

Migrant Networks and Climate Change

Shifrah Aron-Dine
UC Berkeley

Emmanuella Kyei Manu
Stanford
(*Job Market Paper*) *

October 23, 2026
[\[link to latest version\]](#)

Abstract

Migration is a potentially important form of climate adaptation if people choose to relocate to less exposed places. However, using global data on migration and climate, we find that new emigrants follow existing diasporas and that diasporas are concentrated in places with similar climates to the origin country. To assess the implications of these patterns for future adaptation, we build a dynamic spatial model with endogenous migrant networks and spatially correlated climate risk. Networks exacerbate climate losses when they are concentrated in locations whose expected damages and climate shocks positively covary with the origin's. Quantitatively, networks increase the welfare cost of climate change for two-thirds of the world. In poor, climate-exposed countries, they raise consumption-equivalent welfare losses by 2 percentage points, roughly the cost of an additional degree of warming by 2100. Targeted migration policies improve adaptation by building networks in less exposed destinations, and earlier policies generate larger gains.

*Emails: aron-dine@berkeley.edu and kyeimanu@stanford.edu. We are especially grateful to Pete Klenow, Melanie Morten, Patrick Kehoe, and Benny Kleinman for their guidance and support. We would also like to thank Ran Abramitzky, Adrien Auclert, Adrien Bilal, Johannes Boehm, Kirill Borusyak, Tamma Carleton, Doireann Fitzgerald, Kelsey Jack, Chad Jones, Ishan Nath, Elena Pastorino, Monika Piazzesi, Stephen Redding, Martin Schneider, Felix Soliman, Chris Tonetti, Mike Waugh and other participants in the Stanford macro group for helpful feedback and suggestions. Emmanuella Kyei Manu acknowledges support by the Leonard W. Ely and Shirley R. Ely Graduate Student Fellowship through a grant to the Stanford Institute for Economic Policy Research.

1. Introduction

The impacts of climate change are unevenly distributed across the globe: high-latitude economies may benefit from warming, while low-latitude countries face potentially catastrophic losses. International migration has the potential to be a key endogenous adaptation margin by enabling large-scale reallocation toward less climate-exposed locations. Migrant networks may, however, shape where people move. These networks consist of diaspora communities from an origin spread across destinations. A large literature focused on specific migration episodes or narrower geographic settings establishes that a larger diaspora in a destination attracts subsequent migrants from the same origin. In the context of climate change, a large diaspora in a climate-safe destination could facilitate adaptation by lowering migration costs to that destination. However, dependence on the existing migrant network may lead people affected by adverse climate events in their origin country to move to climate-exposed destinations where their diaspora is already large rather than to climate-safe destinations with no established community. Whether migrant networks facilitate or limit migration away from places exposed to climate change therefore depends on where these diaspora communities are located and the climate damages those destinations face.

In this paper, we study how migrant networks shape the geography of climate-induced migration and assess whether they help or hinder adaptation to climate change. We first combine global migration and temperature data to measure how strongly existing diasporas steer the destination choices of new migrants and whether those diasporas are located in places with similar climate exposure to the origin. We also examine how the climate exposure of migrant networks has changed as the world has warmed. We then develop a dynamic spatial model in which larger diasporas attract new migrants and migrant networks evolve endogenously as people move. Climate change is modeled as a warming trend with persistent, spatially correlated climate shocks. We quantify the model to assess the value of migrant networks for future climate adaptation, decompose the channels through which networks interact with climate losses, and study the role of network dynamics for migration policies.

We find that migrant networks shape both migration flows and climate exposure. Destinations with larger existing diasporas receive larger subsequent migration flows, and these diasporas are located in places with similar climates to their origins, although migrant networks have shown some adaptation with warming over our sample period. As climate change worsens over the century, these networks will increase the welfare cost for two-thirds of the world. For countries that lose from climate change, networks add 1.7% in consumption-equivalent welfare losses, roughly equivalent to an additional 1°C of warming. The key determinant is where existing migrant networks are located relative to future climate damages. Finally, migration policy can have persistent effects on climate adaptation through the networks it builds or depletes. Temporary restrictions on migration to the United States and European Union increase climate losses even after those restrictions are lifted, while targeted climate visas can build networks in destinations relatively spared by climate change and facilitate future migration. Because these networks take time to build, delaying such policies

reduces their benefits.

We begin by establishing three empirical facts about global migration and exposure to climate change. First, we show new emigrants follow the existing network. Using United Nations data on global migrant stocks and estimates of migration flows from [Abel and Cohen \(2019\)](#) from 1990 to 2020, we show that a larger diaspora attracts more migrants from the same origin. This pattern holds globally and is robust to controlling for origin and destination fixed effects that absorb all push and pull factors, as well as origin–destination fixed effects that account for time-invariant bilateral determinants of migration. To address concerns about time-varying confounders, we implement an instrumental-variables strategy following [Burchardi et al. \(2019\)](#). The instrument predicts changes in an origin–destination diaspora using the interaction of an origin-specific emigration wave with a destination-specific immigration boom. We construct both components using geographic leave-outs so that the bilateral flow between the origin and destination does not mechanically enter the instrument. The IV estimates show that a 10 percent larger diaspora increases subsequent migration by 7 percent.

Second, given the strong network effect we observe globally we try to understand whether migrant networks facilitate climate adaptation by helping individuals from climate-exposed origin countries move toward safer destinations. We measure climate exposure using temperature data from Berkeley Earth. For each origin country, we compute the climate exposure of its diaspora network as the migrant-weighted average temperature across destination countries. We find that origins and their networks face similar climates. Origin countries exposed to high temperatures have diasporas concentrated in destinations that are similarly exposed, leaving little scope for migration through existing networks to reduce climate exposure.

Third, migrant networks show some adaptation with the warming observed over recent decades. We decompose changes in diaspora-network exposure for each country into network reallocation, holding destination temperature fixed, and changes due to warming at existing destinations, holding the network fixed. We find that warming at destinations increased diaspora exposure, while network reallocation offset about 40% of this increase. This suggests that migrant networks may adjust as climate conditions change, but the data cannot tell us whether that adjustment is enough to limit the welfare losses from future warming.

To quantify the welfare effects of future warming, we develop a dynamic model of migration in which existing diasporas shape destination choices and climate damages vary across locations. Households choose their location each period based on local wages, the size of the diaspora from their home country, and a destination-specific cost of residing away from home. They also receive idiosyncratic taste shocks for each destination, which gives a tractable structure for location choice. Climate change affects local productivity through deterministic warming trends and persistent, spatially correlated shocks, generating changes in migration incentives over time. Because current migration changes the size and geographic distribution of diaspora communities, migrant networks evolve endogenously and shape future migration decisions. The model delivers the positive relationship between migrant stocks and subsequent flows that we document in the

data. When utility is increasing in the local diaspora, a larger migrant stock makes a destination more attractive both to migrants already there and to new migrants from the same origin.

The network elasticity summarizes how strongly past migration shapes current destination choices. It measures how strongly a destination’s attractiveness rises with its migrant stock relative to the dispersion of idiosyncratic location tastes. When the elasticity is greater than one, the feedback is strong enough to generate multiple steady states. Temporary changes in migration incentives such as policy or productivity shocks can permanently alter the geography of migrant networks. When the elasticity is below one, the long-run migrant network is unique given the path of climate. However, network feedback still shapes migration dynamics through persistence and amplification. Persistence comes from the effect of past migration on current destination choices. A temporary shock that sends migrants to a destination increases the diaspora there, which continues to attract migrants after the original shock has disappeared. The effect of the initial shock therefore decays gradually, at a rate governed by the network elasticity. Networks also amplify the migration response to changes in fundamentals. An increase in a destination’s productivity raises migration to that destination and, in turn, expands its diaspora. The larger diaspora then raises subsequent migration further, so the overall response exceeds the response to productivity alone. For a permanent change in fundamentals, a higher network elasticity implies greater long-run amplification of migration.

We next quantify the model to measure how migrant networks affect adaptation to future climate change. The quantification requires three inputs: economic growth absent climate change, the evolution of expected climate damages and climate risk across countries, and the strength of the network in migration decisions. For economic growth absent climate change, we use country-level projections from the Shared Socioeconomic Pathways. For climate, we separately parameterize the warming trend and climate risk. We map projected warming into country-level productivity losses using the temperature response function of [Nath et al. \(2024\)](#). We estimate climate risk from the spatial covariance of weather shocks across one hundred runs of a single climate model under a common emissions scenario. Because all runs use the same climate model and emissions scenario, differences across runs capture variation in weather realizations rather than differences across climate models. As temperatures rise, the same weather fluctuations can generate larger productivity effects, allowing the volatility of climate damages to change over time. Finally, we jointly estimate the network elasticity and residency costs using five-year changes in bilateral migrant stocks from 1995 to 2020. Starting from the observed network in each period, we estimate the parameters so that the model reproduces where migrants from each origin are located five years later. The estimated network elasticity is below one, implying a unique long-run network, but sufficiently close to one that network-driven persistence and amplification are quantitatively important.

We use the quantified model to ask how migrant networks change the welfare costs of climate change. We ask how large climate losses would be if migrants faced the same productivities, residence costs, and climate conditions, but did not choose destinations based on existing

diasporas. This comparison decomposes the welfare cost of climate change into the loss that would arise without network-driven migration and the additional loss or gain attributable to migrant networks. To implement this comparison, we solve four versions of the model: with and without climate change, and with and without the effect of existing diasporas on destination choice. We first measure the welfare cost of climate change when existing diasporas shape migration decisions by comparing welfare with climate change to welfare without climate change. We then repeat the same calculation after removing the effect of existing diasporas on destination choice. The difference between these two climate losses measures how much migrant networks amplify or mitigate the welfare cost of climate change.

We find that network-driven migration raises the welfare cost of climate change for two-thirds of the world. The largest increases are in Sub-Saharan Africa, the Middle East, and South and Southeast Asia, where projected climate damages are high. For origins in the hottest quartile, migrant networks add 2% of permanent consumption to the welfare cost of climate change. By contrast, for origins in the coldest quartile, networks reduce climate losses by 0.35% of consumption on average. The effects also vary with the income level of the origin country. For countries below median GDP per capita, migrant networks add 1.7% of consumption to climate losses, roughly equivalent to the welfare cost of an additional degree of warming in these countries.

The network raises the welfare cost of climate change through two channels. First, it anchors households to the origin and to destinations where their diaspora is already large, changing their exposure to the expected warming trend. Second, it changes how well households can diversify climate risk across destinations, and therefore how much insurance migration provides against climate shocks. We derive both channels from a second-order expansion of expected lifetime utility around the expected climate path. For anchoring, the key moment is the covariance between diaspora shares and damages from the warming trend. Networks raise climate losses when migrants are concentrated in destinations that also experience large damages from warming. For diversification, the relevant moment is whether the network is concentrated in pairs of destinations whose climate shocks are positively correlated. When climate shocks across network destinations move together, alternative destinations provide less protection in a bad climate realization, reducing the option value of migration. Quantitatively, anchoring accounts for nearly all of the network effect on expected welfare. Although diversification has a small effect on expected losses, it matters in the tails of the distribution of climate realizations. In the worst 5% of climate draws, the consumption-equivalent loss from climate shocks is twice as large with the network as without it.

Finally, we show that migration policies can shape climate adaptation by changing the geography of migrant networks. We illustrate this with two counterfactuals. The first temporarily closes the borders of the United States and the European Union to the rest of the world from 2020 to 2060. The restrictions end before the worst decades of climate change, but their effects persist. By reducing migration during the closure, the policy depletes the networks that would otherwise facilitate future migration to these destinations. As a result, 41 percent of the additional climate

losses caused by the restrictions occur after the borders have reopened. The second counterfactual considers climate visas that allow migrants from poor countries facing large climate damages to move to destinations that are spared by climate change. The visa program provides a small number of admissions, which establish diaspora communities in the receiving countries. These networks then attract additional migrants through other migration channels. As those migrants arrive, they further expand the diaspora and increase subsequent migration to the same destinations. Because this process takes time, delaying the policy is costly. A policy introduced in the 2040s gives the network less time to build and delivers about half the protection of a policy introduced in 2020 over the same subsequent years.

Related Literature. Our main contribution is to introduce a dynamic spatial model with endogenous migrant networks. A large empirical literature shows that existing networks shape destination choice and migrant outcomes through a variety of mechanisms (Bartel, 1989; Card, 2001; Edin et al., 2003; Munshi, 2003; Beine et al., 2011; Beaman, 2012; Mahajan and Yang, 2020; Stuart and Taylor, 2021; Blumenstock et al., 2025).¹ We add to this literature by providing global evidence on the relationship between existing migrant stocks and subsequent migration flows across more than 200 countries. Despite the breadth of empirical evidence on migrant networks, these mechanisms are rarely embedded in quantitative models of migration. A notable exception is Carrington et al. (1996), who formalize this network effect by assuming that moving costs fall with the stock of prior migrants, yielding self-reinforcing dynamics and the potential for path dependence. We build on their insight by incorporating endogenous migrant networks into a multi-period, multi-location equilibrium framework that can be solved quantitatively, allowing us to study transition dynamics as well as equilibrium properties.

This framework also connects migrant networks to the development literature on risk-sharing. A central result in this literature is that migration allows households to insure against location-specific shocks. This literature shows that households use family ties across space and migration to insure against local shocks and smooth downturns (Rosenzweig and Stark, 1989; Yagan, 2014; Borusyak et al., 2026; Bryan et al., 2014; Lagakos et al., 2023). At the same time, access to insurance at the origin can reduce incentives to migrate (Munshi and Rosenzweig, 2016; Morten, 2019; Meghir et al., 2022). Closest to our mechanism, Schwartzman (2025) considers how the insurance value of migration depends on the spatial covariance of shocks, showing that migration within the Philippines is directed toward destinations whose income shocks are less correlated with those at home. Our contribution is to formally characterize and quantify how the migrant network shapes the insurance value of migration when losses are spatially correlated and evolve over time.²

Additionally, we build on a growing literature that studies global warming in dynamic spatial

¹See Munshi (2020) for a recent review.

²A related literature studies diversification against spatially correlated shocks in production and trade including plant location within the firm (Castro-Vincenzi, 2024), supplier networks (Castro-Vincenzi et al., 2024), and the policy choice between diversification and reshoring (Grossman et al., 2023; Kleinman et al., 2023). We study diversification through migration, where the set of destinations available to a household is shaped by the migrant network.

equilibrium settings where migration is a central adaptation margin.³ Recent work emphasizes that heterogeneous climate damages can trigger large reallocations of people and economic activity across space, and that mobility frictions are therefore first-order for the welfare effects of climate change (Desmet and Rossi-Hansberg, 2015; Conte et al., 2021; Desmet et al., 2021; Bilal and Rossi-Hansberg, 2023; Conte, 2024; Cruz and Rossi-Hansberg, 2024; Bilal et al., 2025)⁴. In these models, mobility frictions are typically fixed over the transition. In our framework, part of the mobility friction comes from the migrant network, which evolves endogenously as people move. Migration today therefore changes the relative attractiveness of destinations for future migrants, so the mobility frictions also evolve as climate damages change across space.⁵ Because these mobility frictions evolve with migration, they can either reinforce existing migration patterns and slow reallocation toward destinations with lower climate damages, or facilitate reallocation when migrant networks are already established in those destinations.

The remainder of the paper is organized as follows. Section 2 describes the data and documents three motivating facts on the role of diaspora networks in shaping migration destinations and their climate exposure. Section 3 develops a dynamic spatial equilibrium model of migration with endogenous migrant networks and climate change. Section 4 quantifies the model. Section 5 discusses the role of the migrant network in the welfare cost of climate change. Section 6 explores how migration policy affects adaptation through the migrant network. Section 7 concludes.

2. Three Facts about Migrant Networks and Climate Exposure

The first part of the paper documents three empirical patterns that motivate our analysis. First, migrant networks shape international migration destinations. Second, these networks are relevant for climate adaptation because origins and their networks face similar climates. However, our third fact shows that network reallocation over the past decades offset some of the warming that the diaspora would have otherwise faced. Together, these facts suggest that existing migrant networks may shape future climate-induced migration, and motivate a structural model that can assess whether network-driven migration facilitates or limits future adaptation to climate change.

2.1 Data

Migrant stocks and flows. We use data on migrant stocks from the United Nations Global Migration Database (United Nations, Department of Economic and Social Affairs, Population

³A related empirical literature estimates how long-run climate trends (Cattaneo and Peri, 2016; Beine and Parsons, 2017; Jessoe et al., 2018) and extreme events (Hornbeck, 2012; Groeger and Zylberberg, 2016; Deryugina et al., 2018; Mahajan and Yang, 2020; Aron-Dine, 2025) affect migration. See Cattaneo et al. (2019) for a review.

⁴See Desmet and Rossi-Hansberg (2024) for a recent review of these types of models and Carleton et al. (2024) for a broader summary of work on climate adaptation

⁵This relates to Seo and Oh (2024) who show that persistent worker types can make reallocation across industries in response to sector-specific shocks substantially more costly than implied by models without type persistence. In our framework, country of origin specific migrant networks are a persistent state that creates broad heterogeneity in the cost and desirability of relocating across destinations as location-specific climate damages evolve.

Division, 2021). This is the most comprehensive source of international migration statistics available. The database provides bilateral migration matrices that capture the stock of foreign-born populations residing in destination countries at 5-year intervals from 1990 to 2020 across over 200 countries. Our measure of bilateral migration flows comes from [Abel and Cohen \(2019\)](#) pseudo-Bayesian demographic-accounting series which builds on the UN bilateral migrant stock data.⁶ Their approach explicitly accounts for births and deaths in the migrant population and employs a statistical estimation procedure to infer the most likely migration flows consistent with observed stocks.⁷ Appendix B contains a discussion of the data construction and sources of potential systematic measurement error.

Climate data. We use data from the Berkeley Earth Surface Temperature (BEST) dataset. This dataset has daily temperature measurements at $1^\circ \times 1^\circ$ spatial resolution from 1970-2020. It is one of the most comprehensive land-surface temperature datasets available, using data from over 30,000 weather stations worldwide and statistical methods to address common biases in raw station data. To construct country-level temperature exposure measures, we combine the gridded temperature data with high-resolution population data. We calculate daily temperature experienced by an average person in a country. From this, we construct annual measures such as mean temperatures and measures of extreme heat exposure, such as the number of days exceeding specific temperature thresholds. Appendix B provides more detail about the data and additional climate measures we use in robustness exercises.

2.2 Fact 1: New emigrants follow the existing network

The data show that existing diaspora communities play an important role in explaining both the size and the destination of migration flows. Figure 1a illustrates this relationship using global bilateral migration over the last quarter century. For each origin–destination pair (o, d) and period t , we plot a binscatter of the share of origin- o flows going to destination d against the share of origin- o 's pre-existing diaspora located in d five years earlier. There is a strong positive relationship between contemporaneous flows and past migrant stocks. This pattern is robust over time and across levels of economic development (see Appendix A).


The positive correlation in Figure 1a does not by itself identify a causal network effect. Migrant stocks and flows may both reflect origin and destination characteristics, such as wages or political stability, as well as bilateral factors such as language, colonial ties, cultural proximity, distance, or climate similarity.

To address these concerns, we implement an IV strategy following [Burchardi et al. \(2019\)](#) who

⁶Differencing migrant stocks can generate spurious zeros and mismeasure flows because stock changes also reflect deaths, naturalization, and data revisions, and may miss onward migration.

⁷They use a pseudo-Bayesian approach to mix two inference strategies: The first attempts to maximize matrix diagonals (stayers) while still respecting the migrant accounting identity imposed by a start and ending stocks. The second approach attempts to spread flows more smoothly across destinations given the same margins. They find that this approach performs best against the limited sample of reported global flows measures.

Figure 1: Role of Networks in Migration and Climate Exposure



Notes: Panel (a) uses data on migrant stocks from UN Global Migration Database, flows from [Abel and Cohen \(2019\)](#) estimates. The figure shows a binscatter plot of flows and lagged migrant stocks for an origin- o and destination d pair using data from 1990-2020. The orange line plots the line of best fit and the dashed gray line shows the 45 degree line. Panel (b) bins origin-country temperature against the average temperature of its diaspora. The solid line is the line of best fit and the dashed line is the 45-degree line.

study the impacts of ancestry on foreign investment flows across US counties. The instrument exploits the interaction of push and pull factors that coincide in a time-period to channel migrants from an origin- o into a destination- d . When a large emigration wave from origin o coincides with a period in which destination d is attractive to migrants more generally, more migrants from o settle in d . Because migrant stocks accumulate past flows, these historical push-pull interactions generate differences in the distribution of the origin- o diaspora across destinations.⁸ We adapt this logic to bilateral international migration. Our push factor captures an origin country specific emigration wave, while our pull factor captures the contemporaneous attractiveness of a destination country to migrants from other countries.

Using total emigration from o and total immigration into d would mechanically load the $o \rightarrow d$ flow into both components of the instrument. Therefore, we follow the leave-out construction in [Burchardi et al. \(2019\)](#). They construct the origin push factor excluding migration to the destination US county Census region and the destination pull factor excluding migrants from the origin's continent. In our global setting, we use the analogous geographic leave-out at the UN subregion level. The push factor excludes flows from o to the UN subregion containing d , while the pull factor excludes migrants from the UN subregion containing o .

⁸[Burchardi et al. \(2019\)](#) provide an example with nineteenth-century German migration to the Midwest, where a German emigration wave coincided with strong destination pull. We see similar episodes in our data. For example, migration from Pakistan to Turkey increased following the 2010 Pakistan floods and during Turkey's 2010–2015 immigration reforms. Appendix A provides additional case studies.

The identifying assumption is that, conditional on our fixed effects, the interaction of these push and pull shocks affects subsequent o - d migration only through the migrant stock it creates, and is uncorrelated with time-varying bilateral factors that directly affect future migration between o and d . A remaining concern is that the push-pull interaction may be correlated with persistent shocks that drive o - d migration. For this violate the exclusion restriction, the shock would need to be persistent for o and d by driving emigration from o while also keeping d relatively attractive. The geographic leave-out further limits this concern. For a persistent confounder to violate the exclusion restriction, it would need to continue pushing migrants from origin o toward destinations outside d 's subregion, while simultaneously making destination d attractive to migrants from origins outside o 's subregion.

