly, purged of each origin's annual mean and the global year-by-month path by the bilateral regression's fixed effects. To check that this identifying variation is representative, we extend the population-weighted CRU TS country-month series (Harris et al. 2020) back to 1990 and residualise temperature and precipitation on origin-by-calendar-month, origin-by-year, and year-by-calendar-month fixed effects, exactly as in the bilateral regression. Figure A1 compares the distribution of these residual shocks over 2019–2022 against the 1990–2018 baseline. The two are of essentially the same magnitude: the residual temperature shock has a standard deviation of 0.73,°C over 2019–2022 against 0.75,°C over 1990–2018, and the residual precipitation shock is likewise unchanged (1.41 versus 1.40 per 30 mm/month). The window thus carries ordinary weather variation, and the elasticity it identifies is not an artifact of an anomalous weather pattern.

A.4 Representativeness of the 2019–2022 migration variation

The same concern applies to the migration flows themselves, and is equally unfounded. Because no single bilateral-flow series spans both our window and the earlier decades, we compare our flows to the historical record after stripping out the gravity structure: in each source we residualise the log bilateral flow, defined as $\ln(1 + \text{flow})$ so that corridors with no recorded flow are retained, on corridor, origin-year, and destination-year fixed effects. These absorb every country's overall migration level in each period and leave the *bilateral surprise*, the deviation of a corridor from what its origin and destination totals imply. We apply this identically to our 2019–2022 flows (Chi et al. 2025a), aggregated to annual corridor totals, and to the five-yearly 1990–2015 historical flows of Abel & Cohen (2019), over the same developing-origin corridors. Figure A2 compares the two distributions. The residual bilateral shock has a

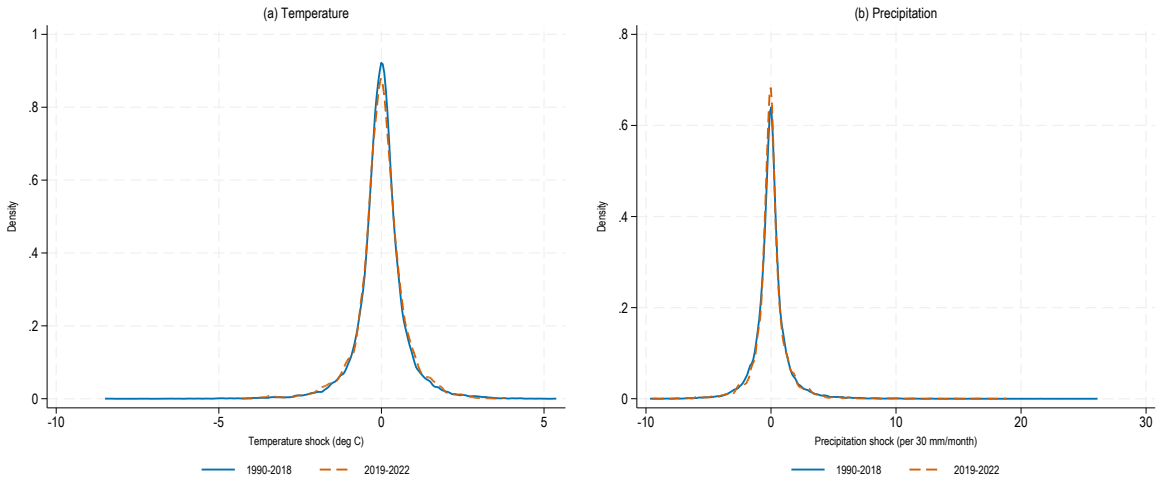

Figure A1: Distribution of the identifying weather variation, 2019–2022 versus a 1990–2018 baseline. Left panel: temperature; right panel: precipitation (per 30 mm/month). Each panel shows the kernel density of the within-origin-month residual shock for the 127 developing origins, purged of origin-by-calendar-month, origin-by-year, and year-by-calendar-month fixed effects (the variation that identifies the bilateral regression).

