Climate migration projections across disciplines: a model inter-comparison for the US

Tom Bearpark[*†1,2], Jordan Rosenthal-Kay[*3], Nic Choquette-Levy[4], Tingyin Xiao[5], and Michael Oppenheimer[5]

[1]University of Exeter
[2]Grantham Research Institute on Climate Change and the Environment, LSE
[3]Federal Reserve Bank of San Francisco
[4]Pennsylvania State University
[5]Princeton University

August 9, 2026

Abstract

How will climate change reshape migration patterns? Different disciplines approach this question with models that vary in their assumptions and find divergent results. Here, we present the first systematic intercomparison of four approaches to modelling the climate migration relationship: causal inference, agent-based, gravity, and general equilibrium. We introduce a unified conceptual framework designed to compare their outputs at a common scale, and illustrate the framework by projecting temperature-driven internal migration across U.S. states, estimating over one hundred models per approach. Despite their structural differences, the approaches agree on the pattern and magnitude of warming-driven redistribution. Climate change causes cold regions to gain population and hot regions to lose it by under one percent of the population over a 25-year future horizon (representing less than three percent of the overall projected migration during that period). The disagreement that remains is dominated by researchers' implementation choices. These contribute as much to variance in results as the choice of approach itself and roughly twice as much as statistical uncertainty, the only source of uncertainty studies conventionally report. More broadly, our framework offers a transparent, reproducible way to identify the characteristics of climate migration that are robust across modelling traditions. Like the intercomparison efforts that transformed climate and crop science, it can be extended to other regions, to hazards beyond heat, and to new modelling frameworks.

[*]These authors contributed equally to this work. The views in this paper do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

[†]Corresponding author: `t.bearpark@exeter.ac.uk`

1 Introduction

Migration is a crucial mechanism through which individuals and communities can adapt to a changing climate (Black et al., 2011). However, mobility barriers may foreclose this option for the most vulnerable (Benveniste et al., 2022), and migration may strain receiving regions' resources and heighten the risk of conflict (Abel et al., 2019). These complexities have drawn the attention of institutions such as the Intergovernmental Panel on Climate Change (IPCC, 2022), the Food and Agriculture Organization (FAO, 2018), and the World Bank (Rigaud et al., 2018; Clement et al., 2021). As global temperatures continue to rise, anticipating the impact on future migration patterns is essential to develop strategies for adaptation, resilience, and international cooperation.

This policy imperative has spurred a large, interdisciplinary literature to develop quantitative models of the climate-migration relationship. Environmental economists employ *causal inference* tools to isolate the effects of climate variables on migration decisions. Geographers, environmental scientists, and demographers use *agent-based* models to simulate individual migration behaviour in response to climate stressors. Statisticians use *gravity* models to predict migration flows based on bilateral factors. Meanwhile, spatial economists use *general equilibrium* models to capture feedbacks between individual behaviours and market outcomes, and their effect on geographic distributions of populations. However, despite their common goal of understanding climate-driven migration, these approaches have largely developed in isolation, making direct comparisons across approaches challenging.

Existing meta-analyses (Cattaneo et al., 2019; Beine and Jeusette, 2021; Hoffmann et al., 2020, 2021; Šedová et al., 2021; Moore and Wesselbaum, 2022; Schewel et al., 2024) document substantial heterogeneity in estimates of climate-induced migration. The studies underlying these reviews differ widely in the magnitude of estimated effects (Moore and Wesselbaum, 2022), and can even reach qualitatively different conclusions about the direction of climatic impacts (Beyer et al., 2023). However, interpreting these differences is difficult because they reflect variation not only in study contexts, datasets, definitions of migration and climate exposure, but also in broad disciplinary traditions, which bundle modelling approaches with specific methodological choices. Published estimates are typically accompanied only by statistical uncertainty; the sensitivity of each estimate to the implementation choices behind it is rarely quantified or reported. Therefore, despite growing interest, we know little about how alternative quantitative approaches compare when applied to the same research question under common assumptions and inputs. This fragmentation limits our ability to extract clear, cross-disciplinary insights for policy and research.

Here, we bridge this gap by developing and implementing a framework that allows us to perform the first systematic comparison across widely used quantitative climate-migration models.[1] We consider four widely used *approaches* to modelling climate migration: causal inference, agent-based, gravity, and general equilibrium.[2] We apply our framework to a real-world case study: estimating and projecting inter-state migration caused by temperature changes in the United States. In total, across two warming scenarios, and four approaches, we estimate more than 500 distinct models to decompose projection uncertainty into that driven by approach versus modelling choices. For each approach, we select a "preferred model" which we further bootstrap across parameter estimates to capture statistical uncertainty. This lets us synthesize results and decompose uncertainty in the climate-migration relationship into contributions from the modelling approach, implementation choices, and statistical estimation.

2 A framework for climate migration model intercomparison

Since our study is the first systematic intercomparison of key approaches to modelling the impacts of climate on migration, a central part of our contribution is that we develop a clear framework for future intercomparison. Figure 1 displays a broad outline of our analysis, illustrating five steps involved in developing a generic quantitative model for projecting climate-induced migration. Each step in this process introduces points of divergence between existing studies, which differ according to (1) the output they aim to produce, (2) the modelling approach they use, (3) the choices made in implementation, (4) the methods used for estimating model parameters, and (5) the way they present their results.

Step 1: A common target output

The four modelling traditions we consider are built to answer different questions and estimate different quantities, often at different scales and in different units. Thus, their findings may not be comparable, and any disagreement among them can not be cleanly attributed to modelling choice. A meaningful comparison must therefore begin by specifying a common output in which these diverse models can be cleanly compared (Fig. 1, Step 1).

We outline a common mathematical language that nests the four approaches (Methods), that allows us to precisely define a common target output: the change in the spatial distribution of population caused by warming, measured at a fixed future horizon (Fig. 2a). Concretely, it is the

[1]While qualitative research has significantly advanced our understanding of the climate migration relationship (see, e.g., Ford et al., 2010; Tebboth et al., 2019; Ou-Salah and Van Praag, 2025), in this paper we restrict our analysis to quantitative models, for the purposes of structured intercomparison.

[2]We define an *approach* as a broad category of models characterised by shared disciplinary conventions (e.g., causal inference, gravity, general equilibrium, and agent-based models), while a *model* refers to a specific implementation within one of these approaches.

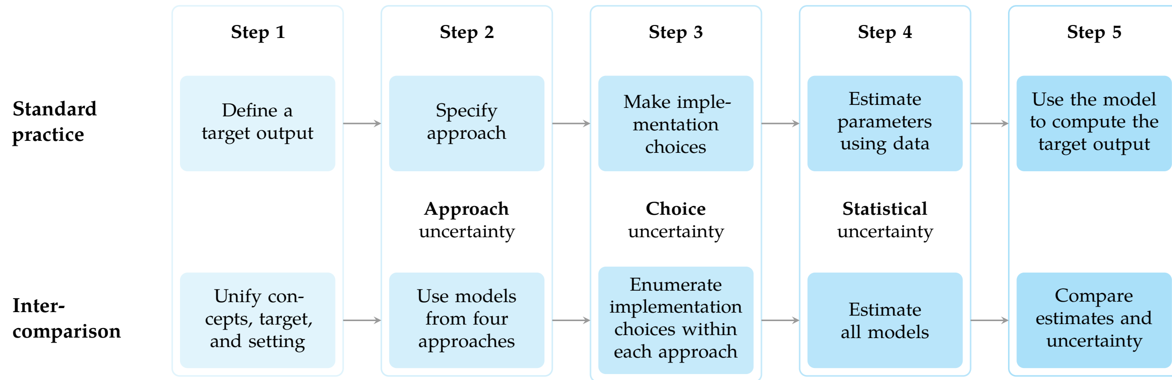

Figure 1: A common framework for comparing climate-migration models.
The top row shows standard practice, in which a researcher commits to a single approach. The bottom row shows our intercomparison. Differences introduced at successive steps correspond to approach, choice, and statistical uncertainty.

difference between a prediction of where people will live by a given year under a given warming scenario, and where they *would have lived* by that same year had the climate instead been held at its baseline.

Because only one of these two worlds is ever realised, the target is an inherently causal, unobserved quantity that no dataset can report directly. For the same reason, the target cannot be validated by simply comparing fit on historic data: reproducing observed migration demonstrates nothing about where people *would have lived* under a counterfactual climate. Estimating the target therefore requires modelling assumptions, and it is these assumptions that the four approaches supply differently. This is what gives intercomparison its purpose: were the target directly observable, model disagreement could be resolved by selecting the best performer. However, models that fit the observed history equally well can imply very different counterfactuals; agreement across models and approaches suggests that the projection robust to alternative modelling decisions.

The choice of this particular target also has two practical virtues. First, fixing the time horizon and warming scenario, all four approaches can be adapted to estimate it (see details in Fig. 2 and Supplementary Section B.2), so disagreement among them does not merely reflect differences in what is measured. Second, our target describes how climate change will reshape demands on housing, infrastructure, and public services, and is therefore essential for developing climate adaptation and resilience policy.

Step 2: Choice of approach

In reality, agents (which could be individuals, households, or even whole communities) choose where to live by selecting the location that best serves their needs, subject to the constraints that they face (see formalisation of this in Methods). Thus, the migration rate between two locations depends on how agents value the features of the origins and destinations available to potential movers, and what forces impede or facilitate migration between locations. We encode this dependence in a generic mapping, from local conditions (including weather) to migration rates, which we call the *migration-rate generating function*. This function aggregates the individual location choices described above, averaged over the many idiosyncratic factors that differ across agents, and summarises the propagation of migration flows across time and space.

If this function were known, projecting the effect of warming would be straightforward. We could alter the temperature inputs and directly calculate the resulting migration flows and population distribution. In practice, however, the function is unknown, and estimating features of it that can be used to estimate the target output requires a sequence of modelling decisions (Fig. 2b, Methods).

The first decision is whether to model the underlying behavioural processes explicitly at all. Causal inference models do not. Instead, they exploit random year-to-year weather variation to identify the local slope of the migration-rate generating function and then extrapolate these estimates to future climates. This slope directly encodes how migration rates respond to a marginal change in temperature, and is usually assumed to be a parametric non-linear function in local weather conditions. The appeal is that this locally identified effect requires limited assumptions about how agents behave. The cost is that turning it into a projection requires strong assumptions elsewhere: that the estimated causal effect of temperature on migration rates is not biased by spillovers, and that the response to short-run weather is informative about the response to long-run climate change.[3]

Each of the other three approaches models the behavioural process in full (Methods). The next modelling choice concerns the degree of heterogeneity to allow among agents. Agent-based models retain the greatest degree of heterogeneity. They can incorporate arbitrarily rich behavioural mechanisms that vary along a wide range of agent characteristics. They can also account for social networks and interactions between small sets of agents. Because these mechanisms are difficult (or impossible) to represent analytically, agent-based models simulate large numbers of individual agents directly. This flexibility allows them to represent behavioural processes that other approaches cannot, but makes them difficult to calibrate tightly to observed data. Moreover, projections may be path-dependent and have multiple equilibria.

The remaining two approaches sacrifice some heterogeneity in exchange for tractability, working with aggregate migration flows rather than individual-level simulations. The distinction between them concerns the importance assigned to equilibrium feedbacks—that is, the extent to which

[3] Frontier implementations sometimes seek to relax these assumptions, e.g. by controlling for spatial lags of the treatment variables to admit local spillovers, or estimating long differences to bring the identifying variation closer to climate.

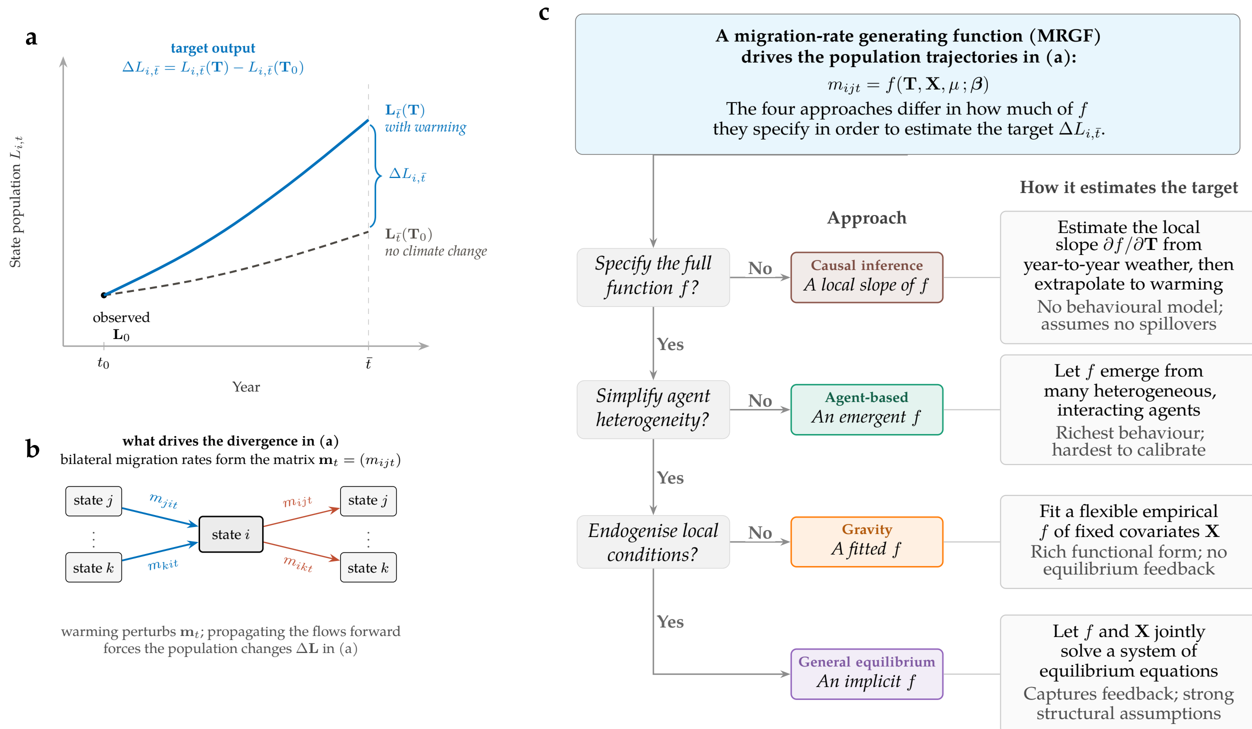

Figure 2: The target output, and the four approaches as different ways of estimating it. **(a)** The target for a single state: two projected population trajectories share the observed baseline population $\mathbf{L}_0$ and then diverge due to different climates: one holding the climate (denoted $\mathbf{T}$) at its baseline ($\mathbf{L}_{\bar{t}}(\mathbf{T}_0)$, dashed), one under warming ($\mathbf{L}_{\bar{t}}(\mathbf{T})$, solid). The target output is the gap between them at the horizon $\bar{t}$, $\Delta L_{i,\bar{t}} = L_{i,\bar{t}}(\mathbf{T}) - L_{i,\bar{t}}(\mathbf{T}_0)$ (Methods, Equation 3). The illustration shows a state that gains population due to climate change. **(b)** The object that generates these trajectories: in each year t, state i receives migrants from every other state at rates m_{jit} (blue) and sends migrants to them at rates m_{ijt} (red). Collected over all pairs of states, these bilateral rates form the migration matrix $\mathbf{m}_t = (m_{ijt})$. Populations evolve by repeated application of $\mathbf{m}_t$, so warming produces the divergence in (a) by perturbing $\mathbf{m}_t$ (Methods). **(c)** Estimating the target requires the *migration-rate generating function* (MRGF), $m_{ij} = f(\mathbf{T}, \mathbf{X}, \mu; \boldsymbol{\beta})$, which maps local conditions ($\mathbf{T}$ and non-climate covariates $\mathbf{X}$) to migration rates (m_{ij}) (Methods, Equation 2); μ collects migration costs and $\boldsymbol{\beta}$ a vector of unknown parameters. The four modelling approaches differ in how much of f they specify. Full specifications are given in Supplementary Sections B.5–B.8.

the population distribution itself changes the conditions that shape future migration decisions.

Gravity models abstract from such feedbacks. Like the causal-inference approach, they are *flow* models: they work directly with the number of people moving between each pair of states. Gravity models assume a specific behavioural model under which each bilateral flow depends only on observable characteristics of the origin, the destination, and the pair. This is sometimes motivated by the empirical regularity that larger and closer places exchange more migrants; an analogy to Newtonian gravity gives the approach its name. Because gravity models treat the state-level characteristics as exogenous, the modeller can concentrate on representing the migration-rate generating function in empirical detail, using rich covariate sets, flexible functional forms, and machine-learning methods. The cost is that migration cannot alter the conditions that shape later migration: in a gravity model, an influx into a cool state does not bid up its housing costs or depress its wages, and so can neither deter further arrivals nor induce a reverse flow. Gravity models are therefore not equipped to capture endogenous adjustment to large-scale climate change.

General equilibrium models place these feedbacks at the centre of the analysis. In these models, constantly re-equilibrating markets endogenously shape wages, housing costs, amenities, and other local conditions, which in turn influence migration decisions. Population distributions emerge from an equilibrium in which individual choices and local conditions are mutually consistent. This framework captures adjustment mechanisms that the other approaches omit, but requires strong structural and parameter assumptions that cannot generally be identified from migration data alone. To remain tractable, general equilibrium models also typically restrict the set of factors influencing migration decisions.

Step 3: Implementation choices

We define a model as a unique estimator of a given target output, so specifying an approach is not enough to pin down a model. Within each approach there is a rich literature of competing variants, and once an approach is chosen many consequential researcher degrees of freedom remain (Fig. 1, Step 3). For example, researchers still must decide on the choice of covariates, functional forms, estimation strategies, and so on.

Comparing the approaches fairly imposes two competing demands on the implementation choices that we make in our analysis. On one hand, each approach must be represented by models that are genuine examples of their tradition. They should be close to the research frontier and faithful to disciplinary conventions. On the other, the models must be mutually comparable: built in our common framework, capable of estimating the common target output, computationally feasible, and estimable from a single shared dataset. The models we implement in this paper are designed to sit at the intersection of these two demands. Meeting both at once required a number of methodological innovations in how each approach is specified and estimated, which we introduce alongside the models themselves in the Methods (Fig. 6).

For each approach we first specify a *preferred model*: a single implementation we judge to be defensible and representative of its approach. Around it we then enumerate a structured set of alternative implementation choices, each a reasonable perturbation of the preferred model (Methods, Fig 6). Comparing across this set lets us diagnose how projections vary both across approaches and across the implementation choices within them.

Importantly, the choice uncertainty we report here is a conservative measure of the true variability across models that could be run for a given approach. The space of models one might bring to this question is far larger, and structurally far more varied, than the grids we enumerate. There are plausible alternatives which embed altogether different mechanisms, or entirely different estimation strategies than we consider here. In this first intercomparison attempt, we vary only modest perturbations of each approach's preferred model, limiting the scope of our intercomparison.[4]

Step 4: Estimation

Projecting the target for any model requires estimating a set of unknown parameters using data from a specific setting (Fig. 1, Step 4). In this paper, we focus on internal migration within the United States. The US provides rich data on migration flows and local economic conditions, allowing all four modelling approaches to be implemented on a common empirical footing. At the same time, relatively little work has examined the effects of gradual climate change on migration in high-income countries.[5] Much of the existing literature instead focuses on rapid-onset climate shocks in lower-income, agriculturally dependent settings (Cattaneo et al., 2019). Thus, our analysis here aims to contribute to our empirical understanding of the relationships in this setting, as well as the intercomparison of the approaches.

We assemble a panel of bilateral state-to-state migration flows and state characteristics from 1991-2010 (see full details in Methods). Migration flows come from the IRS Statistics of Income data. These data record changes in tax-filing locations and the associated number of tax exemptions, which provides a proxy for the number of individuals moving between states. We combine these data with population-weighted annual mean temperatures derived from ERA5 reanalysis, real gross state product, bilateral distances between states, interstate trade flows, and measures of interstate social connectedness derived from Facebook's Social Connectedness Index. For future climate conditions, we use state-level population-weighted temperature projections from the Climate Impact Lab. We consider two warming scenarios, one with moderate warming, and one with extreme warming (Supplementary Fig. A1). We calibrate each model to a historical baseline using observed migration patterns and then use the estimated models to project migration responses over the subsequent 25 years. The full set of datasets, and which approach uses

[4]An example of where this restriction binds is on the implementation choices used when estimating the climate-damage functions that force the agent-based and general equilibrium models. Our historical data only weakly identify the impacts of climate on productivity, and hence reasonable perturbations of the estimation strategy can lead to flipped damage function estimates in which climate change *increases* productivity, and wildly different projections. In this study, we confine ourselves to forms that find results consistent with the established empirical literature (e.g., Burke et al., 2015) that increased temperatures reduce productivity for hot regions. Our choice uncertainty is therefore a conservative lower bound on the disagreement a fuller exploration of the model space would reveal.

[5]Existing studies that examine climate driven migration in the US include reduced form studies (Feng et al., 2012; Leduc and Wilson, 2023; Baylis et al., 2025) and structural studies (Fan et al., 2018; Sinha et al., 2018; Bilal and Rossi-Hansberg, 2023). In addition, recent large scale studies have incorporated US internal migration data (Benveniste et al., 2025).

each, is given in Supplementary Table A1. Variable construction and sources are detailed in Supplementary Section B.1, with summary statistics in Supplementary Table A3.

As every approach requires estimating unknown parameters, each projection inherits statistical uncertainty from the data used to fit it. We quantify this uncertainty across approaches using a bootstrap, re-estimating the target output from repeated draws of each model's fitted parameters, so that every approach carries a comparable measure of statistical error (Methods).

Step 5: Comparing results

The final step collects estimates of the target output across models and approaches, so their projections can be cleanly compared (Fig. 1, Step 5). Our analysis produces a large ensemble of projections. Across the four approaches we estimate 520 models in total — 128 causal-inference, 120 agent-based, 144 gravity and 128 general-equilibrium specifications (Fig. 6b). The preferred model within each approach additionally carries a parametric bootstrap on the key estimation step of how temperature affects migration rate that allows us to quantify statistical uncertainty (Methods). Each computation projects the change in population attributable to warming for every state, under a baseline in which the climate is held fixed, and our two climate change temperature trajectories. Because every projection shares a common target, dataset, and scenario, comparison proceeds on common terms. The spread across projections can be decomposed into the variance contributed by the choice of approach, by implementation choices within an approach, and by statistical estimation, while the qualitative conclusions of each approach can be synthesised into a single ensemble view.

Our four implementations

We instantiate this framework with one preferred model from each tradition, chosen to be representative of its literature while estimating the common target on the shared US dataset (Methods).

The *causal inference* model is a fixed-effects panel regression that reads the temperature-migration slope off conditionally random year-to-year weather variation, similar to Baylis et al. (2025). Our implementation choices span the estimator used (Poisson vs. ordinary least squares), the structure of the time controls, population weights, and assumed lag structure of the impacts (Supplementary Section B.5). Building from a recent literature that shows migrants tend to move to locations with similar climates to their origin (Obolensky et al., 2024), we consider heterogeneity according to the degree of climate similarity between the origin and destination (Methods, Supplementary Section B.5).