We instrument for migrant stocks in year $t - 5$ using the history of push-pull interactions up until year $t - 5$. Appendix A contains more details on constructing the instrument. We then estimate

$$\log(\text{flows}_{od,t}) = \beta \log(\widehat{\text{stocks}}_{od,t-5}) + \delta_o + \delta_d + \delta_{od} + \varepsilon_{od,t} \quad (1)$$

where $\text{flows}_{od,t}$ is the level of the migrant flows and $\widehat{\text{stocks}}_{od,t-5}$ is the instrumented lagged level of the migrant stock. δ_o and δ_d are origin and destination fixed effects that absorb time-invariant push and pull factors common across destinations and origins, respectively. The origin-destination pair fixed effect, δ_{od} , absorbs stable bilateral characteristics such as distance, shared language, and colonial history. The specification including origin-destination pair fixed effects identifies the network effect from within-corridor changes over time. We report heteroskedasticity-robust standard errors clustered at the origin-destination pair level.

Table 1: Elasticity of Flows to Stocks

	OLS			IV	
	(1)	(2)	(3)	(4)	(5)
$\log(\text{stocks}_{od,t-5})$	0.882*** (0.002)	0.830*** (0.003)	0.089*** (0.006)	0.723*** (0.111)	0.723*** (0.112)
o, d		✓	✓		✓
$o \times d$			✓	✓	✓
First stage F -stat				11.8	11.7

Notes: The table reports estimates of the elasticity of bilateral migration flows to lagged migrant stocks from 1990 to 2020. Migrant stocks are from the UN and gross flows from [Abel and Cohen \(2019\)](#). Columns (1)–(3) report OLS estimates; columns (4)–(5) instrument for lagged stocks using the leave-out push-pull instrument of [Burchardi et al. \(2019\)](#). Standard errors are clustered at the origin-destination pair level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The coefficient β measures the elasticity of bilateral migration flows to the lagged diaspora stock. Table 1 reports the estimates. Columns (1)–(3) show the OLS results.⁹ In column (1), which includes no fixed effects, a 10% larger lagged migrant stock is associated with 8.3% higher subsequent flows. The coefficient decreases when we include origin-destination fixed effects in column (3). This is because migrant stocks are highly persistent over our relatively short panel, so o - d fixed effects

⁹The PPML estimates, reported in Table A1 in Appendix A, are similar to the OLS results.

absorb much of the identifying variation. The IV estimates in columns (4) and (5) correct for the attenuation in the OLS estimates.¹⁰ The IV estimate implies that a 10% increase in the migrant stock raises subsequent bilateral flows by 7.2%.¹¹

To illustrate the magnitude of the network effect, we ask how much of the variation in migration flows across destinations is associated with differences in existing diaspora stocks, after removing destination factors that makes some locations more attractive. For each origin-year, we use our elasticity to estimate the share of the variance in migration flows associated with differences in lagged diaspora stocks. The covariance decomposition shows that the network accounts for 59% of the remaining variation in log flow shares across destinations, after accounting for destination fixed effects.

2.3 Fact 2: Origins and their networks face similar climates

Having established that migrant networks shape destination choices, we next examine whether they also shape migrants' climate exposure across destinations. In this section, we focus on climate as measured by temperature. Average temperature is the standard measure in the damage-function literature, which estimates how warming affects economic activity (Dell et al., 2012; Burke et al., 2015; Nath et al., 2024). We report robustness to other measures of climate in Appendix A. The key takeaway is that migrants are located in destinations with similar climate as the origin. Countries that are more exposed to high temperatures also have diasporas concentrated in similarly exposed destinations.

Let T_t^ℓ denote the average temperature of the country ℓ in period t and let $s_t^{o,d}$ denote the share of migrants from origin o residing in destination d at time t ¹². Then, origin o 's diaspora's temperature is

$$\bar{T}_t^o \equiv \sum_{d \neq o} s_t^{o,d} T_t^d. \quad (2)$$

Intuitively, $\bar{T}_t^o$ measures the climate conditions experienced by the average member of origin o 's diaspora at time t . We plot this against the origin country's average temperature in Figure 1b. Appendix A also plots the full distribution of diaspora temperature conditional on origin temperature

Figure 1b shows a strong positive relationship between origin and diaspora temperature, with a correlation coefficient of 0.7. Migrants from hotter countries therefore tend to be concentrated in destinations that are also relatively hot. This suggests that the existing geography of migrant

¹⁰In addition to measurement error, OLS may be biased downward by time-varying bilateral factors, such as labor-market congestion or policy responses to larger migrant stocks.

¹¹This estimate is in line with related work. Beine et al. (2011) find an elasticity of 0.76 for 1990–2000 flows to OECD destinations relative to 1990 stocks, while Beine and Parsons (2015) estimate an OLS coefficient of 0.40. Stuart and Taylor (2021)'s estimates imply roughly 0.5 additional Black migrants and 0.1 additional white migrants per diaspora member over a five-year period during the Great Migration. Our estimate implies about 0.3 additional migrants per diaspora member over the same horizon.

¹²Such that $\sum_{d \neq o} s_t^{o,d} = 1$.

networks provides limited diversification from origin-country climate exposure. Appendix A shows that this relationship is robust across subsamples and alternative measures of climate exposure.

The positive correlation is consistent with gravity. Migration is concentrated toward nearby destinations, and climate is spatially correlated, so origins and diasporas have similar climates. Migrant networks are geographically concentrated near the origin, with differences by origin-country income. Migrants from low-income countries are closer to home: the average migrant from a low-income country lives 2,700 km from the origin, compared with 3,800 km for migrants from high-income countries. Appendix A reports the full distribution of migrant distances by origin-country income group. Nearby destinations are valuable when shocks are country-specific, which is consistent with the strong role of distance in many historical migration episodes, including refugee flows (Moore and Shellman, 2007).¹³ Climate exposure, however, is spatially correlated across neighboring countries. As a result, networks concentrated in nearby destinations may provide limited protection against climate change.

2.4 Fact 3: Migrant networks show some adaptation with warming

Average global temperatures have increased by 1.5°F from 1995 to 2020 (Rohde and Hausfather, 2020). A key question for understanding the role of networks is whether people have adapted to warming to date by migrating to cooler destinations and, in doing so, shifted diaspora networks toward countries with lower exposure. Our central finding is that migration patterns over the past three decades have partially offset the climate exposure of networks. We find that most changes in diaspora-network exposure reflect evolving climate conditions in existing destinations, but a reallocation of migrant networks toward safer destinations mitigated some of that increasing exposure.

To formalize this analysis and distinguish changes in $\bar{T}_t^o$ driven by warming in destinations versus shifts in the migrant network, we decompose the change in the diaspora exposure of an origin country, $\Delta\bar{T}^o = \bar{T}_{2020}^o - \bar{T}_{1995}^o$, as:


$$\Delta\bar{T}^o = \underbrace{\sum_{d \neq o} s_{1995}^{o,d} \Delta T^d}_{\text{warming-at-destinations}} + \underbrace{\sum_{d \neq o} \Delta s_{1995}^{o,d} T_{1995}^d}_{\text{network reallocation}} + \underbrace{\sum_{d \neq o} \Delta s_{1995}^{o,d} \Delta T^d}_{\text{interaction}} \quad (3)$$

The first term is the warming-at-destinations term that isolates how exposure changes if destination exposure evolves while network shares stay fixed at $s_{1995}^{o,d}$. We expect this term to be positive given the warming that has happened over this time period. The second term is a network reallocation term that isolates how exposure changes holding destination temperature at baseline levels, T_{1995}^d , while allowing network shares to shift from $s_{1995}^{o,d}$ to $s_{2020}^{o,d}$. If people are shifting migration patterns toward safer destinations, then this term should be large and negative. The interaction term captures whether the network is reallocating toward destinations whose exposure is changing

¹³Related work shows that the value of proximity depends on the spatial scale of risk; see Rosenzweig and Stark (1989) and Kazianga and Udry (2006).

the most. It is positive when the network shifts toward destinations that are warming more, and negative when they shift toward destinations that are warming less or cooling.¹⁴

Figure 2: Changes in Climate Exposure and Geography of Networks



Notes: Panel (a) decomposes the change in migrant-network climate exposure from 1995 to 2020 into network reallocation (blue), warming at destinations (orange), and their interaction (grey), globally and for origins in the bottom and top temperature quartiles. Panel (b) plots migrants' average distance from the origin over time for the same quartiles. Temperature data are from BEST; bilateral distance is from GeoDist (Mayer and Zignago, 2011).

Figure 2a plots this decomposition globally and for countries in the top and bottom quartile of our climate measure. Globally, the orange bar representing warming-at-destinations dominates. However, the blue bars representing network reallocation are negative and large. This suggests that migration has been to cooler regions. Network reallocation offset about 40% of warming, and this offset is larger for the hottest countries. This pattern could reflect networks becoming geographically dispersed over this time period. Figure 2b plots the average distance of migrants from their origin for countries in the hottest and coldest quartiles. For the hottest quartile, the average migrant is 550 km further from home in 2020 than in 1995, while the coldest quartile had little change.

The empirical evidence shows that migrant networks shape where people move and, in turn, the climate exposure of migrants' destinations. Migrants from hotter origins tend to have diasporas in similarly hot destinations, although there has been some reallocation toward cooler and more distant destinations over time. These facts describe the past role of networks, but our main question is how they will affect adaptation as climate change worsens. Answering this question requires

¹⁴We do not interpret the network-reallocation term as migration caused by climate change. In practice, it is difficult to define a "climate move," since migration may respond to events partly shaped by climate e.g. Kelley et al. (2015) link the 2007–2010 drought in Syria to the civil war. Additionally, a regression of migration on climate shocks may not recover a causal estimate since exposure to those shocks is non-random (Borusyak and Hull, 2023). Therefore, we use the decomposition to ask if the sum of all migration over the period changed the climate exposure of networks.

a model so that we can predict how migration and migrant networks will evolve under future warming, isolate the role of the network by constructing a counterfactual in which destination choices are not shaped by existing diaspora stocks, and translate differences in climate exposure into welfare. We turn to a dynamic model of migrant networks and climate change in Section 3.

3. Model of Migrant Networks and Climate Change

The findings in Section 2 suggest that migrant networks will shape adaptation to climate change. Networks can contribute to adaptation by helping move people out of the most climate-exposed regions. But they may also limit adaptation by directing migrants to destinations facing a similar climate. We build a model to formalize both channels and measure their net effect on welfare as the difference in climate damages between an economy with migrant networks and one without. Which channel dominates depends on how strongly the network pulls migrants, on how correlated damages are between an origin and its network's destinations, and on how the network shifts as warming proceeds.

The empirical findings dictate three new features absent from existing quantitative spatial models of climate change. First, new emigrants follow the existing network. In the model, the diaspora at a destination therefore enters household utility. Second, migrants settle in destinations whose climate resembles that of their origin. The value of a destination therefore depends on how its stochastic climate damages co-move with damages at home. Third, networks have shifted with recent warming. An origin's diaspora is therefore a state variable that updates with migration.

The rest of the economic environment is deliberately sparse. We abstract from trade in differentiated goods, from capital, from any fixed factor of production, from endogenous amenities, and from any market size effects. We also abstract from differences in household income within an origin, and we capture financial constraints on migration through a residence cost that falls with origin income. Including these features would change the magnitude of climate damages at a given location. Our results, however, are driven by how attractive a destination is to one origin relative to others, and a destination-wide force changes the level of damages in the same way with or without the network, leaving the interaction unchanged.

3.1 Set Up

We consider a world with N locations. Time is discrete.

Production. Competitive firms in each location ℓ produce a freely traded good with a linear technology in labor, $Y_t^\ell = A_t^\ell L_t^\ell$. The local wage is $w_t^\ell = A_t^\ell$. Productivity has three components,

$$A_t^\ell = A_0^\ell g_t^\ell z_t^\ell, \quad (4)$$

a time-invariant fundamental A_0^ℓ , a deterministic growth path g_t^ℓ , and proportional climate damages z_t^ℓ .

Climate. The climate at a location is the distribution of its temperature, and climate change is a known shift in that distribution. The temperature at location ℓ in period t is its mean plus a weather shock, $q_t^\ell = \bar{q}_t^\ell + \eta_t^\ell$. Across locations, mean temperature is the vector $\bar{\mathbf{q}}_t = (\bar{q}_t^1, \dots, \bar{q}_t^N)$, and the weather shocks are jointly normal with covariance $\mathbf{\Gamma}_t$. We assume that weather shocks are independent across periods. Climate change is then the sequence of weather distributions $\{\bar{\mathbf{q}}_t, \mathbf{\Gamma}_t\}_{t=0}^T$. We assume that after some date T the climate stabilizes at $(\bar{\mathbf{q}}_T, \bar{\mathbf{\Gamma}}_T)$.

Following the empirical climate literature, we model climate damages as a proportional output loss that depends on temperature (Dell et al., 2012; Burke et al., 2015; Nath et al., 2024). A damage function d maps temperature into log output. Damages may persist beyond the period of the shock (Newell et al., 2021; Nath et al., 2024), so they follow an AR(1) around a trend,

$$\log z_t^\ell = \log \bar{z}_t^\ell + \psi(\log z_{t-1}^\ell - \log \bar{z}_{t-1}^\ell) + \zeta_t^\ell, \quad (5)$$

with persistence $\psi \in (-1, 1)$. The trend is the damage at the mean temperature, $\log \bar{z}_t^\ell = d(\bar{q}_t^\ell)$, and the innovation ζ_t^ℓ is the damage from the weather shock, which we assume is normally distributed with covariance Σ_t .

Households. A household is characterized by its origin o , which is fixed for life. At $t = 0$, there is an initial population distribution for every origin country across locations. Let $L_t^{o,\ell}$ denote the number of households from origin o living in location ℓ at the start of period t . We abstract from population growth so the total population from o is constant over time and equal to L^o . We define $x_t^{o,\ell} = L_t^{o,\ell} / L^o$ as the share of the origin- o population living in ℓ , and call $\mathbf{x}_t^o = (x_t^{o,1}, \dots, x_t^{o,N})$ the migrant network of origin o .¹⁵

Each household supplies one unit of labor inelastically and earns the wage w_t^d in its chosen location d . Households purchase a tradable consumption good at a price normalized to one in all locations. With no saving and no borrowing, households consume their wage income. Households are forward-looking, discount the future at rate $\beta < 1$, and have rational expectations. Flow utility is $\log c_t^d + h(x_t^{o,d})$, so a household values both consumption at its chosen destination and the share of its own origin population living there.¹⁶ The function h is increasing, bounded above, and continuously

¹⁵Because the model is written in shares, any demographic process proportional to $L_t^{o,\ell}$ leaves $\mathbf{x}_t^o$, and hence all decisions, unaffected. We abstract from two demographic channels that could matter in practice. First, differential fertility across origins or with climate exposure (e.g., Doepke et al., 2023; Barreca et al., 2018) would make L^o time-varying and add a location-specific drift to the law of motion (9). Second, the assimilation of o -origin descendants into destination communities (e.g., Abramitzky et al., 2014; Bisin and Verdier, 2011) can be modeled as a per-period hazard with which a household exits the co-ethnic network, which erodes the effective network stock over generations. Both extensions preserve the origin-by-origin structure of the law of motion and neither channel is likely first order at the horizons we study.

¹⁶We assume the climate affects only productivity and not the amenity of living in a location. Cruz and Rossi-Hansberg (2024) estimate both effects and find they contribute about equally to welfare losses from warming. Because flow utility is log consumption, an amenity loss would enter exactly as a productivity loss of the same size in logs, so including it would

differentiable.

In addition, households have idiosyncratic preference shocks for each potential destination, $\epsilon_t^{o,d}$. We assume these shocks are i.i.d. across households, destinations, and time, and drawn from a Gumbel distribution normalized to have mean zero. The shape parameter $\nu > 0$ scales their dispersion and is common across origins. These assumptions imply a common elasticity of migration to utility differences.

Households from origin o also pay a utility cost $\kappa_t^{o,d}$ in every period they live in location d , with $\kappa_t^{o,o} = 0$, as in [Desmet et al. \(2018\)](#) and [Cruz and Rossi-Hansberg \(2024\)](#).¹⁷ We interpret the residency cost as capturing pecuniary and non-pecuniary costs associated with living away from a household's origin county. These include conversing in a foreign language, maintaining legal status, and living in a foreign culture with different institutions. The cost is additive in utility and independent of intermediate moves. This formulate of costs matches our data. The migrant stocks in Section 2 record where a household was born and where it lives now, not where it lived last period. No global data exists that would identify the cost of moving between two foreign countries.

Based on empirical evidence that financial constraints limit migration ([Bryan et al., 2014](#); [Angelucci, 2015](#); [Cattaneo and Peri, 2016](#); [Bazzi, 2017](#); [Aron-Dine, 2025](#); [Barbosa-Alves, 2025](#)), we allow emigration costs to fall with origin country income. We do not model household financial constraints explicitly. Instead we capture them in reduced form by letting residence costs fall uniformly for households from origin o as its economy grows:

$$\kappa_t^{o,d} = \kappa^{o,d} - \lambda \log(A_0^o g_t^o), \quad d \neq o. \quad (6)$$

Here $A_0^o g_t^o$ is origin country income absent climate change. A one percent rise in origin income lowers the cost of living abroad by λ percent of consumption. Because the cost depends on income absent climate change, a climate damage at home lowers the wage but does not make it harder to leave.¹⁸

3.2 Migration Decisions

In the model, migration is the household's only technology for adjusting its exposure to climate change. Where a household chooses to live depends on wages, the cost of living away from home, and the size of its own community at each destination.

raise the level of damages without changing their interaction with the network.

¹⁷[Desmet et al. \(2018\)](#) let the flow cost of moving from r to s factor into an origin and a destination component, $m(r, s) = m_1(r)m_2(s)$, with $m(r, r) = 1$. Because the costs of a household's successive moves multiply, the product becomes a cost of residing in d that depends only on the household's origin o . Our $\kappa^{o,d}$ is that object with the separability relaxed, so it need not be the difference between a destination term and an origin term.

¹⁸If the cost instead depended on realized origin income, a climate shock in o would raise the cost of leaving by λ times the damage, and the response of emigration to a damage would fall from $1/\nu$ to $(1 - \lambda)/\nu$. That is the trapped-population channel as in [Cattaneo and Peri \(2016\)](#) and [Barbosa-Alves \(2025\)](#). We leave it out for two reasons. A period in our model is five years, long enough for a household to save for a move after a bad year in line with findings from [Barbosa-Alves \(2025\)](#) that households save ahead of anticipated warming. Adding this channel would lower adaptation through migration but would not interact with the migrant network.

Based on evidence that the value of a network comes from migrants who are already established, we let flow utility depend on the network at the start of the period, $\mathbf{x}_t^o$. Established migrants provide information about jobs and housing (Munshi, 2003; McKenzie and Rapoport, 2007) and job referrals (Bayer et al., 2008). Migrants who arrive at the same time do not. Beaman (2012) shows this with the random assignment of refugees across US cities. Refugees have better labor market outcomes when assigned to cities with a larger established community from their own country, and worse outcomes the larger the cohort resettled alongside them.

In each period a household observes climate damages $\mathbf{z}_t$ in every location and its preference shocks ϵ_t^o , and chooses the location d where it will live and work. A household from origin o solves

$$v_t^o(\mathbf{x}_t^o, \mathbf{z}_t) = \max_d \left\{ \log c_t^d + h(x_t^{o,d}) - \kappa_t^{o,d} + \nu \epsilon_t^{o,d} + \beta \mathbb{E} \left[v_{t+1}^o(\mathbf{x}_{t+1}^o, \mathbf{z}_{t+1}) \mid \mathbf{x}_t^o, \mathbf{z}_t \right] \right\}. \quad (7)$$

The continuation value is the same at every destination, so it cancels from the Gumbel choice probabilities. The probability that a household from origin o chooses destination d is

$$\pi_t^{o,d} = \frac{\exp([\log c_t^d + h(x_t^{o,d}) - \kappa_t^{o,d}]/\nu)}{\sum_{i=1}^N \exp([\log c_t^i + h(x_t^{o,i}) - \kappa_t^{o,i}]/\nu)}. \quad (8)$$

Migration is reversible. A household pays $\kappa_t^{o,d}$ only as long as it lives in d , and a household that moves to d' pays $\kappa_t^{o,d'}$ instead. A household that leaves o and later returns pays nothing for having left. There is no sunk cost.¹⁹

Every household from origin o solves the same problem, because the choice probabilities do not depend on its current location. Next period's network is therefore the vector of choice probabilities,

$$x_{t+1}^{o,d} = \pi^{o,d}(\mathbf{x}_t^o, \mathbf{z}_t). \quad (9)$$

Households never need to forecast the network their own choices produce. Origins do not interact, because wages do not depend on who lives where and each household values only the share of its own origin at a destination. Households from the same origin do interact, through h , and a household that moves does not take into account that it raises the utility of every compatriot at its destination.

While the household's migration decision in the model is static, the welfare effects of climate change are not. The discount factor cancels from (8) because the continuation value is the same at every destination. The household's value function in (7) still contains the full expected path of wages, climate damages, and the network. Expected lifetime utility before the preference shocks

¹⁹A sunk cost of moving instead of a flow residency cost would add three things. First, households would move in anticipation of future climate damage, since the choice would depend on future as well as current payoffs. To date there is little evidence that households do so (Hoffmann et al., 2020; ?). Second, households that had already paid to move would be more likely to stay in a destination given a bad climate realization. Third, households would wait to see how bad a climate shock is before paying to move. The first would raise the value of the network as insurance and the second would lower it.

are drawn, V_t^o , solves

$$V_t^o = \nu \log \sum_{d=1}^N \exp\left([\log c_t^d + h(x_t^{o,d}) - \kappa_t^{o,d}]/\nu\right) + \beta \mathbb{E}[V_{t+1}^o \mid \mathbf{x}_t^o, \mathbf{z}_t]. \quad (10)$$

The welfare effects of climate change therefore depend on where the network will be as climate damages grow, on how correlated future weather shocks are across the destinations open to a household from origin o , and on how quickly the network adjusts.