standard deviation of 0.42 over 2019–2022 against 0.61 over the 1990–2015 historical record. The historical figure is the larger of the two, but the difference most likely reflects the two sources' sampling frequencies (annual over a short period versus five-yearly over a longer period) rather than unusually quiet migration in our window. The comparison moreover concerns the structure of bilateral variation, not its level: the origin-year and destination-year fixed effects absorb each country's overall migration in every period, including any pandemic-era contraction, so a low level of total migration would not register here. Allowing for the coarser historical sampling, the bilateral variation our window exploits is in line with the historical record rather than peculiar to the period.[1]

A.5 Country-month NTL panel: construction

The country-month nighttime-lights proxy of Section 5.2 spans April 2012 – December 2024 for 136 developing countries ($n = 20{,}808$ country-months, of which 17 have zero land coverage in a given month and are missing by design). Monthly composites are built from NASA's Black Marble VNP46A2 daily product (Román et al. 2018) on Google Earth Engine.

For each 500 m land pixel and each calendar day we drop the retrieval if the mandatory quality flag indicates a gap-filled or borderline view angle (value 2) or no retrieval (value 255); if the snow flag is set; if the pixel sits on a JRC Global Surface Water occurrence above 50

[1]We report the pseudo-Bayesian closed demographic-accounting estimates, which Abel & Cohen (2022) find best reproduce observed flows. The comparison does not depend on the flow-estimation method: across the six methods of Abel & Cohen (2019) the historical standard deviation ranges from 0.30 to 1.06, and our window's 0.42 lies within that range.

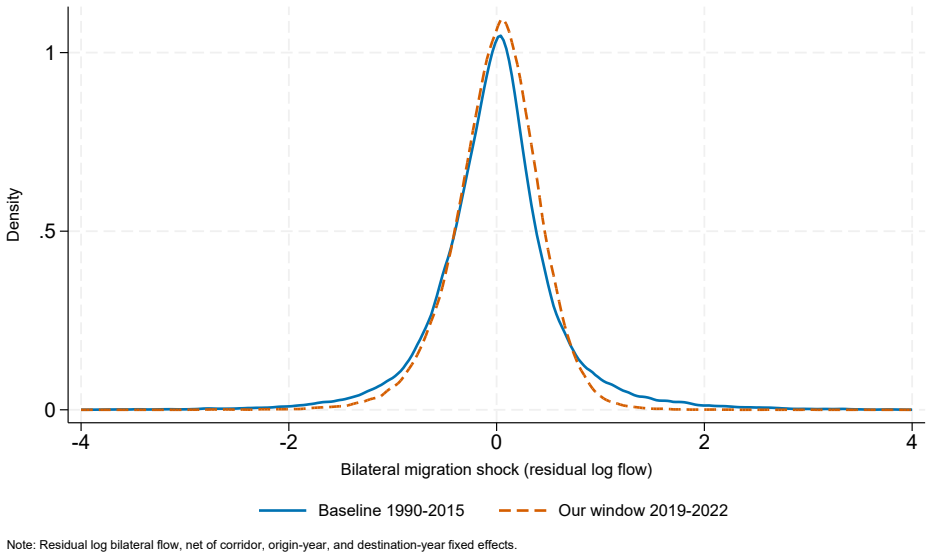

Figure A2: Distribution of the bilateral migration shock, the 2019–2022 window versus a 1990–2015 baseline. The shock is the residual log bilateral flow, net of corridor, origin-year, and destination-year fixed effects, for the developing-origin corridors; the window from Chi et al. (2025a), the baseline from Abel & Cohen (2019).

percent (Pekel et al. 2016); or if it falls within 1.5 km of any point in the EOG VIIRS Nightfire gas-flaring catalogue 2012–2024 (Elvidge et al. 2013, Zhizhin et al. 2021). Surviving daily radiances are clipped at the $0.3\,\mathrm{nW/cm^2/sr}$ noise floor. A pixel's monthly value is the mean of its surviving daily radiances when it has at least three valid daily retrievals; otherwise the pixel is missing.[2]