The *agent-based* model (ABM) simulates $100{,}000$ "individuals" who choose where to live each period. They are pulled by each destination's "amenities", which are influenced by climate. Agents also value the locations of their social ties, on a network anchored to Facebook's Social Connectedness Index which evolves as agents move, linking our approach to agent-based models that account for social networks when modelling climate migration (Kniveton et al., 2012; Entwisle et al., 2016; Bell et al., 2019; Choquette-Levy et al., 2021). A key methodological innovation is how we calibrate it: a gravity approximation of the ABM's underlying decision rule provides an initial estimate of coefficients (e.g., how much agents weigh amenities and social network factors in their location decisions). Then, we iteratively simulate the ABM and compare results to observed migration flows, rescaling coefficient values until simulated flows match observed data (a one-dimensional root-finding exercise with only one period of simulation). This fits the ABM efficiently, transparently, and stably, whereas unbounded search over the parameter space is extremely computationally costly and can be unstable when the model allows for multiple equilibria. [6] Our implementation choices span how damages and social preferences are estimated, and how agents' social networks are constructed (e.g., whether anchored networks are random or small-world, à la Watts and Strogatz 1998) and how they evolve over time (Supplementary Section B.6).

The *gravity* model is a two-stage bilateral-flow regression (Supplementary Section B.7), in the tradition of gravity models that capture complex interdependencies with frontier statistical tools (e.g., Xiao et al., 2022). A key methodological innovation is the division of labour between the two stages: a Poisson (PPML) gravity equation provides a stable, interpretable average response, while a machine-learning layer searches flexibly for heterogeneity, letting the temperature response vary with the other covariates in ways a fixed parametric form would miss (similar in spirit to Benveniste et al. (2025), who use a cross validation exercise to select a model for heterogeneity). Unlike for the other approaches, machine learning is a natural tool here: gravity is the sole approach whose estimation task is direct prediction, seeking the best empirical representation of the migration-rate generating function itself, whereas causal inference requires identification rather than fit, and the ABM and GE approaches require parameters with behavioural meaning. Our implementation choices span how temperature, controls, and fixed effects enter the first stage, and which machine learners are used in the second.

The *general equilibrium* (GE) model is a quantitative spatial model with interstate trade and migration, taken from the menu of models in Redding and Rossi-Hansberg (2017) with damages similar to Cruz and Rossi-Hansberg (2024). In the model, warming damages both local productivity and amenities and the population distribution emerges endogenously as a market

[6] Importantly, this borrowing does not tie the two approaches together: the gravity approximation only initialises the search, and the final parameters are those under which the ABM's own simulated flows reproduce the observed data. The projections are then generated by mechanisms the gravity model lacks (evolving networks and interacting agents) and indeed the two approaches' headline projections differ roughly threefold (Fig. 4a).

equilibrium. The damage functions are estimated with reduced-form panel data regressions. Our model includes a variety of market forces: interstate trade, agglomeration, local housing markets, and roundabout production. The strength of these forces, the calibration method for the trade and migration costs, and structure of the damage function are implementation choices (Supplementary Section B.8).

Full specifications, including the choice grids, are given in the Methods and Supplementary Information (see graphical outline in Fig. 6).

3 The impact of temperature changes on US internal migration

Projected population redistribution

The four approaches broadly agree on the geography of warming-driven redistribution (Fig. 3a). Relative to a world without climate change, each preferred model projects that cold states gain population and hot states lose it (Fig. 3b, c). They agree less on its precise *magnitude*. For example, climate change is projected to cause North Dakota's population to grow between 0.3% (causal inference) and 4% (gravity) over a 25 year horizon, while a hot state, Arizona, is projected to lose between roughly 1.4% and 6% of its population. This redistribution is concentrated: across the choice grid, the few most-affected states account for a large portion of all population moved (Supplementary Fig. A2).

Fig. 3a maps the state-to-state corridors that each approach predicts will have the highest number of climate migrants. The identity of these top corridors is governed by two forces: the pre-existing migration network, which sets which state pairs carry enough movers for warming to redirect many of them, and an approach-specific climate response, which sets how strongly each pair is reshaped. Under causal inference, projected corridor changes are close to a uniform rescaling of observed baseline flows, which explain 83% of the variation across its growing corridors. The gravity and general equilibrium models produce similar corridor maps (rank correlation 0.87): in each, the change in flows primarily depends on the origin and destination climate response and the distance between them; neither the general equilibrium model's trade linkages and endogeneity nor the gravity models' flexible learner reorganize the broad pattern. The agent-based corridors are only loosely correlated with those of the other approaches (all rank correlations below 0.1) and roughly twenty times more dispersed relative to baseline flows. They emerge from the simulated social network and its rewiring rather than from any observed bilateral covariate. Only the Arizona to California corridor ranks among the five largest under every approach (Fig. 3a). Supplementary Fig. A3 contrasts these net corridors with gross flow increases and decreases, all of which reveal the importance of displacement effects: e.g., that increased migration from Arizona to California can lead some California residents to move towards Colorado.

These central estimates carry substantial uncertainty, from both implementation choices and statistical estimation, which we decompose in the next section. The uncertainty is large at the state level: for example, under the gravity approach, Florida's projected population loss ranges from 2% to 10% once we account for statistical error, and that range widens further across implementation choices (Supplementary Fig. A4).

To assess the overall net-migration impact of climate change, we aggregate state-level estimates into a single migration share for each approach (Equation 4, which measures the total number of climate migrants as a share of US population), under our 1°C and 2°C warming scenarios (Fig. 3d). We find broad agreement on the rough magnitude. Across the choice grids, all four approaches project that warming generally relocates under one percent of the population over the period. The preferred causal inference, agent-based, and general equilibrium models clustere near a quarter of a percent, while the gravity model reports a share about threefold larger. For context, the models project that roughly 1.4–1.8% of the population moves state *each year* in the baseline, so climate-induced migration amounts to under three percent of the movement projected to occur anyway over the horizon. The migration share scales approximately linearly with total warming, so each approach's moderate-warming projection is close to half its extreme-warming one (Supplementary Fig. A5). A multi-model ensemble view is shown in Supplementary Fig. A6.

Explaining variance in the results

Our projection of the aggregate migration share caused by climate change is sensitive to implementation choices. For each approach, the spread of the migration share across its choice grid is wide. Gravity, for example, ranges from below 0.1% to about 3% of the population depending only on which specification is used (Fig. 4a; both warming scenarios are shown in Supplementary Fig. A7). For every approach, the choice spread is as wide as, or wider than, the statistical 95% confidence band around the preferred model's target output (grey bars in Fig. 4a) — the band that studies conventionally report.

Implementation choices reshape not only the headline magnitudes but also the geography of the projected redistribution. Comparing each specification's map of state-level changes to that of the preferred model (Fig. 4b), the general-equilibrium map is essentially choice-invariant and the gravity map is fairly stable, but the causal-inference and agent-based maps reorganise substantially across the grid, with some causal-inference specifications even reversing the rank order which states gain vs lose population. The agent-based model is distinctive in where this

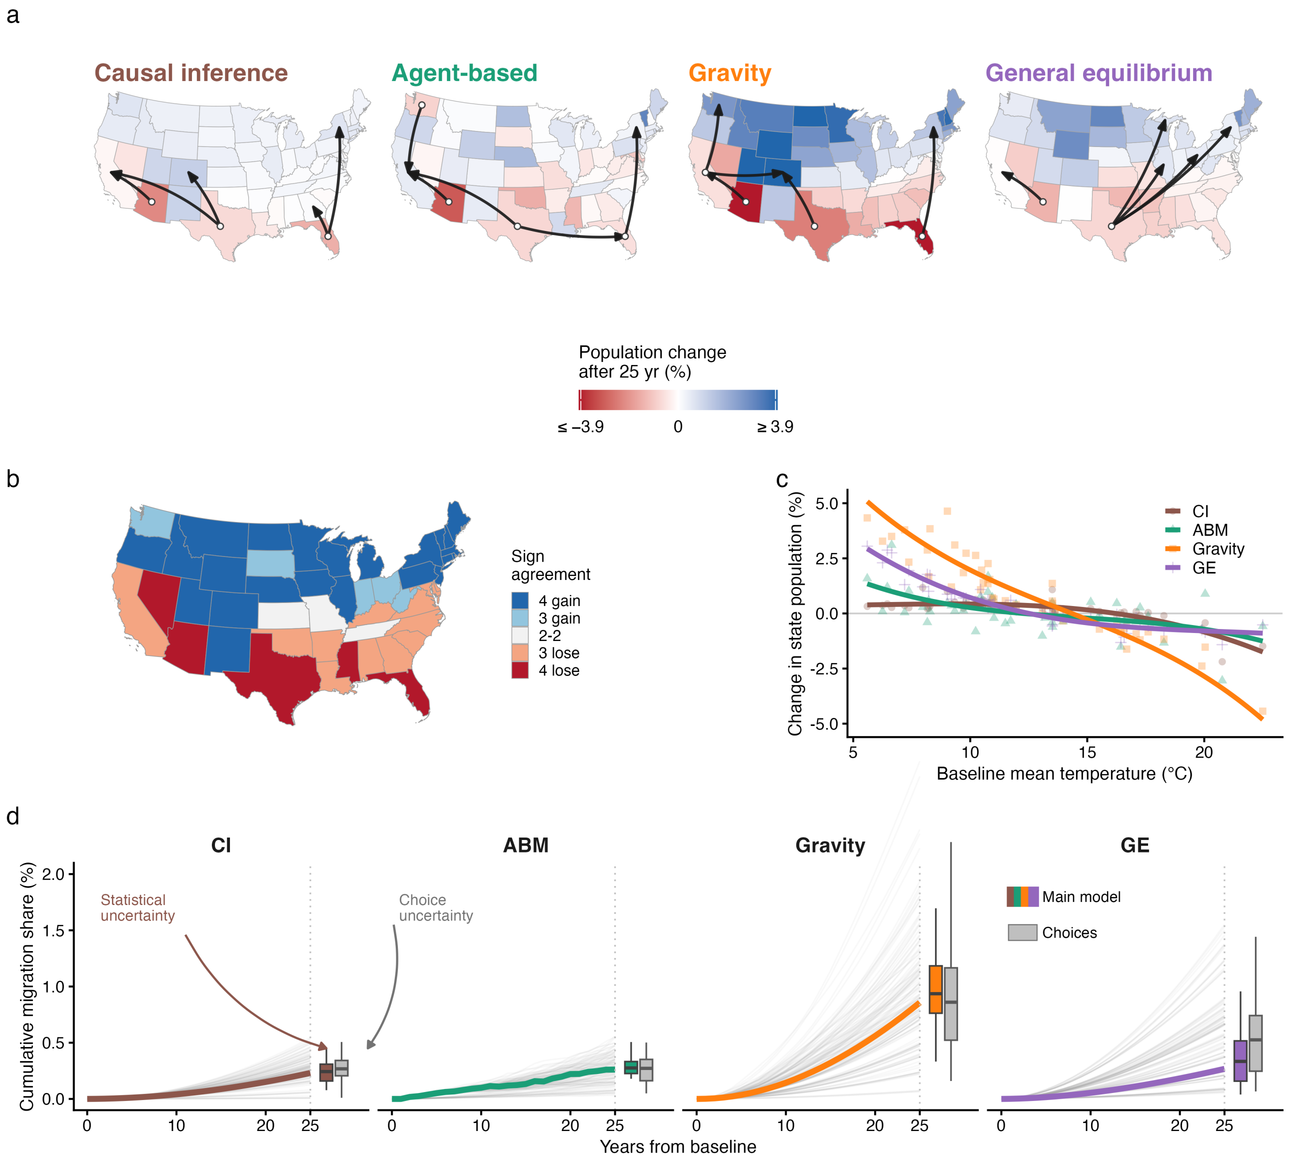

Figure 3: The projected redistribution of the US population under warming.
Each panel compares the four modelling approaches for the high-warming scenario (Supplementary Fig. A1). All changes are relative to a no-climate-change baseline. **(a)** State-level change in population caused by climate change (%) under each approach's preferred model (shading); arrows show each approach's five largest warming-induced net migration corridors, the change in flow between two states net of the change in the reverse direction (Supplementary Fig. A3 contrasts net corridors with gross flow increases and decreases). Per-state uncertainty is in Supplementary Fig. A4. **(b)** For each state, how many of the four approaches' preferred models agree on the *direction* of the projected change. **(c)** Baseline annual mean temperature against the preferred-model population change due to climate change (points) and a cubic fit per approach (lines). **(d)** Cumulative climate-induced migration as a proportion of US population (the common target, Fig. 2a; aggregated in Equation 4) over the 25-year horizon. Thin grey lines trace the implementation-choice grid for each approach (Fig. 6). At the horizon, box plots show *statistical* uncertainty (coloured) and *choice* uncertainty (grey). Boxes show the IQR, whiskers the 2.5–97.5% interval, the central line the median.

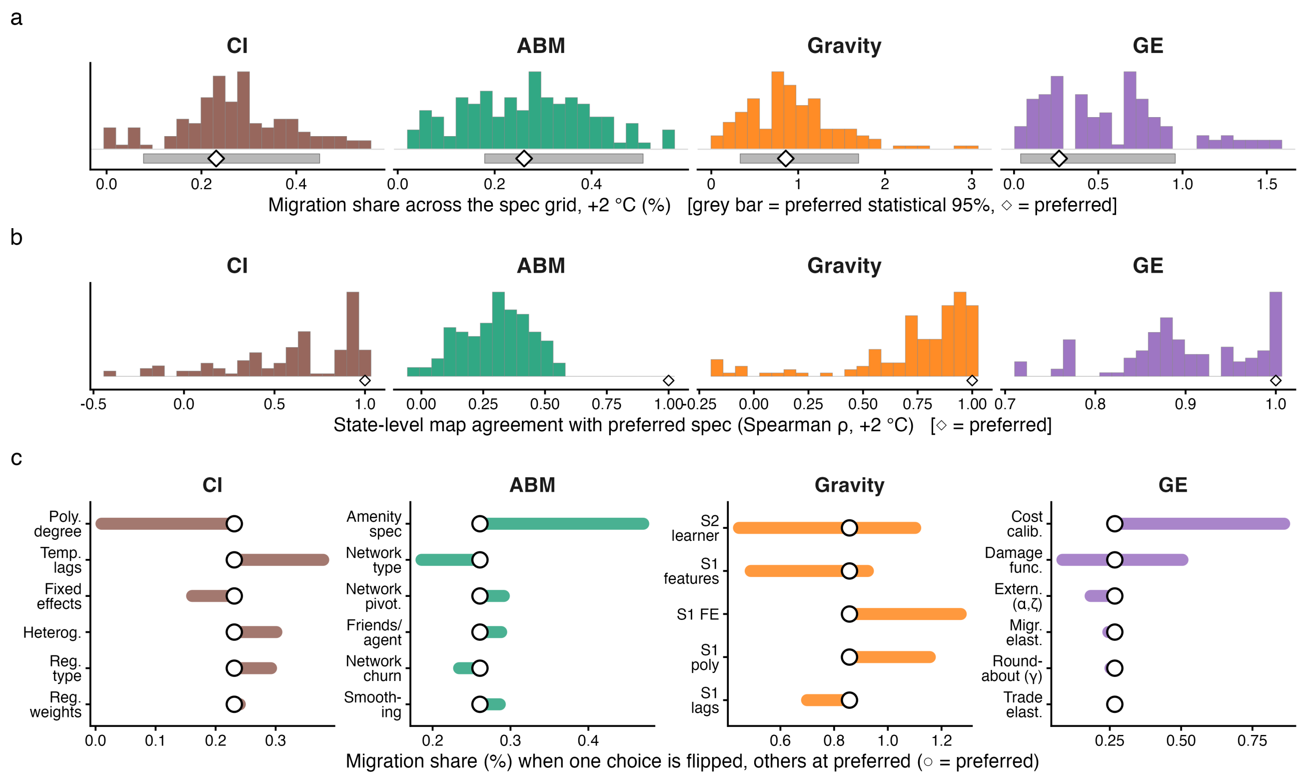

Figure 4: Which implementation choices drive the disagreement.
All panels use the extreme-warming scenario (Supplementary Fig. A1) and the headline *migration share*, i.e., the percentage of the population that relocates due to of warming, relative to the no-climate change baseline (Equation 4); outcomes for each approach are shown on independent horizontal scales. **(a)** Spread of the migration share across each approach's implementation-choice grid (Fig. 6). The grey bar beneath each shows the *statistical* 95% interval of the preferred model and the diamond (⋄) its point estimate. **(b)** How implementation choices reshape geography: the rank correlation (Spearman ρ) of each specification's state-level changes with those of the preferred model (⋄, $\rho = 1$ by construction). Values near 1 mean the map is insensitive to the choice; lower values mean it reorganises. **(c)** For each implementation choice, the range of the migration share as that single choice is varied with all others held at the preferred model (○ = preferred), ordered by range.

variance falls: dynamic social networks make its spatial pattern intrinsically choice-sensitive, so no single specification's map is representative of the others (the lowest cross-specification map agreement of the four approaches; Fig. 4b), even as the aggregate migration share stays comparatively stable.

A few implementation choices tend to drive the majority of the variance in projections across choices within each approach (Fig. 4c; the full per-axis distributions are in Supplementary Fig. A8). For causal inference, it is the polynomial degree of the temperature response; for the agent-based model, the amenity specification together with the structure of the social network; for gravity, the stage-2 machine-learning learner and the stage-1 feature set. For general equilibrium, the wide range of the target output is driven by the cost calibration choice and damage-function estimator, while the structural elasticities most debated in that literature — the trade and migration elasticities — barely move it. The choices that dominate the spread are therefore the ones where better empirical evidence would most sharpen projections, and our decomposition points to them approach-by-approach.

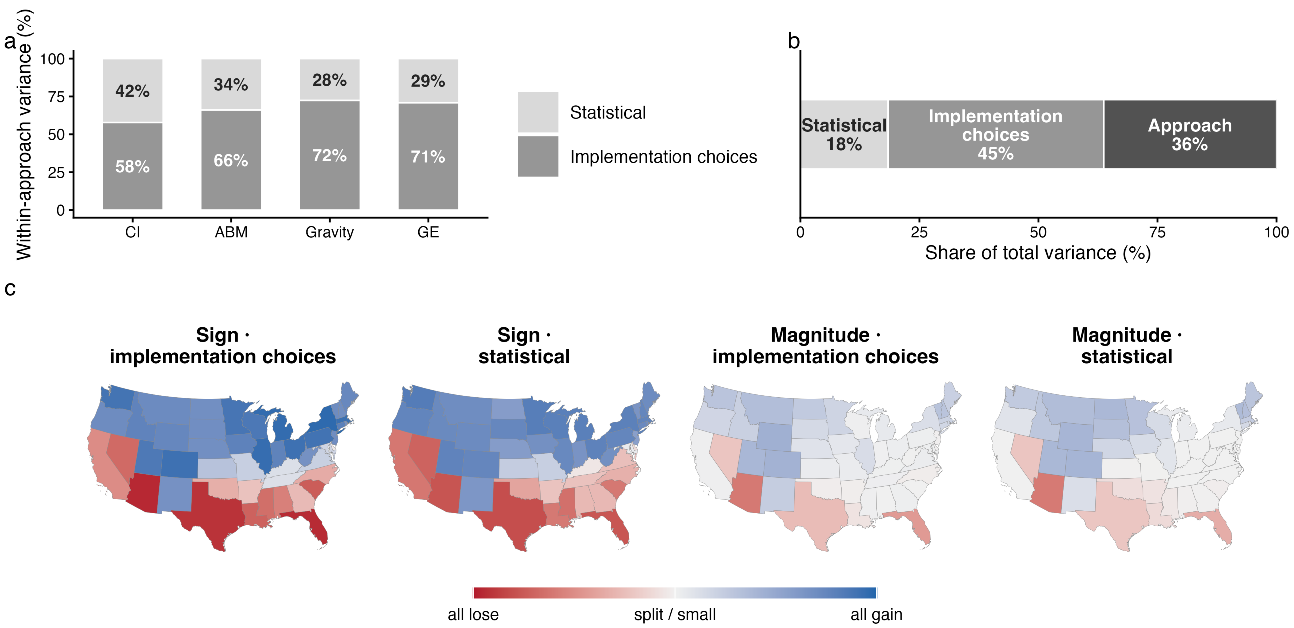

Figure 5: What the approaches agree on, and where their disagreement comes from. All panels use the extreme-warming scenario (Supplementary Fig. A1). (**a**) Within each approach, the share of the variance in the migration share attributable to statistical estimation versus implementation choice. The split is each component's variance as a share of their sum: implementation variance is taken across the choice grid at the parameter point estimates, statistical variance across the preferred model's bootstrap draws. (**b**) Decomposition of the total variance in the migration share across all approaches into three sources: the choice of approach, implementation choice within an approach, and statistical estimation. The approach component is the variance of the four preferred models' migration shares; the implementation and statistical components are the averages across approaches of the two variances defined in (a); each source's share of total variance is its component divided by the sum of the three. (**c**) Cross-approach agreement on the geography of the projected change, pooled with equal weight on the four approaches. The two left-hand maps measure agreement on the *direction* of each state's change; the two right-hand maps, agreement on whether it exceeds $\pm 2\%$ in *magnitude*. Within each pair the first map varies the implementation choice (across the choice grid) and the second varies the preferred model's statistical draw.

Decomposing the variance confirms that implementation choice is a first-order source of disagreement. Within approaches (Fig. 5a), choice uncertainty dominates statistical uncertainty for all approaches (causal inference, 58% vs 42%; agent-based, 66% vs 34%; gravity, 72% vs 28%; and general equilibrium, 71% vs 29%). Aggregating across all approaches (Fig. 5b), implementation choice and the choice of approach contribute almost equally, while statistical estimation accounts for just 18% of the cross-model variance in the projected migration share. In other words, the implementation choices researchers make *within* an approach explain as much of the overall disagreement as the choice of approach itself, and more than twice as much as statistical uncertainty, the only source of uncertainty studies conventionally report.

Beneath this disagreement over magnitudes is a robust qualitative consensus. Pooling across approaches, implementation choices and statistical draws, the ensemble agrees on the *direction* of each state's change — cold states gain, hot states lose — and this agreement is largely insensitive to both implementation choices and statistical error. The geography of where each source of uncertainty dominates is mapped in Supplementary Fig. A4.

Robust conclusions

Our case study highlights key areas of agreement and divergence among different modelling approaches regarding climate-induced migration patterns. The evaluation of qualitative statements across multiple modelling approaches is a key advantage of our approach intercomparison. To illustrate this, we qualitatively consider the extent to which our model results are consistent with four policy-relevant statements (summary in Supplementary Table A2).

First, there is strong consensus across every specification in our ensemble that 1–2°C of warming alone is unlikely to drive extreme population redistribution in the United States: all four approaches project that a maximum of a few percent of the population is displaced by warming over the 25-year horizon. As noted above, this is under three percent of the baseline interstate migration projected to occur anyway over the period. This does not mean that non-climatic drivers can be neglected when estimating climate-induced migration — our target nets them out by construction, as the difference between two worlds that share all drivers except the climate. Rather, it means the mid-century population map in the US will be shaped primarily by non-climatic forces, with warming acting as a relatively small but persistent and directional pull. Second, the approaches agree on the direction of redistribution (colder regions gain population and hotter regions lose it) and on the character of the dominant corridors, which mostly run from hot origins to cold destinations, although the specific top corridors differ across approaches (Fig. 3a, Supplementary Fig. A3).