Dividing (8) for destination d by the same expression for the origin country and taking logs delivers a gravity equation for migrant stocks,

$$\log\left(\frac{x_{t+1}^{o,d}}{x_{t+1}^{o,o}}\right) = \frac{1}{\nu} \left[\log\left(\frac{w_t^d}{w_t^o}\right) - \kappa_t^{o,d} + h(x_t^{o,d}) - h(x_t^{o,o}) \right]. \quad (11)$$

Three forces determine the size of the diaspora at d relative to the population still in the origin, and the migration elasticity ν converts each of them from utility into log odds. The first two are the standard gravity forces, the wage ratio w_t^d/w_t^o as the return to migration and $\kappa_t^{o,d}$ as the cost of living away from the origin. The third is the network.

3.3 Equilibrium

Given the structure of the model, the equilibrium path is unique independent of the strength of network preferences. However, network preferences determine the long-run population distribution that the economy reaches.

Definition 1 (Equilibrium). Given initial networks $\mathbf{x}_0^o$ and residence costs κ^o for all o , initial damages $\mathbf{z}_0$, fundamental productivities $\mathbf{A}_0$, growth paths $\{\mathbf{g}_t\}_{t \geq 0}$, climate paths $\{\bar{\mathbf{q}}_t, \Gamma_t\}_{t \geq 0}$, and a sequence of weather shocks $\{\boldsymbol{\eta}_t\}_{t \geq 1}$, an equilibrium is a sequence of damages $\{\mathbf{z}_t\}_{t \geq 1}$, wages $\{\mathbf{w}_t\}_{t \geq 0}$, and networks $\{\mathbf{x}_t^o\}_{t \geq 1}$ and choice probabilities $\{\pi_t^o\}_{t \geq 0}$ for each o such that in every period damages follow the climate process (5), wages equal local productivity (4), households optimize so that choice probabilities satisfy (8), and the population distribution is consistent with household choices, $\mathbf{x}_{t+1}^o = \pi_t^o$ for every origin.

Proposition 1 (Existence and Uniqueness). Given initial networks and damages, an equilibrium exists and is unique for any residence costs, any realization of weather shocks, and any strength of network preferences.

The full proof can be found in Appendix C. Intuitively, this period's wages and the inherited networks determine the household choice probabilities, and the choice probabilities determine next period's networks. Each step has a single solution. The unique equilibrium path can therefore be constructed period by period.

3.4 Network Dynamics, Steady States, and Stationary Distributions

The key object for the dynamics of the network is the response of migration to the size of the existing network. Differentiating (11), holding the share at home fixed, gives the elasticity of migration to the network,

$$\varepsilon_t^{o,d} \equiv \frac{\partial \log(x_{t+1}^{o,d}/x_{t+1}^{o,o})}{\partial \log x_t^{o,d}} = \frac{x_t^{o,d} h'(x_t^{o,d})}{\nu}. \quad (12)$$

The numerator is the local strength of the network preference. A one percent larger diaspora in d raises the utility of living there for households from o by $x_t^{o,d} h'(x_t^{o,d})$. The denominator is the dispersion of location tastes. When ν is large, a household's idiosyncratic attachment to a location outweighs differences in wages and networks, so few households move in response to a change in relative utility.

The migrant network makes the migration response to a shock persistent. Consider a one-time shock δ to the log relative wage of d that fades after one period. The direct response is δ/ν in the first period. Because this migration increases the size of the diaspora in d , it draws additional migrants in subsequent periods even after the shock has faded. After t periods, the log odds of living in d are still higher by $(\delta/\nu) \varepsilon^t$. When $\varepsilon < 1$ the effect of a one-time shock fades and the diaspora returns to the size supported by fundamentals. When $\varepsilon \geq 1$ the effect does not fade and so the size of the diaspora depends on the full history of past shocks rather than on fundamentals.

The migrant network also amplifies the migration response to a permanent change in fundamentals. Consider a permanent change δ in the log relative wage of d . The direct response is again δ/ν in the first period, but the change is now permanent, so it raises the log odds of living in d by δ/ν in every period. The direct response increases the diaspora in d and draws additional migrants in every period. The long-run change in the log odds of living in d is the sum $\sum_{t=0}^{\infty} (\delta/\nu) \varepsilon^t$. When $\varepsilon < 1$ the sum converges and the network multiplies the direct response to a permanent change by $1/(1 - \varepsilon)$. When $\varepsilon \geq 1$ the sum does not converge, the diaspora in d continues to grow.

Persistence and amplification matter for the effectiveness of migration as an adaptation strategy and for the role of migration policy. Amplification through the migrant network determines the magnitude of the migration response to climate change. Persistence determines whether a temporary policy, such as a border closure, has a permanent effect on where people live. Both results hold only locally, since they fix the elasticity at the current network. To characterize the long-run location of the migrant network, we first describe the steady states of the law of motion (9) in the absence of weather shocks.

Definition 2 (Steady State). Hold the climate fixed at its mean, $\eta_t = 0$ for all $t > T$, and the growth path for productivity in every location at $\bar{g}$. A steady state for origin o is a network $\mathbf{x}^o$ that solves $\mathbf{x}^o = \pi^o(\mathbf{x}^o, \bar{z}_T)$.

Proposition 2 (Unique Steady State). A steady state exists for every origin. If the elasticity of migration to the network is below one, $\varepsilon_t^{o,d} < 1$ for every $x_t^{o,d} \in [0, 1]$, the steady state is unique and the network converges to it from any initial population distribution. If at a steady state $\varepsilon^{o,d}(1 - x^{o,d}) >$

1 for some destination d , that steady state is unstable and (9) has more than one steady state.

The full proof can be found in Appendix C. Intuitively, when the elasticity is below one everywhere, two economies with the same fundamentals but different initial networks converge to the same steady state, at a rate equal to the elasticity. When the elasticity is above one at a steady state, two economies with the same fundamentals but slightly different initial networks move apart, and converge to different steady states. In this case, a small group of migrants who move away from the steady state generates sustained additional migration, so a network near that steady state moves away from it rather than converging back.

We now characterize the long-run location of the migrant network with weather shocks. Even after the climate has stabilized, weather shocks continue to generate migration. For example, a dry year in Ghana induces Ghanaians to migrate to Nigeria. The larger Ghanaian community makes Nigeria a more attractive destination in subsequent years even after the drought ends, so the share of Ghanaians in Nigeria falls only gradually. In a stable climate with weather shocks, the network never settles at a steady state. Instead, the distribution of the network converges. The analogous question is whether that distribution is unique.

Definition 3 (Stationary Distribution). Hold the climate fixed at $(\bar{q}_T, \Gamma_T)$ for all $t > T$ and the growth path at $\bar{g}$. A stationary distribution for origin o is a distribution μ^o over networks such that if $\mathbf{x}_t^o$ is drawn from μ^o and damages follow (5), then $\mathbf{x}_{t+1}^o$ is also distributed according to μ^o .

Proposition 3 (Unique Stationary Distribution). A stationary distribution exists for every origin. If the elasticity of migration to the network is below one, $\varepsilon_t^{o,d} < 1$ for every $x_t^{o,d} \in [0, 1]$, the stationary distribution is unique and the distribution of the network converges to it from any initial population distribution.

The full proof can be found in Appendix C. Intuitively, when the elasticity is below one everywhere, two economies with the same fundamentals and the same sequence of weather shocks but different initial networks converge to the same network. The distribution of the network after many periods therefore does not depend on the initial network, and the stationary distribution is unique. When the elasticity exceeds one at some steady state, the same two economies do not necessarily converge to the same network distribution. Each economy remains near the closest stable steady and only a sequence of unusually large weather shocks moves it to another. The distribution of the network can therefore depend on the initial network.

Propositions 2 and 3 are the counterparts of the uniqueness results in the spatial equilibrium literature. Allen et al. (2020) and Kucheryavyy et al. (2023) give conditions under which a static spatial equilibrium is unique by comparing the strength of agglomeration with the dispersion of tastes. Allen and Donaldson (2020) provide the dynamic version for a single perturbation to a deterministic system. None of these economies have persistent risk, so the object their conditions pin down is a single allocation of households across locations. In our model with climate risk, the object is a distribution of networks, and the analogous condition determines whether that distribution is unique. We return to the policy implications in Section 6.

4. Model Quantification

To quantify the interaction between climate change and the migrant network we need to discipline three objects, economic growth absent warming, the spatial distribution of climate risk, and the strength of the network in driving migration decisions. We take the growth path from the no-warming baseline of the shared socioeconomic pathways. We parameterize the mean and the spatial covariance of the damage process by combining the temperature response function of [Nath et al. \(2024\)](#) with output from many runs of a climate model. Finally, we estimate the migration elasticities and residency costs to match the evolution of migrant networks from 1995 to 2015. Table 2 below summarizes the model inputs. Additional details can be found in Appendix D.

Table 2: Inputs to the Quantitative Model

Parameter	Description	Value	Source
<i>Panel A. Economic growth absent climate change</i>			
A_0^ℓ	productivity level		GDP per capita in 2019, Penn World Table
g_t^ℓ	growth path		SSP3, OECD ENV-Growth
<i>Panel B. Climate damages</i>			
$\bar{q}_t^\ell$	expected temperature path		MPI-GE mean, observed level
Γ_t	weather shock covariance		MPI-GE deviations from the mean
$d'(\bar{q}_t^\ell)$	marginal damage, percent per degree		Nath et al. (2024)
$\bar{z}_t^\ell$	expected damage path		Nath et al. (2024) , integrated along $\bar{q}_t^\ell$
ψ	damage persistence	0.66	Nath et al. (2024)
<i>Panel C. Migration elasticities and costs</i>			
β	discount factor, per period	0.90	2 percent annual rate
$1/\nu$	migration wage elasticity	0.70	Ortega and Peri (2013)
θ_1	strength of the network preference	0.90	Section 4.3
θ_0	network floor	$1e - 08$	Section 4.3
$\kappa_t^{o,d}$	residence costs		Section 4.3

4.1 Economic growth absent global warming


To understand how migrant networks interact with climate damages we first need to predict how migrant networks would evolve absent climate change. Countries will grow at different rates whether or not the climate changes, and economic growth will change relative wages between countries and generate migration. Some households in the countries forecast to lose from climate change would have migrated to the countries forecast to gain regardless of climate change. Without a growth path in the counterfactual, the model would classify migration that would have happened anyway as adaptation.

We take the growth projections from the shared socioeconomic pathways. These are the scenarios underlying much of the climate literature ([Burke et al., 2015](#); [Carleton et al., 2022](#); [Nath et al., 2024](#)). The five pathways differ in their assumptions about population growth, investment in

education, and how fast poorer countries catch up to richer ones (Riahi et al., 2017).²⁰ We use the growth rates implied by the income projections of the OECD’s ENV-Growth model (Dellink et al., 2017; International Institute for Applied Systems Analysis, 2025), which map directly into g_t^ℓ in (4). We use SSP3, a slow growth and catch-up scenario, as our baseline.²¹

Under these projections the return to migration falls over the century, even absent climate change. Figure 3 plots the distribution of the destination wage relative to the origin wage over every pair of countries, in 2019 and in 2095, splitting origins at the observed median of GDP per capita in 2019. For the poorer half of origins the median wage ratio falls from 2.7 to 2.0, and the ninetieth percentile by more, from 15 to 7.9. Under SSP3 poorer countries grow at 1.6 percent a year, compared to 0.7 percent for richer ones. A key driver of migration in the model is relative growth, not the average rate of global economic growth.

Figure 3: The Evolution of the Distribution of Relative Wages



Notes: The figure plots histograms of the destination wage relative to the origin wage, over all pairs of countries. The vertical axis counts country pairs. Wages are the no-warming path $g_t^\ell A_0^\ell$ of (4) under SSP3. A_0^ℓ comes from observed per-capita GDP in 2019 from Penn World Tables and g_t^ℓ the projected growth from SSP3. Origins are split by median of GDP per capita in 2019. The horizontal axis is on a log scale, the dashed vertical line marks a ratio of one.

4.2 Climate damages

We forecast climate damages by combining an empirical damage function with the output of a climate model.

Deriving the damage process from weather. A damage function $d(q)$ maps a country’s temperature q into the climate component of its productivity, $\log z = d(q)$. The slope of the damage function, $d'(q)$, is the output loss in percentage points from an additional degree of warming.

²⁰There are two main alternative approaches to long-run growth prediction. The first is probabilistic. Müller et al. (2022) fit a Bayesian factor model to a century of growth in 113 countries and report a distribution instead of a set of pathways. However, a century of data can only weakly identify century-long persistence, so the dispersion of that distribution depends heavily on the prior about how persistent growth is. The second is structural. Models in the tradition of Desmet et al. (2018) generate growth inside the model rather than imposing an external path, but that path then depends heavily on assumptions about local innovation and diffusion. Every approach to predicting economic growth through the end of the century must make strong assumptions.

²¹Appendix D.1 compare growth and catch-up across the five pathways. We report our main results under the four alternative growth pathways in the Appendix.

Projections for the path of warming for each country come from a climate model. Climate models divide the globe into a grid and simulate how heat, moisture and momentum move between cells of the atmosphere and the ocean over time, under a given path for greenhouse gas emissions. The climate system is chaotic, meaning that two model runs with the same emissions path but started from infinitesimally different initial conditions produce different weather paths for the same warming. We therefore write temperature in location ℓ and period t as an expected path and a weather shock around it,

$$q_t^\ell = \bar{q}_t^\ell + \eta_t^\ell, \quad (13)$$

where the shocks η_t are mean zero with spatial covariance Γ_t and independent across periods.²²

Both moments of the damage process follow from applying the damage function to (13). Because (5) is written in deviations from trend, the expected path of climate damages is the trend, $E_0[\log z_t^\ell] = \log \bar{z}_t^\ell$, for every t and any persistence ψ . Trend damages in any location are simply the damage function evaluated at the expected temperature,

$$\log \bar{z}_t^\ell = d(\bar{q}_t^\ell). \quad (14)$$

This is the standard measure of climate damages, a country's GDP per capita under warming relative to its GDP per capita without warming. Global warming moves a country along its damage function, so the marginal damage from the next degree of warming, $d'(\bar{q}_t^\ell)$, is evaluated at its current temperature.

There are many empirical estimates of this object for a specified warming path (Burke et al., 2015; Carleton et al., 2022; Bilal and Känzig, 2024; Nath et al., 2024), as well as of the ambiguity in those estimates across economic (Newell et al., 2021; Rennert et al., 2022) and climate models (Burke et al., 2018). However, there are no estimates of climate uncertainty given a specific climate model and warming path. More precisely, for our setting, we need to quantify the variance of damages around their expected path as well as the covariance of those deviations across countries. Current empirical approaches cannot deliver an estimate of Σ_t . The standard panel specification identifies the response of a country's output to its own weather, and its year fixed effects absorb any component common across countries, which is exactly the comovement we need. Identifying damages from a single global temperature shock, as in Bilal and Känzig (2024), recovers comovement, but any two countries' damages are then perfectly correlated.²³ We therefore construct Σ_t from climate model output rather than estimating it from historical data.

The spatial covariance of damages, Σ_t , follows from the spatial covariance of weather, Γ_t , and

²²Temperature is persistent and spatially correlated in levels. That correlation structure is embedded in the expected temperature path $\bar{q}_t^\ell$. What remains is the deviation from the mean warming path, η_t^ℓ . In observed temperatures over 1970 to 2019, the deviation of a country's annual temperature from its trend has a lag-one autocorrelation of 0.14, and a process with that persistence implies a correlation between adjacent five-year means of 0.03. At the frequency of the model, weather shocks are independent across periods.

²³Dingel et al. (2019) show that comovement is identifiable from ENSO events, which change the spatial distribution of weather anomalies, and use it to show that comovement in agricultural productivity strengthens under warming. Their measure summarizes comovement across all countries in a single index, while the value of a migrant network depends on the covariance between a specific origin and each of its destinations.

the slope of the damage function in each location. Substituting (13) into the damage function and taking a first order expansion around the expected temperature, the shock to climate damages in (5) is then simply the marginal damage times the weather shock, $\zeta_t^\ell = d'(\bar{q}_t^\ell) \eta_t^\ell$.²⁴ Since the slope is known at $t-1$, the damage shock inherits the properties of the weather shock. It is mean zero, and its covariance across locations i and j is the product of the two marginal damages and the covariance of their weather,

$$\Sigma_t^{ij} = E_{t-1}[\zeta_t^i \zeta_t^j] = d'(\bar{q}_t^i) d'(\bar{q}_t^j) \Gamma_t^{ij}. \quad (15)$$

Both slopes are evaluated at the expected temperature in period t , so the covariance of damages changes over the century as countries warm, even if the covariance of weather does not.

Correlated weather does not imply correlated damages. The correlation of climate damages for any two locations i and j is a function of their weather correlation and the slopes of their damage functions:

$$\rho_t^{ij} = \frac{\Sigma_t^{ij}}{\sqrt{\Sigma_t^{ii} \Sigma_t^{jj}}} = \frac{d'(\bar{q}_t^i) d'(\bar{q}_t^j)}{|d'(\bar{q}_t^i)| |d'(\bar{q}_t^j)|} \cdot \frac{\Gamma_t^{ij}}{\sqrt{\Gamma_t^{ii} \Gamma_t^{jj}}} = \text{sign}\left(d'(\bar{q}_t^i) d'(\bar{q}_t^j)\right) \cdot \frac{\Gamma_t^{ij}}{\sqrt{\Gamma_t^{ii} \Gamma_t^{jj}}}. \quad (16)$$

When the marginal damage from warming has the same sign in both locations, a shared weather shock moves their output in the same direction and the correlation of damages equals the correlation of weather. Because the marginal damage is linear in temperature, the damage function has a single peak, and two locations on the same side of it have marginal damages of the same sign, so a warm year is a cost to both or a gain to both. Locations on opposite sides of the peak have marginal damages of opposite sign, so the same weather shock raises output in one and lowers it in the other, and their damages are negatively correlated even when their weather is highly correlated. A location with zero marginal damage suffers no economic impact from a weather shock, so its damages are uncorrelated with those of all other locations.

Mapping to data. We take our damage function, $d(\bar{q}_t^\ell)$, from [Nath et al. \(2024\)](#), who estimate the response of output to temperature as a function of a country's baseline temperature. Their specification is the empirical counterpart of our climate process specified in (5). They identify the effect of an annual temperature deviation, the weather shock in (13), on GDP per capita using local projections, and find that a temperature shock lowers GDP per capita for at least a decade. Because of their relatively short sample, they use a simple model of catch-up to recover the persistence of damages by indirect inference. In their model a country's productivity is a weighted average of the world frontier and its own past, so a temperature shock lowers productivity below the frontier and each year the country closes a fraction of that gap equal to the weight on the frontier. They choose the weight to match long-run growth across countries and their local projection estimates.

²⁴To first order, $d(q_t^\ell) = d(\bar{q}_t^\ell + \eta_t^\ell) \approx d(\bar{q}_t^\ell) + d'(\bar{q}_t^\ell) \eta_t^\ell$. Households know $\bar{q}_t^\ell$ at $t-1$ and $E_{t-1}[\eta_t^\ell] = 0$ by (13), so taking expectations, $E_{t-1}[d(q_t^\ell)] \approx d(\bar{q}_t^\ell)$, which is the trend (14). The shock is the realized value less its expectation, $\zeta_t^\ell = d(q_t^\ell) - E_{t-1}[d(q_t^\ell)] \approx d'(\bar{q}_t^\ell) \eta_t^\ell$. It is mean zero, and since $d'(\bar{q}_t^\ell)$ is known at $t-1$, $E_{t-1}[\zeta_t^i \zeta_t^j] = d'(\bar{q}_t^i) d'(\bar{q}_t^j) E_{t-1}[\eta_t^i \eta_t^j]$, which is (15). The damage function we use is linear in the weather shock, so the expansion is exact.

Taking their estimate implies that the deviation from trend in (5) decays by 0.92 each year and by $\psi = 0.92^5 = 0.66$ each period in our model.

The expected damage path integrates their damage function along the expected temperature path. They estimate a marginal damage that is linear in a country’s temperature, $d'(q) = \beta_0 + \beta_1 q$. Their estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ imply that the marginal damage is positive in cold countries and negative in hot ones, so the damage function has a single peak. Both the expected temperature path and the weather covariance come from many runs of the same climate model under the same emissions path. We use the Max Planck Institute Grand Ensemble, which runs the MPI-ESM1.1 climate model 100 times under the RCP 4.5 emissions path from slightly different initial conditions (Maher et al., 2019). The expected temperature path, $\bar{q}_t^\ell$, is simply the average across the N runs. Since the marginal damage depends on the level of temperature and climate models are known to have level bias, we set the initial level of $\bar{q}_t^\ell$ from observed temperature and take only its changes from the ensemble estimates. We then cumulate damages in 0.1°C increments along that path, evaluating the marginal damage at each increment, and because a damage persists, the marginal damage that enters (14) is their cumulative response to a permanent degree of warming, 5.05 times their annual response.²⁵

The weather covariance comes from the spread of the same runs around their average. The weather shock in each run is its deviation from that average, and the weather covariance is the sample covariance of those deviations across countries,

$$\eta_t^{\ell,n} = q_t^{\ell,n} - \bar{q}_t^\ell, \quad \mathbf{\Gamma}_t = \frac{1}{N-1} \sum_{n=1}^N \eta_t^n (\eta_t^n)^\top, \quad (17)$$


where $q_t^{\ell,n}$ is the five-year mean temperature in run n . A multi-model ensemble would not work, because each model contributes one run and the spread would then mix weather shocks with disagreement between models about the climate process. We pool the ensemble into two eras, every period to 2050 and every period after it, and estimate one weather covariance for each. A single period’s estimate from 100 runs is too noisy to track the weather covariance over time, and fifty years is the shortest window over which the one change the ensemble shows, a fall in the mean weather correlation of about 0.02 after mid-century, stands out from the noise. The marginal damage that enters (15) is their annual response times 4.26, the accumulated effect of a five-year shock of which the first year’s damage has partly faded by the fifth.