The country aggregate is the sum of pixel-level monthly radiances over GAUL 2015 admin-0 polygons; the *coverage share* is the fraction of country pixels surviving every mask with a non-missing monthly value. We derive three log-transformed dependent variables: total country brightness ($\ln$ sum radiance), the extensive margin ($\ln$ lit area in $\mathrm{km^2}$, "lit" meaning above $1.0\,\mathrm{nW/cm^2/sr}$), and the intensive margin ($\ln$ mean radiance per valid pixel). Climate exposures are the population-weighted CRU TS v4.09 (Harris et al. 2020) country-month series of the bilateral panel, with GPW v4 (Center for International Earth Science Information Network (CIESIN), Columbia University 2018) weights. The regression uses the $\{0,\ 1\text{-}4,\ 5\text{-}8,\ 9\text{-}12\}$ grouped-lag windows of the bilateral regression, country-by-year and year-month fixed effects, country-clustered standard errors, and coverage-share observation weights, so country-months in which more of the country was seen contribute more.

[2]The three-retrieval floor is permissive by design. Across the 136 countries and 153 months, the mean pixel carries 14.7 valid daily retrievals per month (median 14.6), roughly five times the floor, and only 2.6 percent of country-months fall below it. A stricter seven-retrieval floor would discard 18 percent of country-months, concentrated in the cloud-covered wet tropics (the Congo Basin, maritime South-East Asia, the Amazon fringe) and polar-night Russia: a weather-correlated sample selection in the places where climate exposure is greatest. The three-retrieval floor avoids that selection; residual variation in monthly data quality is instead absorbed smoothly through the coverage-share analytic weight. This is consistent with the minimum-observation guidance in NASA's Black Marble user guide and Román et al. (2018).

A.6 Census-microdata arrival panel: construction

Source and flows. The microdata are IPUMS International census extracts covering 366 census samples and 894.6 million person records. A bilateral arrival is a person born in origin o, enumerated in destination d at census c, who reports having immigrated in calendar year y (harmonised variables BPLCOUNTRY and YRIMM). Summing sampling weights within (o, d, c, y) cells yields directly measured annual arrival flows: each census contributes a retrospective annual series for every origin. We retain arrival years within ten years of the census, so that arrivals are reported close to the event; the recall window limits bias from return migration, mortality and misremembered years, and the estimates are insensitive to tightening it to five years or removing it. We further restrict to arrival years from 1990 onward, covering roughly the last three decades; the resulting panel spans arrival years 1990 to 2019, drawing on 45 destination-census pairs from 18 distinct census years between 1991 and 2020. Given the recall window, the 1990 floor drops fifteen further destination-census pairs from the 1970s and 1980s and trims pre-1990 arrival years from six 1990s censuses, retaining 97 percent of weighted arrivals within the recall window. The year-of-immigration question is asked almost exclusively by immigration societies: of the 28 destination countries that ever field it, Argentina and Venezuela do so only in censuses whose recall windows predate 1990, and Italy records birthplace only at the continental level, leaving 177 origin countries and 25 destinations (the United States, Canada, most of Latin America, Spain, Greece, Switzerland, South Africa, Kenya and Malaysia). Birthplace codes that do not identify a single country (9.4 percent of weighted arrivals) are dropped, as are persons born in the census country.

Sampling weights. Each IPUMS sample is a draw from the full census returns, typically 5 or 10 percent of households. Every person record i carries an expansion weight w_i equal to the inverse of its sampling probability. The bilateral flow on corridor $o \to d$ in arrival year y, as recorded by census c, is

$$\text{flow}_{odcy} \;=\; \sum_{i \in S(d,c)} w_i \, \mathbf{1}\{\text{BPL}_i = o\} \, \mathbf{1}\{\text{YRIMM}_i = y\}, \tag{A6}$$

where $S(d, c)$ is the set of sampled persons in census c of destination d, BPL_i is birthplace, and YRIMM_i the reported year of immigration. The cell is therefore an estimate of the total population of qualifying arrivals, not a count of sampled records.

Zeros. Aggregation creates a cell only where the census sampled at least one arrival, so corridor-years in which nobody came are absent rather than zero; a regression on observed cells alone would estimate the intensive margin only, and a weather shock that suppresses migration entirely produces exactly such an empty cell. We therefore rectangularise: for every destination-census, we form the full grid of all 177 origins and all recall years, coding as zero

every cell with no sampled arrival (75.5 percent of the resulting 71,114 cells).