We then consider the scope of uncertainty both within and across approaches. Accounting for both statistical and choice uncertainty, every approach admits substantially larger relocation than its central estimate (by a factor of roughly two to five) even though such outcomes remain unlikely. Finally, the balance of these two uncertainties is roughly consistent across approaches, despite each approach's choice grid enumerating a distinct set of implementation choices.

4 Discussion

We have brought together four approaches to modelling the climate-migration relationship that have until now largely been deployed in isolation. By recasting each as a different way of handling a common migration-rate generating function, we made each model estimate the same target output so that their disagreements become measurable. The resulting ensemble of more than 500 model specifications, each projected under two warming scenarios, produces internally consistent, state-level projections. This lets us decompose the uncertainty in a climate-migration projection into three sources: the choice of approach, the implementation choices made within it, and statistical estimation.

Our first result is that there is substantial agreement across the four approaches. Despite resting on different theories of behaviour and different computational machinery, our preferred models converge on a common picture. Warming over the coming decades redistributes a modest share of the US population; climate change causes states like Florida, Texas, and Arizona shrink, while Utah, Colorado, and Wyoming grow. That these distinct approaches agree on both the spatial pattern and the rough magnitude is a strong robustness check, and suggests that the broad climate-migration signal we recover is a property of the data rather than of a single modelling assumption. Part of this agreement may reflect common structure. The four approaches estimate the same migration-rate generating function from the same data, and in special cases collapse into one another (e.g., the GE model without market forces collapses into a gravity model, as does the ABM without network effects), so some directional agreement is unsurprising.

Agreement on the overall magnitude and pattern of climate migration masks how implementation and statistical uncertainty vary across the approaches. Residual uncertainty is dominated not by statistical estimation but rather by within-approach implementation choices: the often-unreported decisions about functional forms, covariates, and estimators.

These results have implications for how climate-migration evidence should inform policy. The robust conclusions are the kind decision-makers can build on, since they hold across methods rather than resting on a single contestable estimate. The precise magnitudes are less certain. For a given hot state, population losses due to climate change reach roughly 6% of baseline population in the most affected states, while the coldest states grow by up to roughly 6%, across the four approaches' preferred models, and the full uncertainty range admits movements two to five times the central estimate. Thus, policy design should be robust across this range rather than optimised to a single estimate: the directional findings identify which states should plan for additional arrivals or departures due to climate change, while housing, infrastructure, and service commitments can be scaled as realised flows are observed.

More broadly, our framework aims to provide the climate-migration field the infrastructure that reshaped climate and crop science: a common target and a transparent, reproducible ensemble into which any model can be placed and compared. Intercomparison projects such as CMIP for climate models (Eyring et al., 2016) and AgMIP for agricultural models (AgMIP, 2023) converted fragmented, hard-to-reconcile literatures into cumulative, policy-relevant science, and the same opportunity is now open here.[7] Seizing it will require the community to treat implementation choice as a key object of analysis, which is studied and documented rather than implicit, and to standardise targets and benchmark data so that new models are compared, and not merely accumulated.

One immediate dividend of the infrastructure we have built here is the field's first multi-model ensemble product: averaging the four preferred models yields an ensemble-mean map of projected redistribution, with the pooled spread across all 520 specifications as its accompanying uncertainty (Supplementary Fig. A6), the analogy of a CMIP ensemble mean. Our results also identify the key drivers of implementation-choice uncertainty within each approach (Fig. 4c), and the states where the approaches disagree most, directing future methodological work toward the choices, and the cases, that matter most. Future work could evaluate hybrid models against a common target, asking whether projections using combinations of approaches (e.g. reduced-form estimates disciplining structural models) can reduce implementation sensitivity.

Several limitations of our study are worth highlighting. We isolate the effect of annual mean temperature, yet other hazards (such as drought, windspeed, and flooding) also shape mobility and would broaden the relevant uncertainties. We examine internal migration in a high-income country with exceptional data; repeating the comparison in lower-income, agriculturally dependent settings, where climate-migration links are believed strongest (Cai et al., 2016), is an important direction for future research.[8] Because the framework holds the models and the target fixed, the same design extends directly to other settings, and replications elsewhere would test whether these findings are properties of the methods or of this application.

A further limitation is heterogeneity. The drivers of climate migration differ by occupation, age, and income (Benveniste et al., 2025), yet all four of our implementations capture heterogeneity only through aggregate state characteristics that we fix at baseline, not through the demographic composition of movers: the general equilibrium model distinguishes agents only through idiosyncratic location tastes, and the agent-based model's heterogeneity lies in social network position rather than demographics. Similarly, boundedly rational mechanisms common in the agent-based literature, such as memory, information sharing, and place attachment (Thober et al., 2018; Groeneveld et al., 2017; Schwarz et al., 2020), enter our implementation only implic-

[7] https://environmenthalfcentury.princeton.edu/research/2024/researchers-around-world-attend-climate-migration-workshop-princeton-0

[8] A first check on temporal stability is available within our own panel: we re-estimate each approach's preferred model independently on the two halves of the sample, and find that this leaves the direction of redistribution and the sub-one-percent magnitudes intact (Supplementary Fig. A9).

itly through the idiosyncratic shocks, and so cannot be varied systematically. Extending the intercomparison to richer, more flexible models is a priority for future work.

Several well-documented drivers of migration do not enter our approaches explicitly. These include cultural and social factors such as family ties and place attachment, political and institutional barriers, and demographic structure such as the concentration of moves in particular life stages. Because all four models are calibrated to observed data, the average influence of these factors is absorbed into the fitted baseline, so their imprint on who moves where today is retained. What our framework does not do is vary them. It is built to trace how one driver, temperature, reshapes migration. It is not designed to evaluate how migration would respond to climate and also simultaneous changes in, say, immigration policy or an ageing population. Extending the common target to other drivers is a natural direction for future intercomparison.

Further, our estimates of choice and approach uncertainty are themselves conservative. The grid of implementations we enumerate here is necessarily incomplete. The specifications within it are structurally similar perturbations of a common preferred model. The true spread of defensible models is wider than the choice uncertainty we report. Comparisons of choice uncertainty across approaches should also be treated with caution: we cannot verify that the four grids span comparably representative ranges of reasonable specifications, so differences in choice bands partly reflect differences in grid construction. Similar reasoning applies to statistical uncertainty: because sampling error is evaluated only at each approach's preferred model, the statistical bands are conditional on those same implementation choices, and may understate sampling error elsewhere in the grid (see Methods).

Across four disciplines and hundreds of model runs, the results here agree on the shape of climate migration in the United States. The largest remaining source of disagreement about is driven by the implementation choices researchers make and seldom disclose. Bringing those choices into the open, and comparing models on common terms, is the route to developing climate-migration projections that policymakers can trust.

Methods

Our comparison is organised around a framework that allows to consistently describe the four modelling traditions, so that each is adapted to report one common summary output, and they differ in their stated assumptions (Fig. 1). Here, we first set out that framework and the target quantity, then the data, then the four approaches and the implementation choices within each, and finally how we quantify and decompose uncertainty.

A common framework

Agents, indexed by a, live in locations $i \in \{1, \ldots, N_i\}$ and in each period $t \in \{1, \ldots, N_t\}$ decide whether to stay or move. The measure of agents residing in i at t is L_{it}. Agents choose the location that maximises utility, with the chosen location j^*_{ait} determined by

$$j^*_{ait} = \mathrm{argmax}_j \, U_{aijt}(\mathbf{T}_{jt}, \mathbf{X}_{jt}, \mu_{ijt}, \epsilon_{ajt}) \tag{1a}$$

$$\mathbf{X}_{jt} = \mathbf{X}_{jt}(\mathbf{T}, \mathbf{L}) \tag{1b}$$

$$L_{jt} = \sum_a 1(j^*_{ait} = j) \tag{1c}$$

$$\epsilon_{ajt} \sim g(\cdot). \tag{1d}$$

Four kinds of variable enter the location decision (1a). Weather $\mathbf{T}$ (here temperature) affects utility directly. A vector of non-weather variables $\mathbf{X}$ (economic and social conditions) affects it too, and may itself depend on weather and on the population distribution (1b), allowing for spatial linkages, dynamic feedback and adaptation. Migration costs μ_{ijt} capture the tangible and intangible costs of moving $i \to j$, and an idiosyncratic shock component $\epsilon_{ajt} \sim g(\cdot)$, which captures factors specific to the agent, such as personal preferences, life-course events, and social or cultural ties not represented elsewhere in the model. The remaining two equations make the system dynamic and self-contained: the accounting identity (1c) sums individual choices into the location populations $\mathbf{L}$ — the object whose warming-induced change $\Delta\mathbf{L}$ is our target — while the shock distribution g (1d) converts the deterministic argmax (1a) into smooth migration probabilities, the bridge to the migration-rate generating function below. Because $\mathbf{X}$ depends on $\mathbf{L}$ (1b), aggregation feeds back: today's moves reshape the conditions that drive tomorrow's. Climate forcing thus acts both directly, through $\mathbf{T}$, and indirectly, through $\mathbf{X}$.

The target output and the migration-rate generating function

Because ϵ_{ajt} is random, the probability that an agent moves $i \to j$ in period t can be written as

$$m_{ijt} = P\Big(U_{aijt} \geq \max_k U_{aikt}\Big) = f(\mathbf{T}, \mathbf{X}, \mu; \boldsymbol{\beta}), \tag{2}$$

where f is the *migration-rate generating function* (MRGF), $\mu := \{\mu_{ijt}\}$ collects migration costs, and $\boldsymbol{\beta}$ is a vector of unknown parameters.[9] We write $m_{ijt}(\mathbf{T}) = f(\mathbf{T}, \mathbf{X}, \mu; \boldsymbol{\beta})$ for the migration rate under climate $\mathbf{T}$, leaving the other arguments implicit. The function f is generic and unknown. Each approach is a different way of modelling, estimating, and projecting it (Fig. 2).

Collecting the rates into the (column-stochastic) matrix $\mathbf{m}_t(\mathbf{T}) = [m_{ijt}(\mathbf{T})]_{ji}$ (column i holds the rates out of i) and letting $\mathbf{L}_t(\mathbf{T})$ be the population vector under climate $\mathbf{T}$ with $\mathbf{L}_0(\mathbf{T}) = \mathbf{L}_0$ observed at baseline, the population distribution evolves as $\mathbf{L}_t(\mathbf{T}) = \mathbf{m}_t(\mathbf{T})\,\mathbf{L}_{t-1}(\mathbf{T})$. Our common *target* is the warming-induced change in the population distribution at a fixed horizon $\bar{t}$,

$$\Delta\mathbf{L}_{\bar{t}}(\mathbf{T}) = \mathbf{L}_{\bar{t}}(\mathbf{T}) - \mathbf{L}_{\bar{t}}(\mathbf{T}_0) = \left(\prod_{t=1}^{\bar{t}} \mathbf{m}_t(\mathbf{T}) - \prod_{t=1}^{\bar{t}} \mathbf{m}_t(\mathbf{T}_0)\right)\mathbf{L}_0 \tag{3}$$

the difference between where people are projected to live under warming and where they would have lived under a baseline climate, with the annual transition matrices applied in chronological order, the most recent on the left (see a conceptual illustration in Fig. 2a).

The national *migration share* aggregates these state-level changes. Because the population is closed (no births, deaths, or international migration), total population is conserved, so aggregate gains equal aggregate losses, so summing absolute state-level changes counts each move twice; we report half the summed absolute change as a percentage of the baseline national population,

$$\text{migration share} = \frac{100}{2}\,\frac{\sum_i |\Delta L_{i,\bar{t}}|}{\sum_i L_{i,0}}, \tag{4}$$

where $\Delta L_{i,\bar{t}}$ is the state-i component of $\Delta\mathbf{L}_{\bar{t}}$ and $L_{i,0}$ its baseline population. This migration share is the headline metric reported in Figs. 3d, 4 and 5.

Data and setting

We estimate the target for internal migration between the 48 contiguous US states. Bilateral state-to-state flows are derived from IRS Statistics of Income data, which record changes in tax-filing location and the associated exemption counts (a proxy for individuals); the IRS panel

[9] The parameter vector $\boldsymbol{\beta}$ is not primitive: it is induced by the structure of the decision problem in (1a), so specifying the utility model determines both the functional form of f and the content of $\boldsymbol{\beta}$.

covers above 95% of the tax-filing universe (Hauer and Byars, 2019). We restrict to 1990–2010, before a documented change in data quality (DeWaard et al., 2022), and combine the flows with population-weighted annual mean temperatures from ERA5 reanalysis, real gross state product, bilateral distances, interstate trade flows, and Facebook's Social Connectedness Index. Future temperatures come from Climate Impact Lab projections under the SSP3-7.0 scenario, considered at two warming scalings (moderate, $\approx +1\,°\mathrm{C}$, and extreme, $\approx +2\,°\mathrm{C}$, at the horizon; Supplementary Fig. A1 and Section B.3). All four approaches project from a common launch point: an observed baseline cross-section of populations, with the warming anomaly phased in linearly over the 25 years to the horizon (Section B.3). Estimated parameters and exogenous non-climate inputs are held at their baseline values, while model-internal responses (equilibrium wages and prices in the general-equilibrium model, network evolution in the ABM) update endogenously; each approach applies its own annual update rule, chronologically from the baseline.

The full set of datasets, and which approach uses each, is documented in Table A1. Variable construction and sources are detailed in Supplementary Methods Section B.1, with summary statistics in Table A3.

We define the migration rate as $M_{ijt} = L_{ijt}/\sum_k L_{ki,t-1}$. Here L_{ijt} counts the agents who move from i to j in year t, and the denominator $\sum_k L_{ki,t-1}$ is the population resident in i at $t-1$; the sum over prior locations k includes $k = i$, so stayers are counted. For example, taking i as Arizona and j as California, M_{ijt} is the number of tax filers moving from Arizona to California in year t as a share of Arizona's filer population in year $t-1$. This definition matches the timing of our model: M_{ijt} is a random variable with $E(M_{ijt}) = m_{ijt}$, providing a basis for estimating $\boldsymbol{\beta}$ in any approach (Supplementary Methods, Section B.2). As M_{ijt} is non-negative, the model also permits working with log migration flows, under which the heavily skewed empirical rates are roughly normal.

The four approaches

Although all four approaches estimate the same target output, they differ in how they identify it. Causal inference and gravity are *reduced-form*: they identify the net response of the population distribution to warming, without attributing it to underlying channels. The agent-based and general equilibrium models are *structural*: they separately identify how warming alters the relative attractiveness of each location — through amenities, and, in the general-equilibrium case, wages — from the mobility parameters governing how readily agents respond to it. The approaches differ, finally, in whether warming in one location is permitted to affect migration elsewhere: only the general-equilibrium and agent based models admit such cross-state spillovers, through market and network linkages. Supplementary Section B.4 formalises these distinctions.

Figure 6 sets out each approach as a staged pipeline and explains the modelling choices that branch off each stage. Here, we briefly describe each approach in turn, deferring the full specifications to the Supplementary Information.

Causal inference

Causal-inference (CI) models do not model behaviour explicitly; they use random year-to-year weather variation to identify the local slope of f with respect to temperature and then extrapolate. We estimate

$$\log E(M_{ijt}) = h(\mathbf{T}_{it}; \boldsymbol{\beta}^{CI}) + \gamma^{CI} \cdot \mathbf{W}_{it}, \qquad \forall\, i \neq j, \tag{5}$$

where h is a polynomial in (lagged) origin temperature and $\mathbf{W}_{it}$ collects fixed effects (including the always-included origin–destination pair effect), origin- and destination-specific trends, and a quadratic precipitation control, which together render the remaining temperature variation as good as randomly assigned. The OLS variant of the choice grid instead models the rate M_{ijt} in levels. Identification is local, in the sense that the recovered $d\log f/d\mathbf{T}$ estimates the in-sample, direct, *causal* effect of a temperature shock on migration. To use this estimate to project the target output requires the assumption that conditional on the fixed effects and time controls, one origin's out-migration is unaffected by other locations' temperatures (a stable-unit-treatment-value assumption) and that the in-sample slope applies to counterfactual climate regimes. We project by anchoring on the in-sample average rate and applying a first-order temperature adjustment, $\widehat{m}_{ijt} = \bar{m}_{ij}\,[\,1 + h(\mathbf{T}_{it}; \hat{\boldsymbol{\beta}}^{CI}) - h(\bar{\mathbf{T}}_i; \hat{\boldsymbol{\beta}}^{CI})\,]$, where $\bar{m}_{ij}$ is the in-sample average rate and $\bar{\mathbf{T}}_i$ the in-sample average origin temperature, with negative off-diagonal rates floored at zero and the diagonal fixed by the column-sum constraint (the OLS variant updates additively). The full form is in Section B.5, with the estimated response in Fig. A10). Six choices make up the CI grid (Table A4, choices 1–6 in Fig. 6b): regression type (OLS or Poisson), the fixed-effect/trend block, regression weights, the polynomial degree, the temperature lag structure, and directional heterogeneity (a separate temperature response by whether the destination is colder than, similar to, or hotter than the origin).

Agent-based

Our agent-based model (ABM) retains the most agent-level heterogeneity and endogenises agents' social networks. It is an amenity-only multinomial-logit model in which climate enters *only* through a destination-temperature amenity. This focuses the experiment on a single behavioural channel and distinguishes the ABM from the GE model, which also damages produc-

a

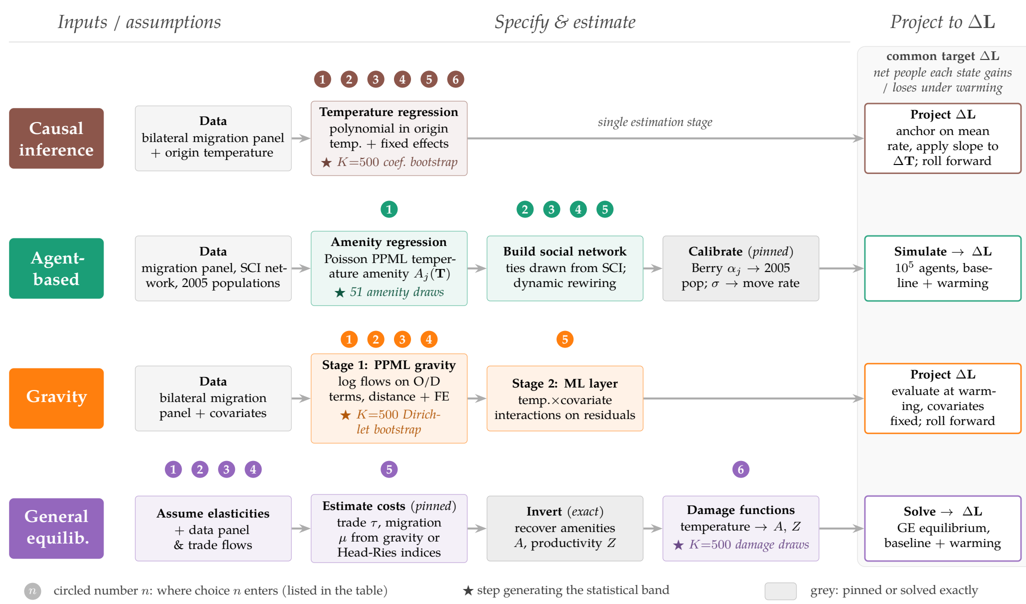

b

Causal inference (= 128 models)

No.	Choice	*What it controls*	Options (***preferred*** *bold*)
1	Regression type	functional form for flows; treatment of zero flows	OLS **Poisson PPML**
2	Fixed effects & trends ($\mathbf{W}_{it}$)	which confounding variation is absorbed (identification)	pair + year **+ origin- & dest.-specific trends**
3	Temperature polynomial	curvature of the temperature response	1 2 **3** 4
4	Regression weights	whether populous states dominate the fit	unweighted **population-weighted**
5	Temperature lags	delayed response to past temperature	**contemporaneous** + one-year lag
6	Direction-of-move interaction	whether the response depends on destination climate	**none** climate-similarity buckets

Agent-based (= 120 models)

No.	Choice	*What it controls*	Options (***preferred*** *bold*)
1	Amenity specification	how the climate signal (temperature amenity) is estimated	**pop.-weighted Poisson amenity** unweighted Poisson pop.-weighted + destination trend
2	Network type	structure of social ties that guide moves	**small-world (Facebook SCI)** random
3	Friends per agent	density of each agent's social network	**10** 25
4	Within-state smoothing	how strongly friend locations are smoothed	**0.01** 0.10
5	Network-rewiring regime	whether ties evolve as agents move (amplifies the signal)	static dynamic (σ_0, κ): $(.3, .5)$, $(.3, 1.5)$, $(.5, .5)$, $(\mathbf{.5}, \mathbf{1.5})$

Gravity (= 144 models)

No.	Choice	*What it controls*	Options (***preferred*** *bold*)
1	Stage-1 fixed effects	which confounding variation is absorbed	full FE drop-origin FE **drop-origin + state trends**
2	Stage-1 feature set	which covariates the flow model uses	**base** + inequality & precipitation + SCI & log-population
3	Temperature polynomial	curvature of the temperature response	**quadratic** cubic
4	Temperature lags	delayed response to past temperature	1 **2**
5	Stage-2 ML layer	richness of nonlinear temp.×covariate effects	none **lasso** neural network gradient boosting

General equilibrium (= 128 models)

No.	Choice	*What it controls*	Options (***preferred*** *bold*)
1	Migration elasticity (ε)	how readily people relocate for a given gain	1 **3**
2	Trade elasticity ($\sigma - 1$)	substitutability of goods across states	**4** 8
3	Externalities (α, ζ)	agglomeration / congestion feedbacks	**off** $(0, 0)$ on $(0.3, 0.05)$
4	Roundabout share (γ)	input–output linkages that amplify shocks	**0** 0.5
5	Cost calibration ($\{\mu\}, \{\tau\}$)	the size and pattern of trade and migration costs	**smoothed (estimated on distance)** Exactly fit (Head-Ries indices)
6	Damage-function spec.	how temperature maps to productivity & amenity damage	**quadratic, levels-with-lag** quadratic first-difference cubic levels-with-lag quadratic pop.-weighted

Figure 6: The four modelling approaches and the implementation choices within each.
(**a**) Each approach runs left to right, converging on the shared target output. Circled numbers mark the step at which each implementation choices enter; ★ marks the step whose resampling generates that approach's statistical-uncertainty band. (**b**) For each approach a single *preferred* model (values in bold) is varied across a grid of alternatives. Each numbered choice matches the identically numbered step in (a).

tivity. Here 'amenity' is a reduced-form measure of a location's desirability: it absorbs income, wages, cost of living and other attributes that make a place attractive, all of which are pinned down in the baseline calibration.