Forecasting damages. Most countries are forecast to lose from global warming, and the hottest countries today are forecast to lose the most. Figure 4(a) maps the expected damage $\bar{z}_t^\ell$ in 2095. The

²⁵Nath et al. (2024) estimate long-run impacts of climate change by integrating their damage function along a warming path constructed by rescaling a multi-model RCP 8.5 projection to RCP 4.5. This is a common approach when only one scenario is available, but it does not provide the weather variation we need to construct Σ_t . We therefore integrate their damage function along our own warming path, so that the expected damage path and the covariance of damages are built on the same climate. Applied to their warming path, our construction matches their projected damage per degree of warming. Appendix D.2 provides additional details.

median country loses 13.0 percent of output, with losses of up to 28 percent across the tropics and gains of up to 29 percent at high latitudes. The countries that lose most are the ones that are hot already, because the marginal damage increases with temperature. Figure 4(b) plots every country's path $\bar{z}_t^\ell$ from 2020 to 2095. The median loss is 4.3 percent by 2050 and 13.0 percent by 2095. Warming widens the gap between winners and losers without changing who they are.

Figure 4: Expected Damages



Notes: Both panels plot $\bar{z}_t^\ell$, the expected damage of (14), in percent of GDP per capita relative to the no-warming path. Panel (a) shows a choropleth map of the level in 2095. Panel (b) plots every country's path from 2020 to 2095 in grey and the median across all countries in blue. Warming paths come from the ensemble mean of the MPI Grand Ensemble under RCP 8.5 and the damage function is that of Nath et al. (2024). The map uses the Equal Earth projection of Šavrič et al. (2019).


Weather risk to output rises with the marginal damage, so it is largest in the hottest countries. Figure 5(a) maps the standard deviation of the five-year damage shock, $\sqrt{\Sigma_t^{\ell\ell}}$, in 2095. For the median country it is 0.8 percent of GDP, and it is 1.4 percent in the hottest decile of countries and 0.3 percent in the coldest. Figure 5(b) plots the same standard deviation against the expected damage. Risk is highest among the countries forecast to lose the most, because a steeper damage function turns the same warming into a larger expected loss and the same weather into a larger shock.

Our forecast losses from climate change are large. The population-weighted world loss in 2095 is 14.2 percent of GDP and the loss in Sub-Saharan Africa is 18 percent.²⁶ For the median country the loss from climate change equals 12 years of growth under SSP3. The countries that lose the most in expectation also face the most risk. The weather risk is 0.8 percent of GDP over five years for the median country, but 1.4 percent in the hottest decile and up to 2.9 percent.

Climate damages are weakly correlated across countries, and the correlation falls with distance. Figure 6(a) plots the distribution of the correlation of the five-year damage shock across all pairs of countries in 2095. The mean is 0.14, 37 percent of pairs are negatively correlated, and only 14 percent have a correlation above one half. Figure 6(b) plots the mean correlation in twenty equal-count bins of the distance between the two countries. It is 0.45 in the closest bin, falls to 0.06 by 4 thousand kilometres, and stays there until it rises for the most distant pairs, which sit on

²⁶Under the same emissions scenario Nath et al. (2024) project a loss of 16 percent for the world and 21 percent for Sub-Saharan Africa. The gap comes from the warming path, not the damage function. They cumulate damages from 2010 along a linear path to 4.3 degrees of warming for the median country, while we start in 2020 and trace the ensemble mean to 3.2 degrees. Because the marginal damage steepens as a country warms, the last degree is the most costly one.

Figure 5: The Standard Deviation of Damages




Notes: Both panels plot $\sqrt{\Sigma_t^{\ell\ell}}$, the standard deviation of the five-year damage shock of (15), in percent of GDP, in 2095. Panel (a) shows a choropleth map. Panel (b) plots it against the expected damage $\bar{z}_t^\ell$ of Figure 4, with a dashed line at zero. The weather covariance is that of the second era of the MPI Grand Ensemble under RCP 8.5, 2050 to 2095, and the damage function is that of Nath et al. (2024). The map uses the Equal Earth projection of Šavrič et al. (2019).

opposite sides of the earth.

The correlation of damages is inherited from the correlation of weather. Neighboring countries are often governed by the same weather systems, so a hot year in one country is a hot year in its neighbors. The weather correlation is 0.75 in the closest distance bin and decays with distance. The sign of the damage correlation depends on the damage function as well as the weather correlation. A hot year lowers output in a country above the peak of the damage function and raises it in a country below. Two countries with positively correlated weather on opposite sides of the peak therefore have negatively correlated damages. This is why 37 percent of pairs have negatively correlated damages although only 33 percent have negatively correlated weather.

Figure 6: The Correlation of Damages



Notes: Both panels plot the correlation of the five-year damage shock, $\Sigma_t^{ij} / \sqrt{\Sigma_t^{ii} \Sigma_t^{jj}}$ from (15), in 2095, over every pair of countries. Panel (a) plots the distribution of correlations in 2095. Panel (b) splits the pairs into 20 equal-count bins by the great-circle distance between the two countries' population-weighted centroids and plots, at each bin's mean distance, the mean correlation of damages and the mean correlation of weather across the pairs in the bin.

4.3 Migration elasticities and costs

Three objects govern the household's migration decision, the dispersion of location tastes ν , the network preference, and the bilateral residence costs κ . We estimate the network preference and the bilateral residence costs by fitting the model's law of motion (9) to bilateral migrant stocks observed every five years, and set the dispersion of location tastes from the literature. We parameterize the network preference as

$$h(x) = \theta_1 \log(x + \theta_0), \quad (18)$$

where θ_1 governs the strength of the network preference and θ_0 the size below which a diaspora community stops mattering. Given ν and the network parameters, the gravity equation (11) recovers the residence cost of every migration corridor from the observed migrant stocks. We therefore select the network parameters to best predict each period's stocks from the previous period's.

Identification. The elasticities of migration to wages and to the network can only be identified from variation within a migration corridor over time. Anything about a corridor that is constant over time is absorbed by its residence cost $\kappa^{o,d}$, independent of the values of θ_0 , θ_1 , and ν . Therefore, in a single cross-section, a migration corridor that is large because its residence cost is small is indistinguishable from one that is large because its diaspora community is large. Over time, however, the two corridors differ, because all else equal a larger diaspora community draws more migrants while a smaller residence cost does not. This response of a corridor's migrant network to the existing migrant stock is what identifies the network parameters. In the model, that response is the network elasticity (12), which under (18) equals $\theta_1 x / (x + \theta_0) / \nu$, constant at θ_1 / ν for corridors larger than θ_0 and falling to zero below it. The elasticity of large corridors identifies θ_1 / ν , and whether the smallest corridors respond less than large ones bounds θ_0 .

Wages do not vary enough within a migration corridor over time to identify the elasticity of migration to wages. In our relatively short data sample, we only have five observations per migration corridor. Over those observations, destinations that are getting richer are also the ones whose diaspora communities are growing. A rising wage pulls migrants from every origin into a destination by the same amount, while a growing diaspora community pulls migrants only from its own origin, in proportion to its size. The data therefore identify the network parameters but not the dispersion of location tastes, which is what governs the response of migration to a change in wages. We instead take $1/\nu = 0.70$ from [Ortega and Peri \(2013\)](#), who estimate the elasticity of bilateral migration flows to destination income per capita over a longer panel.

Estimation. Given ν , the network parameters, and the observed migrant networks and wages, the gravity equation (11) pins down the residence cost of every open migration corridor. For a corridor observed in t and $t + 1$, the residence cost is

$$\kappa^{o,d} = \log \frac{w_t^d}{w_t^o} + \theta_1 \log \frac{x_t^{o,d} + \theta_0}{x_t^{o,o} + \theta_0} - \nu \log \frac{x_{t+1}^{o,d}}{x_{t+1}^{o,o}}, \quad (19)$$

with $\kappa^{o,o} = 0$. The residence cost is constant over time in the model, so we average (19) over the five transitions we observe for each corridor. An empty migration corridor only provides a lower bound on its residence cost, $\kappa^{o,d}$. We treat a corridor that is empty in 2020 as closed in the counterfactuals.²⁷

We estimate the network parameters by minimizing the distance between the migrant stocks the model predicts one period ahead and the stocks we observe. For a candidate pair of parameter values, (θ_0, θ_1) , we recover residence costs and form the choice probabilities from the observed migrant stocks and wages in t . We use the Poisson deviance as our distance metric to compare predicted, $\hat{x}_{t+1}^{o,d}$, and observed, $x_{t+1}^{o,d}$, migrant stocks:

$$\mathcal{D}(\theta_0, \theta_1) = 2 \sum_{o,d,t} \left[x_{t+1}^{o,d} \log \left(\frac{x_{t+1}^{o,d}}{\hat{x}_{t+1}^{o,d}} \right) - (x_{t+1}^{o,d} - \hat{x}_{t+1}^{o,d}) \right]. \quad (20)$$

A distance metric based on squared errors in levels would be dominated by the few largest corridors, and one based on squared errors in logs would overweight the smallest and drop every corridor that empties. Under the Poisson deviance each corridor enters weighted by its size. A ten percent miss on a corridor of a million people therefore costs a thousand times more than the same miss on a corridor of a thousand, and a corridor that empties stays in the sample.²⁸ We search over θ_1 on a grid and repeat for values of θ_0 from 10^{-9} to 10^{-2} , since θ_0 only matters for corridors of the same order of magnitude.

Estimates. We estimate a migrant network elasticity of 0.63. Given our choice of functional form, the elasticity does not depend on the size of the migrant stock for any corridor larger than θ_0 . The data support this assumption. Figure 7(a) splits migration corridors into five groups by size and, within each group, regresses a corridor's migrant stock on its own stock five years earlier with corridor and origin-year fixed effects. The regression is the gravity equation (11), with the residence cost $\kappa^{o,d}/\nu$ as the corridor fixed effect and the origin's wage and home network in the origin-year fixed effect, so the coefficient on the lagged stock is the network elasticity θ_1/ν . The model estimate is within the confidence interval across several orders of magnitude of corridor size.

An elasticity below one means that a shock to a migration corridor fades over time. Consider a migration corridor whose migrant stock is one percent above its steady state, the size its wages and residence cost support. The larger diaspora community draws more migrants in the next period, but only enough to keep the corridor θ_1/ν percent above its steady state, and the period after that $(\theta_1/\nu)^2$ percent. Given our estimate, the half-life of a shock to the network is 7 years, and 15 percent of a one percent shock remains after twenty years. For the same reason, no corridor is in the region of Proposition 2, where the elasticity exceeds one and more than one long-run migrant network


²⁷Corridor formation is small in the data. Of the 23,303 corridors that are empty in 1995, 1,008 hold any migrants by 2020, with a median of 22 people and 0.4 percent of the migrant stock between them.

²⁸This metric is twice the difference between the Poisson log-likelihood at the best possible fit, $\hat{x} = x$, and at the model's prediction. Near the best fit, each term is $(x_{t+1}^{o,d} - \hat{x}_{t+1}^{o,d})^2 / (2\hat{x}_{t+1}^{o,d})$, a squared error weighted by one over the predicted stock. It is the criterion behind PPML, the standard estimator for gravity equations of trade and migration (Santos Silva and Tenreyro, 2006; Beine et al., 2016).

distribution exists.

Although a one-time shock to the network fades over time, the network is a source of amplification. A permanent change in a corridor's wages or residence cost moves its migrant stock nearly three times as much as it would in a model without the network. Figure 7(b) plots the long-run size of a migration corridor against its residence cost net of the wage gap, $\kappa^{o,d} - \log(w_t^d/w_t^o)$, the corridor's fundamentals. In the model without the network the slope is $-1/\nu$, the direct response of the corridor's migrant stock to its cost. In the baseline model with the network effect the slope is instead $-1/(\nu - \theta_1)$. A permanent bad shock to a corridor's fundamentals shrinks its migrant stock directly, and the smaller diaspora community then draws fewer migrants in every subsequent period. The migrant stock declines each period by θ_1/ν of its decline in the prior period, so the long-run impact is $1/(1 - \theta_1/\nu) = 2.7$ times the direct decline.

Figure 7: The Network Elasticity



Notes: Panel (a) plots the network elasticity by quintile of corridor size. Each point is the coefficient on a corridor's lagged log odds against home, $\log(x_t^{o,d}/x_t^{o,o})$, in a regression of its log odds five years later on the lagged log odds interacted with the size quintile, the wage gap, and corridor and origin-year fixed effects, on the 44,912 corridor-transitions of 1995 to 2020. Size is the corridor's mean share of the origin population over the sample. Bars are 95 percent confidence intervals with standard errors clustered by origin and by destination. The orange line is the network elasticity θ_1/ν at the estimated parameters. Panel (b) plots the model's long-run relation between the size of a corridor, $\log(x_t^{o,d}/x_t^{o,o})$, and its residence cost net of the wage gap, $\kappa^{o,d} - \log(w_t^d/w_t^o)$, through the mean of the 8,877 corridors open in 2020. The blue line has slope $-1/(\nu - \theta_1)$ and the orange line, the model with the network held fixed, slope $-1/\nu$.

Our estimates of the network elasticity and of the residence costs are reasonable. The network elasticity of 0.63 sits within the range that gravity regressions of bilateral migration on the diaspora stock report (Beine et al., 2011; Beine and Parsons, 2017). The residence costs rise with distance and fall with every tie between two countries. Table 3 projects the recovered residency costs on bilateral frictions with origin and destination fixed effects. A corridor's cost rises by 0.67 log points with each log point of distance, falls by 0.95 for a colonial tie, and rises by 0.34 where a visa is required, and the frictions explain 0.88 of the variation in costs across corridors. Through the amplification of the previous paragraph the distance coefficient implies a distance elasticity of migrant stocks of 1.3, in the range the gravity literature reports.

Table 3: The Residence Costs and Bilateral Frictions

	(1)	(2)
Log distance	0.67 (0.01)	0.64 (0.01)
Contiguous	−0.54 (0.05)	−0.50 (0.05)
Common official language	−0.45 (0.03)	−0.44 (0.03)
Colonial tie	−0.95 (0.05)	−0.94 (0.05)
Common coloniser	−0.69 (0.04)	−0.61 (0.04)
Same country in the past	−0.29 (0.06)	−0.27 (0.06)
Visa required		0.34 (0.03)
Migration corridors	8,560	8,415
R^2	0.88	0.88

Notes: Each column is a regression of the residence cost $\kappa^{o,d}$ recovered by (19), in log wage points, on bilateral frictions, over the migration corridors open in 2020. Distance is the population-weighted bilateral distance of CEPII. Contiguity, common official language, colonial tie, common coloniser and same country in the past are the indicators of Mayer and Zignago (2011). Visa required is the share of the 1995 to 2010 vintages of the DEMIG visa database in which nationals of the origin needed a visa to enter the destination, and column (2) drops corridors it does not cover. Both columns include origin and destination fixed effects, which are not shown. Standard errors are in parentheses.

5. Migrant Networks and Climate Losses

In this section, we show that migrant networks increase the welfare cost of climate change for most countries. This is because networks anchor households to their origin and to destinations where their diaspora is already large, which are places that experience substantial damages from warming. Networks also reduce the insurance value of migration because they tilt destination choices to locations whose climate shocks move together. We establish these results in three steps. We first define the effect of the migrant network by comparing climate losses with and without the influence of existing diasporas on migration decisions. We then quantify how this effect varies across countries and show that it is largest for poor origins facing high climate damages. Finally, we decompose the network effect into an anchoring channel, which operates through exposure to the warming trend, and a diversification channel, which operates through the covariance of climate shocks across destinations.

5.1 Measuring the effect of the migrant network on climate losses

We measure the effect of the migrant network on climate losses by comparing the welfare cost of climate change in economies with and without network preferences. Even absent climate change, household utility differs across the two economies, so we measure climate losses within each economy and then compare the losses. Let $W^o \equiv V_0^o(\mathbf{x}_0^o, \mathbf{z}_0)$ be the expected lifetime utility (10) of an origin- o household at the initial network and climate. In an economy where households value their migrant network, the welfare loss of climate change is

$$D^o \equiv (1 - \beta) (W_{\dagger}^o - W^o), \quad (21)$$

where the dagger ($\dagger$) denotes the no climate change benchmark. The benchmark holds the trend and the spatial covariance of the climate process (5) at their initial values, $\bar{z}_t = \bar{z}_0$ and $\Sigma_t = \Sigma_0$ for all t . We define climate losses the same way in a no-network economy where households do not value their migrant network ($h \equiv 0$),

$$\tilde{D}^o \equiv (1 - \beta) \left(\tilde{W}_\dagger^o - \tilde{W}^o \right), \quad (22)$$

where the tilde ($\sim$) denotes the no-network economy.

In the no-network economy migration is driven by fundamentals alone. Households do not value their migrant network, $h \equiv 0$, and therefore choose where to live based on wages and residence costs. We hold every other parameter fixed across the two economies, so a household from origin o faces the same wages and residence cost in every location as in the network economy, but no longer values living among others from its origin, at home or abroad. Without the network, households leave destinations where their diaspora is concentrated for destinations with strong fundamentals. The reallocation is large, because we estimate that the network accounts for most of the wedge that holds a corridor at its observed size. The no-network economy is not a prediction but a benchmark for valuing network preferences.

With log utility, D^o is exactly the consumption-equivalent welfare loss from climate change. It is the fraction of consumption that a household from origin o would be willing to give up in every location and every period to avoid living in a world with climate change.²⁹ The same holds for $\tilde{D}^o$ in the no-network economy. Because D^o is a proportional loss, it does not depend on the initial income level of the origin. Because it is a difference within an economy, it does not depend on the level of welfare in that economy. The losses can therefore be compared across origins and across the two economies with and without network preferences.

The effect of the migrant network on climate losses is the difference between the two:

$$D^o - \tilde{D}^o = (1 - \beta) \left[(W_\dagger^o - \tilde{W}_\dagger^o) - (W^o - \tilde{W}^o) \right]. \quad (23)$$

It measures the interaction between the migrant network and climate change, not the value of network preferences or the welfare losses from climate change. The network effect is a double difference over the four welfare levels. The first term in parentheses, $W_\dagger^o - \tilde{W}_\dagger^o$, is how much better off a household from origin o is from living in an economy with network preferences when the climate does not change. The second term, $W^o - \tilde{W}^o$, is the same value with global warming.


Equation (23) isolates the interaction between the migrant network and climate change. The first difference compares welfare with and without the network while holding the climate fixed. Under the no-warming benchmark, $W_\dagger^o - \tilde{W}_\dagger^o$ is the value of the migrant network to a household from origin o . Under climate change, the corresponding value is $W^o - \tilde{W}^o$. The network effect on climate

²⁹Suppose a household from origin o loses the fraction $1 - e^{-D^o}$ of its consumption in every location and every period. With log utility its flow utility falls by D^o . Because the fall is the same proportion in every location, its migration decisions remain unchanged. Lifetime utility falls by $D^o/(1 - \beta)$. This equals the welfare loss from climate change, $W_\dagger^o - W^o$, by (21).

losses is the difference between these two values. A positive value of $D^o - \tilde{D}^o$ means that the migrant network increases the welfare cost of climate change, while a negative value means that the network reduces it.

We compute the four components of the effect of migrant networks on climate damages by solving the model forward from the observed 2020 network, averaging over 10,000 Monte Carlo draws of the climate process. All four economies start from the 2020 migrant stocks $\mathbf{X}_0$ and the current climate $\mathbf{z}_0$ and follow the growth path of Section 4 through . In the two economies with climate change, the trend and the covariance of the climate process follow the projections of Section 4. In the two economies without global warming, both are held at their 2020 values. The same draws of the underlying shocks are used in all four economies, so that differences across them are not driven by sampling. We assume that after the trend and the covariance of the climate process are fixed at their terminal values.

Figure 8: The Effect of Migrant Networks on Climate Losses




Note: This figure maps the effect of the migrant network on climate losses, $D^o - \tilde{D}^o$. It is the consumption equivalent welfare losses defined in (23), computed from the observed 2020 network $\mathbf{X}_0$ under the projected climate path. The difference compares the present discounted cost of climate change with and without the network. The counterfactual removes network preferences ($h \equiv 0$), holds the initial network and all other parameters fixed, and re-solves the economy. Numbers above zero (red) indicate that the network amplifies climate losses, and below zero (blue) that it contributes to adaptation.

Our quantification shows that migrant networks raise climate losses for two thirds of the world. Figure 8 maps the effect of the migrant network on climate losses, $D^o - \tilde{D}^o$. They are lifetime welfare losses evaluated in 2020 under the projected climate path, so the difference compares the present discounted cost of climate change with and without the network. The network amplifies losses for most origins, with the largest values concentrated in Sub-Saharan Africa, the Middle East, and South and Southeast Asia.

Migrant networks amplify climate losses most for the origins already facing the largest climate damages. Figure 9 plots the network effect in consumption units, $(D^o - \tilde{D}^o)$, against origin's projected percent loss in 2095 GDP per capita from climate change in panel (a) and origin GDP per capita in panel (b). The network effect is the change in the permanent consumption cost of climate change attributable to the network, so a positive value means the network raises the share of consumption a household would give up to avoid climate change. For origins expected to lose from climate change, the network raises damages. For the hottest quartile, the network adds a 2% consumption loss. The largest increase is for Qatar which has its diaspora in the Gulf, on the order of 6.7% of consumption. For the coldest countries, the network lowers welfare losses from climate change on the order of 0.35% of consumption. For Mongolia, which has a diaspora primarily in cold destinations, the network lowers climate losses by 5.3% of consumption. The network effect on climate losses falls with income. For countries below the median GDP per capita, the network adds 1.7% of consumption, whilst mitigating 0.1% of consumption for countries above the median.