Weather and controls. Weather is the population-weighted CRU TS country mean of temperature and precipitation (per 30 mm/month, as in the main text), with weights from the GPW v4 2015 population grid applied consistently across all arrival years. The relevant timing for both origin and destination is the year the household leaves the origin, since that is when the migration decision is taken. Census microdata, however, record only the arrival year: each destination census asks its residents when they came, not when they left somewhere else. We therefore use the arrival year y throughout, which for the typical international journey, completed within a calendar year, coincides with the year of departure. The cumulative rows of the table additionally enter $y-1$ and report the summed effect, capturing conditions over the twelve months preceding departure. All specifications control for log origin population, log destination population, and log destination GDP per capita (Feenstra et al. 2015), and the estimation sample is held fixed across columns. Column (3) adds spatial-lag (SLX) terms, row-standardised inverse-distance averages of the origin's neighbouring-country weather within 2,000 km, mirroring the spatial robustness specification of the main text; column (4) adds log origin GDP per capita, the income channel itself.

Validation. The panel reproduces external benchmarks at every level. Summing person weights to the country level reproduces the Feenstra et al. (2015) annual mid-year population to within 5 percent across seven benchmark censuses (Argentina 1970 and 2010, Brazil 2010, the United States 2000, Mexico 2010, Spain 2011, and Indonesia 2010); the residual reflects timing (census-day versus mid-year) and post-enumeration adjustments. The Mexico-born population in the United States 2000 census is 9.25 million against the published 9.18 million. The arrival-year dimension itself traces known histories, with the retrospective Mexico–United States series peaking at 647,000 in 1999 — the documented all-time peak of Mexican migration — and the Morocco–Spain series collapsing from roughly 50,000 per year before 2008 to 21,000 by 2010.

A.7 Compound climate events: definitions and construction

Events are identified at the 0.5° cell-month level using the standardised climate-index thresholds of the WMO (World Meteorological Organization 2012). The Standardised Temperature Index (STI) is the calendar-month normal score of monthly mean temperature: the observed value is ranked within the cell's calendar-month-specific 1981–2010 benchmark distribution, and the resulting percentile is mapped through the inverse standard normal. The Standardised Precipitation-Evapotranspiration Index (SPEI) of Vicente-Serrano et al. (2010) is the monthly water balance $P - PET$ at a one-month timescale, standardised by fitting a Generalised Logistic distribution (Hosking 1990) to the cell's 1981–2010 benchmark; P is monthly precipitation and PET is monthly potential evapotranspiration, both from

CRU TS v4.09 (Harris et al. 2020), which derives PET using the FAO-56 Penman-Monteith formulation.

The base layer flags a cell-month for a *heatwave* whenever $\mathrm{STI} \geq +1.0$, for a *coldwave* whenever $\mathrm{STI} \leq -1.0$, for a *drought* whenever $\mathrm{SPEI} \leq -1.0$, and for a *flood* whenever $\mathrm{SPEI} \geq +1.0$. Because each index is standardised to a standard normal, a unit threshold marks one standard deviation from the long-run norm, equivalently the 16th/84th percentile; this is the WMO "moderate" cutoff. Compound events are cell-month intersections of two single-event flags: *hot+dry* requires $\mathrm{STI} \geq +1.0$ and $\mathrm{SPEI} \leq -1.0$ jointly, *hot+wet* requires $\mathrm{STI} \geq +1.0$ and $\mathrm{SPEI} \geq +1.0$, and *cold+dry* and *cold+wet* analogously. The country-month exposure for each event is the share of country population in cells flagged for the event in the month, using GPW v4 (Center for International Earth Science Information Network (CIESIN), Columbia University 2018) 2015 gridded population as the within-country weight. The severity layers used in robustness ($|\mathrm{index}| \geq 1.5$ and ≥ 2.0, that is, 1.5 and 2 standard deviations) raise the bar to the 7th/93rd and 2nd/98th percentiles respectively and are otherwise constructed identically.