Each agent maximises

$$U_{aijt} = \gamma^{ABM} \log G_{aj,t-1} + A_{jt}(\boldsymbol{\beta}^{ABM}, \mathbf{T}_{jt}) + \alpha_j^{ABM} - \delta^{ABM}(d_{ij}) + \sigma^{ABM}\, \epsilon_{aijt} \qquad \epsilon_{aijt} \sim \text{Gumbel}(0, 1), \tag{6}$$

where $G_{aj,t-1}$ is the smoothed share of the agent's social ties resident in j (the smoothing weight is a swept implementation choice), $\boldsymbol{\beta}^{ABM}$ the amenity coefficients common across states, A_{jt} a temperature amenity, α_j^{ABM} a location effect absorbing j's time-invariant attractiveness, such as its labour market, cost of living, and non-climate amenities and $\delta^{ABM}(d_{ij})$ a distance cost (the full specification is in Section B.6). The behavioural parameters are estimated *outside* the simulation by a Poisson amenity regression; the location effects α_j^{ABM} are then re-fit by way of a "Berry inversion" (Berry et al., 1995) so the model reproduces observed 2005 populations, and the logit scale σ^{ABM} is calibrated to the observed year-1 move rate (Section B.6.4, Fig. A12). Networks are dynamic: ties are re-sampled as agents move, with a rich-get-richer rule. We then simulate a population of $100{,}000$ agents forward under baseline and warming temperatures, averaging outcomes over 10 independent simulation runs. Five choices make up the ABM grid (Table A5): the amenity specification, the network type, the number of ties per agent, a within-state smoothing parameter, and the network-rewiring regime (Fig. 6b). The network rewiring introduces path-dependence, suggesting that without conditioning on a sequence of shocks, the model may not have a unique steady-state.

Gravity

Our gravity model treats migration flows as a function of observable origin, destination, and bilateral characteristics. While it abstracts from equilibrium feedback, it pairs a standard estimator with a machine-learning layer to allow for flexible nonparametric heterogeneity. A Type-I extreme-value assumption on ϵ_{ajt} delivers a multinomial-logit migration probability whose ratio to the stay probability gives the estimating equation

$$\frac{E(M_{ijt})}{E(M_{iit})} = \exp\Big(\gamma^{Gr,d}\mathbf{X}_{jt} + h(\mathbf{T}_{jt}; \boldsymbol{\beta}^{Gr,d}) - \big(\gamma^{Gr,o}\mathbf{X}_{it} + h(\mathbf{T}_{it}; \boldsymbol{\beta}^{Gr,o})\big) - D(d_{ij}; \delta^{Gr})\Big), \tag{7}$$

where $D(d_{ij}; \delta^{Gr})$ proxies moving costs. We fit this in two stages: first, we estimate a Poisson pseudo-maximum-likelihood (PPML) regression equation that is log-linear in time-varying origin and destination covariates, distance, and origin fixed effects (Stage 1). In Stage 2, we fit a regularised machine learner by way of cross-validation, treating the Stage-1 fit as an offset.

This procedure allows the estimated temperature response to interact flexibly with the other features. The machine learning step introduces no new variables; its purpose is only to introduce data-driven flexibility in the response function.

We use machine learning in the gravity approach because prediction is that tradition's estimation objective. In causal inference, a regularised prediction layer would change the estimand[10], while the agent-based and general equilibrium approaches require parameters with behavioural interpretations.[11] What learner we use is one of our implementation choices, so its contribution to uncertainty is measured as choice uncertainty within the gravity approach. All learners share the same training panel and a single year-blocked five-fold partition for cross-validation so implementation choice variation reflects the learners rather than the data splits.[12]

Projection in the gravity approach plugs counterfactual temperatures into the fitted model, holding non-climate covariates fixed (full specification in Section B.7, with the pipeline in Fig. A15). Five choices make up the gravity choice grid (Table A6): the Stage-1 fixed-effect block, the Stage-1 feature set, the temperature polynomial degree, the temperature lag structure, and the Stage-2 learner (none, lasso, neural net, or gradient boosting; Fig. 6b).

General equilibrium

Our general equilibrium (GE) model places market feedback at the centre: it is a dynamic quantitative spatial model with bilateral migration costs (Desmet and Rossi-Hansberg, 2024; Bryan and Morten, 2019). Agents choose locations on real wages w_j/P_j and amenities A_{jt} with Fréchet-distributed idiosyncratic shocks; goods are produced from local productivity Z_{jt} and population, and traded subject to iceberg shipping costs τ_{ij}^{GE}; and productivities and amenities respond to temperature through damage functions, so the population distribution emerges from an equilibrium in which choices and local conditions are mutually consistent. Absent input-output linkages, the equilibrium of our model is guaranteed to be unique following the theorems of Allen and Donaldson (2020). Five structural parameters are not identified from migration data and must be assumed: the migration elasticity ε^{GE}, the elasticity of substitution σ^{GE} (whose value less one is the trade elasticity), the congestion (α^{GE}) and agglomeration (ζ^{GE}) externalities, and the roundabout intermediate-input share γ^{GE}. Conditional on these, we recover migration and trade costs from either structural gravity regressions (which project costs on distance, our preferred specification) or exact identification of the costs from the migration

[10]Although it is possible to use ML tools to search for heterogeneity in the response function; see Benveniste et al. (2025).

[11]In is possible to use machine learning to calibrate migration and trade costs in the general equilibrium setting; see Dingel and Tintelnot (2026).

[12]The lasso and gradient-boosted trees are cross-validated over the fixed data partitions, while the neural network trains on the first four blocks and uses the final block for early stopping.

and trade flows alone (using Head and Ries 2001 indices, which assume no measurement error in flows), invert the model for amenities and productivities (exactly, hence without statistical uncertainty), and estimate temperature damage functions; we then re-solve the equilibrium each period under baseline and warming temperatures (full specification in Section B.8, with the estimation pipeline in Fig. A16 and the estimated damage functions in Fig. A17). Six choices make up the GE grid (Table A7): ε^{GE}, σ^{GE}, the externality regime $(\alpha^{GE}, \zeta^{GE})$ (varied jointly, off or on), the roundabout share γ^{GE}, the damage-function specification, and the choice of how costs are calibrated, a known source of implementation uncertainty (Dingel and Tintelnot, 2026); see Fig. 6b.

Preferred models and the choice grid

Specifying an approach does not pin down a unique model: within each there remain many defensible implementation choices — covariates, functional forms, estimators — and these are not innocuous. For each approach we therefore define a single *preferred model*, chosen to be representative of its tradition, and enumerate a structured grid of reasonable alternatives around it (a total of 520 models across the four approaches; Fig. 6 and Section B.3). The preferred model is designed to satisfy two competing sets of objectives. On the one hand, it must fit with *intercomparison* objectives, in that it must fit in the common framework, estimate the common target, be computationally feasible, and be estimable from the shared dataset. On the other, it must satisfy *model-specific* objectives that keep it faithful to its own literature's norms. Our preferred model is subjectively chosen to satisfy both of these criteria. Varying the model across the grid, while holding the target, data and scenario fixed, isolates how much projections move with implementation choice.

Uncertainty

For every approach we report three kinds of uncertainty, defined identically across the four so they are comparable (Section B.3). None is the preferred model's point estimate. Statistical uncertainty reflects sampling error in the fitted parameters: we draw repeatedly from the sampling distribution of the *preferred model only*. We run $K = 500$ draws for causal inference, gravity and general equilibrium; 51 for the ABM where a single draw incorporates 10 runs using different draws of the shock and network parameters. The draws are taken using parametric bootstrap of the regression coefficients (CI), a Dirichlet bootstrap of the two-stage fit (gravity), draws from the damage-function coefficient covariance (GE), or draws of the amenity coefficients evaluated on the run-averaged population vector (ABM) — and re-estimate the target for each draw. Choice uncertainty reflects implementation choice: we re-estimate the target across the full grid (Fig. 6) at the parameter point estimate.

In each case the bootstrap targets the single step that estimates how temperature enters the model. This step varies by approach, it is the migration-rate response (causal inference, gravity), the destination-temperature amenity (agent-based), or the productivity and amenity damage functions (general equilibrium). Parameters that are calibrated or estimated outside this step, such as the trade and migration elasticities of the general equilibrium model, are held fixed across draws.

Two caveats apply to this decomposition. First, because statistical uncertainty is evaluated only at the preferred model, and sampling error varies across the grid, the reported statistical bands may understate statistical uncertainty elsewhere in the grid. Second, the boundary between statistical and choice uncertainty is itself a classification judgement: a setting such as the agent-based model's friends-per-agent count could be viewed either as an implementation choice or as an approximation to a quantity that better data would pin down. The decomposition should therefore be read as conditional on the definitions used here (see full model details in Supplementary Section B).

Acknowledgements

We thank Andrew Bell, Hélène Benveniste, Emile Esmaili, Filiz Garip, Peter Howard, Ishan Nath, Guglielmo Zappalà, the workshop participants at Princeton University's Center for Policy Research on Energy and the Environment (C-PREE), the Climate Migration Modeling Intercomparison Workshop, the Eastern Economics Association conference, the Managed Retreat conference, the Scenarios Forum, and the World Congress of Environmental and Resource Economists for their comments, suggestions and advice.

T.B gratefully acknowledges funding from C-PREE and the Dragon Capital chair at the University of Exeter.

Declarations

The authors declare no conflicts of interest relevant to this study.

Disclaimer

The views in this paper do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

Author contribution statement

M.O and N.C.L conceived of the study. All authors refined the study concept. T.B, J.R.K, N.C.L and T.X designed an initial version of the models and framework. T.B. and J.R.K refined and implemented the models and framework, and wrote the first draft. All authors contributed to revising the draft.

Data and code availability

All data used in this study is compiled from publicly available data sources. A full replication package that generates all figures and results used in the paper is available at `https://zenodo.org/records/21840473`. This also contains documentation on the source of all data used.

AI declaration

The authors used Claude to assist with writing code, and for help checking the draft and refining language. The authors take full responsibility for all content.

Supplementary Materials

Contents

A Supplementary Tables and Figures

A.1 Supplementary Tables

			Used by			
Dataset	Source / provider	Variable(s) and role	CI	ABM	Grav.	GE
IRS Statistics of Income migration	US IRS (via Hauer and Byars, 2019 replication)	Bilateral state-to-state migration flows; tax-exemption counts as the population/agent base (1990–2010)	•	•	•	•
ERA5 / ERA5-Land reanalysis	ECMWF / Copernicus (Hersbach et al., 2020; Muñoz-Sabater et al., 2021)	Population-weighted annual mean temperature (and precipitation) by state-year	•	•	•	•
Climate Impact Lab US projections	Climate Impact Lab (SSP3-7.0)	Projected state temperature under warming; the future-climate forcing	•	•	•	•
State economic activity	Mohaddes et al. (2023) replication package	Real gross state product			•	•
Social Connectedness Index	Meta / Facebook Data for Good (Oct. 2021)	State-pair social connectedness; anchors the friend network / network covariate		•	•	
Centers of Population 2010	US Census Bureau	State population centroids $\rightarrow$ bilateral great-circle distance (moving-cost proxy)		•	•	•
Commodity Flow Survey 2012 (PUMF)	US Census Bureau	Interstate trade flows $\rightarrow$ trade costs τ_{ij}				•
State inequality / income shares	Frank et al. (2015)	Inequality covariates (relative mean deviation)			•	

Table A1: Datasets used in the analysis by approach.

	CI	ABM	Gravity	GE	Ensemble
What the models imply about climate-driven migration					
Moderate warming (+1 °C) alone relocates more than 3% of the population	Low	Low	Low	Low	Low
Warming raises the populations of cold states and lowers those of hot states	High	High	High	High	High
Accounting for the full uncertainty, relocation could be at least double the central estimate	Medium	Medium	High	High	Medium
Where the disagreement comes from					
Implementation choices contribute as much uncertainty as statistical estimation	Medium	Medium	High	High	High

Table A2: Synthesis: strength of evidence for policy-relevant statements, by approach and pooled. Each cell reports the strength of evidence (Low / Medium / High) that the statement holds, evaluated over that approach's full ensemble of implementation choices (Fig. 6) and the preferred-model bootstrap. The Ensemble column pools across approaches. The first row is evaluated at the $\times 1$ scaling ($\approx$1 °C); all others at $\times 2$, A1). Levels are computed directly from the model ensemble: for the first two statements, the share of specifications supporting it (Low $< 1/3$, Medium $< 2/3$, High otherwise); for the third, the factor by which the full uncertainty range (choice grid or preferred-model bootstrap, 97.5th percentile) exceeds the central estimate (Low $< 1.5\times$, Medium $< 2.5\times$, High otherwise); for the fourth, the share of within-approach variance attributable to implementation choice rather than statistical estimation.

731 ## A.2 Supplementary Figures

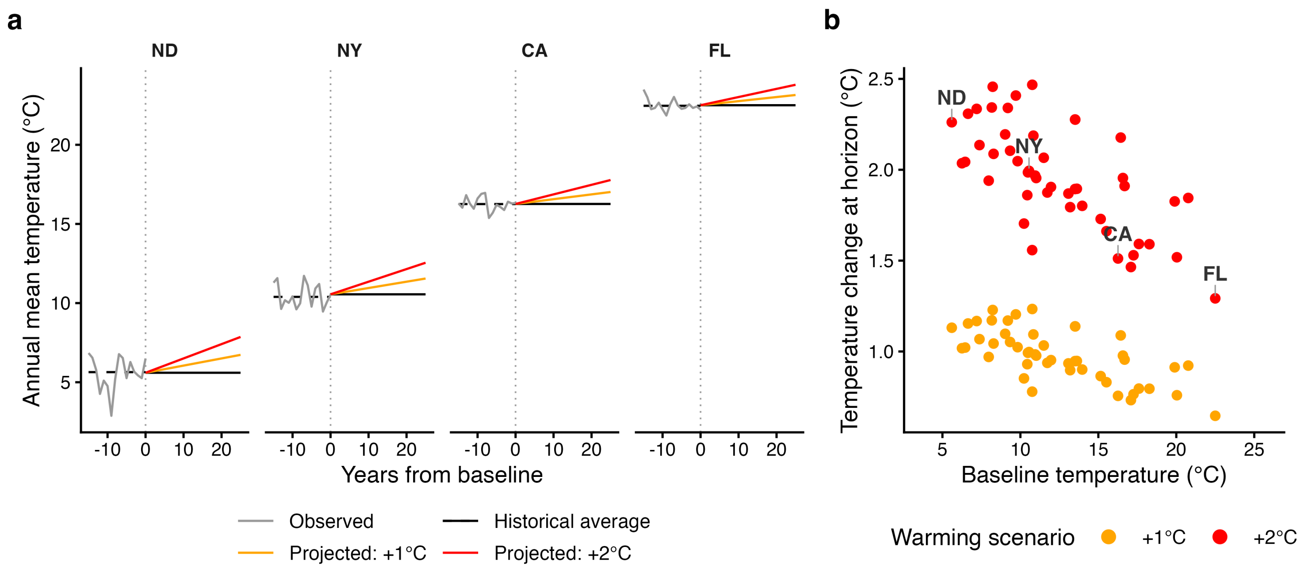

Figure A1: Temperature forcing used in the projections.
(**a**) Projected annual mean temperature at four representative states (North Dakota, New York, California, Florida) under the SSP3-7.0 scenario, at warming scalings of $\times 1$ ($\approx +1\,°C$) and $\times 2$ ($\approx +2\,°C$) by the projection horizon, shown against the observed series and the historical baseline mean. (**b**) Baseline state level annual mean temperature against the projected temperature change at the horizon, under both warming scenarios.

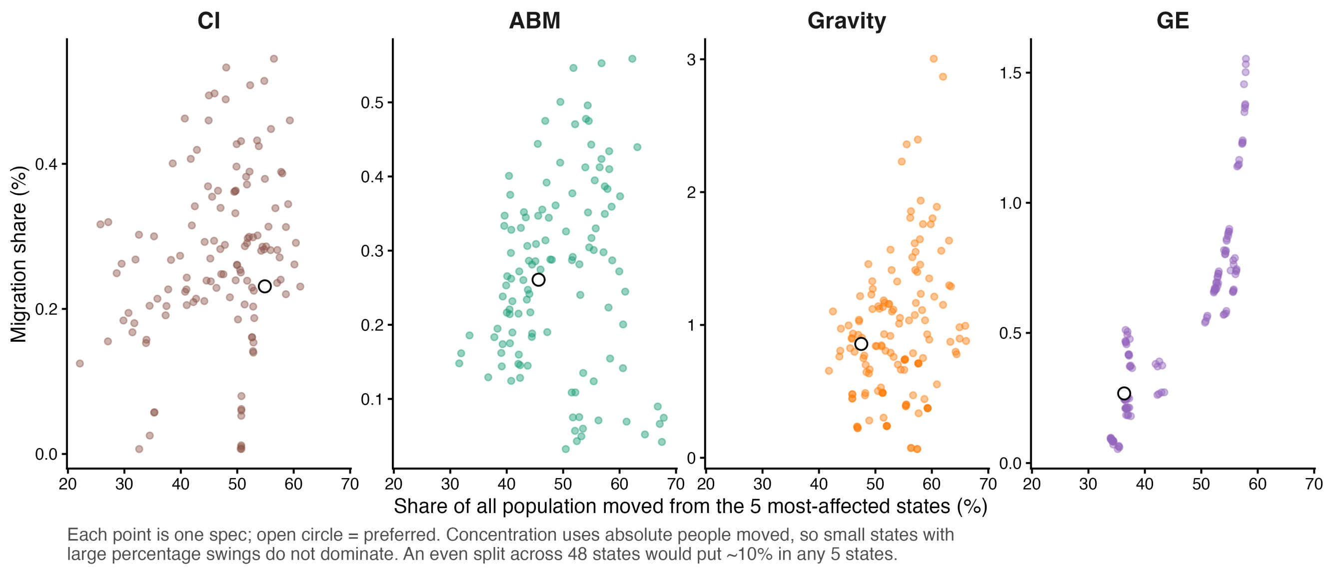

Figure A2: Geographic concentration of the redistribution.
Each point is one specification (extreme-warming scenario): the share of all population moved that is accounted for by the five most-affected states (horizontal) against the aggregate migration share (vertical); open circle = preferred. Concentration is measured in absolute people moved, so small states with large percentage swings do not dominate; an even split across 48 states would place ≈10% in any five. The redistribution is robustly concentrated in a handful of states across the entire choice grid.

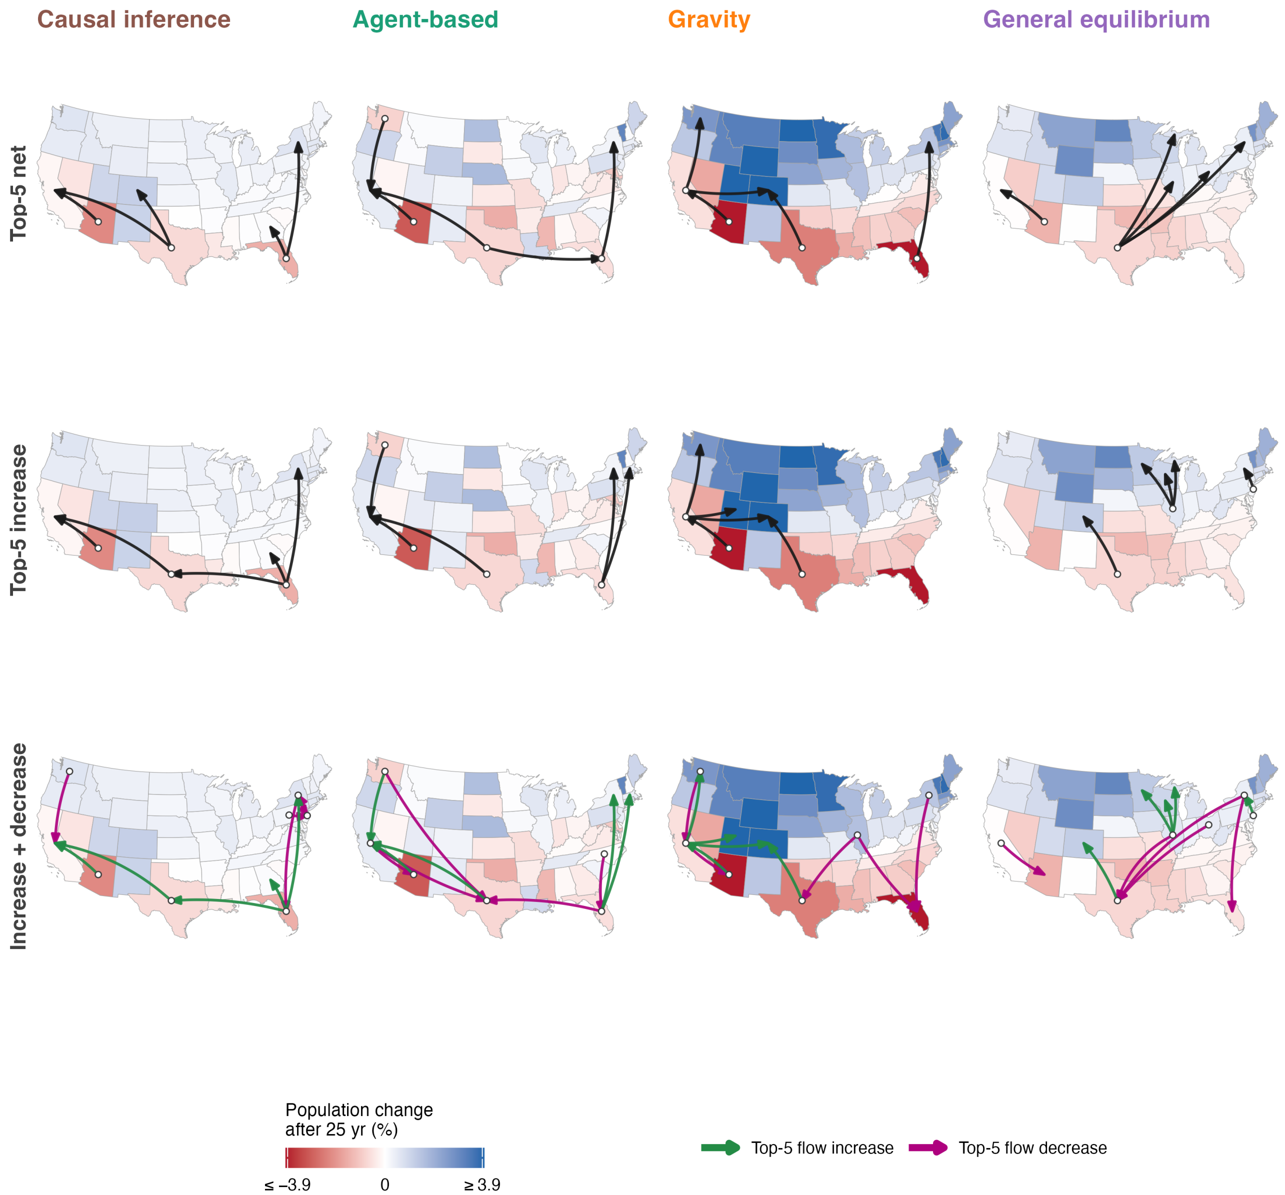

Figure A3: Corridors of climate-induced migration, by approach.
For each approach's preferred model under the extreme scenario (≈2,°C at the 25-year horizon), the three rows contrast three definitions of the top migration corridors. Top row: the five largest net migration corridors (the warming-induced change in cumulative migration flow from one state to another, net of the change in the reverse direction), as in Fig. 3a. Middle row: the five largest increases in gross directional migration flows. Bottom row: the five largest gross increases (green) together with the five largest gross decreases (magenta); decrease arrows are drawn along the diminishing origin-to-destination flow. Arrows summarise state-to-state flows. States are shaded by the projected population change as in Fig. 3. As the shading integrates over all corridors, a state can receive a top corridor yet lose population on net (e.g., California).