Figure 9: The Incidence of the Network Effect Across Origins



Note: This figure plots the network effect on climate damages against origin characteristics. Each point is an origin country, and the network effect is $(D^o - \tilde{D}^o)$ in log points of permanent consumption. They are lifetime welfare losses evaluated in 2020 under the projected climate path, computed as in Figure 8. Panel (a) plots the effect against origin climate damages, measured as projected percent loss in 2095 GDP per capita from climate change. Panel (b) plots the effect against origin GDP per capita in 2019, PPP-adjusted, on a log scale. The orange line is the least squares fit weighted by origin population.

5.2 Understanding the Effect of the Network in Climate Losses

The value of a location reflects both the expected future path and the variability of the climate. Expanding W_0^o , the ex-ante value of an origin- o household, to second order in the climate shocks

around the expected climate path, $\bar{z}$, decomposes the value of a location into a trend and risk term:

$$W_0^o \approx V_0^o(\mathbf{x}_0^o, \bar{z}) + \underbrace{\frac{1}{2} \text{tr}(\nabla_z^2 V_0^o \cdot \Sigma)}_{\equiv \Omega^o}. \quad (24)$$

The first-order term of the expansion vanishes because by definition climate shocks are mean zero. The trend term, V_0^o , is the household's welfare if climate change follows its projected path exactly. The risk term Ω^o adjusts this value for the shocks around that path, scaling their variance by the curvature of the value function. With log utility, the flow utility is linear in the log climate shock. Therefore, a household that does not migrate is risk neutral in the shock. The risk matters only because households choose where to migrate to after seeing the shock. This makes the risk term Ω^o an option value. The value of moving depends on how differently destinations are hit because moving is valuable only when one place a household might choose does better than another. Hence, this option value is always nonnegative: it rises when options are risky, falls when they move together, and is zero when every option is exposed the same so that there is nowhere safe to move.

We therefore decompose the network's effect on climate losses into a trend term and a variability term. We separate the two effects because they are economically distinct. The first would be present even if the climate followed its projected path exactly, and the second arises from the variability of realized climate around that path. Because the expansion in (24) applies to each of the four values in (23), we can isolate the interaction between the network and each moment. Since the welfare losses are, to second order, the sum of the trend and variability terms, the network effect is the sum of the difference for each term:

$$D^o - \tilde{D}^o \approx (1 - \beta) \left[\underbrace{\Delta V^o - \Delta \tilde{V}^o}_{\text{anchoring}} + \underbrace{\Delta \Omega^o - \Delta \tilde{\Omega}^o}_{\text{diversification}} \right] \quad (25)$$

We call the losses from the trend term the anchoring channel, and the losses from the risk term the diversification channel.

The network affects both terms of the value of a location by changing where households choose to live. With network preferences additively separable from consumption, we have that:³⁰

$$\pi_t^{o,d} \propto \tilde{\pi}_t^{o,d} e^{h(x_t^{o,d})/\nu} \quad (26)$$

where $\tilde{\pi}_t^{o,d}$ represents migration choice probabilities in an economy without network preferences. Network preferences tilt migration flows toward destinations with larger diasporas. The network therefore changes the probability that a household lives in a particular destination in the future, and this can change its exposure to the climate trend. The network also changes the destinations a household would move to after a shock, and the value of being able to move.

The network amplifies trend damages if it distorts migration decisions towards more climate exposed destinations. Applying the decomposition for the choice probabilities above, we can

³⁰Full derivation is in Appendix C.5

express the anchoring term as

$$\Delta V^o - \Delta \tilde{V}^o = \sum_t \beta^t \frac{\text{Cov}_{\tilde{\pi}}(e^{h(x_t^{o,d})/\nu}, \Delta \log \bar{z}_t^d)}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]} \quad (27)$$

A formal derivation can be found in Appendix C.5, but the key steps are straightforward. First, logit with an additive destination shifter implies an exponential tilting of choice probabilities. Second, by the envelope theorem, ΔV^o is, to first order, an expectation of trend damages $\Delta \log \bar{z}^d$ under the choice probabilities π . Combining these observations, the network changes ΔV^o only by reweighting destinations d by the factor $e^{h(x^{o,d})/\nu}$, which yields the covariance representation above.

A sufficient statistic is the sign of covariance of diaspora shares and trend losses. The migrant network enters the household value function like a taste shifter $h(x^{o,d})$. In this multinomial logit setting, the network reweights the baseline choice probabilities $\tilde{\pi}^{o,d}$ into $\pi^{o,d} \propto \tilde{\pi}^{o,d} e^{h(x^{o,d})/\nu}$, pulling households toward destinations where the migrant stock is large. When the destinations with a large diaspora are also facing large trend damages, then the covariance is positive and the network raises the welfare cost of climate change. On the other hand, if the network is located in locations that gain from warming, then the covariance is negative and the network mitigates the welfare cost of climate change.


Our quantification shows that the network raises trend damages by keeping households at home and by directing those who leave toward losing destinations. With the network, 95% of households remain in their origin, compared with 40% in the no-network economy. By staying at home, these households face trend losses of 12.3% of consumption, compared with 8.8% at the locations they would choose without network preferences. The 5% of households who emigrate follow the network into a narrow set of destinations whose climates resemble home's. Figure 10a plots the diaspora weighted trend damages as a percent of 2095 GDP by the origin country's damages with and without the network. The network steepens the positive correlation between the origin and diaspora damages by concentrating the diaspora in fewer geographically close destinations. The Herfindahl–Hirschman Index (HHI) of the diaspora with network preferences is 0.19, compared to an HHI of 0.04 without the network. For this channel, the network amplifies losses for countries in Africa and the MENA region, and mitigates losses for former Soviet countries which have their diasporas in other cold countries.

The diversification channel works through the option value to move, and the network reduces this value when it is concentrated in co-moving places. Applying the decomposition of the choice probabilities above to pairs of destinations, we can express the diversification term as

$$\Delta \Omega^o - \Delta \tilde{\Omega}^o = \sum_t \beta^t \frac{1}{2\nu} \frac{\text{Cov}_{\tilde{\pi}}(e^{(h(x_t^{o,i})+h(x_t^{o,j}))/\nu}), \Delta \Sigma_t^{ii} - \Delta \Sigma_t^{ij})}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,i})/\nu}] \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,j})/\nu}]},$$

where the covariance is taken over ordered pairs of destinations (i, j) weighted by $\tilde{\pi}$, and $\Delta \Sigma^{ii}$ and $\Delta \Sigma^{ij}$ are the changes under climate change in the variance of destination i 's shock and in its

Figure 10: Decomposing the Role of the Network in Climate Losses



Note: In Figure 10a, each dot is an origin country. The x-axis plots the climate damages from the warming trend as a percent of 2095 GDP per capita, and the y-axis plots the diaspora-weighted damages of the destinations where the diaspora lives. Orange represents the diaspora in 2095 for the economy without network preferences, and blue represents the diaspora in 2095 for the economy with network preferences. The lines are the least squares fit lines. Figure 10b plots the diversification channel, $\Delta\Omega^o - \Delta\tilde{\Omega}^o$, across origin countries. For each origin, we rank 10,000 climate draws by the realized loss at home, group them into deciles, and report the distribution of the diversification channel within each decile.

covariance with destination j 's. A formal derivation can be found in Appendix C.5, but the key steps are straightforward. First, the exponential tilting of choice probabilities (26) applies to each member of a pair, so the network reweights the pair weights $\tilde{\pi}^{o,i} \tilde{\pi}^{o,j}$ by the product factor $e^{(h(x^{o,i}) + h(x^{o,j}))/\nu}$. Second, the option value is an expectation over pairs of destinations, $\Omega^o = \frac{1}{2\nu} E_{\tilde{\pi}\tilde{\pi}} [\Sigma^{ii} - \Sigma^{ij}]$, in which each destination contributes the variance of its own shock net of the shock's co-movement with the alternative. Combining these observations, the network changes the option value only by reweighting destination pairs (i, j) , which yields the covariance representation above.

The key moment driving the diversification channel is the covariance between the size of stocks in a pair of destinations and the co-movement of the pair's shocks. The network reweights the baseline pair weights in proportion to their stocks, so the pairs that matter for the option value are those where the migrant stocks are large. A pair of destinations provides insurance when its climate shocks are negatively correlated. Migration is the option to move to a good draw, and that option is worth most across weakly correlated places. The network reweighting erodes the option value in two ways. It concentrates weight on fewer destinations, and the destination pairs that keep weight are similar places whose shocks co-move. A diaspora clustered in similar places offers nowhere to move. When the network is concentrated in destination pairs whose shocks co-move, the covariance term is positive, so the network raises climate losses through lower diversification. When the network

places more weight on destination pairs with offsetting shocks, the covariance term is negative, increasing the option value of migration.

Our quantification shows that the diversification channel is close to zero in expectation but large in the tail of climate draws. The expected option value is small to begin with. With log utility, this channel only matters through the ability to re-sort after the shock. However, the shocks themselves are small, with a standard deviation of roughly 0.8% of GDP in each period, so Ω^o is modest for every origin. The network then removes most of the option that remains. With the network, 95% of households stay at the origin and never use the migration option. The diversification channel becomes important in bad climate realizations. Figure 10b shows the distribution of realized losses across climate draws after removing the anchoring channel. When the origin gets a bad draw, some destinations are still doing relatively well. Without the network, households are more able to move toward those destinations. With the network, they are more likely to stay at home or move to destinations experiencing similar shocks. In the worst 5% of climate draws, the consumption-equivalent loss from climate shocks is twice as large with the network as without it.

6. Counterfactual Migration Policies

The findings in Section 5 present policymakers with an additional instrument for climate adaptation. The migrant networks that exist when the worst of climate change arrives determine who can leave the most exposed regions, and whether they leave for destinations that climate change spares. Although policymakers cannot directly determine these networks, they can influence their development through immigration policy by deciding who can enter, from where, and at what cost. This section quantifies how migrant networks respond to these policy choices and whether they mitigate or amplify the welfare costs of climate change.

We use two policy experiments to show that migration policy is climate policy, and that the effects of these policies are amplified by the migrant networks it creates or destroys. The first experiment closes the borders of Europe and the US from 2020 to 2060, and then reopens them. Despite reopening the borders during the worst decades of climate change, the lost networks from the temporary closure worsen climate damages. This loss falls on countries whose networks would have facilitated emigration. In the second experiment, ten climate-spared destinations extend climate visas to ten poor countries most exposed to warming. We show that such a policy can mitigate the impact of climate damages on these countries. The migrant network plays a key role in this policy: the visas admit migrants who make the destination more attractive for the next round of migrants.

6.1 Temporary tightening of migration policy


The first experiment tightens the borders of the EU, the UK, and the US against the rest of the world for forty years, while allowing migration within the bloc. We use this counterfactual to ask whether migration policies that are not explicitly designed to address climate change can still affect climate

adaptation, and who bears this cost. The experiment is motivated by historical precedent. In 1924, the United States introduced restrictions that greatly reduced immigration admissions. This system remained in place until 1965 (Abramitzky and Boustan, 2017). Our experiment quantifies the costs of a similarly prolonged migration restriction in the context of climate change.

We implement the tightening of migration policy by increasing residence costs for households from outside the bloc who reside in the US, the UK, or the twenty-six EU member countries in our sample. Beginning in 2020, we increase these costs each period, reaching their maximum just before the policy ends in 2060. We calibrate the increases so that, by 2060, migrant stocks in the bloc from the rest of the world fall to 5 percent of their no-policy levels. Higher residence costs reduce new arrivals and gradually reduce existing migrant stocks as households return to their origin countries or relocate elsewhere. Figure 11a shows the decline in the bloc's migrant networks during the four decades of tightening. We leave residence costs within the bloc and for migration from the bloc to the rest of the world unchanged. After 2060, we restore all residence costs to their original levels.

The policy increases the welfare costs of climate change for countries excluded from the bloc. We measure the impact as the difference between the countries' welfare with and without climate change, compared across the policy and no-policy worlds. Therefore, the number isolates the impact of the tightened policy on climate damages, without confounding it with the impact of the policy restricting migration to richer countries. By this measure, the temporary policy decreases consumption equivalent welfare by 0.010 percent, increasing the climate damages by 0.4 percent.

Figure 11: Impact of the Temporary Tightening of Migration Policy




Notes: Panel (a) shows migrant stocks in the bloc from non-bloc origins as a percentage of their no-policy levels. The shaded area marks the tightening period, during which residence costs increase from 2020 until their removal in 2060. Panel (b) plots the increase in climate damages from the temporary policy for each climate-affected origin, expressed as a percentage of permanent consumption, against its 2020 migrant stock in the bloc as a share of its population (log scale).

The effects are heterogeneous across countries and depend on the size of their existing migrant

networks in the bloc. For example, Türkiye loses 0.030 percent of permanent consumption, representing 5 percent of its climate burden. In contrast, the policy has no effect on countries such as Djibouti, which initially had only 0.01 percent of its population in the bloc. Figure 11b plots the impact of the tightened migration policy against the country's initial population share in the bloc. For most countries, the policy has no effect. Losses emerge only for countries with more than 1 percent of their population in the bloc and increase sharply with the size of their initial migrant communities.³¹

The costs of the tightened migration policy persist after its removal. Figure 12a plots the damage from the tightened policy in each period. Damages increase throughout the four decades of tightening as migration restrictions intensify and climate change worsens, reaching their maximum just before the policy ends in 2060. After we remove the restrictions, damages decline gradually but remain substantial into the 2080s. Although the policy ends before the most severe decades of climate change, 41 percent of its total damages occur after 2060.

Figure 12: Damages During and After the Temporary Tightening



Notes: Panel (a) shows the increase in climate damages from the temporary tightening relative to the no-policy world, expressed as a percentage of consumption per period and aggregated across climate-affected origins outside the bloc using population weights. The dotted line marks the policy's removal in 2060. Panel (b) compares damages from the temporary tightening, removed in 2060, with those from a permanent immigration ban beginning in 2060, over 2065–2100. Both are measured relative to the same no-policy world. The curve shows the temporary policy's damages as a percentage of the permanent ban's damages in each period. Damage shares reported in the text are discounted, while the plotted series are undiscounted.

The damages persist after reopening because reopening restores the migration costs to their

³¹Two origin countries benefit from the closure: Albania and Georgia. Both have their largest migrant communities in bloc destinations that experience substantial climate damage: Albania in Greece and Italy, and Georgia in Cyprus and Greece. Without the closure, existing migrant networks would direct migration to these destinations, where climate change reduces welfare relative to remaining at home. By restricting this migration, the policy increases welfare in both origin countries.

original levels but does not restore the migrant networks that would otherwise facilitate migration into the bloc. The right half of Figure 11a shows that a decade after reopening, migrant stocks in the bloc remain below 31 percent of their no-policy levels. By the end of the century, they recover to 93 percent, but remain below their no-policy levels. These smaller networks continue to restrict migration during the most severe decades of climate change, leaving more households in climate-damaged origin countries even after the policy ends.

To quantify the costs of these reduced networks, we run an additional counterfactual where the bloc keeps its borders open until 2060 and then bans immigration for the rest of the century. From 2065 to 2100, both counterfactuals face the same climate conditions but differ in their migration restrictions and existing networks. The first has open borders but depleted migrant networks, while the second has closed borders but initially retains its existing networks. We compare the damages from the temporary tightening with those from the immigration ban to measure the costs of depleted networks relative to a complete restriction on immigration. Figure 12b shows this ratio for each period. In the first years after reopening, the temporary policy generates 52 percent of the damages from the immigration ban. This share declines as migrant networks recover. Over the full period, the damages from the temporary policy represent 13 percent of those from the immigration ban, demonstrating that temporary migration restrictions can impose substantial costs long after their removal by reducing migrant networks.

6.2 Climate visa issuance

The second experiment studies climate visa issuance. Under this policy, the destinations that climate change spares offer dedicated visa slots to the countries that it damages most. The case for such a policy comes from the results in Section 5. The origins facing the largest climate damages in the poorest regions of the world are those without diaspora communities in the destinations that warming spares. Therefore, a policy that lowers the cost of migration by creating these pathways acts directly on that constraint.

Examples of such climate visa policies already exist. In 2023, Australia signed a treaty with the Pacific island of Tuvalu granting 280 Tuvaluans per year permanent residence in Australia. With a population of 11,000, this comes to 2.5 percent of the population annually. The treaty states that the reason for this pathway is rising sea levels from climate change, which are expected to submerge most of the country in the coming decades. These visas are in addition to other migration pathways that exist for Tuvaluans into Australia. Migrants under the climate visa are chosen by lottery from an open ballot ([Department of Foreign Affairs and Trade, Australian Government, 2024](#)). Demand in the first year revealed that the policy had relaxed a binding constraint. More than a third of Tuvalu's population entered the first ballot, and by its close over 80 had applied ([Agence France-Presse, 2025](#)). Our experiment asks what this instrument does at the scale of the climate problem, for regions with millions of people.

We run a similar policy in our model. Ten destinations that climate change spares offer climate visas to ten of the poorest countries it damages most, so that each treated origin holds a visa

pathway to each of the ten destinations. The destinations are the largest climate winners in our sample. The targeted origin countries are the ten poorest countries with the largest climate damages, ranging from 3.5 percent to 5.1 percent of permanent consumption.³² The model has no separate visa channel, so we implement the policy through the residence costs, κ . Specifically, we set costs on the visa corridors so that in the policy's first period, each origin sends approximately 500 additional migrants per year into the destination bloc as a whole, roughly 50 per year to each of its ten destinations.³³ The policy is small by design. Sudan's pathway to Iceland, for instance, admits approximately hundred additional people per year, and Mauritania's about ten.

The network amplifies the migration response to the policy. The climate visas issued build communities in the destinations, and the communities lower the cost of following for future migrants through other channels outside of the climate visas. Although the climate visas guarantee immigration for 0.003 percent of the origin country's population, over the century, the policy relocates approximately 2 percent of the treated origins' population into the destinations that climate change spares. We show the value of this migration response for climate adaptation in Table 4.

Table 4: Impact of Climate Visa Policy on Ten Treated Countries

	Change in climate damages vs. no-policy world	
	Consumption-equivalent welfare (%)	Share of no-policy world climate damages
Climate visa policy in 2020	0.149	3.5%
<i>Role of the endogenous network:</i>		
visa policy only	0.001	0.0%
+ endogenous network	0.148	3.5%

Notes: The table reports the decrease in climate damages relative to the no-policy world. The first column reports the reduction in the welfare cost of climate change as a percentage of permanent consumption. The second expresses this reduction as a share of the treated origins' climate damages without the policy, equivalent to 4.25 percent of permanent consumption. The "visa policy only" row holds migrant networks fixed at their 2020 levels, isolating the policy's effect from the network feedback effect. The "+ endogenous network" row reports the additional impact of the endogenous network. All entries are population-weighted across the ten treated origins.


The migration response to the policy reduces climate damages for the ten treated countries. As before, we isolate the effect of climate visas on climate damages by comparing the welfare effects of climate change with and without the policy, separating the policy's effect on climate adaptation from the welfare gains associated with migration from poorer to richer countries. The policy removes damages worth 0.15 percent of permanent consumption for the treated group,

³²The ten targeted origin countries are Benin, Burkina Faso, Côte d'Ivoire, Ghana, Guinea, Mali, Mauritania, Niger, Sudan, and Chad. The ten targeted climate-spared destination countries are Austria, Canada, Switzerland, Estonia, Finland, Iceland, Lithuania, Norway, Russia, and Sweden.

³³The policy sets a uniform rate: each origin's guaranteed increment is 0.003 percent of its population per year, divided across its ten corridors. The 500 is the mean across origins but the policy scales with the origin countries' population, from roughly 90 additional admissions per year for Mauritania to 1,050 for Sudan.

representing 3.5 percent of the damages that climate change imposes on them (Table 4). The impact varies widely across the treated countries. Figure 13a shows the reduction in each origin's climate damages as a share of its damages without the policy. The policy reduces Benin's climate damages by 8.5 percent, compared with 0.5 percent for Ghana. Differences in initial migrant networks help explain this variation. Benin's initial diaspora share in the destination bloc is one-sixteenth of Ghana's. Climate visas therefore create new migration opportunities for Benin, where few migrants previously resided in the bloc. In contrast, Ghana already has established migration corridors and gains less from the additional migration opportunities provided by the policy.

Figure 13: Impact of Climate Visa Policy



Notes: Panel (a) shows the reduction in climate damages from the 2020 policy for each treated origin, expressed as a percentage of its climate damages without the policy. Panel (b) shows, for each start year, the percentage of the 2020 policy's adaptation benefits forgone by delaying implementation, comparing both policies over the same remaining years along deterministic trend paths. We weight origins by population when aggregating. Table 5 provides the corresponding decomposition of the costs of delay.

Because the endogenous network amplifies the migration response to the policy, removing that channel decreases the adaptation value of the policy. To isolate the impact of the policy independent of the network amplifier, we run the climate visa policy in a model where we hold the network at its 2020 value throughout the transition. Agents make migration decisions as if the network stays at its 2020 value. This counterfactual removes the feedback of migrants increasing the attractiveness of a destination and inducing more migration. In this counterfactual world, the policy removes 0.001 percent of climate damages in permanent consumption, less than 1 percent of the full effect (Table 4). The policy on its own guarantees admissions for 500 migrants from each origin country per year, and the counterfactual exercise shows that these are too few to change the impacts of climate change. The value of the policy comes from its interaction with the endogenous network. Section 5 shows that, without this policy, the network alone would increase climate damages for these countries. With the policy, however, the climate visas admit the first migrants, these migrants

make the destinations attractive for future migrants, and the effect compounds into the 3.5 percent of climate damages that the policy removes.

Since the impacts of climate change are backloaded, one could argue for introducing the policy later in the century. Table 5 evaluates this argument by introducing the same policy in 2035, 2050, and 2070. Delay is costly at every horizon. The 2020 policy removes climate damages worth 0.15 percent of permanent consumption. The same policy started in 2035 removes 0.089 percent, and a late launch in 2070 removes only 0.019 percent. Waiting 50 years costs nearly 90 percent of the adaptation value of the policy.