The eight regressors used in Section 4.3.1 are mutually exclusive partitions of the population-share space, built by subtraction. The heat-only share is the heatwave share net of the hot-dry and hot-wet shares, so it isolates a heatwave that does not coincide with concurrent drought or flood; the cold-only, drought-only, and flood-only categories are built analogously; and the four compound categories enter as their joint-event shares directly. Each rescaled regressor lies in $[0, 1]$.

A.8 Conflict risk and origin nighttime lights

Table A1 reports the income-channel result for the non-climate shock of Section 5.4. It applies the country-month nighttime-lights specification of Appendix A.5 with the origin domestic conflict-risk index entered as grouped distributed lags: alone (column 1), against the weather lags alone (column 2), and jointly (column 3), all on a common sample. The eight-month cumulative effect of conflict risk on origin brightness is negative and statistically significant, and remains negative when the weather lags are included, confirming that the non-climate shock erodes origin economic activity at the monthly frequency. The conflict coefficient attenuates once weather is added (column 1 versus column 3), while the temperature coefficient is essentially unchanged (column 2 versus column 3), a pattern consistent with climate stress contributing to the conflict-risk signal; we do not attempt to identify that channel here. Because conflict timing is less plausibly exogenous than weather, we read the conflict result as corroborative of the income channel rather than as a clean causal effect.