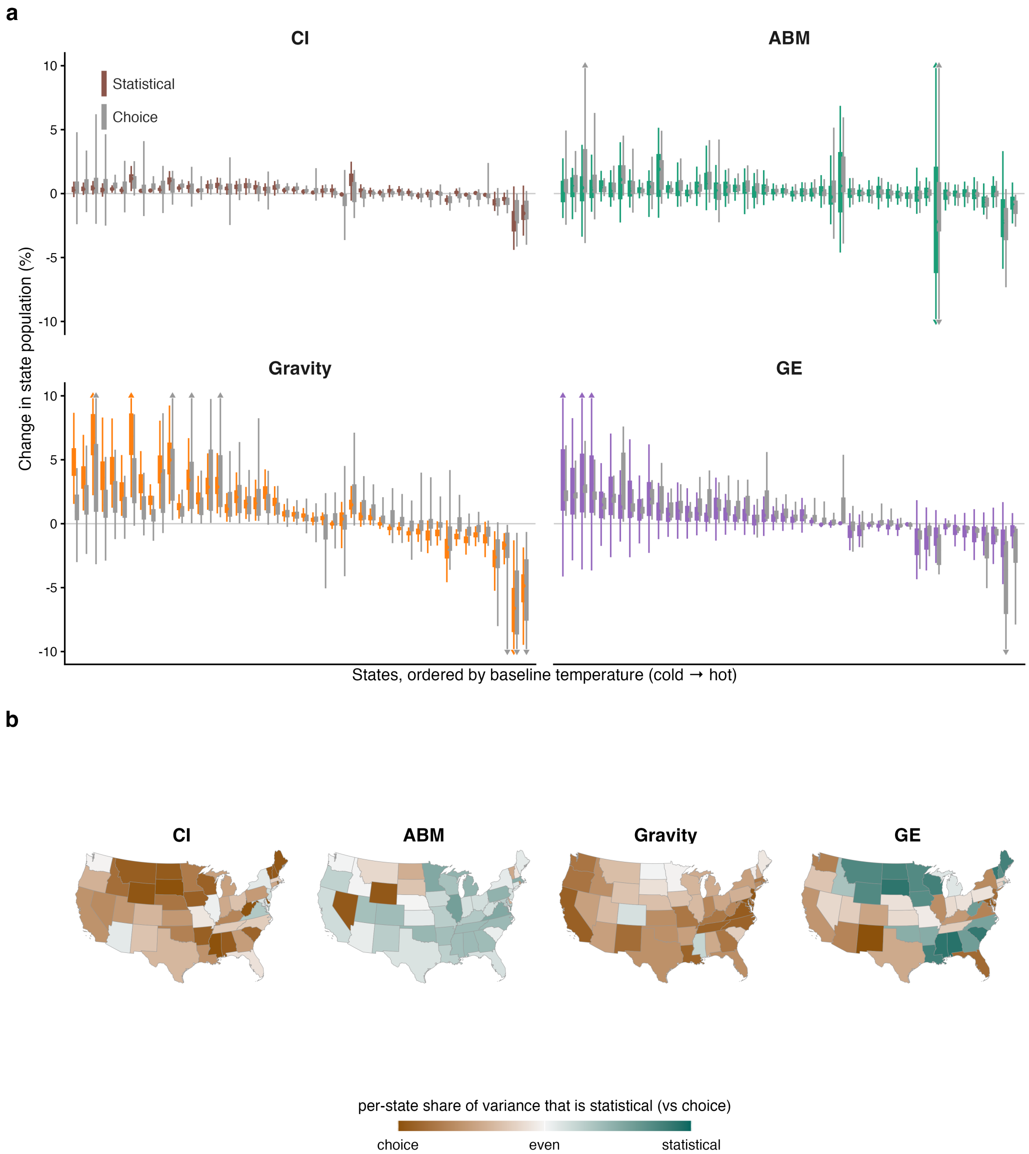

Figure A4: State-level statistical versus choice uncertainty in the projected population change, by approach.

Both panels use the extreme-warming scenario ($\approx +2\,°\mathrm{C}$ at the 25-year horizon) and the four approaches (causal inference, agent-based, gravity, general equilibrium). **(a)** Projected change in each state's population (%), with two adjacent error bars per state: *statistical* uncertainty (approach colour, left) for the preferred model, and *choice* uncertainty (grey, right) across the implementation grid. For each bar the thick segment is the interquartile range and the thin segment the 95% interval (2.5–97.5%); the point marks the median. States are ordered along the horizontal axis by baseline temperature (cold to hot), equally spaced and in the same order in every panel. The vertical axis is clipped to $\pm 10\%$ and bars extending beyond the axis are flagged with a triangle (chiefly under the gravity model, whose state-level responses are largest). **(b)** The share of the per-state variance in the population change that is statistical (preferred-model bootstrap) rather than implementation choice across approaches, $\mathrm{var}_{\mathrm{stat}}/(\mathrm{var}_{\mathrm{stat}} + \mathrm{var}_{\mathrm{choice}})$. This is not directly comparable to the aggregate choice-versus-statistical split in Fig. 5a as it is at the state level and variance shares are not easily aggregable.

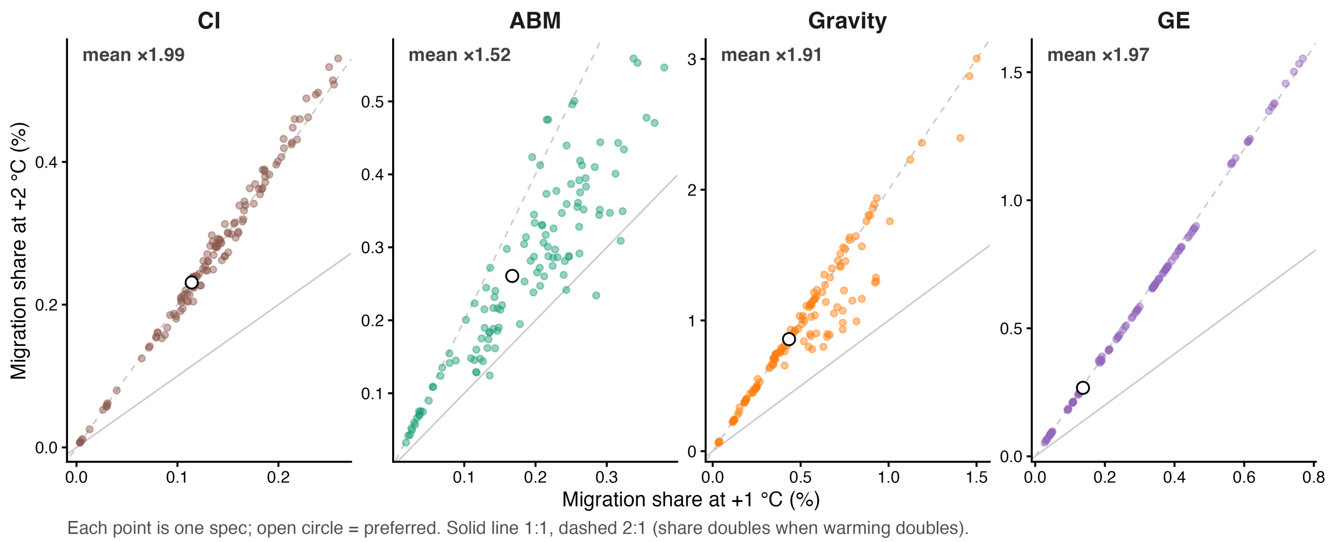

Figure A5: Linearity of the migration response in warming intensity.
Each point is one specification: its migration share under moderate warming (+1 °C, horizontal) against extreme warming (+2 °C, vertical); open circle = preferred model. The solid line is 1:1 and the dashed line 2:1, so a cloud lying along the 2:1 line means doubling the warming anomaly doubles the migration share.

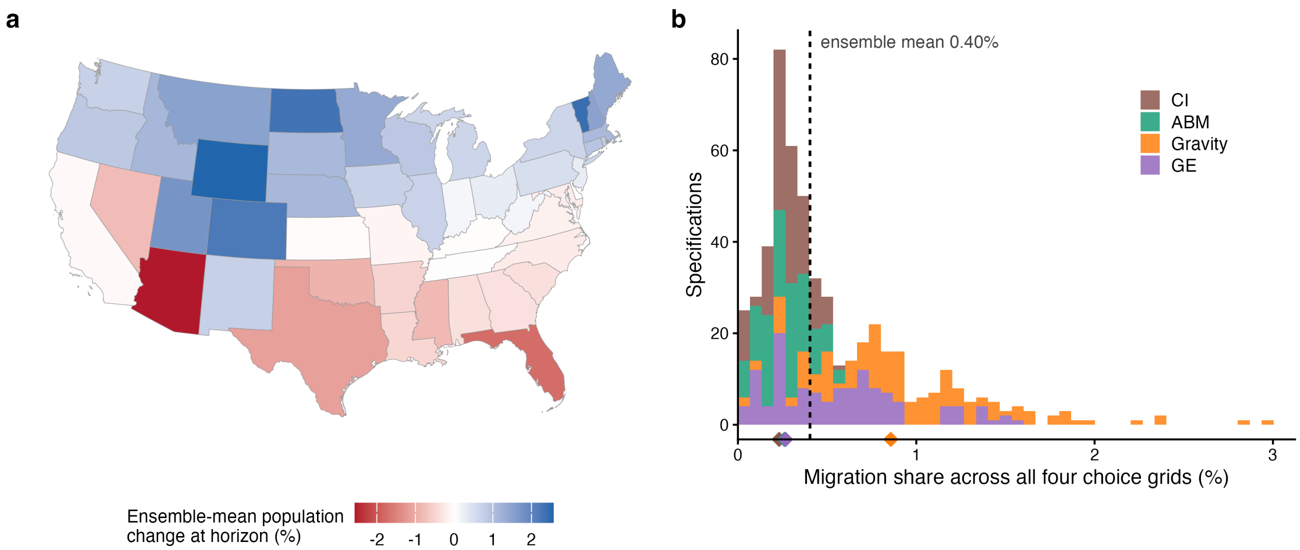

Figure A6: Multi-model ensemble view of the projected redistribution.
(**a**) The equal-weight ensemble mean of the four preferred models' projected population change at the 25-year horizon under the extreme scenario (≈2 °C) for each state. (**b**) The ensemble migration share: the distribution of the migration share pooled across all four approaches' implementation choices (520 specifications). Diamonds below the axis mark each approach's preferred model. The dashed line is the equal-weight mean of the four preferred models (0.40%).

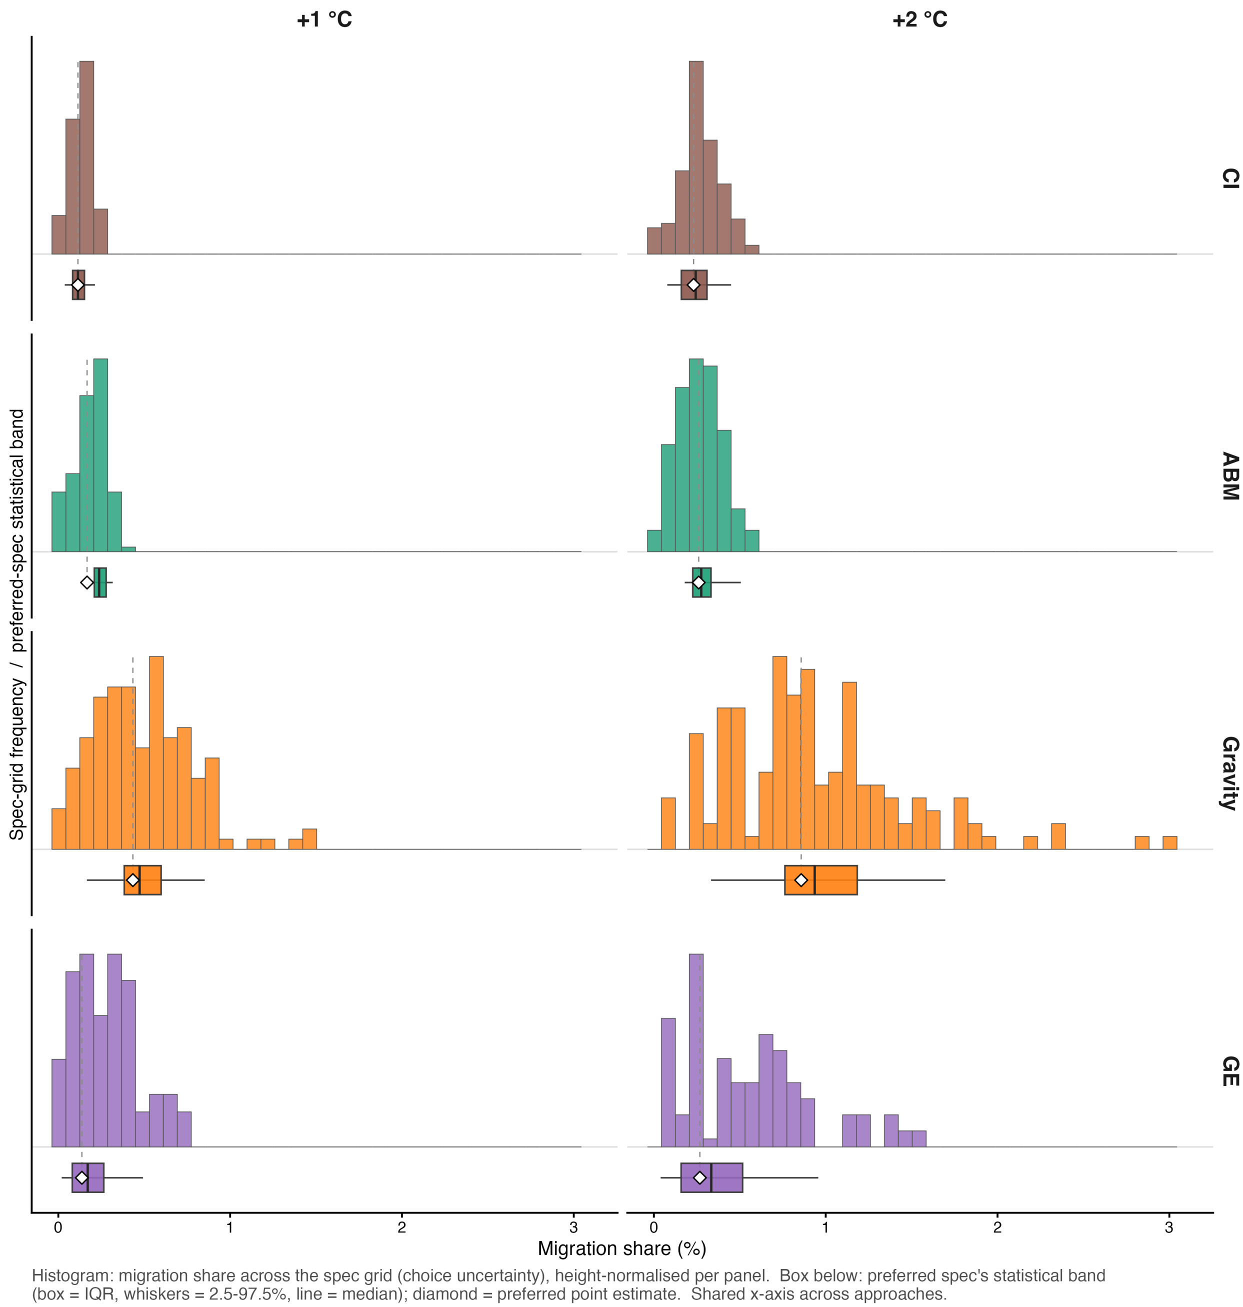

Figure A7: Distribution of the migration share across the choice grid, with the preferred model's statistical band.
Rows are approaches; columns are the moderate (+1 °C) and extreme (+2 °C) warming scenarios. In each panel the histogram reflects the distribution of the migration share across the full implementation choice grid (choice uncertainty), while the box plot directly below it reflects the preferred model's statistical uncertainty (box = interquartile range, whiskers = 2.5–97.5%, line = median; diamond = preferred point estimate). Choice uncertainty is comparable to or wider than statistical uncertainty for every approach in both scenarios. This is the two-scenario companion to the extreme-warming panel in Fig. 4a.

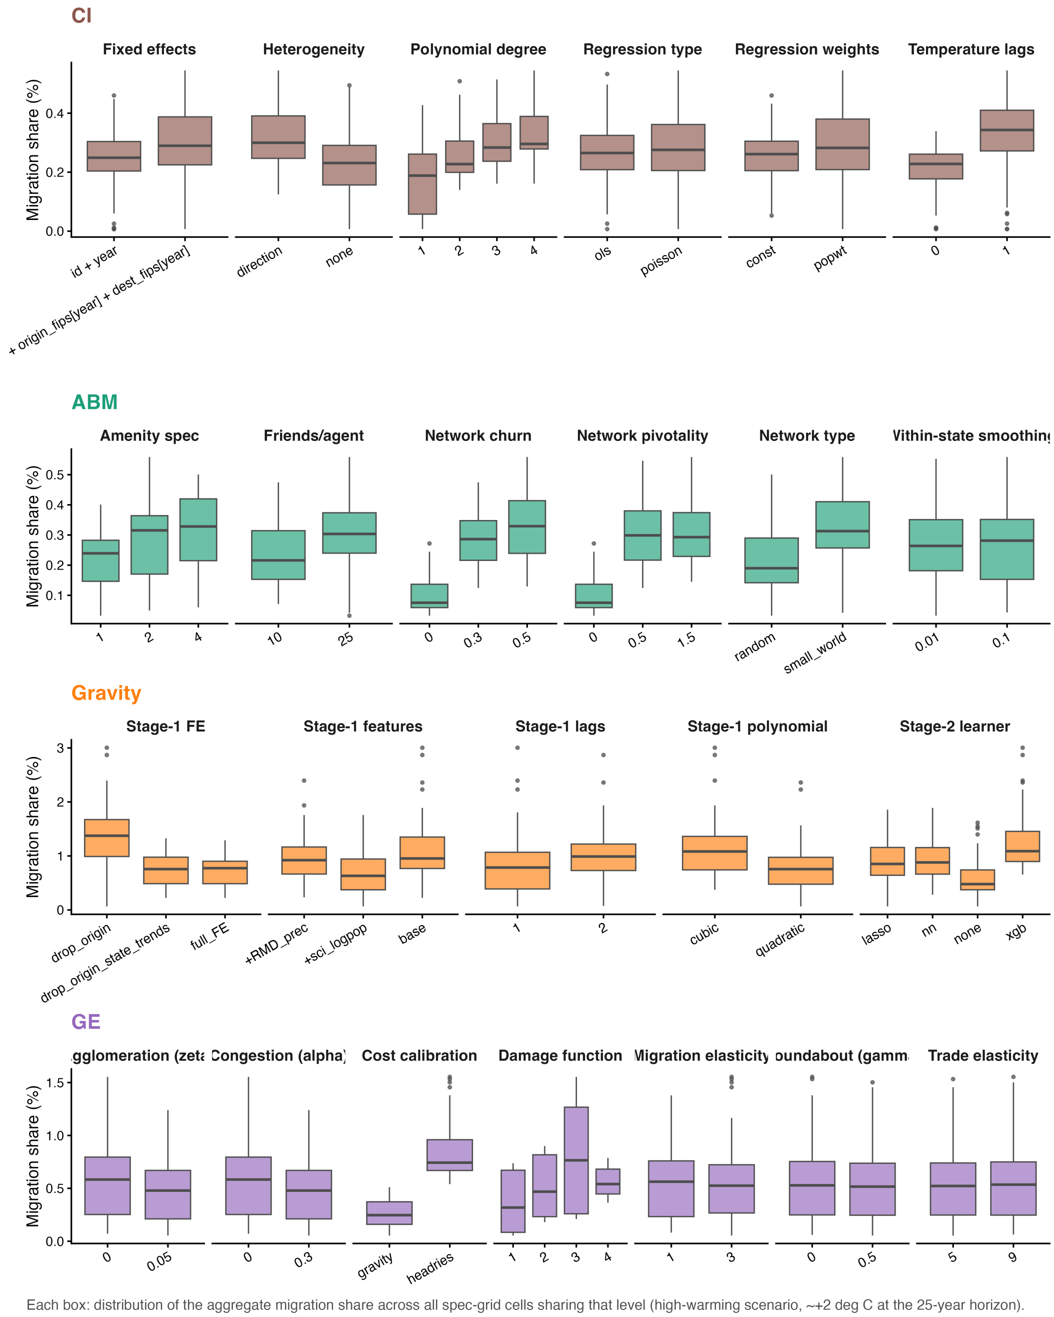

Figure A8: Marginal effect of each implementation choice on the migration share. One row per approach (causal inference, agent-based, gravity, general equilibrium), for the extreme-warming scenario ($\approx +2\,°$C at the 25-year horizon). Within each panel, a box plot shows the distribution of the aggregate migration share (Equation 4) across every specification in the choice grid that shares a given level of that axis. This is the full per-axis view underlying the tornado summary in Fig. 4c.