Table 5: Cost of Delay in Climate Visa Policy

	Decrease in climate damages (% ce)	Loss to later start (% ce)	Accumulated damages before policy begins (% ce)	Lost adaptation from compounding network effect
2020	0.149			
2035	0.089	0.060	0.002	39%
2050	0.047	0.101	0.036	58%
2070	0.019	0.130	0.099	62%

Notes: Each start date introduces the same climate visa policy. The first column reports the reduction in climate damages relative to the no-policy world, expressed as a percentage of permanent consumption. The second column reports the total cost of delay, measured as the difference in climate damages avoided by the 2020 policy and the policy introduced in the indicated year. The third column measures the portion of this cost incurred before the delayed policy begins. The final column measures the cost of delayed network formation. For each start date, we calculate the reduction in climate damages under both policies over the years following the delayed launch and express the difference as a percentage of the reduction achieved by the 2020 policy over those same years. All calculations use discounted utility flows along the deterministic trend path and are population-weighted across the ten treated origins.

Policy delay is costly for two reasons. The first is that prior to issuing climate visas, people who would have migrated remain in the damaged origins, absorbing the losses the policy would have prevented. The third column of Table 5 measures it as the damages the 2020 policy removes in the years before the later policy begins. These accumulated damages grow with the start date, and they grow faster later in the century, when the climate burden is largest.

The second reason is that a later start date limits the compounding network effect. The policy's value comes from the network that the visas build. A policy that starts in 2070 gives the network less time to build, and it faces the worst of climate change with communities still near the beginning of their growth. The fourth column of Table 5 measures this by comparing the two policies over the same calendar years, after the later one has started. A policy that starts in 2035 delivers 39 percent less protection than the 2020 treaty over those same years, and one introduced in 2070, 62 percent less.

Figure 13b shows how the cost of delay from the network channel varies with the policy's start date. These costs increase sharply during the first two decades of delay. A policy that starts in the 2040s delivers only half the protection that the 2020 policy provides over the same remaining years. The costs then increase more slowly through mid-century as two effects offset each other. A later start gives migrant networks less time to develop, but also shifts the comparison window forward,

excluding years in which the 2020 policy already benefits from its established networks. After 2075, the costs increase sharply again because the later policy has too little time to build migrant networks and benefit from the additional migration they facilitate. A policy that starts in the 2080s delivers less than 10 percent of the protection that the 2020 policy provides over the same remaining years.

7. Conclusion

Whether migration will help people adapt to climate change depends not only on how many move but on where they go. Our estimates show that migrant networks, which play an important role in determining migration destinations, raise the welfare cost of climate change for most origin countries. Larger diasporas draw larger subsequent flows, and diasporas live in climates similar to their origins, so the network steers new migrants toward destinations that face the same changing climate as where they came from. The network also offers little insurance against climate shocks, because a bad year at the origin is likely to be a bad year in the diaspora's destinations as well. The network amplifies losses from climate change for the hottest and poorest countries, who have geographically closer diasporas, and the smallest networks in countries that will be harmed the least by global warming.

The same force that raises the cost of climate change makes migration policy a more effective tool for climate adaptation. Because new migrants follow the diaspora, and the diaspora is the accumulation of past migrants, networks amplify the effects of immigration policy on migration flows. At the parameters we estimate, a temporary policy does not have permanent effects, but its effects persist for decades after the policy itself has ended. A small climate visa program can therefore reduce climate losses substantially. The IPCC recognizes migration as an adaptation to climate change ([Cissé et al., 2022](#)), the Global Compact for Migration asks states to open regular pathways for people displaced by it ([United Nations General Assembly, 2018](#)), and a few governments have begun to do so, but no framework exists for choosing which origins to admit, to which destinations, and how soon. Our framework provides a sufficient statistic that allows policymakers to compare the reduction in climate losses across potential visa recipients. We forecast that the gain from a visa program would be largest when the connected origin is expected to experience worse impacts from climate change, when origin and destination climate shocks are uncorrelated, and when the initial origin diaspora in the destination is small.

Our results point to two directions for future research. First, understanding how networks respond to climate shocks requires more frequent and more detailed data on international migration than statistical agencies currently collect. Migrant stocks observed at a five year frequency show that migrant networks are adapting slowly to global warming, but they cannot show whether the network facilitates short run migration after a large climate shock such as a flood. Bilateral flows at an annual frequency, with information on return and onward moves, would help researchers better understand the value of the diaspora as insurance after a shock. Second, many of the climate shocks that drive migration are extreme events rather than changes in

average temperature. Hurricanes, droughts, and floods displace people ([Hornbeck, 2012](#); [Groeger and Zylberberg, 2016](#); [Mahajan and Yang, 2020](#); [Aron-Dine, 2025](#)). The insurance value of a network may therefore differ when damages come from extreme events rather than from gradual warming. Estimating a damage function for such climate shocks as well as its spatial covariance remains an open problem.

References

- Abel, Guy J. and Joel E. Cohen, "Bilateral international migration flow estimates for 200 countries," *Scientific Data*, June 2019, 6, 82.
- Abramitzky, Ran and Leah Boustan, "Immigration in American Economic History," *Journal of Economic Literature*, 2017, 55 (4), 1311–1345.
- , Leah Platt Boustan, and Katherine Eriksson, "A Nation of Immigrants: Assimilation and Economic Outcomes in the Age of Mass Migration," *Journal of Political Economy*, 2014, 122 (3), 467–506.
- Agence France-Presse, "More than 80% of Tuvalu Seeks Australian Climate Visa," Agence France-Presse, 23 July 2025 July 2025. By Laura Chung, Sydney.
- Allen, Treb and Dave Donaldson, "Persistence and Path Dependence in the Spatial Economy," Working Paper 28059, National Bureau of Economic Research November 2020. Revised November 2022.
- , Costas Arkolakis, and Yuta Takahashi, "Universal Gravity," *Journal of Political Economy*, 2020, 128 (2), 393–433.
- Angelucci, Manuela, "Migration and Financial Constraints: Evidence from Mexico," *The Review of Economics and Statistics*, 2015, 97 (1), 224–228.
- Aron-Dine, Shifrah, "Rebuild or Relocate? Recovery after Natural Disasters," 2025. Working paper.
- Baldwin-Edwards, Martin and Albert Kraler, "REGINE: Regularisations in Europe," Technical Report, International Centre for Migration Policy Development (ICMPD), Vienna 2009. Edited volume/report on regularisation practices across Europe; PDF hosted by ICMPD.
- Barbosa-Alves, Maurício, "Climate Change and International Migration," 2025. Working paper.
- Barreca, Alan, Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S. Shapiro, "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century," *Journal of Political Economy*, 2016, 124 (1), 105–159.
- , Olivier Deschênes, and Melanie Guldi, "Maybe Next Month? Temperature Shocks and Dynamic Adjustments in Birth Rates," *Demography*, 2018, 55 (4), 1269–1293.
- Bartel, Ann P., "Where Do the New U.S. Immigrants Live?," *Journal of Labor Economics*, 1989, 7 (4), 371–391.
- Bayer, Patrick, Stephen L. Ross, and Giorgio Topa, "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes," *Journal of Political Economy*, December 2008, 116 (6), 1150–1196.

- Bazzi, Samuel, “Wealth Heterogeneity and the Income Elasticity of Migration,” *American Economic Journal: Applied Economics*, 2017, 9 (2), 219–255.
- Beaman, Lori A., “Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the U.S.,” *The Review of Economic Studies*, January 2012, 79 (1), 128–161.
- Beine, Michel and Christopher Parsons, “Climatic Factors as Determinants of International Migration,” *The Scandinavian Journal of Economics*, 2015, 117 (2), 723–767.
- and Christopher R. Parsons, “Climatic factors as determinants of international migration: Redux,” *CESifo Economic Studies*, 2017, 63 (4), 386–402.
- , Frédéric Docquier, and Çağlar Özden, “Diasporas,” *Journal of Development Economics*, 2011, 95 (1), 30–41.
- , Simone Bertoli, and Jesús Fernández-Huertas Moraga, “A Practitioners’ Guide to Gravity Models of International Migration,” *The World Economy*, 2016, 39 (4), 496–512.
- Bilal, Adrien and Diego R. Känzig, “The Macroeconomic Impact of Climate Change: Global vs. Local Temperature,” Working Paper 32450, National Bureau of Economic Research May 2024.
- and Esteban Rossi-Hansberg, “Anticipating Climate Change Across the United States,” Working Paper 31323, National Bureau of Economic Research 2023. Conditionally accepted, *Econometrica*.
- , FIRSTNAME Franco, and Esteban Rossi-Hansberg, “Climate Risk in Dynamic Spatial Equilibrium,” Working Paper 33377, National Bureau of Economic Research 2025.
- Bisin, Alberto and Thierry Verdier, “The Economics of Cultural Transmission and Socialization,” in Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, eds., *Handbook of Social Economics*, Vol. 1A, North-Holland, 2011, pp. 339–416.
- Blumenstock, Joshua E., Guanghua Chi, and Xu Tan, “Migration and the Value of Social Networks,” *The Review of Economic Studies*, January 2025, 92 (1), 97–128.
- Borusyak, Kirill and Peter Hull, “Nonrandom Exposure to Exogenous Shocks,” *Econometrica*, 2023, 91 (6), 2155–2185.
- , Rafael Dix-Carneiro, and Brian K. Kovak, “Understanding Migration Responses to Local Shocks,” Working Paper 35921, National Bureau of Economic Research October 2026.
- Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak, “Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh,” *Econometrica*, 2014, 82 (5), 1671–1748.
- Burchardi, Konrad B., Thomas Chaney, and Tarek A. Hassan, “Migrants, Ancestors, and Foreign Investments,” *The Review of Economic Studies*, 2019, 86 (4), 1448–1486.

- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel, “Global non-linear effect of temperature on economic production,” *Nature*, October 2015, 527, 235–239.
- , W. Matthew Davis, and Noah S. Diffenbaugh, “Large Potential Reduction in Economic Damages under UN Mitigation Targets,” *Nature*, 2018, 557 (7706), 549–553.
- Card, David, “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration,” *Journal of Labor Economics*, 2001, 19 (1), 22–64.
- Carleton, Tamma A., Esther Duflo, B. Kelsey Jack, and Guglielmo Zappalà, “Adaptation to climate change,” in Lint Barrage and Solomon Hsiang, eds., *Handbook of the Economics of Climate Change*, Vol. 1, Elsevier, 2024, pp. 143–248.
- Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E. Kopp, Kelly E. McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang, “Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits,” *The Quarterly Journal of Economics*, 2022, 137 (4), 2037–2105.
- Carrington, William J., Enrica Detragiache, and Tara Vishwanath, “Migration with Endogenous Moving Costs,” *American Economic Review*, September 1996, 86 (4), 909–930.
- Castro-Vincenzi, Juanma, “Climate Hazards and Resilience in the Global Car Industry,” 2024. Working paper.
- , Gaurav Khanna, Nicolas Morales, and Nitya Pandalai-Nayar, “Weathering the Storm: Supply Chains and Climate Risk,” Working Paper 32218, National Bureau of Economic Research March 2024. Revised March 2025.
- Cattaneo, Cristina and Giovanni Peri, “The Migration Response to Increasing Temperatures,” *Journal of Development Economics*, 2016, 122, 127–146.
- , Michel Beine, Christiane J. Fröhlich, Dominic Kniveton, Inmaculada Martinez-Zarzoso, Marina Mastrorillo, Katrin Millock, Etienne Piguet, and Benjamin Schraven, “Human Migration in the Era of Climate Change,” *Review of Environmental Economics and Policy*, 2019, 13 (2), 189–206.
- Center for International Earth Science Information Network (CIESIN), Columbia University, United Nations Food and Agriculture Programme (FAO), and Centro Internacional de Agricultura Tropical (CIAT), “Gridded Population of the World, Version 3 (GPWv3): Population Count Grid,” 2005. Accessed [your download date].
- Chen, Xiaoguang and Lu Yang, “Temperature and industrial output: Firm-level evidence from China,” *Journal of Environmental Economics and Management*, 2019, 95, 257–274.

- Cissé, Guéladio, Robert McLeman, Helen Adams, Paulina Aldunce, Kathryn Bowen, Diarmid Campbell-Lendrum, Susan Clayton, Kristie L. Ebi, Jeremy Hess, Cunrui Huang, Qiyong Liu, Glenn McGregor, Jan Semenza, and Maria Cristina Tirado, “Health, Wellbeing, and the Changing Structure of Communities,” in Hans-Otto Pörtner, Debra C. Roberts, Melinda Tignor, Elvira S. Poloczanska, Katja Mintenbeck, Andrés Alegría, Marlies Craig, Stefanie Langsdorf, Sina Löschke, Vincent Möller, Andrew Okem, and Bardhyl Rama, eds., *Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK and New York, NY: Cambridge University Press, 2022, pp. 1041–1170. Cross-Chapter Box MIGRATE: Climate-related migration.
- Conte, Bruno, “Climate change and migration: the case of Africa,” Working Paper 1411, Barcelona School of Economics 2024.
- , Klaus Desmet, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg, “Local Sectoral Specialization in a Warming World,” *Journal of Economic Geography*, 2021, 21 (4), 493–530.
- Cruz, José-Luis and Esteban Rossi-Hansberg, “The Economic Geography of Global Warming,” *The Review of Economic Studies*, 2024, 91 (2), 899–939.
- Delforge, Damien, Valentin Wathelet, Regina Below, Cinzia Lanfredi Sofia, Margo Tonnelier, Joris A. F. van Loenhout, and Niko Speybroeck, “EM-DAT: The Emergency Events Database,” *International Journal of Disaster Risk Reduction*, 2025, p. 105509.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, “Temperature Shocks and Economic Growth: Evidence from the Last Half Century,” *American Economic Journal: Macroeconomics*, 2012, 4 (3), 66–95.
- , —, and —, “What Do We Learn from the Weather? The New Climate-Economy Literature,” *Journal of Economic Literature*, 2014, 52 (3), 740–798.
- Dellink, Rob, Jean Château, Elisa Lanzi, and Bertrand Magné, “Long-Term Economic Growth Projections in the Shared Socioeconomic Pathways,” *Global Environmental Change*, 2017, 42, 200–214.
- Department of Foreign Affairs and Trade, Australian Government, “Explanatory Memorandum: Falepili Union between Tuvalu and Australia,” <https://www.dfat.gov.au/geo/tuvalu/australia-tuvalu-falepili-union-treaty/explanatory-memorandum-falepili-union-between-tuvalu-and-australia> 2024. Accessed September 2026.
- Deryugina, Tatyana, Laura Kawano, and Steven Levitt, “The Economic Impact of Hurricane Katrina on Its Victims: Evidence from Individual Tax Returns,” *American Economic Journal: Applied Economics*, 2018, 10 (2), 202–233.

- Deschênes, Olivier and Michael Greenstone, “Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US,” *American Economic Journal: Applied Economics*, 2011, 3 (4), 152–185.
- Desmet, Klaus and Esteban Rossi-Hansberg, “On the Spatial Economic Impact of Global Warming,” *Journal of Urban Economics*, 2015, 88, 16–37.
- and —, “Climate Change Economics over Time and Space,” Working Paper 32197, National Bureau of Economic Research 2024.
- , Dávid Krisztián Nagy, and Esteban Rossi-Hansberg, “The Geography of Development,” *Journal of Political Economy*, 2018, 126 (3), 903–983.
- , Robert E. Kopp, Scott A. Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H. Strauss, “Evaluating the Economic Cost of Coastal Flooding,” *American Economic Journal: Macroeconomics*, 2021, 13 (2), 444–486.
- Dingel, Jonathan I., Kyle C. Meng, and Solomon M. Hsiang, “Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate,” Working Paper 25447, National Bureau of Economic Research 2019. Revised.
- Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt, “The Economics of Fertility: A New Era,” in Shelly Lundberg and Alessandra Voena, eds., *Handbook of the Economics of the Family*, Vol. 1, North-Holland, 2023.
- Edin, Per-Anders, Peter Fredriksson, and Olof Åslund, “Ethnic Enclaves and the Economic Success of Immigrants—Evidence from a Natural Experiment,” *The Quarterly Journal of Economics*, February 2003, 118 (1), 329–357.
- for International Earth Science Information Network (CIESIN), Center, “Low Elevation Coastal Zone (LECZ) Urban-Rural Population and Land Area Estimates, Version 3,” Dataset, Version 3.00 2021. Accessed [your download date].
- Gasparrini, Antonio, Yuming Guo, Masahiro Hashizume, Eric Lavigne, Antonella Zanobetti, Joel Schwartz et al., “Mortality risk attributable to high and low ambient temperature: a multicountry observational study,” *The Lancet*, 2015, 386 (9991), 369–375.
- Gover, Mary, “Mortality during Periods of Excessive Temperature,” *Public Health Reports*, July 1938, 53 (27), 1122–1143.
- Groeger, Armando and Yanos Zylberberg, “Internal Labor Migration as a Shock Coping Strategy: Evidence from a Typhoon,” *American Economic Journal: Applied Economics*, 2016, 8 (1), 268–289.
- Grossman, Gene M., Elhanan Helpman, and Hugo Lhuillier, “Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?,” *Journal of Political Economy*, 2023, 131 (12), 3462–3496.

- Hersbach, Hans, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla et al., “The ERA5 Global Reanalysis,” *Quarterly Journal of the Royal Meteorological Society*, 2020, 146 (730), 1999–2049.
- Hoffmann, Roman, Anna Dimitrova, Raya Muttarak, Jesús Crespo Cuaresma, and Jonas Peisker, “A Meta-Analysis of Country-Level Studies on Environmental Change and Migration,” *Nature Climate Change*, 2020, 10 (10), 904–912.
- Hornbeck, Richard, “The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe,” *American Economic Review*, 2012, 102 (4), 1477–1507.
- International Institute for Applied Systems Analysis, “SSP Scenario Database,” <https://data.ece.iiasa.ac.at/ssp> 2025. Version [fill in version], accessed [your download date].
- International Monetary Fund, “IMF Executive Board Approves USD451 Million Disbursement in Emergency Natural Disaster Assistance for Pakistan,” Press Release No. 10/344 September 2010. Accessed 2026-01-19.
- , “Zimbabwe: 2011 Article IV Consultation—Staff Report; Staff Supplement; Public Information Notice on the Executive Board Discussion; and Statement by the Executive Director for Zimbabwe,” IMF Country Report 11/135, International Monetary Fund May 2011. Staff report completed May 12, 2011; Executive Board discussion June 1, 2011.
- Jessoe, Katrina, Dale T. Manning, and T. Robert Taylor, “Climate change and labour allocation in rural India: implications for welfare and gender,” *The Economic Journal*, 2018, 128 (608), 230–261.
- Jones, T. S., A. P. Liang, E. M. Kilbourne, M. R. Griffin, P. A. Patriarca, S. G. Wassilak, R. J. Mullan, R. F. Herrick, Jr. Donnell H. D., K. Choi, and S. B. Thacker, “Morbidity and Mortality Associated With the July 1980 Heat Wave in St Louis and Kansas City, Mo,” *JAMA*, 1982, 247 (24), 3327–3331. Published June 25, 1982.
- Kazianga, Harounan and Christopher Udry, “Consumption Smoothing? Livestock, Insurance and Drought in Rural Burkina Faso,” *Journal of Development Economics*, 2006, 79 (2), 413–446.
- Kelley, Colin P., Shahrzad Mohtadi, Mark A. Cane, Richard Seager, and Yochanan Kushnir, “Climate Change in the Fertile Crescent and Implications of the Recent Syrian Drought,” *Proceedings of the National Academy of Sciences*, 2015, 112 (11), 3241–3246.
- Kleinman, Benny, Ernest Liu, Stephen J. Redding, and Motohiro Yogo, “Neoclassical Growth in an Interdependent World,” Working Paper 31951, National Bureau of Economic Research 2023. Revised February 2026.

- Kucheryavyy, Konstantin, Gary Lyn, and Andrés Rodríguez-Clare, “Grounded by Gravity: A Well-Behaved Trade Model with Industry-Level Economies of Scale,” *American Economic Journal: Macroeconomics*, 2023, 15 (2), 372–412.
- Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E. Waugh, “The Welfare Effects of Encouraging Rural–Urban Migration,” *Econometrica*, 2023, 91 (3), 803–837.
- Mahajan, Parag and Dean Yang, “Taken by Storm: Hurricanes, Migrant Networks, and US Immigration,” *American Economic Journal: Applied Economics*, April 2020, 12 (2), 250–277.
- Maher, Nicola, Sebastian Milinski, Laura Suarez-Gutierrez et al., “The Max Planck Institute Grand Ensemble: Enabling the Exploration of Climate System Variability,” *Journal of Advances in Modeling Earth Systems*, 2019, 11 (7), 2050–2069.
- Mayer, Thierry and Soledad Zignago, “Notes on CEPII’s Distances Measures: The GeoDist Database,” Working Paper 2011-25, CEPII December 2011.
- McKenzie, David and Hillel Rapoport, “Network Effects and the Dynamics of Migration and Inequality: Theory and Evidence from Mexico,” *Journal of Development Economics*, 2007, 84 (1), 1–24.
- Meghir, Costas, A. Mushfiq Mobarak, Corina Mommaerts, and Melanie Morten, “Migration and Informal Insurance: Evidence from a Randomized Controlled Trial and a Structural Model,” *The Review of Economic Studies*, 2022, 89 (1), 452–480.
- Moore, Will H. and Stephen M. Shellman, “Whither Will They Go? A Global Study of Refugees’ Destinations, 1965–1995,” *International Studies Quarterly*, 2007, 51 (4), 811–834.
- Morten, Melanie, “Temporary Migration and Endogenous Risk Sharing in Village India,” *Journal of Political Economy*, 2019, 127 (1), 1–46.
- Müller, Ulrich K., James H. Stock, and Mark W. Watson, “An Econometric Model of International Growth Dynamics for Long-Horizon Forecasting,” *The Review of Economics and Statistics*, 2022, 104 (5), 857–876.
- Munshi, Kaivan, “Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market,” *The Quarterly Journal of Economics*, May 2003, 118 (2), 549–599.
- , “Social Networks and Migration,” *Annual Review of Economics*, 2020, 12 (1), 503–524.
- and Mark Rosenzweig, “Networks and Misallocation: Insurance, Migration, and the Rural–Urban Wage Gap,” *American Economic Review*, 2016, 106 (1), 46–98.
- Nath, Ishan B., Valerie A. Ramey, and Peter J. Klenow, “How Much Will Global Warming Cool Global Growth?,” Working Paper 32761, National Bureau of Economic Research July 2024.