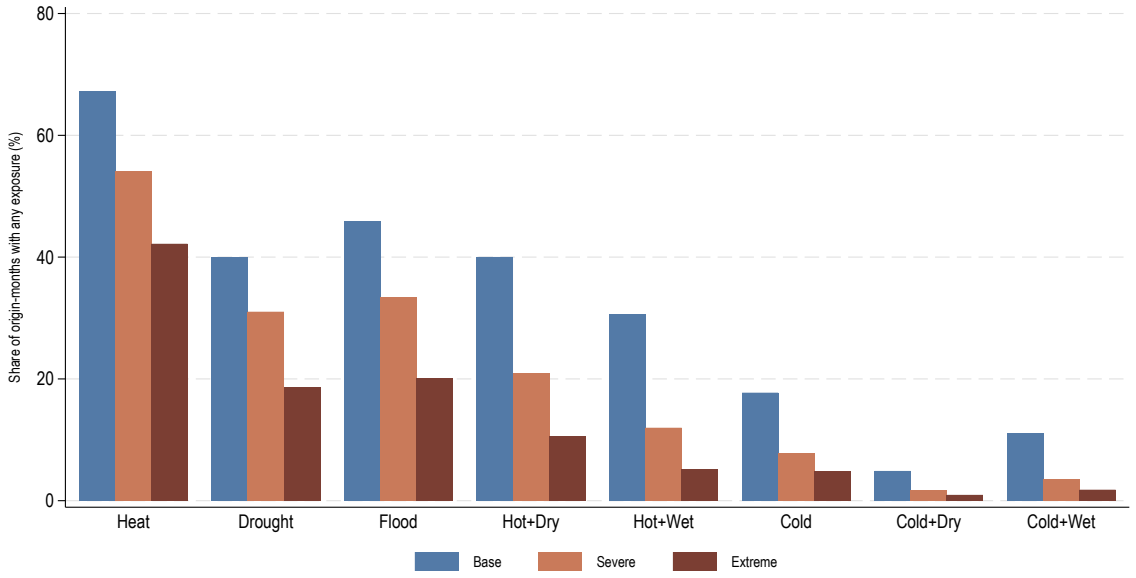

(a) Event frequency

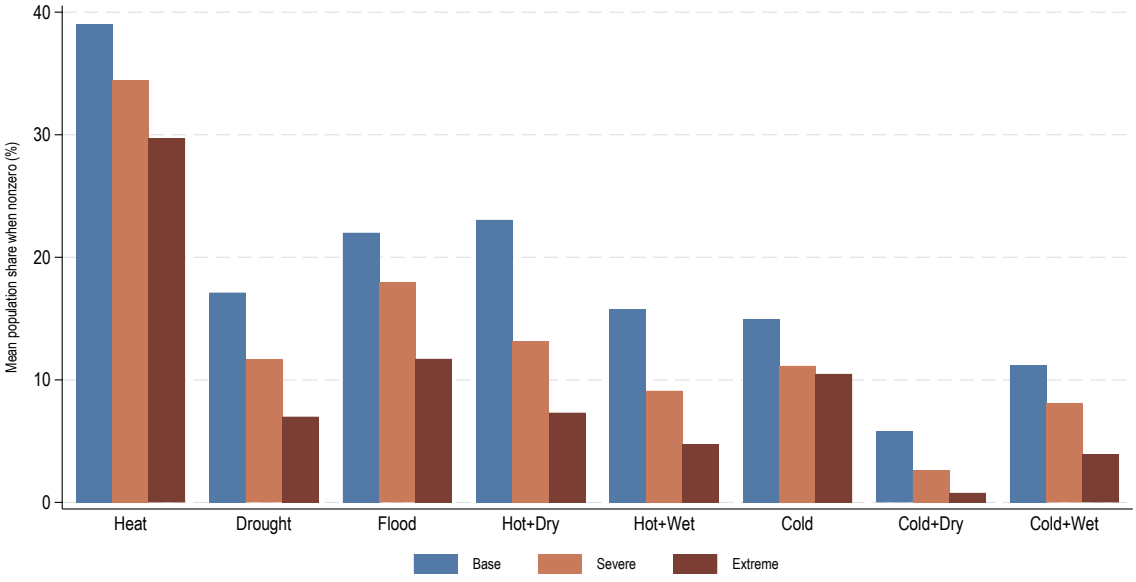

(b) Conditional intensity

Figure A3: Frequency and conditional intensity of the eight compound climate-event categories on the developing-origin sample, 2019–2022, by WMO severity threshold. Panel (a): share of origin-months with any positive exposure to each category. Panel (b): mean share of country population exposed in origin-months with positive exposure (conditional intensity). Base layer: $|\text{index}| \geq 1.0$ (WMO "moderate" threshold, 16th/84th percentile). Severe layer: $|\text{index}| \geq 1.5$ (7th/93rd percentile). Extreme layer: $|\text{index}| \geq 2.0$ (2nd/98th percentile). Heat-only events dominate frequency and conditional intensity at every layer; cold-related events are present but rare, especially the compound cold categories at higher severity.

Table A1: Domestic conflict risk, climate, and origin economic activity.

	(1) Conflict only	(2) Climate only	(3) Conflict + climate
Conflict risk, cumulative 8-month	-0.229^{**}		-0.175^{*}
	(0.096)		(0.095)
Conflict risk, cumulative 12-month	-0.241^{*}		-0.186
	(0.134)		(0.133)
Temperature, cumulative 8-month		-0.035^{***}	-0.036^{***}
		(0.012)	(0.012)
Temperature, cumulative 12-month		-0.053^{***}	-0.054^{***}
		(0.014)	(0.014)
Conflict-risk lags	Yes	No	Yes
Climate lags	No	Yes	Yes
Country-by-year fixed effects	Yes	Yes	Yes
Year-month fixed effects	Yes	Yes	Yes
Observations	19,392	19,392	19,392

Notes: Dependent variable is log total country brightness (sum of nighttime-light radiance) in the country-month panel of Appendix A.5. Conflict risk is the GDELT-based calibrated domestic conflict-risk index (in $[0, 1]$); weather is origin temperature and precipitation. All regressors enter as grouped distributed lags $\{0, 1\text{-}4, 5\text{-}8, 9\text{-}12\}$, and reported coefficients are cumulative sums over the 8- and 12-month horizons. The three columns share a common estimation sample. Precipitation is included in columns (2) and (3) on the same grouped-lag structure but its cumulative coefficients are not statistically significant in either column and are omitted from the display for legibility; the temperature coefficients are essentially identical whether precipitation is included or excluded. All regressions weight by the monthly coverage share and absorb country-by-year and year-month fixed effects; standard errors in parentheses are clustered by country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.