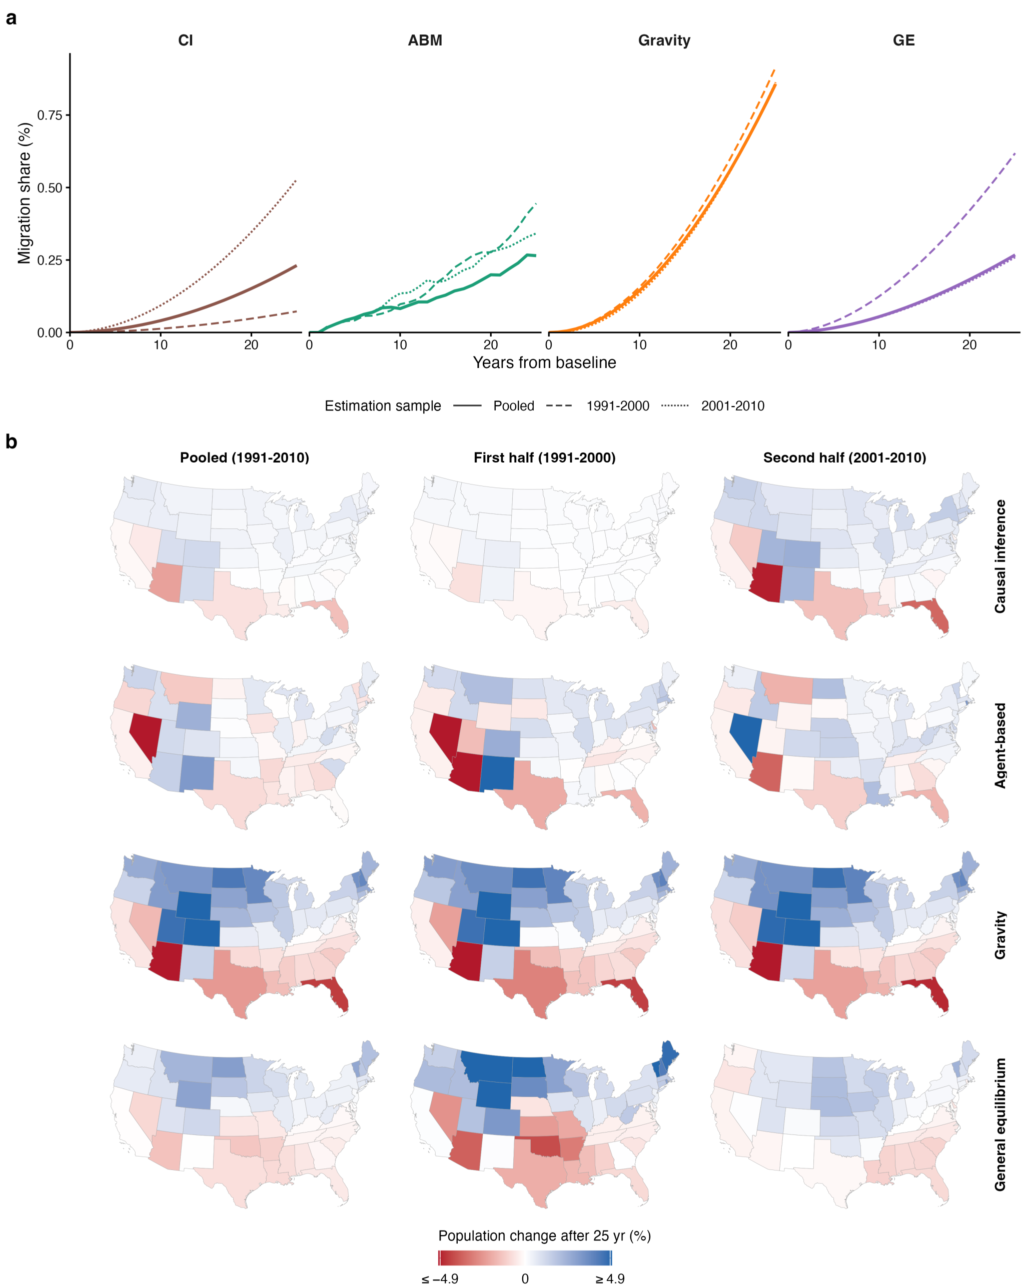

Figure A9: Temporal stability of the preferred models.
Each approach's preferred model is re-estimated independently on the two halves of the migration panel and on the pooled panel. Inputs with no within-panel time variation (e.g., interstate trade flows and the social-connectedness network), the projection baseline, and the projection itself (25-year horizon, extreme scenario, $\approx 2\,^{\circ}$C) are held fixed identical across samples. (**a**) Migration-share trajectories over the horizon for each estimation sample. (**b**) State-level projections at the horizon, by approach (rows) and estimation sample (columns). All values are point estimates. Splitting the panel in two roughly doubles sampling uncertainty, so part of the cross-era spread in our estimates here reflects expected statistical noise rather than true instability.

B Supplementary Methods

This Supplementary Section details the methods used in the four-approach climate-migration model intercomparison analysis.

- Section B.1 describes the datasets and the construction of the warming forcing;
- Section B.2 shows that the observed migration rate is an unbiased estimator of the underlying migration probability that every approach targets;
- Section B.3 details the common structure of all four approaches;
- Sections B.5, B.6, B.7, and B.8 give the full specification of the causal-inference, agent-based, gravity, and general equilibrium approaches respectively.

B.1 Data and forcing construction

This appendix documents the datasets that enter the intercomparison and how we construct each variable. All four approaches are estimated on a common panel of the 48 contiguous US states (i.e., Alaska, Hawaii, and the District of Columbia are excluded).

Migration. Bilateral migration flows come from the IRS county-to-county migration series, measured as the number of persons represented on a tax return. We aggregate the county-to-county flows to state-to-state flows. These flows data are the same as those used in Hauer and Byars (2019). We transform this data to a complete bilateral panel and generate zeros for state-pair-years in which no migration is observed. To protect confidentiality, the IRS suppresses county-to-county flows below a disclosure threshold (fewer than ten movers), pooling them into a residual "other" destination. We drop this unattributable before we collapse the county data to states, rather than reassigned it. The bilateral migration rate is $M_{ijt} = L_{ijt}/N_{it}$, where N_{it} is the origin's filer *base* (its total exemptions in the previous year) while L_{ijt} are the interstate *flows* between origin state i, destination state j, in year t. Using the state tax bases N_{it}, we impute and retain "self-flows" ($i = j$), since they supply the "stayer" counts used by the gravity and agent-based normalisations and are used in migration cost estimation by the general equilibrium model. The usable data span 1991–2010 (1990 is excluded from estimation as our self-flow imputation requires data at year $t-1$).

Temperature (in sample). The key treatment variable is each state's population-weighted annual mean near-surface ($2\,\mathrm{m}$) air temperature, in °C. To calculate this, we weight gridded daily temperature reanalysis data by gridded population to a daily state series and averaged within the calendar year. The reanalysis product is the ERA5-Land (Muñoz-Sabater et al., 2021), generated using Copernicus Climate Change Service information from the ERA5 reanalysis (Hersbach et al., 2020). We use gridded population data from LandScan Global 2000 (Dobson et al., 2000). The polynomial specifications use powers T_{it}, T_{it}^2, $\ldots$ of this annual mean, and one- and two-year lags are constructed for the lag choice.

Precipitation. Origin precipitation enters the causal-inference and gravity specifications as a control: each state's population-weighted annual *total* precipitation (and its powers and lags), from the same ERA5-Land pipeline as temperature (in metres of water equivalent).

Climate projection (warming forcing). The forcing is the SSP3-7.0 projected annual mean temperature. We take the median projected 2020–2040 mean for each state and convert it to

°C. We convert its anomaly relative to the historical reference period into a per-year increment between the two periods' midpoints (1995 and 2030) and apply that increment linearly from the baseline year, so the 25-year horizon carries ≈ 0.98 °C under WS$\times 1$ scaling. This increment is added to each state's in-sample baseline temperature $\bar{T}_i$ to form the projected path, and the warming-scaling factor (WS$\times 1$/WS$\times 2$) multiplies the anomaly as in Section B.3. These projected temperatures are the state-level SSP3-7.0 series from the Climate Impact Lab Impact Map (Climate Impact Lab, 2023), a probabilistic projection derived from the CMIP6 climate-model ensemble.

Other covariates. Real gross state product per capita is taken from the state-level macroeconomic panel of Mohaddes et al. (2023); the bilateral social network is the Facebook Social Connectedness Index (Bailey et al., 2018); bilateral distance is the great-circle distance between 2010 population-weighted state centroids, in miles; bilateral state-to-state trade flows, used by the general equilibrium approach to calibrate trade costs, are from the 2012 Commodity Flow Survey (U.S. Census Bureau and Bureau of Transportation Statistics, 2015); the relative mean deviation of income (an inequality measure entering a gravity specification) is from Frank (2009); and state population (the IRS filer base) defines both the population weights and the baseline population vector $\mathbf{L}_0$ used in projection.

Summary statistics. Table A3 reports summary statistics for all the data variables used across the four approaches.

Table A3: Summary statistics for the data variables used across the four approaches, over the full 48×48 state grid. We split summary statistics for each bilateral variable across two rows: first, statistics are shown for the self-pairs ($i = j$; e.g., stayers for migration, within-state links for the SCI, intra-state trade, and zero distance); the second row contains summary statistics for off-diagonal pairs ($i \neq j$). The panel covers years 1991–2010. The self-flow migration rate can exceed 100% because its denominator is the prior-year tax filer base. Origin climate and economic variables are summarised over origin state-years, and the projected warming over the 48 states at the 25-year horizon.

Variable	Structure	Mean	SD	Min	Max	N
Bilateral migration rate, M_{ij} (%)	$i{=}j$, 48×20 yr	97.429	1.391	87.363	104.161	960
	$i{\neq}j$, 48×47×20 yr	0.029	0.075	0.000	2.456	45,120
Bilateral migration flow (exemptions)	$i{=}j$, 48×20 yr	4,460,810	4,680,796	374,212	28,681,489	960
	$i{\neq}j$, 48×47×20 yr	1,291	3,612	0	84,614	45,120
Bilateral distance (miles)	$i{=}j$, 48 pairs	0	0	0	0	48
	$i{\neq}j$, 48×47 pairs	1,033	600	37	2,645	2,256
Log Social Connectedness Index	$i{=}j$, 48 pairs	13.40	0.97	11.01	15.09	48
	$i{\neq}j$, 48×47 pairs	8.49	0.79	7.22	12.03	2,256
Bilateral trade flow, CFS 2012 ($m)	$i{=}j$, 48 pairs	770.5	1,598.5	30.8	10,036.2	48
	$i{\neq}j$, 48×47 pairs	20.2	49.1	0.0	1,043.4	2,256
Origin temperature, T_{it} (°C)	48 states × 20 yr	12.4	4.2	2.9	23.0	960
Origin annual total precipitation	48 states × 20 yr	0.49	0.16	0.06	0.90	960
Real GSP per capita (USD)	48 states × 20 yr	41,282	8,677	23,269	71,155	960
Origin population (IRS filers)	48 states × 20 yr	4,548,116	4,720,806	393,084	28,832,510	960
Income inequality (rel. mean dev.)	48 states × 20 yr	0.814	0.047	0.721	0.978	960
Projected warming, WS×1 (°C)	48 states	0.98	0.14	0.65	1.23	48
Projected warming, WS×2 (°C)	48 states	1.96	0.29	1.29	2.47	48

B.2 Linking observed migration rates to migration probability

Since we assume the idiosyncratic utility component ϵ_{aijt} is a random variable, as in the standard random-utility model (McFadden, 1974), we can denote the probability that agent a moves from i to j in period t as:

$$m_{aijt} = P(U_{aijt} \geq \max_k U_{aikt}). \tag{8}$$

Taking the expected value of m_{aijt} over the agents in a given location gives an expression for the expected migration rate between i and j in period t. Denote the population residing in i at the start of period t as: $\mathbf{a}^*_{it} = \{a \mid j^*_{a,i,t-1} = i\}$, and $N_{\mathbf{a}^*_{it}}$ as the number of such agents. Then,

$$m_{ijt} := E(m_{aijt}) = \frac{1}{N_{\mathbf{a}^*_{it}}} \sum_{\mathbf{a}^*_{it}} m_{aijt} \tag{9}$$

Let L_{ijt} be the flow of migrants from $i \to j$ observed in the data, i.e.,

$$L_{ijt} = \sum_{\mathbf{a}^*_{it}} 1(j^*_{ait} = j)$$

Then, the observed migration rate is,

$$M_{ijt} = \frac{L_{ijt}}{\sum_k L_{ki,t-1}} \tag{10}$$

Note, since i and t here are generic, this definition is consistent with short-term, low distance migration (for example by setting i equal to towns, and t equal to weeks), or long distance permanent migration (let i be countries, and t be years or decades).

We now show that, conditional on the realised origin population $\mathbf{a}^*_{it}$, the observed rate M_{ijt} is an unbiased estimator of m_{ijt}. Conditioning on $\mathbf{a}^*_{it}$ fixes both the count $N_{\mathbf{a}^*_{it}}$ and the agent-level probabilities $\{m_{aijt}\}$, so that the only remaining source of stochasticity is agents' period-t destination draws. Then

$$\begin{aligned}
E(M_{ijt} \,|\, \mathbf{a}^*_{it}) &= \frac{1}{N_{\mathbf{a}^*_{it}}} \sum_{\mathbf{a}^*_{it}} E(1(j^*_{ait} = j) \,|\, \mathbf{a}^*_{it}) && \text{(def. of } M_{ijt}\text{; } N_{\mathbf{a}^*_{it}} \text{ fixed)} \\
&= \frac{1}{N_{\mathbf{a}^*_{it}}} \sum_{\mathbf{a}^*_{it}} m_{aijt} && (E \text{ of an indicator} = \text{move probability}) \\
&= m_{ijt}. && \text{(def. of } m_{ijt}\text{)}
\end{aligned}$$

Averaging over realised populations gives the unconditional $E(M_{ijt}) = E(m_{ijt})$. When the residents of i are ex-ante identical — the reading used by the causal-inference, gravity, and general equilibrium approaches, for which $m_{aijt} = m_{ijt}$ is a location-pair parameter — the conditioning is immaterial and $E(M_{ijt}) = m_{ijt}$ holds unconditionally; the population-averaging matters only for the heterogeneous-agent (agent-based) case.

This result is what licenses the four approaches below to treat the observed bilateral migration rate M_{ijt} as an unbiased estimator (conditional on the origin population) of the migration probability m_{ijt} that each method targets. The approaches differ in how they specify and project the migration-rate generating function f that produces m.

B.3 Common structure of the four implementations

The four approaches share a common reporting structure, instantiating the framework of Section 4.

Shared forcing and notation. Climate enters through the SSP3-7.0 projected temperature, reduced to a state-year scalar T_{it}^{raw} (each state's population-weighted annual mean). Letting $\bar{T}_i$ represent the baseline temperature, the temperature path applied in projection is simply,

$$T_{it}^{\mathrm{proj}} = \bar{T}_i + \left(T_{it}^{\mathrm{raw}} - \bar{T}_i\right) \cdot \mathrm{WS},$$

where the warming-scaling factor is $\mathrm{WS} = 1$ (the native SSP3-7.0 anomaly, $\approx +1\,°\mathrm{C}$ at the horizon) or $\mathrm{WS} = 2$ (twice the anomaly, $\approx +2\,°\mathrm{C}$). The no-warming counterfactual holds temperature at its baseline ($\mathrm{WS} = 0$). As the population distribution is not in steady state across the approaches, even absent forcing the population distribution evolves. The projection runs 25 years forward from a baseline cross-section. We index time by years since that baseline so that $t = 0$ is the baseline and $t = 25$ is our horizon. The baseline is calibrated to a historical reference period — the 2005 cross-section in our data.

WS$\times 1$ is the SSP3-7.0 median temperature projection at the horizon (a US-average warming of $\approx +1\,°\mathrm{C}$ from baseline, at the scenario's $\approx 0.4\,°\mathrm{C}$/decade rate), whereas WS$\times 2$ ($\approx +2\,°\mathrm{C}$) is the warming the same projection reaches around mid-century (2040–2059). We feed the same forcing to every approach. The agent-based model applies it with a one-period lag, so its terminal year reflects warming through the penultimate year.

Each approach designates a single *preferred model* (Methods) and embeds it in a *choice grid* that enumerates the principal discretionary modelling choices (the researcher degrees of freedom). For each warming scenario we then report three quantities, all computed on the common migration-share metric (defined in each appendix's Projection subsection, and identical across all four approaches):

None. the preferred-model point estimate;

Statistical. the sampling uncertainty of the preferred model, obtained by resampling its estimated parameters and propagating each draw through the projection, reported as the $\{q_{2.5}, q_{25}, q_{75}, q_{97.5}\}$ quantiles of the resulting migration-share values;

Choice. the dispersion induced by the modelling choices, obtained by re-projecting across the full choice grid (the same four quantiles).

Approach-specific parameters carry an approach superscript (^{CI}, ^{ABM}, ^{Gr}, ^{GE}); objects of the common framework (L, M, m, T, t, and the idiosyncratic shock ϵ) are unsuperscripted and shared. The general equilibrium appendix (Section B.8) omits the ^{GE} superscript after each parameter's first definition, as flagged there.

B.4 What each approach identifies

Although the four approaches estimate $f(\cdot)$ differently, climate enters the migration-rate generating function (2) through a single scalar per location, which clarifies what each can and cannot separate. Let $v^c_j(\mathbf{T})$ denote the *climate-relevant value* of location j: i.e., the part of utility affected by warming that is common across agents (the superscript c distinguishes it from the real wage v_{jt} of the GE inversion, Section B.8). The ABM carries the temperature amenity only, $v^c_j = A_j$, while the GE model adds real income, $v^c_j = \log A_j + \log(w_j/P_j)$. Because climate acts *only* through v^c, the rate is a composition $m_{ij}(\mathbf{T}) = f\big(v^c(\mathbf{T}), \mathbf{X}, \mu; \boldsymbol{\beta}\big)$, where μ is the migration-cost vector (Section B.8), and, holding the other locations' values fixed, the partial derivative of a row of f with respect to it own-temperature separates into a behavioural and a preference component,

$$\frac{\partial m_{ij}}{\partial T_j} = \underbrace{\frac{\partial f}{\partial v^c_j}}_{\text{mobility}} \underbrace{\frac{\partial v^c_j}{\partial T_j}}_{\text{marginal value}}, \tag{11}$$

where the mobility term $\partial f/\partial v^c_j$ — how strongly flows respond to a unit of value — is governed by the dispersion of the shocks idiosyncratic shocks. The approaches sit at different points of this factorisation: CI and gravity fit migration directly and so identify only the composite $f \circ v^c$, reporting the product in (11) without separating its factors; the ABM, a calibrated multinomial logit, recovers $v^c = A(T)$ and f separately, isolating the marginal amenity $\partial A/\partial T_j$; the GE model likewise separates v^c (amenity and real income) from f, but only by *assuming* a migration elasticity that migration data cannot pin down.[13] If f is increasing in value (i.e., $\partial f/\partial v^c_j > 0$), then the *location* of any peak in the damage function is robust, since $\partial m_{ij}/\partial T_j = 0 \Leftrightarrow \partial v^c_j/\partial T_j = 0$. That is, the reduced form pins down the value-maximising temperature without any elasticity, whereas the *size* of the slope is not: a steep response can mean either that warming hurts or that people are mobile, and separating the two requires the migration elasticity. This own, ceteris-paribus effect is also distinct from the response that warming realises. Equation (11) holds the rest of the temperature vector $\mathbf{T}_{-j}$ fixed and its identifying variation is purely *relative* (j hotter than k); this is a no-interference (stable-unit-treatment-value) assumption — that one origin's migration is unaffected by other locations' temperatures — which the reduced-form approaches require to project (Section B.5) and the structural models instead relax by

[13]Identification of the amenity component is conditional on the migration elasticity in the GE model.

explicitly modeling and re-solving. Warming shifts both the partial and spillover effects, with the warming-scaling factor WS of the shared forcing above (so $\mathbf{T} = \mathbf{T}(\mathrm{WS})$, $\mathrm{WS} = 0$ at baseline and $\dot{T}_j \equiv dT_j^{\mathrm{proj}}/d\mathrm{WS} = T_j^{\mathrm{raw}} - \bar{T}_j$), the realised response is the total derivative

$$\frac{dm_{ij}}{d\mathrm{WS}} = \underbrace{\frac{\partial m_{ij}}{\partial T_j}\,\dot{T}_j}_{\text{own (reported)}} + \underbrace{\sum_{k \neq j} \frac{\partial m_{ij}}{\partial T_k}\,\dot{T}_k}_{\text{rivals warm too}}\,. \tag{12}$$

As migration is relative — f is normalised over destinations, $\sum_j m_{ij} = 1$ — the part of the shift common to all locations partly cancels, so the realised cold-gain, hot-lose pattern is the *differential* response across the temperature distribution (the heterogeneity of $\partial v_j^c/\partial T_j$), not the own coefficient: warm j alone and it simply loses, but warm everyone and a cold j can gain. The ABM and GE models perform this field shift natively, re-normalising the whole choice over the updated field; the reduced-form approaches report the own effect (11) directly, and turning it into a projection composes that building block over the field under the no-interference and stable-slope assumptions the reduced-form projections require.

B.5 Causal inference

This appendix gives the full specification of our causal inference (CI) implementation. The preferred model is a Poisson pseudo-maximum-likelihood (PPML) bilateral migration regression with origin-by-year and destination-by-year linear trends, a cubic polynomial in origin temperature, and population weights. We accompany the preferred model with a 128-choice grid that enumerates the principal researcher degrees of freedom, and we propagate two distinct sources of uncertainty: *statistical*, from a parametric bootstrap of the regression coefficients of the preferred model, and *choice*, from re-running projections across the full 128-choice grid.

B.5.1 Approach overview

This implementation is close to Baylis et al. (2025), who estimate the temperature–migration relationship on US *county*-level data and project it forward. Our causal-inference approach works at the state level, adds the directional-heterogeneity channel, and propagates a full choice grid rather than a single specification; but our central estimates are broadly consistent with theirs.

The CI implementation works directly with the migration-rate generating function f introduced in Section 4 and attempts to statistically recover its semi-elasticity with respect to origin temperature. We exploit conditionally random year-to-year fluctuations in state temperature, after partialling out a flexible set of fixed effects and trends, to identify the function $h(\mathbf{T}_{it}; \boldsymbol{\beta}^{CI})$ in

$$\log E(M_{ijt}) = h(\mathbf{T}_{it}; \boldsymbol{\beta}^{CI}) + \gamma^{CI} \cdot \mathbf{W}_{it}, \qquad \forall i \neq j, \tag{13}$$

where $\mathbf{W}_{it}$ collects fixed effects, trends, and a quadratic precipitation control, and M_{ijt} is the observed bilateral migration rate (origin i to destination j in year t). The estimand of interest is the local derivative of $\log f$ with respect to origin temperature.

To see why the regression estimates that derivative, write the migration rate as $m_{ijt} = f(\mathbf{T}_{it})$ (suppressing the non-temperature arguments). Taking a first-order Taylor expansion of $\log f$ in temperature about the in-sample origin mean $\bar{\mathbf{T}}_i$, we have that,

$$\log m_{ijt} \;\approx\; \log f(\bar{\mathbf{T}}_i) \;+\; \left.\frac{\partial \log f}{\partial \mathbf{T}}\right|_{\bar{\mathbf{T}}_i} \big(\mathbf{T}_{it} - \bar{\mathbf{T}}_i\big) \;+\; \cdots\,,$$

where the omitted higher-order terms are captured by the polynomial $h()$. The leading term $\log f(\bar{\mathbf{T}}_i)$ is origin-specific and time-invariant, so it is therefore absorbed by the origin fixed effect. The second term is the local semi-elasticity. Provided temperature is conditionally exogenous given $\mathbf{W}_{it}$ — that is, $E[\epsilon_{ijt} \mid \mathbf{T}, \mathbf{W}] = 0$ (where conditioning here can be interpreted as partialling out the fixed effects and time trends), the identifying assumption standard in empirical climate-impact studies (no spillovers, conditional random assignment of temperature; see e.g. Kolstad and Moore, 2020) — regressing $\log E(M_{ijt})$ on $h(\mathbf{T}_{it})$ and $\mathbf{W}_{it}$ consistently estimates $d \log f/d\mathbf{T}$ at the in-sample mean.