- National Legislative Bodies / National Authorities, “Türkiye: Regulation on Temporary Protection,” Refworld October 2014. Accessed 2026-01-20.
- Newell, Richard G., Brian C. Prest, and Steven E. Sexton, “The GDP-Temperature Relationship: Implications for Climate Change Damages,” *Journal of Environmental Economics and Management*, 2021, 108, 102445.
- Ortega, Francesc and Giovanni Peri, “The Effect of Income and Immigration Policies on International Migration,” *Migration Studies*, 2013, 1 (1), 47–74.
- Park, R. Jisung, Joshua Goodman, Michael Hurwitz, and Jonathan Smith, “Heat and Learning,” *American Economic Journal: Economic Policy*, 2020, 12 (2), 306–339.
- Presidency of Migration Management (Republic of Türkiye), “Law on Foreigners and International Protection,” n.d. Accessed 2026-01-20.
- Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon et al., “Comprehensive Evidence Implies a Higher Social Cost of CO₂,” *Nature*, 2022, 610 (7933), 687–692.
- Republic of Türkiye, “Türkiye: Law No. 6458 of 2013, Law on Foreigners and International Protection [Amended],” Refworld April 2013. Accessed 2026-01-20.
- Riahi, Keywan, Detlef P. van Vuuren, Elmar Kriegler, Jae Edmonds, Brian C. O’Neill, Shinichiro Fujimori, Nico Bauer, Katherine Calvin et al., “The Shared Socioeconomic Pathways and their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview,” *Global Environmental Change*, 2017, 42, 153–168.
- Rohde, Robert A. and Zeke Hausfather, “The Berkeley Earth Land/Ocean Temperature Record,” *Earth System Science Data*, 2020, 12, 3469–3479.
- Rosenzweig, Mark R. and Oded Stark, “Consumption Smoothing, Migration, and Marriage: Evidence from Rural India,” *Journal of Political Economy*, 1989, 97 (4), 905–926.
- Šavrič, Bojan, Tom Patterson, and Bernhard Jenny, “The Equal Earth Map Projection,” *International Journal of Geographical Information Science*, 2019, 33 (3), 454–465.
- Schlenker, Wolfram and Michael J. Roberts, “Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change,” *Proceedings of the National Academy of Sciences*, 2009, 106 (37), 15594–15598.
- Schwartzman, Matthew, “Risk Management and Aggregate Migration Flows,” 2025. Working paper.
- Schwarzwald, Kevin and Kerrie Geil, “xagg: A Python Package to Aggregate Gridded Data onto Polygons,” *Journal of Open Source Software*, 2024, 9 (104), 7239.

- Seo, Jaeun and Ryungha Oh, “Sectoral Shocks and Labor Market Dynamics: A Sufficient Statistics Approach,” December 2024. Working paper. Conditionally accepted at the *Review of Economic Studies*.
- Silva, J. M. C. Santos and Silvana Tenreyro, “The Log of Gravity,” *The Review of Economics and Statistics*, 2006, 88 (4), 641–658.
- Stuart, Bryan A. and Evan J. Taylor, “Migration Networks and Location Decisions: Evidence from US Mass Migration,” *American Economic Journal: Applied Economics*, July 2021, 13 (3), 134–175.
- United Nations, Department of Economic and Social Affairs, Population Division, “International Migrant Stock 2020: Destination and Origin,” Dataset, POP/DB/MIG/Stock/Rev.2020 2021. Accessed 27 July 2026.
- United Nations General Assembly, “Global Compact for Safe, Orderly and Regular Migration,” Resolution A/RES/73/195, adopted 19 December 2018 December 2018. Objective 5: Enhance availability and flexibility of pathways for regular migration.
- Vicedo-Cabrera, Ana M., Noah Scovronick, Francesco Sera et al., “The burden of heat-related mortality attributable to recent human-induced climate change,” *Nature Climate Change*, 2021, 11, 492–500.
- WHO Regional Office for the Eastern Mediterranean, “Annual report of the Regional Director, 2010: Executive summary,” Technical Report, World Health Organization, Regional Office for the Eastern Mediterranean 2010. Accessed 2026-01-19.
- Yagan, Danny, “Moving to Opportunity? Migratory Insurance over the Great Recession,” 2014. Working paper.

Supplemental Appendix

A	Additional Tables and Figures	56
A.1	Additional Evidence on Flows–Stocks Relationship	56
A.2	Additional Evidence on Diaspora Climate Exposure	56
B	Data	61
B.1	Data Construction and Limitations	61
B.2	Constructing the Instrument	63
B.3	Examples of Variation Underlying the Instrument	64
C	Model Proofs and Derivations	67
C.1	Proof of Proposition 1: Existence and Uniqueness of Equilibrium	67
C.2	Proof of Proposition 2: Unique Steady State	67
C.3	Proof of Proposition 3: Unique Stationary Distribution	67
C.4	Persistence and Amplification with N Destinations	68
C.5	Second Order Expansion of Welfare	68
D	Quantification and Computation	71
D.1	Differences across Shared Socioeconomic Pathways	71
D.2	The Damage Process	72
D.3	Comparison with Cruz & Rossi-Hansberg (2024) Quantification	72

A. Additional Tables and Figures

A.1 Additional Evidence on Flows–Stocks Relationship

This appendix provides additional evidence on the relationship between migrant stocks and subsequent migration flows documented in Section 2. We first show that the positive relationship is visible in the raw data and is similar across alternative transformations, time periods, and country income groups. We then report alternative estimates that account for the large number of zero bilateral flows.

Figure A1 provides additional evidence on the flows–stocks relationship. Figure A1a plots bilateral flow shares against lagged stock shares. One concern is that this relationship could be driven by the normalization used to construct the shares. Figure A1b therefore plots migrant flows against lagged stocks in levels and shows the same positive relationship. Another concern is that the quality and coverage of the migration data vary across countries and over time. The aggregate relationship could therefore be driven by particular parts of the sample. Figures A1c and A1d plot the relationship separately by year and by UN income group. The positive relationship is present across these subsamples, suggesting that it is not driven by a particular time period or income group.

We also show that our results are robust to including corridors with zero migration flows. Our main specification is estimated in logs using OLS so that the estimates are directly comparable to the 2SLS estimates in Table 1. A limitation of this specification is that bilateral flows are frequently zero and are therefore dropped when taking logs. As a robustness check, we estimate the relationship using Poisson pseudo–maximum likelihood (PPML), which retains zero-flow observations and allows the conditional variance of flows to vary with the covariates.


Table A1 reports the results. As in the OLS specifications, lagged migrant stocks are strongly positively associated with subsequent flows. The estimated elasticity is 0.77 without fixed effects and 0.73 with origin and destination fixed effects, compared with 0.88 and 0.83 in the corresponding OLS specifications. The coefficient falls once origin–destination fixed effects are included, again reflecting the limited within-corridor variation in migrant stocks, but remains positive and precisely estimated. Overall, the positive flows–stocks relationship is not driven by dropping corridors with zero migration flows.

A.2 Additional Evidence on Diaspora Climate Exposure

This appendix provides additional evidence on the relationship between origin-country climate exposure and the exposure of its diaspora. We first show the distributions underlying the averages reported in the main text. We then examine whether the relationship differs across time periods and country income groups, given differences in the coverage and quality of the migration data. Finally, we repeat the analysis using alternative measures of climate exposure.

Figure A2 shows the distributions behind the averages reported in Section 2. Figure A2a plots the distribution of diaspora temperature conditional on origin-country temperature. As origin

Figure A1: Migrant Flows by Migrant Stocks



Notes: Migrant stock data are from the UN Global Migration Database and migrant flows are from [Abel and Cohen \(2019\)](#). Panel (a) plots bilateral migrant flow shares against lagged migrant stock shares. Panel (b) plots the same relationship in levels. Panels (c) and (d) report bincatters of flow shares against lagged stock shares separately by year and by UN income group, respectively. The dashed gray line denotes the 45-degree line.

Table A1: Elasticity of Flows to Stocks, PPML Estimates


	(1)	(2)	(3)	(4)	(5)
$\log(\text{stocks}_{od,t-5})$	0.770*** (0.014)	0.731*** (0.015)	0.789*** (0.011)	0.211*** (0.034)	0.183*** (0.027)
o, d		✓		✓	
$o \times t, d \times t$			✓		✓
$o \times d$				✓	✓

Notes: The table reports PPML estimates of the elasticity of bilateral migration flows with respect to lagged migrant stocks from 1990 to 2020. PPML is estimated in levels and includes observations with zero flows. Migrant stocks are from the UN and gross flows from [Abel and Cohen \(2019\)](#). Standard errors are clustered at the origin–destination pair level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.


temperature rises, the distribution of diaspora temperature shifts toward higher temperatures. Figure A2b reports the distribution of migrants' distance from their origin by income group. The distribution shifts toward larger distances as origin-country income rises: migrants from low-income countries are concentrated much closer to home, while migrants from middle- and high-income countries are spread over a wider range of distances. Migrants from poorer countries are systematically closer to home across the distribution.

Figure A2: Distribution of Climate Exposure and Geography of Migrant Networks

(a) Climate exposure of origins and their diasporas



(b) Spatial concentration of migrant networks




Notes: Panel (a) plots the distribution of diaspora average temperature by origin-country temperature. Origins are grouped by their average temperature, and each vertical histogram shows the distribution of diaspora temperatures for origins in that group. Diaspora temperature is the migrant-stock-weighted average temperature across an origin's destination countries. Colors indicate diaspora temperature in degrees Fahrenheit. Panel (b) plots distribution of migrants' distance from their origin by UN income group. Temperature is from BEST and bilateral distance from GeoDist ([Mayer and Zignago, 2011](#)).

The migration data vary in coverage and quality across countries and over time, raising the

possibility that the aggregate relationship is driven by particular parts of the sample. Figure A3 therefore plots the origin–diaspora exposure relationship separately by year and by UN income group. The positive relationship is present across periods and income groups.


Figure A3: Diaspora Network and Origin Country Climate Exposure



Notes: The figures show a binscatter of the relationship between diaspora exposure against origin exposure. Panel (a) plots relationship by year. Panel (b) plots the same relationship by the UN classifications of income groups. The gray dashed line is the 45-degree line.

Our baseline measure of climate exposure is average temperature. Temperature is measured more consistently across space and time than many other climate variables, particularly precipitation, which varies more sharply across space and is harder to interpolate in areas with sparse station coverage (Dell et al., 2014). Figure A4 shows that the origin–diaspora relationship is not specific to average temperature. We find similar positive relationships using measures of extreme heat, coastal exposure, extreme rainfall, and natural disasters. Appendix B describes the construction and sources of these measures.

Figure A4: Origin–Diaspora Exposure Relationship Under Alternative Climate Measures



Notes: Each panel relates origin-country climate exposure to the migrant-weighted average exposure of its diaspora. Lines show OLS fitted values using the 1990 network (orange) and 2015 network (red); shaded areas are 95% confidence intervals. Migrant stocks are from the UN Global Migration Database ([United Nations, Department of Economic and Social Affairs, Population Division, 2021](#)). See Appendix B for variable definitions and data sources.

B. Data

B.1 Data Construction and Limitations

Migrant stocks. The UN global bilateral migration data are constructed through a systematic harmonization of national censuses and population registers across over 200 countries and territories, spanning multiple decades. The harmonization process attempts to address inconsistencies in national data collection methods, definitions of migrants (including distinctions between foreign-born and foreign nationals), and reporting periods, thereby enabling cross-country comparisons and longitudinal analysis.

However, the data have several important limitations and sources of potential systematic measurement error. First, countries vary significantly in how they define and measure migrants for example, birthplace versus citizenship. This distinction is particularly problematic when comparing countries with and without birthright citizenship policies. For instance, a second-generation immigrant born in the United States would not be counted as a migrant under a birthplace definition, whereas the child of immigrants born in Germany (prior to recent reforms) might be counted as foreign even if they have never lived elsewhere. This implies much greater persistence over time in the stock of people from origin-*o* in Germany than in the United States even if actual migration patterns are similar. Such systematic differences could bias our estimates of migration persistence and network effects across destination countries.

Second, the data are typically drawn from five-year or decennial censuses, which may miss temporary migration episodes and provide limited information on the precise timing of migration. Third, undocumented migrants are systematically undercounted in official statistics, with the degree of undercount likely varying across countries based on enforcement regimes, survey methodology, and migrants' willingness to participate in census activities. Finally, data quality and availability remain particularly limited for developing countries and earlier time periods, potentially biasing our sample toward wealthier destinations and more recent decades. To address concerns about sample composition and data quality, we demonstrate that our main results are robust across different time periods and country income classifications.

Temperature data. We construct country-level temperatures from the Berkeley Earth daily land temperature product ([Rohde and Hausfather, 2020](#)), which reports daily average temperature on a one-degree grid as an anomaly relative to the 1951–1980 climatology. We recover the daily temperature of each grid cell by adding the anomaly to the climatology. We then aggregate grid cells to the country level using the `xagg` package of [Schwarzwald and Geil \(2024\)](#), weighting each cell by the population living in the part of the cell that falls within the country's borders. Country boundaries come from Esri's World Countries shapefile. Population weights come from the Gridded Population of the World version 3 count grid for 1990 ([Center for International Earth Science Information Network \(CIESIN\), Columbia University et al., 2005](#)). This gives daily population-weighted temperature series for each country from 1970 to 2019, representing the

temperature experienced by an average person in the country. From these data, we construct three measures of temperature exposure:

- *Average temperature*, our headline measure, is the mean daily temperature within a country-year.
- *Temperature of the 10 hottest days* captures extreme heat and is defined as the average temperature across the ten hottest days of each country-year.
- *Number of hot days* counts the number of days in a country-year with temperatures above 30°C.³⁴

A limitation of the Berkeley Earth data is that temperatures are interpolated from weather-station records, and station coverage is sparse in parts of Sub-Saharan Africa and South Asia. Temperature estimates in these regions are therefore measured with more error. Our approach of using annual averages of the population-weighted temperature series addresses this concern by smoothing out much of the day-to-day measurement error.

Other climate variables. We show robustness to the following alternative measures of climate exposure. These measures use different data sources and capture dimensions of climate other than average temperature.

- *Temperature of hottest day, ERA5*: We use the temperature of the hottest day of the year. The data come from the ERA5 reanalysis ([Hersbach et al., 2020](#)), using country-level, area-weighted aggregates from the World Bank Climate Change Knowledge Portal. Unlike Berkeley Earth, ERA5 combines a global weather model with satellite and station observations and is therefore less dependent on local weather-station coverage. This measure captures exposure to extreme heat.
- *Low elevation coastal population*: We use the share of a country's population living in the Low Elevation Coastal Zone, defined as land at or below ten meters above sea level that is contiguous with the coast. The measure comes from Version 3 from [for International Earth Science Information Network \(CIESIN\)](#), which combines elevation data with gridded population counts. It provides a measure of exposure to sea-level rise and coastal flooding.
- *Extreme 1-day rainfall*: We use the largest one-day precipitation total in each year, corresponding to the Rx1day index. These data also come from ERA5 through the World Bank Climate Change Knowledge Portal and are aggregated by area. This measure captures exposure to extreme rainfall and flooding.

³⁴There is an extensively documented negative relationship between extreme high temperatures and many outcomes related to welfare. Mortality risk rises non-linearly with high temperatures ([Gover, 1938](#); [Jones et al., 1982](#); [Deschênes and Greenstone, 2011](#); [Gasparrini et al., 2015](#); [Barreca et al., 2016](#); [Vicedo-Cabrera et al., 2021](#); [Carleton et al., 2022](#)). Agricultural yields suffer ([Schlenker and Roberts, 2009](#)), worker productivity declines ([Chen and Yang, 2019](#)), schooling and learning is disrupted ([Park et al., 2020](#)), among many other negative effects.

- *Natural disaster count:* We count the number of floods, storms, heatwaves, droughts, and wildfires recorded in each country-year by EM-DAT, the international disaster database maintained by the Centre for Research on the Epidemiology of Disasters (Delforge et al., 2025). EM-DAT includes events that kill at least ten people, affect at least one hundred people, or lead to a declaration of emergency or a call for international assistance. Unlike the other measures, this variable records realized disaster events rather than meteorological conditions. Because reporting is less complete in earlier periods and in poorer countries, we use it as a complementary measure rather than as our main measure of climate exposure.

B.2 Constructing the Instrument

We follow the construction in Burchardi et al. (2019), who motivate their instrument from a reduced-form law of motion for migrant stocks. In their framework, the stock of migrants from origin o in destination d evolves with migration flows, which depends on the interaction of origin-specific push and destination-specific pull factors, as well as the existing migrant stock. Iterating this law of motion forward expresses the migrant stock in period t as a function of the history of push–pull interactions and higher-order interactions generated by the dependence of migration on past stocks. We follow the same logic and estimate a first stage that projects $\text{stocks}_{od,t}$ on the history of leave-out push–pull interactions, principal components summarizing the higher-order terms, and the initial migrant stock. We refer the reader to Burchardi et al. (2019) for the full derivation.

Similar to Burchardi et al. (2019), we define a leave-out push factor as the total outflow from o to destinations excluding all destinations in the UN sub-region where d is located, denoted by $r(d)$:

$$\text{push}_{o,t}^{-r(d)} \equiv \sum_{d' \notin r(d)} \text{flows}_{od',t}. \quad (\text{B1})$$

This construction ensures the push measure does not mechanically load on o – d migration. Analogously, we define a leave-out pull factor as destination d ’s share of inflows from origins outside of all origins in o ’s sub-region, $r(o)$:

$$\text{pull}_{d,t}^{-r(o)} \equiv \frac{\sum_{o' \notin r(o)} \text{flows}_{o'd,t}}{\sum_{o' \notin r(o)} \sum_{d''} \text{flows}_{o'd'',t}}. \quad (\text{B2})$$

That is, pull is measured using migrants from other UN sub-regions, so it captures destination-year attractiveness that is not mechanically driven by flows from o . For each period $\tau \in \{1995, 2000, 2005, 2010, 2015\}$, the resulting push–pull interaction is

$$Z_{od,\tau} \equiv \text{push}_{o,\tau}^{-r(d)} \times \text{pull}_{d,\tau}^{-r(o)}. \quad (\text{B3})$$

For our first stage for migrant stocks we use the history of push–pull interactions to instrument

for stocks in year t , $\text{stocks}_{od,t}$:

$$\log(\text{stocks}_{od,t}) = \delta_o + \delta_d + \delta_{od} + \sum_{\tau=1995}^t \alpha_{\tau} Z_{od,\tau} + \sum_{k=1}^K \lambda_k \text{PC}_k + u_{od,t}. \quad (\text{B4})$$

where $u_{od,t}$ is the error term and δ_o , δ_d , and δ_{od} are the fixed effects. The interaction terms $Z_{od,\tau}$ are the leave-out push-pull terms that capture quasi-random variation in the allocation of migrants across destinations generated by the coincidence of origin-specific emigration pushes and destination-specific immigration pulls. Following [Burchardi et al. \(2019\)](#), we also include principal components of the higher-order interaction terms implied by the forward solution, PC_k . These terms allow past push–pull shocks to affect later migrant stocks through the evolution of the migrant network. Rather than include the full set of higher-order interactions, we summarize them using their first five principal components, which account for 96% of their variation.

Our preferred specification includes origin–destination fixed effects. Identification therefore comes from changes over time in the push–pull interactions and their higher-order terms within a corridor.

B.3 Examples of Variation Underlying the Instrument

This section provides two examples of the push–pull variation underlying the instrument. In each case, a large origin-specific push coincided with a destination-specific pull, temporarily increasing migration from the origin to the destination.


Zimbabwe and Spain + Italy Example

Zimbabwe's economy began deteriorating from 1998 to 2000, where political turmoil and economic mismanagement generated large output losses ([International Monetary Fund, 2011](#)). This generated strong push incentives for outward migration from Zimbabwe. This period also coincided with pull-side policy changes in key Southern European destinations: Spain implemented large-scale regularization programmes around 2000–2001, and Italy's 2002 framework facilitated legalization, especially for domestic and care work, raising the expected payoff to migrating to these destinations ([Baldwin-Edwards and Kraler, 2009](#)). Figure B1 shows that Spain and Italy became relatively important destinations for Zimbabwean migrants during this period, but do not stand out in later periods.

Pakistan and Turkey Example

Pakistan experienced a major push shock with the floods that began in late July 2010 and affected more than 20 million people ([WHO Regional Office for the Eastern Mediterranean, 2010](#); [International Monetary Fund, 2010](#)). Over the same period, Turkey overhauled its migration governance by adopting a law which established a centralized framework for entry, residence,


Figure B1: Zimbabwe and Italy + Spain Example



Notes: Panels (a)–(b) plot Zimbabwean flows against lagged migrant stocks across destinations, highlighting Spain and Italy. Panels (c)–(d) plot the leave-out push factor for Zimbabwe and the leave-out pull factor for Spain and Italy. Migrant stocks are from the UN Global Migration Database and flows from [Abel and Cohen \(2019\)](#).

and international protection. These institutional upgrades reduced administrative uncertainty and transaction costs, raising the expected payoff from selecting Turkey as a destination.³⁵ Figure B2 shows that Turkey became a relatively more important destination for Pakistani migrants during this period than in the preceding period.

Figure B2: Pakistan and Turkey Example



Notes: Panels (a)–(b) plot Pakistani flows against lagged migrant stocks across destinations, highlighting Turkey. Panels (c)–(d) plot the leave-out push factor for Pakistan and the leave-out pull factor for Turkey. Migrant stocks are from the UN Global Migration Database and flows from [Abel and Cohen \(2019\)](#).

³⁵For more details, see [Republic of Türkiye \(2013\)](#), [Presidency of Migration Management \(Republic of Türkiye\)](#), and [National Legislative Bodies / National Authorities \(2014\)](#).