B.5.2 Preferred model

Our preferred CI model is a Poisson PPML regression with population weighting, full bilateral fixed effects, and a cubic polynomial in contemporaneous origin temperature:

$$\log E(M_{ijt}) = \sum_{p=1}^{3} \beta_p^{CI}\, T_{it}^p + \sum_{q=1}^{2} \eta_q^{CI}\, R_{it}^q + \alpha_{ij}^{CI} + \alpha_t^{CI} + \delta_{1,i}^{CI} \cdot t + \delta_{2,j}^{CI} \cdot t, \qquad \forall i \neq j, \tag{14}$$

The fixed-effect block $\mathbf{W}_{it} = \{\alpha_{ij}^{CI}, \alpha_t^{CI}, \delta_{1,i}^{CI} \cdot t, \delta_{2,j}^{CI} \cdot t\}$ comprises bilateral pair effects α_{ij}^{CI} (absorbing all time-invariant features of origin-destination pairs, including distance and historical network density), year effects α_t^{CI} (absorbing nationwide trends in migration intensity), and *origin-* and *destination-specific linear trends* ($\delta_{1,i}^{CI} \cdot t$ and $\delta_{2,j}^{CI} \cdot t$). The state-by-time trends absorb persistent low-frequency state-level demographic drift that would otherwise be spuriously correlated with the secular trend in temperature. The quadratic in contemporaneous origin precipitation R_{it} (coefficients η_q^{CI}, shown explicitly in Equation 14) is included in every implementation as a continuous control but is not the object of interest. The estimator is Poisson pseudo-MLE using population weights (equal to the time-average of origin-state tax filers). The Poisson specification estiamtes parameters with the standard semi-elasticity interpretation, and accommodates zero migration flows without ad-hoc transformations (e.g., log plus one Wooldridge, 2010; Chen and Roth, 2024). We drop self-flows ($i = j$) from the estimation sample, and observations with a missing migration rate (opposed to zero) are excluded.

We select the polynomial degree $P = 3$ as the preferred value because it retains the negative correlation between origin temperature and projected population change at WS$\times$2 while allowing more curvature than a quadratic ($P = 2$); our choice uncertainty includes $P \in \{1, 2, 3, 4\}$ so sensitivity to this choice is reported explicitly.

Using the fitted model to perform projections requires some care, for which statistical theory is our guide. In the model, the high-dimensional bilateral fixed effects α_{ij}^{CI} are incidental parameters: in short panels (like our 19-year panel) they are not consistently estimated, so the fitted *levels* $\widehat{E}(M_{ijt})$ inherit that inconsistency. We therefore do not project levels. Instead, we anchor our projection on the observed in-sample mean rate $\bar{\mathbf{m}}_{ij}$ and add only the temperature-driven *deviation* implied by the consistently-estimated slope of $h()$ (Equation 16); the inconsistent fixed effects then cancel in this difference. Thus the object we project is the target output, that is, it represents a causal deviation from an observed baseline.

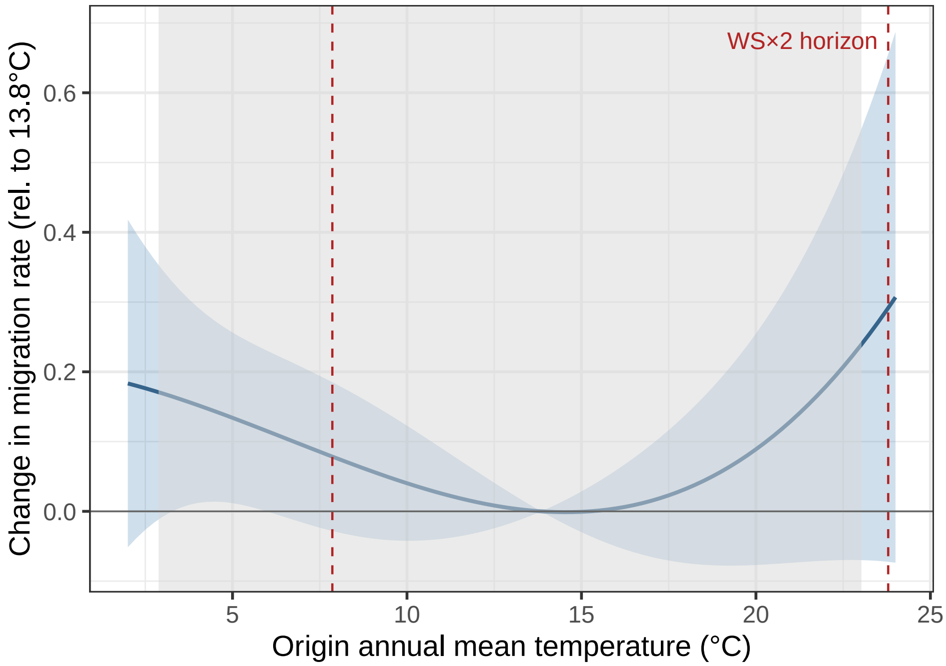

Figure A10: Preferred-model temperature response function. The target parameter from the estimated response $h(T; \hat{\boldsymbol{\beta}}^{CI})$ of the preferred (cubic-PPML) specification: the proportional change in the migration rate relative to the in-sample population weighted mean origin temperature (13.8C), with delta-method 95% confidence bands. The grey band marks the in-sample temperature support; the dashed lines show the bounds of the WS$\times$2-projected temperatures at the 25-year horizon across the 48 states.

B.5.3 Choice grid

We enumerate 128 specifications (Table A4).

Table A4: CI: The 128-choice grid.

Choice	Values	Justification	$\lvert\cdot\rvert$
Estimator	{OLS, PPML}	PPML is the correct model when there are zero-flow cells; OLS remains a common choice in the literature	2
Fixed effects	{full, pair + year}	full $= \alpha_{ij}^{CI} + \alpha_t^{CI} + \delta_{1i}^{CI} t + \delta_{2j}^{CI} t$; pair + year retains $\alpha_{ij}^{CI} + \alpha_t^{CI}$ and drops the state-by-time trends	2
Temp. poly degree	{1, 2, 3, 4}	degree of the polynomial in T_{it}	4
Weighting	{population, equal}	population vs. equal weights	2
Temperature lags	{0, 1}	contemporaneous vs. + one-year lag of T	2
Directional heterogeneity	{none, by direction}	separate response by origin–destination climate gap	2
Cartesian product			128

The implementation choices were chosen for the following reasons: first, polynomial degrees up to $P = 4$ are conventional in the temperature-response literature (see, e.g., Carleton et al., 2022). We include $P = 1$ as a sensitivity row even though linear-only specifications are known to underestimate nonlinear responses. Time fixed effects range from demanding (state-specific flexible time trends) to a conventional year fixed effect, which produces a wider standard-error band but the same sign and similar point estimate (see the per-state response panels in the main text). The two weighting schemes are the two modal choices in the empirical migration literature: the equal-weights specification gives equal weight to small and large states and serves. The OLS implementations use the level of the migration rate M_{ijt} as an outcome variable rather than on its log, so zero flows enter directly as $M_{ijt} = 0$ with no logarithm or ad-hoc offset (the PPML regressions accommodates zeros directly); the OLS models constitute exactly half of the 128-choice grid. The temperature lag choices are $\{0\}$ vs $\{0, 1\}$.

The directional heterogeneity implementation choices split the response function polynomial in temperature into three buckets according to the relative climate of origin and destination. We assign each origin and destination state its average annual temperature, and we classify every bilateral pair by the destination-origin temperature gap: *to-colder* $(C_{ij} = 1)$ if the destination is more than $2°$C cooler than the origin, *similar* $(C_{ij} = 2)$ if within $\pm2°$C, and *to-hotter* $(C_{ij} = 3)$ if more than $2°$C warmer. Our choice of cutoff, $2°$C, is roughly half a cross-state standard deviation and splits the panel into approximately equal-mass groups. These labels are assigned from in-sample panel-average temperatures and we hold these labels fixed in our projections. That is, warming rescales the within-bucket temperature response but never reclassifies a pair across the $2°$C boundary. This interacted specification is motivated by recent findings that climate-induced movers "climate match" by selecting destinations with climates similar to their origin (Obolensky et al., 2024; Leduc and Wilson, 2024):

$$\log E(M_{ijt}) = \sum_{c=1}^{3} \sum_{p=1}^{P} \mathbf{1}(C_{ij} = c) \cdot \beta_{c,p}^{CI}\, T_{it}^{p} + \mathbf{W}_{it}. \tag{15}$$

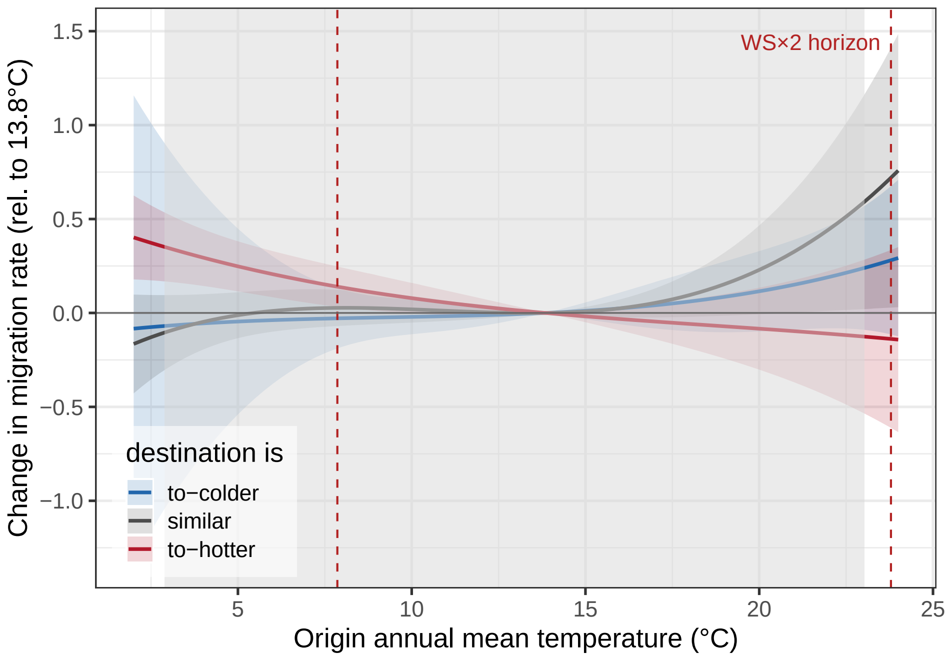

Figure A11: Direction-interacted temperature response functions. Bucket-specific responses from the "climate matching" specification (Equation 15), estimated separately for pairs whose destination is colder than (*to-colder*), similar to (*similar*), or hotter than (*to-hotter*) the origin state. Effects are plotted relative to $13.8\,°\mathrm{C}$ temperature, with delta-method 95% confidence bands. Shading and dashed lines are the same as in Figure A10.

Pinned choices. The following choices are held fixed across all 128 implementations. The treatment variable is population-weighted annual average origin temperature (a single state-year scalar; we do not enumerate over alternative climate indices such as growing-season temperature, which are the subject of a complementary literature). The outcome variable is the gross bilateral migration rate M_{ijt}; we exclude same-state non-movers from the regression but recover the diagonal post-estimation through the column-sum constraint. Standard errors are clustered at the origin level to permit arbitrary serial dependence within an origin's bilateral panel. The projection baseline is the in-sample average migration rate $\bar{\mathbf{m}}_{ij}$ (the within-pair time mean of the *observed* rate over 1991–2010; the migration rate in year t is built from the $t-1 \to t$ flow, so the first in-sample rate is for 1991), consistent with the framework described in Section 4; this choice fixes the linear-extrapolation anchor point.

B.5.4 Estimation

Each model is fit either by Poisson pseudo-maximum-likelihood (via PPML-'HDFE') or by ordinary least squares (OLS). We cluster standard errors at the origin-state level. For the preferred model, we retain the estimated variance-covariance matrix of the temperature coefficients $\hat{\boldsymbol{\beta}}^{CI}$

(denoted $\hat{\Sigma}$) to perform a parametric bootstrap and obtain our estimates' of the approach's statistical uncertainty.

For the parametric bootstrap, we draw $K = 500$ coefficient vectors $\boldsymbol{\beta}^{(k)} \sim \mathrm{MVN}(\hat{\boldsymbol{\beta}}^{CI}, \hat{\Sigma})$ and propagate each through the projection. To be clear, our bootstrap is a *coefficient resample*, not a data refit. Our bootstrap captures sampling uncertainty in the estimated coefficients, conditional on the model specification. It is computationally cheap because no re-estimation is required per draw. Following the cross-approach convention, the headline propagates the $K = 500$ draws through the preferred model alone, while every other model is evaluated at its coefficient mean to estimate the implementation choice uncertainty.

B.5.5 Projection

For each estimated model and each coefficient draw $\boldsymbol{\beta}^{CI,(k)}$, we construct the projected migration rate from the in-sample baseline plus a first-order expansion through the model's polynomial h. The functional form of the update depends on the estimator. For PPML choices the update is *multiplicative*:

$$\widehat{m_{ijt}^{CI,(k)}}(\mathbf{T}) = \bar{\mathbf{m}}_{ij} \left[1 + h\Big(\mathbf{T}_{it}; \hat{\boldsymbol{\beta}}^{CI,(k)}\Big) - h\Big(\bar{\mathbf{T}}_i; \hat{\boldsymbol{\beta}}^{CI,(k)}\Big)\right], \tag{16}$$

which respects the semi-elasticity interpretation of the PPML coefficients. For OLS choices the update is instead *additive*, $\widehat{m_{ijt}^{CI,(k)}} = \bar{\mathbf{m}}_{ij} + \big[h(\mathbf{T}_{it}; \hat{\boldsymbol{\beta}}^{CI,(k)}) - h(\bar{\mathbf{T}}_i; \hat{\boldsymbol{\beta}}^{CI,(k)})\big]$, reflecting that the OLS coefficients are derivatives of the level of M_{ijt} rather than of its log. Here $\bar{\mathbf{m}}_{ij}$ is the within-pair time-average of the observed bilateral rate, $\bar{\mathbf{T}}_i$ is the 1991–2010 origin mean temperature, and $\mathbf{T}_{it}$ is the SSP3-7.0 projected origin temperature in year $t \in \{0, \ldots, 25\}$ (years since the baseline) at warming scaling factor WS$\in \{1, 2\}$ (the projection is anchored at the base value and scaled, $\mathbf{T}_{it}^{\text{proj}} = \bar{\mathbf{T}}_i + (\mathbf{T}_{it}^{\text{raw}} - \bar{\mathbf{T}}_i) \cdot \text{WS}$). For the preferred model ($P = 3$, no interaction, no lag), $h(\mathbf{T}; \boldsymbol{\beta}) = \sum_{p=1}^{3} \beta_p^{CI} T^p$; for directional-heterogeneity models, the projection applies the bucket-specific coefficients to each bilateral pair according to its C_{ij} label, and for one-lag models we sum the temperature difference over the contemporaneous and one-year-lagged terms. We recover the diagonal $\widehat{m_{iit}^{CI,(k)}}$ by using the column-sum constraint $\sum_j m_{ijt} = 1$ (Equation 1c). For negative predicted rates, we clip the projection at 0 before applying the constraint. The implied population reallocation is then

$$\widehat{\Delta \mathbf{L}}_{25}^{CI,(k)} = \left(\prod_{\tau=1}^{25} \widehat{\mathbf{m}}_{\tau}^{(k)} - \prod_{\tau=1}^{25} \bar{\mathbf{m}} \right) \mathbf{L}_0,$$

where each product runs over the 25 annual transition matrices from the baseline to the horizon and $\mathbf{L}_0$ is the baseline origin-filer vector. This gives a horizon population vector under WS$\times 1$ and WS$\times 2$ for each (model, draw) pair. We compute the aggregated target output (migration share at WS$\times 2$) from the resulting L_0 (no-warming) and L_1 (with-warming) trajectories via the migration-share metric used elsewhere in this paper (Equation 4). Here, $L_{1,i,25}$ and $L_{0,i,25}$ are the with-warming and no-warming populations of state i at the 25-year horizon, and $L_{i,0}$ is its baseline population. Total population is conserved over the horizon.

B.5.6 Uncertainty quantification

We report three types of uncertainty for each warming scenario:

None. The point estimate from the preferred model at the coefficient mean $\hat{\boldsymbol{\beta}}^{CI}$. The other 127 models contribute their own point estimates, which enter only the Choice band.

Statistical. The $\{q_{2.5}, q_{25}, q_{75}, q_{97.5}\}$ quantiles of the $K = 500$ projected migration-share values for the preferred model, capturing sampling uncertainty in $\hat{\boldsymbol{\beta}}^{CI}$ at the chosen specification.

Choice. The same four quantiles of the 128 projected migration-share values across the full implementation choice grid, each evaluated at each model's coefficient mean. This captures choice (researcher-degree-of-freedom) uncertainty.

B.5.7 Limitations and robustness

Four caveats deserve emphasis. *First*, the linear extrapolation in Equation 16 assumes that the local derivative $d \log f / d\mathbf{T}$ estimated in-sample applies at the projected WS$\times 2$ temperature for every state. Every one of the 48 states has its horizon WS$\times 2$ temperature outside its 1991–2010 sampled range — on average $+0.8$ °C beyond the hottest year observed in sample, and $+1.3$ °C at most (distinct from the full WS$\times 2$ anomaly relative to the average year, $\approx +2$ °C, reported in Table A3): the projection is mechanically out-of-sample at the state level, so projection requires believing the estimated polynomial represents migration behaviour at the extrapolated temperature. We include the temperature-polynomial-degree choices to expose this sensitivity: as P increases, the variance of the projected migration share grows because higher-order polynomials extrapolate more aggressively, even though they fit better within sample, a classic bias-variance trade-off. This is most acute for the choices that combine the quartic ($P = 4$) temperature response polynomial with directional heterogeneity. We include such models among our 128, but these models have poor fit (substantial estimated parameter uncertainty) and projections should be read accordingly. *Second*, the estimator is bilateral and partial: it recovers an origin-temperature semi-elasticity pair by pair, but migration is a choice among *all* destinations, so a state's in- and out-flows depend on the relative attractiveness of every alternative — the multilateral resistance to migration (Bertoli and Fernández-Huertas Moraga, 2013) — which a local, pair-by-pair slope does not capture. *Third*, migration is a long-run decision, whereas identification here comes from short-run, year-to-year temperature fluctuations; the recovered response is therefore a short-run elasticity that may understate the long-run adjustment to a permanent change in climate. *Fourth*, the approach recovers only the slope of the migration-rate generating function at a single point; it cannot speak to nonlinear adaptation, equilibrium feedback, or social-network effects — which motivate the gravity, GE, and ABM approaches respectively.

B.6 Agent-based

This appendix gives the full specification of our agent-based-model (ABM) implementation. The preferred ABM is an amenity-only multinomial-logit migration model with social-network rewiring: a population-weighted amenity regression, a small-world friendship network of mean degree 10, within-state smoothing $\nu^{ABM} = 0.01$, and dynamic rewiring with churn $\sigma_0^{ABM} = 0.5$ and pivotality $\kappa^{ABM} = 1.5$. We accompany the preferred ABM with a 120-choice grid over the amenity-coefficient specification, network topology, average friends per agent, the smoothing parameter, and the rewiring regime, and we propagate two distinct sources of uncertainty: *statistical*, from 51 draws of the amenity-coefficient posterior of the upstream amenity regression; and *choice*, from re-running the projection across the full 120-choice grid.

B.6.1 Approach overview

Our ABM is a discrete-time multinomial-logit migration model that simulates the location decisions of $100{,}000$ representative agents across the 48 contiguous US states over the 25-year projection horizon (from the baseline, $t = 0$, to $t = 25$). Climate enters utility through the destination-temperature amenity A_{jt}. The strength of this channel is calibrated externally (Section B.6.4); the contribution of this approach is to embed it in a model that includes agent-level social networks.

The ABM is stochastic, and it is the only one of the four approaches that produces *realisation-level* migration counts rather than expected rates. We therefore run it $N_{\text{sim}} = 10$ times under independent shock (and network) realisations and average the per-state population trajectories across the 10 runs *before* forming the migration-share metric — the metric is a convex sum of absolute differences, so averaging afterwards would bias it upward by Jensen's inequality. We use common random numbers across choices, so that differences between models reflect their parameters rather than simulation noise. Aggregate statistics like our target output are then computed by averaging across 10 runs per model implementation or parameter draw.

B.6.2 Preferred model

Each agent a resident in state i at the end of year $t-1$ chooses a destination j in year t by maximising a Gumbel-shocked utility,

$$U_{aijt} = \underbrace{\gamma^{ABM} \log \widetilde{G}_{aj,t-1}}_{\substack{\text{social-network} \\ \text{pull}}} + \underbrace{A_{jt}\big(\boldsymbol{\beta}^{ABM}, T_{jt}\big)}_{\substack{\text{destination} \\ \text{temperature amenity}}} + \underbrace{\alpha_j^{ABM}}_{\substack{\text{location} \\ \text{fixed effect}}} - \underbrace{\delta^{ABM}(d_{ij})}_{\substack{\text{distance} \\ \text{cost}}} + \sigma^{ABM}\,\epsilon_{aijt}, \qquad \epsilon_{aijt} \overset{\text{iid}}{\sim} \text{Gumbel}(0,1). \tag{17}$$

We define each term in turn.

Social-network pull, $\gamma^{ABM} \log \widetilde{G}_{aj,t-1}$. Agents are pulled toward states where their friends already live. The term acts on a within-state-smoothed *friend share* rather than a raw count: with $g_{aj,t-1}$ the number of a's friends resident in j at the close of $t-1$ and N_a the agent's total number of friends,

$$\widetilde{G}_{aj,t-1} = (1 - \nu^{ABM})\, \frac{g_{aj,t-1}}{N_a} + \nu^{ABM}\, \bar{S}_j, \tag{18}$$

where $\bar{S}_j$ is the (sum-normalised) average SCI weight to destination j and ν^{ABM} is the smoothing parameter ($= 0.01$ in the preferred model). The "random surfer" smoothing parameter $\nu^{ABM}\bar{S}_j$ keeps $\widetilde{G}_{aj,t-1} > 0$ for every destination, so the logarithm of G is always defined (i.e., never zero) and each state carries a small information weight even when agent a has no friends there. The coefficient γ^{ABM} is the coefficient on $\log \mathrm{SCI}$ in the gravity regression (Section B.6.4); it multiplies the log smoothed share, not a raw log count.

Why the network forces an agent-level simulation. The friend share $\widetilde{G}_{aj,t-1}$ is specific to each individual agent a. Two agents residing in the same origin state generally have friends in different destinations, so their social-pull terms differ and consequently their behavior will differ. The network enters utility as a genuine agent-level variable that is not iid across agents, periods, or states. It has no "sufficient statistic" (i.e., appropriate aggregate variable) that can be described at the state-pair-year level. In other words there is no representative-agent reduction that preserves it. This is exactly why the ABM must be *simulated agent-by-agent*. Unlike the gravity and GE approaches, flows in the agent-based model cannot be represented by a closed-form bilateral gravity equation at the state-pair level or solved as a fixed point equation in state-pair aggregates: doing so would average away the individual information sets that drive the dynamics.