C. Model Proofs and Derivations

C.1 Proof of Proposition 1: Existence and Uniqueness of Equilibrium

We construct the equilibrium forward. Given $\mathbf{z}_{t-1}$ and the weather shock η_t , the climate process (5) determines $\mathbf{z}_t$ and (4) determines $\mathbf{w}_t$. Given $\mathbf{w}_t$ and the inherited network $\mathbf{x}_t^o$, the choice probabilities (8) are an explicit function of both, and the law of motion (9) determines $\mathbf{x}_{t+1}^o$. Each step has one solution, so the equilibrium exists and is unique. The network enters (8) only as a predetermined argument, so nothing in the construction restricts h .

The value function (10) is well defined along this path. Flow utility is $\nu \log \sum_d \exp([\log w_t^d + h(x_t^{o,d}) - \kappa_t^{o,d}]/\nu)$, in which h is bounded, κ and g are deterministic, and $\log z_t^d$ is normal with variance bounded by the stationary variance of (5). Its expectation is therefore finite in every period, and the discounted sum converges for $\beta < 1$.

C.2 Proof of Proposition 2: Unique Steady State

A steady state exists because $\pi^o(\cdot, \bar{\mathbf{z}}_T)$ is a continuous map from the simplex to itself.

For uniqueness, measure the distance between two networks $\mathbf{x}$ and $\mathbf{x}'$ by the range of their log ratio across destinations, $\max_d \log(x^d/x'^d) - \min_d \log(x^d/x'^d)$, which is zero only when the networks coincide. By (8), the log choice probabilities differ at destination d by $[h(x^d) - h(x'^d)]/\nu$ plus a term common to all destinations, which drops out of the range. As a function of $\log x$, h has slope $xh'(x) = \nu\varepsilon$, so $h(x^d) - h(x'^d)$ has the sign of $\log(x^d/x'^d)$ and is at most $\nu \max_x \varepsilon$ times its size. Since both networks sum to one, the log ratio is positive somewhere and negative somewhere, so the range of the utility differences is at most $\max_x \varepsilon$ times the range of the log ratio. One step of (9) therefore shrinks the distance between any two networks by the factor $\max_x \varepsilon$. When this is below one the map is a contraction, so the steady state is unique and the network converges to it from anywhere.

For instability, the Jacobian of (9) at a steady state $\mathbf{x}$ is $(\text{diag}(\mathbf{x}) - \mathbf{x}\mathbf{x}^\top) \text{diag}(h'(\mathbf{x})/\nu)$, a positive semidefinite matrix times a positive diagonal one. Its eigenvalues are real and non-negative, and the largest is at least the largest diagonal entry, $\max_d \varepsilon^{o,d}(1 - x^{o,d})$. If that exceeds one, a small perturbation of $\mathbf{x}$ along the corresponding eigenvector grows, and the steady state is unstable. The map takes a convex compact set into itself, so the fixed-point indices of its steady states sum to one. A steady state with one unstable direction has index -1 , so another steady state exists. With k unstable directions the index is $(-1)^k$, and the argument covers every odd k .

C.3 Proof of Proposition 3: Unique Stationary Distribution

After T the state is the network and the deviation of damages from trend, $\log \mathbf{z}_t - \log \bar{\mathbf{z}}_T$. The deviation is a Gaussian AR(1) with $|\psi| < 1$, so it has a stationary distribution, and the network lives on the simplex. The state is therefore tight and its transition is continuous, so a stationary distribution exists: any limit of the time averages of the transition probabilities from a fixed starting point is one. Integrating out the deviation gives a stationary distribution of the network.

For uniqueness, take two economies with the same weather shocks but different initial states. Their deviations differ by ψ^t times the initial difference, which vanishes. Their networks are updated each period by the same map, and the argument above shows that map shrinks the distance between any two networks by $\max_x \varepsilon$ whatever the damages, since damages enter both economies identically. When $\max_x \varepsilon < 1$ the two networks converge for every sequence of shocks. The distribution of the network after many periods therefore does not depend on the initial state, so the stationary distribution is unique and the distribution of the network converges to it.

C.4 Persistence and Amplification with N Destinations

The persistence and amplification results in Section 3 perturb one corridor holding the rest of the network fixed. They extend to the full system. In log shares, the linearized law of motion is

$$\frac{\partial \log x_{t+1}^{o,d}}{\partial \log x_t^{o,k}} = (\mathbf{1}\{d = k\} - \pi_t^{o,k}) \varepsilon_t^{o,k},$$

so a corridor's own response is $\varepsilon_t^{o,d}(1 - \pi_t^{o,d})$, which is approximately $\varepsilon_t^{o,d}$ at the small destination shares in the data. The contraction argument in the proof of Proposition 2 bounds the decay of any perturbation of the whole network by $\max_x \varepsilon$ per period for any level of wages, so the persistence rate and the amplification factor $1/(1 - \varepsilon)$ hold for the system. Under the condition of Proposition 2 they hold at every network configuration.

C.5 Second Order Expansion of Welfare

In this appendix, we derive the expressions that decompose the role of the network in climate losses into the anchoring term and the diversification term. First, define the expected lifetime utility of an origin- o household $W^o \equiv V_0^o(\mathbf{x}_0^o, \mathbf{z}_0)$. We define $\tilde{W}^o$ as expected lifetime utility in the corresponding economy without network preferences, setting $h(x) = 0$ for all x .

Taking a second-order expansion of W^o around the expected climate path $\{\bar{\mathbf{z}}_t\}$, with period- t deviations $\log \mathbf{z}_t - \log \bar{\mathbf{z}}_t \sim (\mathbf{0}, \Sigma_t)$,³⁶

$$V_0^o(\mathbf{x}_0^o, \{\mathbf{z}_t\}) \approx V_0^o(\mathbf{x}_0^o, \{\bar{\mathbf{z}}_t\}) + \sum_t \beta^t \nabla_z v_t^o \cdot (\log \mathbf{z}_t - \log \bar{\mathbf{z}}_t) + \frac{1}{2} \sum_t \beta^t (\log \mathbf{z}_t - \log \bar{\mathbf{z}}_t)' \nabla_z^2 v_t^o (\log \mathbf{z}_t - \log \bar{\mathbf{z}}_t)$$

Because the shocks are mean zero, the first order term is zero in expectation. The second order term captures the curvature of the value in each pair of shocks times their covariance

$$\frac{1}{2} \sum_{i,j} \frac{\partial^2 v_t^o}{\partial \log z_t^i \partial \log z_t^j} \Sigma_t^{ij} = \frac{1}{2} \text{tr}(\nabla_z^2 v_t^o \cdot \Sigma_t)$$

³⁶Because choices are static, the value function is additively separable across periods in the climate path: $V_0^o = \sum_t \beta^t \mathbb{E}[\nu \log \sum_d \exp([\log w_t^d + h(x_t^{o,d}) - \kappa_t^{o,d}]/\nu)]$, up to an additive constant from the Gumbel taste shocks, with $\mathbf{z}_t$ entering through wages. The second derivative of V_0^o with respect to the climate path is therefore block-diagonal by period, so cross-period covariances of the shocks drop out of the expansion.

Therefore, we can write the second order effect as

$$W^o \approx \underbrace{V_0^o(\mathbf{x}_0^o, \{\bar{\mathbf{z}}_t\})}_{\text{trend effect}} + \underbrace{\sum_t \beta^t \frac{1}{2} \text{tr}(\nabla_z^2 v_t^o \cdot \Sigma_t)}_{\text{uncertainty effect} \equiv \Omega^o} \quad (\text{C1})$$

where the first term, the trend effect, is the welfare along the expected climate path $\bar{\mathbf{z}}$, and the second order effect, defined as Ω^o , is the welfare from the climate shocks. Now, we use that with the idiosyncratic Gumbel taste shocks, direct differentiation gives

$$\frac{\partial v_t^o}{\partial \log z_t^i} = \pi_t^{o,i}, \quad \frac{\partial^2 v_t^o}{\partial \log z_t^i \partial \log z_t^j} = \frac{1}{\nu} \left(\pi_t^{o,i} \mathbf{1}\{i=j\} - \pi_t^{o,i} \pi_t^{o,j} \right)$$

Substituting this into the uncertainty term gives

$$\Omega^o = \sum_t \beta^t \Omega_t^o, \quad \Omega_t^o = \frac{1}{2\nu} \sum_{i,j} \pi_t^{o,i} \pi_t^{o,j} (\Sigma_t^{ii} - \Sigma_t^{ij}).$$

Ω^o captures the option value of moving under log utility and that option is larger when locations' shocks are negatively correlated.

We define the network effect on climate losses as $D^o - \tilde{D}^o = (1-\beta) [(W_{\dagger}^o - \tilde{W}_{\dagger}^o) - (W^o - \tilde{W}^o)]$ where $\dagger$ represents a no-climate-change economy. The first difference in the square bracket is the network value with no warming, and the second difference is the network value with warming. Now, before deriving how these two terms enter the network effect, let us first show that network preferences tilt choices toward large diasporas. Let $\tilde{\pi}_t^{o,d}$ be migration choices in period t when $h \equiv 0$. Then,

$$\begin{aligned} \pi_t^{o,d} &= \frac{\exp([\log w_t^d + h(x_t^{o,d}) - \kappa_t^{o,d}]/\nu)}{\sum_k \exp([\log w_t^k + h(x_t^{o,k}) - \kappa_t^{o,k}]/\nu)} = \frac{\exp([\log w_t^d - \kappa_t^{o,d}]/\nu) e^{h(x_t^{o,d})/\nu}}{\sum_k \exp([\log w_t^k - \kappa_t^{o,k}]/\nu) e^{h(x_t^{o,k})/\nu}} \\ &= \frac{\tilde{\pi}_t^{o,d} e^{h(x_t^{o,d})/\nu}}{\sum_k \tilde{\pi}_t^{o,k} e^{h(x_t^{o,k})/\nu}} = \frac{\tilde{\pi}_t^{o,d} e^{h(x_t^{o,d})/\nu}}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]} \end{aligned}$$

That is, h tilts migration choices toward destinations with larger diasporas. Using the second order expansion, the network effect can be written as

$$D^o - \tilde{D}^o \approx (1-\beta) [\Delta V^o - \Delta \tilde{V}^o + \Delta \Omega^o - \Delta \tilde{\Omega}^o] \quad (\text{C2})$$

where V^o and $V_{\dagger}^o$ denote welfare along the trend paths with and without climate change, so that by separability $\Delta V^o = \sum_t \beta^t \Delta v_t^o$ with $\Delta v_t^o \equiv v_t^o(\bar{\mathbf{z}}_t^{\dagger}) - v_t^o(\bar{\mathbf{z}}_t)$, and $\Delta \Omega^o = \sum_t \beta^t \Delta \Omega_t^o$ with $\Delta \Omega_t^o \equiv \Omega_{\dagger,t}^o - \Omega_t^o$.

Now, we derive each term. First, let us start with the losses from the expected climate path. Recall

$\Delta \log \bar{z}_t^d \equiv \log \bar{z}_t^{\dagger,d} - \log \bar{z}_t^d$, the trend damages of destination d . Then

$$\begin{aligned}
\Delta v_t^o - \Delta \tilde{v}_t^o &\approx \sum_d \pi_t^{o,d} \Delta \log \bar{z}_t^d - \sum_d \tilde{\pi}_t^{o,d} \Delta \log \bar{z}_t^d \\
&= \frac{\sum_d \tilde{\pi}_t^{o,d} e^{h(x_t^{o,d})/\nu} \Delta \log \bar{z}_t^d}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]} - \mathbb{E}_{\tilde{\pi}}[\Delta \log \bar{z}_t^d] \\
&= \frac{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,d})/\nu} \Delta \log \bar{z}_t^d] - \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,d})/\nu}] \mathbb{E}_{\tilde{\pi}}[\Delta \log \bar{z}_t^d]}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]} \\
&= \frac{\text{Cov}_{\tilde{\pi}}(e^{h(x_t^{o,d})/\nu}, \Delta \log \bar{z}_t^d)}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]}.
\end{aligned}$$

The first line expands each per-period value to first order in trend damages, using that $\partial v_t^o / \partial \log \bar{z}_t^d = \pi_t^{o,d}$. The same $\Delta \log \bar{z}_t^d$ appears in both economies because they share wages and residence costs. The second line substitutes $\pi_t^{o,d} = \tilde{\pi}_t^{o,d} e^{h(x_t^{o,d})/\nu} / \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]$. The third line multiplies and divides the second term by $\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]$. The last line is the definition of covariance. Summing with discount weights gives the lifetime anchoring term,

$$\Delta V^o - \Delta \tilde{V}^o = \sum_t \beta^t \frac{\text{Cov}_{\tilde{\pi}}(e^{h(x_t^{o,d})/\nu}, \Delta \log \bar{z}_t^d)}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,k})/\nu}]}$$

The second order term follows similar steps. Recall $\Delta \Sigma_t^{ij} \equiv \Sigma_t^{\dagger,ij} - \Sigma_t^{ij}$, the benchmark-minus-climate-change difference in the covariance of destinations i and j 's shocks. Then

$$\begin{aligned}
\Delta \Omega_t^o - \Delta \tilde{\Omega}_t^o &= \frac{1}{2\nu} \sum_{i,j} \pi_t^{o,i} \pi_t^{o,j} (\Delta \Sigma_t^{ii} - \Delta \Sigma_t^{ij}) - \frac{1}{2\nu} \sum_{i,j} \tilde{\pi}_t^{o,i} \tilde{\pi}_t^{o,j} (\Delta \Sigma_t^{ii} - \Delta \Sigma_t^{ij}) \\
&= \frac{1}{2\nu} \left[\frac{\sum_{i,j} \tilde{\pi}_t^{o,i} \tilde{\pi}_t^{o,j} e^{(h(x_t^{o,i})+h(x_t^{o,j}))/\nu} (\Delta \Sigma_t^{ii} - \Delta \Sigma_t^{ij})}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,i})/\nu}] \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,j})/\nu}]} - \mathbb{E}_{\tilde{\pi}\tilde{\pi}}[\Delta \Sigma_t^{ii} - \Delta \Sigma_t^{ij}] \right] \\
&= \frac{1}{2\nu} \frac{\text{Cov}_{\tilde{\pi}\tilde{\pi}}(e^{(h(x_t^{o,i})+h(x_t^{o,j}))/\nu}, \Delta \Sigma_t^{ii} - \Delta \Sigma_t^{ij})}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,i})/\nu}] \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,j})/\nu}]}.
\end{aligned}$$

The first line uses that Ω_t^o is linear in the covariance matrix given the weights, so the difference across climate scenarios replaces Σ_t with $\Delta \Sigma_t$. The second line substitutes $\pi_t^{o,i} \pi_t^{o,j} = \tilde{\pi}_t^{o,i} \tilde{\pi}_t^{o,j} e^{(h(x_t^{o,i})+h(x_t^{o,j}))/\nu} / \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,i})/\nu}] \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,j})/\nu}]$. The last line multiplies and divides the second term by $\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,i})/\nu}] \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,j})/\nu}]$, and applies the definition of covariance. Summing with discount

weights gives the lifetime diversification term,

$$\Delta\Omega^o - \Delta\tilde{\Omega}^o = \sum_t \beta^t \frac{1}{2\nu} \frac{\text{Cov}_{\tilde{\pi}\tilde{\pi}}\left(e^{(h(x_t^{o,i})+h(x_t^{o,j}))/\nu}, \Delta\Sigma_t^{ii} - \Delta\Sigma_t^{ij}\right)}{\mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,i})/\nu}] \mathbb{E}_{\tilde{\pi}}[e^{h(x_t^{o,j})/\nu}]}$$

These expressions are to build intuition. In the quantitative exercises, we solve each economy fully and allow migration choices and the migrant network to respond endogenously to climate change.

D. Quantification and Computation

D.1 Differences across Shared Socioeconomic Pathways

The growth paths g_t^ℓ used in the quantification come from the Shared Socioeconomic Pathways (SSPs), which are five scenarios for global development over the century (Riahi et al., 2017; Dellink et al., 2017). We use SSP3 as our baseline. Migration incentives in the model depend on cross-country wage differences, and SSP3 preserves large wage gaps over the century. The 90/10 income ratio remains above 9 in 2095, compared with 3.3 under SSP5, so a fast-convergence scenario would substantially reduce migration incentives even in the absence of climate change. SSP3 is also consistent with the high-warming environment used in the quantitative exercises.³⁷ SSP4 generates similar cross-country income dispersion but substantially lower emissions, making it less consistent with the warming path we study.

Table D1: Growth and Convergence Across the SSPs

	SSP1	SSP2	SSP3	SSP4	SSP5
<i>Growth in GDP per capita, 2020–2095, % per year</i>					
poorest third	3.6	3.2	1.7	2.2	4.3
middle third	2.5	2.3	1.2	2.2	3.1
richest third	1.4	1.3	0.5	1.3	2.0
<i>Dispersion of GDP per capita in 2095</i>					
sd of log GDP per capita	0.51	0.60	0.87	0.97	0.48
90/10 ratio	3.6	4.6	9.4	11.7	3.3

Notes: GDP per capita projections across the five Shared Socioeconomic Pathways. Growth rates are annualized over 2020–2095. Thirds are defined by GDP per capita in 2020, and countries are weighted equally. The 90/10 ratio is the ratio of the 90th to the 10th percentile of GDP per capita across countries in 2095.

³⁷In the SSP scenario matrix, SSP5 is the only baseline pathway that reaches radiative forcing of 8.5 W/m², while SSP3 reaches approximately 7.0 W/m². We follow the climate-impacts literature in pairing weak development with a high-warming path.

D.2 The Damage Process

This appendix details how we construct the trend damages (14) and the damage shocks entering (15) from the annual estimates of [Nath et al. \(2024\)](#). They estimate the response of GDP per capita to an annual temperature shock, with a marginal damage linear in the country's temperature, $d'_a(q) = \hat{\beta}_0 + \hat{\beta}_1 q$, and they find that the damage persists, decaying at an annual rate of 0.92. Our model runs at five-year periods and separates permanent warming from transitory weather, so their annual response enters the trend and the shock with different multipliers.

Trend damages. Warming is permanent, so each degree of the trend generates an accumulated effect. We scale the annual response by the ratio of the cumulative response to the annual response in [Nath et al. \(2024\)](#), 5.05 at their ten-year horizon, so the marginal damage from permanent warming is $d'(q) = 5.05 \times d'_a(q)$. We then integrate this marginal damage along the expected temperature path,

$$\log \bar{z}_t^\ell = 5.05 \int_{\bar{q}_0^\ell}^{\bar{q}_t^\ell} (\hat{\beta}_0 + \hat{\beta}_1 q) dq,$$

We compute the integral in 0.1°C increments, evaluating marginal damages at each step along the temperature path. Since marginal damages are linear in temperature, this is equivalent to using the closed-form integral up to a negligible discretization error. Integrating along the path matters because the implied damage function is quadratic, so the damage from an additional degree depends on the temperature level at which it occurs.

Shock damages. The weather shock η_t^ℓ is the deviation of five-year average temperature from its expected path. If this shock persists for the full five-year period, it affects output in each year. Because the effect of each year's temperature shock decays at the annual rate 0.92, the total damage remaining at the end of the period is

$$\sum_{j=0}^4 0.92^j = \frac{1 - 0.92^5}{1 - 0.92} = 4.26$$

times the annual response. The marginal damage entering (15) is therefore $4.26 \times d'_a(\bar{q}_t^\ell)$. Any damage that carries into the next five-year period persists at rate $\psi = 0.92^5 = 0.66$, as in (5). Thus, both the accumulation of damages within a period and their persistence across periods are implied by the same annual persistence estimate.

D.3 Comparison with Cruz & Rossi-Hansberg (2024) Quantification

[Cruz and Rossi-Hansberg \(2024\)](#) is the closest quantitative model of migration and climate change. This appendix summarizes the main differences between the two quantifications.

The first difference is how migration costs are modeled. In [Cruz and Rossi-Hansberg \(2024\)](#), the cost of living away from home is the product of an origin and a destination component, $m_o \times m_d$,

and is fixed over time. In our model, the effective cost of moving from o to d is


$$\kappa_t^{o,d} - [h(x_t^{o,d}) - h(x_t^{o,o})].$$

The first term is a residence cost, while the second depends on the size of the migrant community at the destination relative to the community left at the origin. Effective migration costs therefore change endogenously as migrant stocks evolve. Panel (a) of Figure D1 illustrates this for migration from Pakistan to Canada. In Cruz and Rossi-Hansberg (2024), the effective cost is constant over time. In our model, it falls as the Pakistani community in Canada grows.

The second difference is the path of income in the absence of climate change. In Cruz and Rossi-Hansberg (2024), growth is endogenous. Innovation is concentrated in more productive locations and diffuses across space, generating convergence over time. Their model also starts from greater spatial income dispersion, in part because the economy is defined over $1^\circ \times 1^\circ$ grid cells and therefore captures within-country inequality. In our model, growth follows the SSP3 path, under which cross-country income differences remain large. Panel (b) of Figure D1 compares the 90/10 ratio of income per capita under the two baselines. Since migration incentives depend on wage differences across locations, the two growth paths imply different incentives to migrate later in the century.

There are several other differences between the two frameworks. First, Cruz and Rossi-Hansberg (2024) model the world at the $1^\circ \times 1^\circ$ grid-cell level, while we model 161 countries. Their framework therefore allows substantial adjustment through within-country migration, whereas our model focuses on international migration, where migrant networks, borders, and migration policy are directly relevant. Second, their climate path is deterministic, while our model includes weather shocks through Σ_t . This additional uncertainty generates an option value of migration and the diversification channel studied in Section 5. Finally, the migration elasticities differ. Their elasticity of migration with respect to income is 2, compared with 0.7 in our calibration based on Ortega and Peri (2013), so migration responds more strongly to a given wage difference in their model.

Figure D1: Comparison with Cruz and Rossi-Hansberg (2024)



Notes: Panel (a) plots the effective cost of the Pakistan–Canada corridor over the century. Panel (b) plots the 90/10 ratio of income per capita in the baseline without climate change.