Destination temperature amenity, $A_{jt}(\boldsymbol{\beta}^{ABM}, T_{jt})$. Climate enters utility only through the amenity term. We specify that the amenity is a quadratic in the destination's own temperature, $A_{jt} = \beta_1^{ABM} T_{jt} + \beta_2^{ABM} T_{jt}^2$, with coefficients common across states, so cross-state variation in A_{jt} comes from temperature levels rather than state-specific slopes. The coefficients are estimated in the gravity regression below (Section B.6.4). Warming shifts each state's A_{jt} heterogeneously, spurring climate migration.

Location fixed effect, α_j^{ABM}. The 'intercept' of the amenity is a time-invariant index of destination attractiveness, absorbing everything about j, like size, wages, housing, and other unmodelled amenities not already captured by temperature, social networks, or geographic position (by way of the distance costs). The α_j^{ABM} are *not* read off the regression: to obtain them, we perform a

"Berry inversion" (Section B.6.4) so that the ABM's choice rule reproduces the observed baseline population *shares* (Berry et al., 1995). This is critical, because the bilateral-flow PPML returns destination effects that are contaminated by its own fixed-effect block. The Berry inversion inverts the observed shares using the choice probabilities (with the estimated coefficients on temperature, network, and distance variables) to recover location fixed effects that rationalized the observed data as the outcome of the choice process.

Distance cost, $\delta^{ABM}(d_{ij})$. We fit an additional bilateral moving cost, parameterized as a function of distance terms (including home-bias) in the gravity regression (Section B.6.4). Because it depends only on the origin-destination pair, it carries no agent-level heterogeneity, in contrast to the network term.

Idiosyncratic taste shock, $\sigma^{ABM}\epsilon_{aijt}$. Each agent draws a vector of taste shocks distributed i.i.d. Gumbel across destinations. Thus, the choice is a multinomial logit, similar to the GE and Gravity approaches. The Gumbel scale parameter σ^{ABM} governs the level of idiosyncratic noise relative to deterministic utility, and hence the overall propensity to move. We calibrate σ^{ABM} to match the observed move rate (Section B.6.4).

Network construction. The preferred social-network specification is a *small-world graph* (Watts and Strogatz, 1998) anchored to the Facebook Social Connectedness Index (SCI), whose connectedness falls steeply with bilateral distance (Figure A13a). It is constructed in three stages from the SCI matrix G (with a fixed random seed shared with the calibration step, so the inverted α_j^{ABM} is solved against the realised network). First, a *ring lattice* of degree equal to the mean degree $\bar{d}$ is built *within each state*, connecting each agent to its $\lfloor \bar{d}/2 \rfloor$ nearest neighbours on each side of the ring (the high local clustering that gives the graph its small-world property); an optional within-state softening parameter (set to 0 throughout, so the within-state lattice is left intact) would rewire within-state edges. Second, each state i is assigned a cross-state rewiring probability $\phi_i^{ABM} = 1 - G_{ii}/\sum_j G_{ij}$, the SCI-implied share of i's connectedness that points outside the state; the ϕ_i^{ABM} are then rescaled by a common factor so that the realised cross-state edge fraction matches the target fraction implied by the SCI-weighted feasible-pair counts $G_{ij}N_iN_j$ (capping at 1). Third, each edge is rewired to a cross-state destination with probability ϕ_i^{ABM}, the destination state $j \neq i$ drawn with probability $\propto G_{ij}N_j$ and a random agent in j selected as the new endpoint. The realised average degree is therefore approximate (states with few agents have lower effective degree, and rewiring conserves edge count): the preferred mean degree of 10 targets roughly 5×10^5 undirected edges across the 10^5-agent simulation.

Dynamic rewiring. The network is *dynamic*: at the end of each year, for every agent that *moved*, friendships are re-sampled. An agent with current degree d_a loses each existing friend independently with probability $\sigma_-^{ABM} = \sigma_0^{ABM}/(1 + d_a/\bar{d})$ (so high-degree agents are stickier), and gains new friends in its destination state in number $\sim \mathrm{Poisson}\big(\sigma_+^{ABM} \cdot n_a^{\mathrm{lost}}\big)$ with the pivotality multiplier $\sigma_+^{ABM} = \max\big(1 + \kappa^{ABM}(d_a/\bar{d} - 1), 0\big)$, where $\bar{d}$ is the reference (mean) degree. The multiplier σ_+^{ABM} holds network size stable in expectation when $d_a = \bar{d}$ (gains replace losses one-for-one) and implements a rich-get-richer mechanism when $d_a > \bar{d}$ (high-degree agents gain more than they lose). New friends in the destination are drawn uniformly from agents currently resident there. The preferred values are $\sigma_0^{ABM} = 0.5$ (the network "churn" parameter) and $\kappa^{ABM} = 1.5$ (the "pivotality" parameter). Taken as a single five-component implementation choice, this rewiring regime is the largest driver of implementation choice uncertainty across the ABM's models: the static regime ($\sigma_0^{ABM} = \kappa^{ABM} = 0$) sits well below the median of the 120-choice migration-share distribution, with a group mean about a third of the fully dynamic regime's. The one-at-a-time sensitivity in Fig. 4c instead attributes the largest *local* spread around the preferred model to the amenity specification, because it varies the churn (σ_0^{ABM}) and pivotality (κ^{ABM}) parameters separately and so never reaches the static regime, which requires both to be zero.

B.6.3 Choice grid

Our ABM choice grid enumerates 120 models across the five implementation choices that we view as consequential discretionary modelling choices for our ABM approach (Table A5).

Table A5: ABM: The 120-choice grid.

Choice	Values	Justification	$\lvert\cdot\rvert$
Amenity regression	{pop.-weighted, unweighted, dest.-trend}	upstream temperature-amenity specification (cubic excluded)	3
Network topology	{small-world, random}	SCI-anchored vs. degree-matched random graph	2
Mean degree	{10, 25}	friends per agent	2
Within-state smoothing	{0.01, 0.10}	smoothing parameter ν^{ABM}	2
Rewiring regime	5 levels	static + 2×2 (churn $\times$ pivotality)	5
Cartesian product			120

The three amenity-regression specifications span the modal Poisson-PPML choices in the gravity literature: population-weighted, unweighted, and with destination-time trends. These differ in the moment of the data they target and in their out-of-sample extrapolation behaviour.

The five rewiring regimes are: *static* ($\sigma_0^{ABM} = 0$, $\kappa^{ABM} = 0$); *low-churn / low-pivotality* ($\sigma_0^{ABM} = 0.3, \kappa^{ABM} = 0.5$); *high-churn / low-pivotality* ($\sigma_0^{ABM} = 0.5, \kappa^{ABM} = 0.5$); *low-churn / high-pivotality*

($\sigma_0^{ABM} = 0.3, \kappa^{ABM} = 1.5$); *high-churn / high-pivotality* ($\sigma_0^{ABM} = 0.5, \kappa^{ABM} = 1.5$, the preferred model). The 2×2 design over churn and pivotality is motivated by the distinct roles these parameters play: σ_0^{ABM} scales the magnitude of per-mover friendship turnover, κ^{ABM} governs how aggressively high-degree agents accumulate ties relative to low-degree agents. Holding the two implementation choices separate prevents confounding the magnitude and the heterogeneity of network rewiring. When $\sigma_0^{ABM} = 0$ the rewiring machinery is disabled and κ^{ABM} is set to 0, so the static regime is a single choice rather than two.

The two mean-degree values (10 and 25) reflect two assumptions on neighbourhood size. The two smoothing values ($\nu^{ABM} = 0.01$ vs $\nu^{ABM} = 0.10$) govern how much we let the agent-level dynamics matter in Equation 18 (when $\nu^{ABM} = 1$, the model can be represented in aggregates); smoothing also keeps the log-share term defined for agents with no friends in a destination. The network topology implementation choices contrast the SCI-anchored small-world topology with the behaviour induced by a degree-matched random graph (sampled directly from the SCI-weighted block-pair probabilities $G_{ij}N_iN_j$, without the ring-lattice clustering step designed to fit the small-world graph); the choice of a random graph serves as a null in which the local clustering structure of friendships is removed.

Pinned choices. Population is $100{,}000$ agents in all implementations (a balance between simulation noise and tractability; for implementations with $N_{\text{agents}} = 10^6$ in pilot runs the central tendency was unchanged). The simulation runs over the 25-year horizon, indexed $t \in \{0, \ldots, 25\}$ (years since the baseline). The number of independent shock realisations per implementation is $N_{\text{sim}} = 10$. We calibrate the Gumbel scale parameter σ^{ABM} to match the empirical year-1 move rate of 1.5%, separately for each (amenity regression, network topology, mean degree, smoothing) combination. The destination-fixed effects α_j^{ABM} estimated using the Berry inversion to match the baseline population vector, again per choice combination. We do not enumerate choices over σ^{ABM} or α_j^{ABM} as both are pinned by external moment conditions.

B.6.4 Estimation and calibration

The ABM is calibrated in three steps, all of which are performed prior to simulation and run *sequentially*; i.e., the two moment conditions are matched one at a time, not jointly. Figure A12 summarises the pipeline.

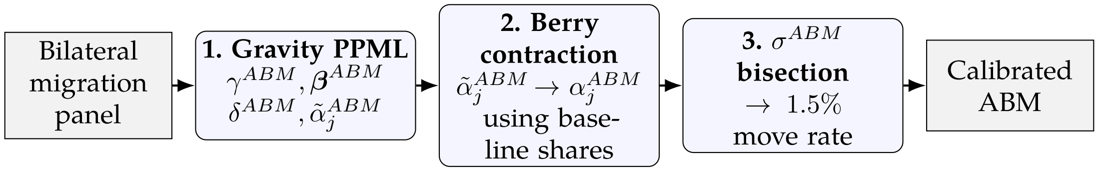

Figure A12: The ABM calibration pipeline. We use a Poisson gravity regression to recover the behavioural elasticities and damage function parameters; then, we use a Berry contraction to invert the destination amenities to match observed baseline population shares given those parameters; finally, we use a one-dimensional bisection to find the logit scale parameter that matches the historical avereage mobility rate. No step requires searching the joint parameter space or simulating the full ABM.

From the utility model to a gravity regression. The structural parameters are recovered *before* any simulation, by exploiting the fact that the agent-level logit aggregates to a gravity equation. Under the choice rule (Equation 17) the probability that an agent in i selects j is proportional to $\exp(V_{aij}/\sigma^{ABM})$, where the deterministic utility is V_{aij}. Replacing the agent-specific friend share with its origin-level average — using the aggregate SCI matrix SCI_{ij} in place of the realised $\widetilde{G}_{aj}$, which holds in expectation at the calibration baseline — makes V separable into origin, destination, and bilateral pieces, so the expected bilateral flow takes the log-linear gravity form

$$\begin{aligned}\log E(m_{ijt}) = {} & \gamma^{ABM} \log \mathrm{SCI}_{ij} + \delta_1^{ABM} \log d_{ij} + \delta_2^{ABM} (\log d_{ij})^2 \\ & + \beta_1^{ABM} (T_{jt} - T_{it}) + \beta_2^{ABM} (T_{jt}^2 - T_{it}^2) + \xi_{it}^{ABM} + \tilde{\alpha}_j^{ABM}, \qquad i \neq j.\end{aligned} \tag{19}$$

with origin-by-year effects ξ_{it}^{ABM} and destination effects $\tilde{\alpha}_j^{ABM}$ (an origin-specific self-flow term carries the diagonal). It is estimated by population-weighted Poisson pseudo-maximum-likelihood, clustered on origin. Because ξ_{it}^{ABM} absorbs the origin component of the temperature gap, the amenity coefficients $(\beta_1^{ABM}, \beta_2^{ABM})$ are identified from *destination* temperatures, which is what licenses reading off the destination amenity $A_{jt} = \beta_1^{ABM} T_{jt} + \beta_2^{ABM} T_{jt}^2$ (concave, peaking near $13\,^{\circ}$C; Figure A13c); the distance terms give the cost $\delta^{ABM}(d_{ij})$ and the log-SCI coefficient gives γ^{ABM}. The simulation then *reintroduces* the agent-level network heterogeneity that this aggregation averaged away.

Step 1 (*gravity calibration*). The coefficients $(\boldsymbol{\beta}^{ABM}, \gamma^{ABM}, \delta^{ABM}, \tilde{\alpha}_j^{ABM})$ are estimated by the population weighted PPML fit of Equation 19 on the bilateral migration panel, clustered on origin state; this is the same regression family used in the gravity approach (Appendix B.7). The specifications differ in their fixed-effect block and weighting. The bilateral-flow contribution is mapped to a structural distance cost $\delta^{ABM}(d_{ij})$, so δ^{ABM} is positive off-diagonal (a cost of moving) and negative on the diagonal (a home-bias bonus).

Step 2 (*Berry inversion*). The gravity regression above gives us $\tilde{\alpha}_j^{ABM}$, which are not the structural coefficients we need. To recover α_j^{ABM}, we invert observed migration probabilities under the ABM's choice rule (Equation 17), holding gravity coefficients fixed at their estimated value, to recover α_j^{ABM} such that the simulated baseline population vector reproduces observed shares. We do so with a contraction mapping,

$$\alpha_j^{(\ell+1)} = \alpha_j^{(\ell)} + \left[\log s_j^{\text{obs}} - \log S_j(\alpha^{(\ell)})\right],$$

with $\alpha^{(\ell+1)}$ mean-centred after each iteration ℓ to fix the level (as only differences in α_j^{ABM} are identified; the ABM tag is suppressed in the iteration superscripts). The target s_j^{obs} is the baseline observed population-share vector; $S_j(\alpha^{(\ell)})$ is the mean over agents of the per-agent multinomial choice probability at the realised network. The iteration is run to a tolerance of 10^{-6} on $\max_j |\Delta\alpha_j|$.

Step 3 (σ^{ABM}-calibration). The Gumbel scale is calibrated *separately*, after Step 2, to match the empirical year-1 cross-state move rate of 1.5% (tolerance $1e$-3). We take advantage of the fact that the move rate is monotone increasing in σ^{ABM}. We calibrate by first estimating migration probabilities over a grid, $\sigma^{ABM} \in \{0.1, 0.3, 0.5, 0.7, 1.0\}$, followed by bisection within the bracketing interval; each candidate σ^{ABM} is evaluated by a short three-year ($t = 0$ to $t = 2$) ABM run with the Berry-inverted α_j^{ABM}, assuming that the year-1 move rate is the no-warming baseline. Calibrated scales are $\sigma^{ABM,*} \approx 0.55$–$0.87$ across implementations (0.74 at the preferred model). Both the inverted α_j^{ABM} and the calibrated σ^{ABM} are computed once per amenity regression $\times$ network topology $\times$ mean degree $\times$ smoothing choice combinations and shared across the five rewiring values, since the rewiring choices affects neither the baseline population nor the year-1 move rate (calibration is performed on a static, year-1 network).

In other words, we calibrate the model to exactly reproduce the baseline choice probabilities and migration rate for almost every parameter configuration. We hold the Gumbel draws fixed across choices, so that differences between choice combinations reflect parameters differences rather than pure simulation noise induced by the taste shocks.

Why targeted moment-matching. A natural alternative would be to calibrate the parameters by brute force. Examples of this include, e.g., Approximate Bayesian Computation, or a grid search that tries many parameter combinations and keeps those whose simulated moments match the data. Such procedures are infeasible in our setting. Avoiding these brute force calibration methods is one of the methodological contributions of our ABM approach. The difficulty of the brute force approach to calibration is that the model must reproduce 48 baseline population shares and the aggregate move rate, but the free parameters that do so (the 48 destination effects α_j^{ABM} and the scale σ^{ABM}) live in a high-dimensional space, and *every* candidate parameter vector requires a full stochastic simulation (averaged over N_{sim} shocks) to evaluate; a search dense enough to pin 48 effects would need an astronomical number of such simulations. We instead exploit the structure of the problem to avoid joint estimation and simulation. Instead, each parameter is matched to a moment it governs: the gravity PPML recovers the bilateral and amenity parameters in a single estimation; the Berry inversion is a fast, guaranteed fixed-point map that inverts the 48 destination effects *exactly* given the estimated PPML parameters, without using any model simulation; and the logit scale is found by searching for the root of one-dimensional monotone function, which we do with bisection method in around eight one-year model simulations. Figure A13b shows the outcome: the calibrated baseline shares sit on the $45°$ line, whereas the uncalibrated gravity destination effects do not. The whole calibration runs in seconds per implementation, rather than the simulation-days a comparable parameter search would demand.

Our approach to calibration is facilitated by our choice of the taste shock vector. The actual distribution of idiosyncratic taste shocks, without substantially more individual level data is unobserved and must be assumed. i.i.d. Gumbel draws allows for aggregation, and conditioning on the realized network allows us to estimate individual-level parameters from aggregate moments. Thus, in our approach to ABM, the model can be calibrated in aggregate even though it must be simulated at the agent level.

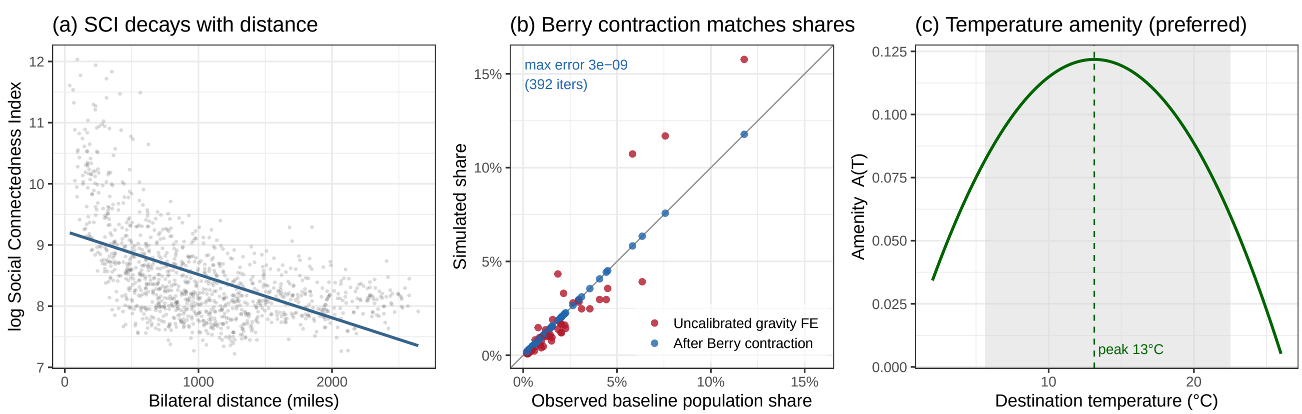

Figure A13: ABM diagnostics. (**a**) Facebook Social Connectedness falls steeply with bilateral distance — the locality that motivates anchoring the friendship network in a small-world graph. (**b**) The Berry contraction drives the simulated baseline population shares onto the observed shares (the $45°$ line; maximum error 3×10^{-9}), whereas the raw, uncalibrated gravity destination fixed effects do not reproduce them. (**c**) The preferred-model destination temperature amenity $A(T) = \beta_1^{ABM} T + \beta_2^{ABM} T^2$ is concave, peaking near $13\,°\mathrm{C}$; the grey band is the in-sample range of state baseline temperatures.

B.6.5 Projection

For each model, two scenarios are simulated: a no-warming baseline (all T_{jt} held at their baseline values) and the SSP3-7.0 warming counterfactual at scaling factors WS$\times$1 and WS$\times$2. The ABM's output is a tensor of state-level population counts $L_{jt}^{(s)}$ indexed by year $t \in \{0, \ldots, 25\}$ (years since the baseline) and shock realisation $s \in \{1, \ldots, 10\}$. We average across the 10 shock realisations (population is a linear functional of the per-period $\arg\max$ counts, so this is the unbiased estimate of $E[L]$) before computing the per-implementation migration share (Equation 4). Agent counts are converted to population units by scaling by the ratio of total population to the number of simulated agents (a constant factor that cancels in the migration-share ratio), giving the per-state population trajectories that enter the headline cross-model comparison.

B.6.6 Uncertainty quantification

We report three uncertainty types for the ABM, in parallel with the other three approaches:

None. The preferred model point estimate at the amenity-coefficient mean.

Statistical. The $\{q_{2.5}, q_{25}, q_{75}, q_{97.5}\}$ quantiles of the migration-share values from the preferred model, one per amenity-coefficient draw. The Statistical band uses 51 draws: the mean coefficient plus 50 random draws sampled from the *origin-clustered* coefficient covariance of the population-weighted amenity regression, over the two temperature coefficients $(\beta_1^{ABM}, \beta_2^{ABM})$. The other structural parameters — γ^{ABM}, the distance cost, the destination effects α_j^{ABM}, and σ^{ABM} — are held at their calibrated values. Fewer draws are used than the $K = 500$ of the other approaches because each draw requires a full set of N_{sim} stochastic simulations (10 per draw); 51 are sufficient to stabilise the reported quantiles. The band is computed on the *per-state run-averaged* population vector (averaging across the $N_{\text{sim}} = 10$ shock realisations within each amenity-coefficient draw before the migration-share aggregation); the rationale is that the migration-share metric involves a sum of absolute differences, which is convex, so per-$(\text{draw}, \text{run})$ sampling would bias the band upward by Jensen's inequality. Run-averaging recovers the (L_0, L_1) pair for each amenity-coefficient draw. Aggregate-level shock variance is small (state-level shocks cancel in expectation; per-state shock noise is larger but does not materially affect the aggregate band).

Choice. The same four quantiles across the 120 implementations at the amenity-coefficient mean, capturing choice (researcher-degree-of-freedom) variance.

B.6.7 Limitations and robustness

Several features of the ABM warrant explicit caveat.

First, The ABM is not designed to capture income or productivity-side responses to climate; it captures the amenity channel only. Cross-channel responses are the subject of the GE implementation (Appendix B.8), in which $\bar{A}$ and $\bar{Z}$ are both functions of T.

Second, in small-population states (NV, WY, VT, ND, SD, DE) the simulation collapses to fewer than ≈ 100 agents by the 25-year horizon; a swing of a few agents produces $\pm 10\%$ per-state ΔL, which is small-N simulation noise rather than meaningful climate response.

Third, and most importantly, our ABM sits at the simple end of the migration agent-based-modelling literature. Our model has a single agent "type," one behavioural channel (a temperature amenity), and a parsimonious friendship network. We omit features that richer migration ABMs may incorporate: for example, bounded rationality, heuristic or satisficing decision rules in place of exact utility maximisation; rich agent heterogeneity by age, income, education, or housing tenure; liquidity, credit, and housing-market frictions; learning and adaptive expectations, and imperfect information beyond our social network; and endogenous local feedbacks such as congestion, wages, and prices that our general equilibrium model can capture. Our minimalism isolates the social-network channel and keeps the ABM comparable to our other three deliberately stylised approaches on all margins *except* for the one that only the ABM can incorporate: the social network. This means that our ABM approach should be read as a controlled experiment in the role of one key mechanism that operates at the agent level. Adding these features is the natural direction for future work, and the choice grid here is built so that further channels could be layered on.

B.7 Gravity

This appendix gives the full specification of our gravity implementation. Our preferred gravity model is a two-stage estimator: in Stage-1, we regress rescaled bilateral migrations flows onto a destination fixed effect, a data-driven self-flow-by-destination anchor, and per-origin linear time trends with a Poisson pseudo-maximum-likelihood (PPML) estimator. In the second stage, we augment the Stage-1 results with a cross-validated lasso interaction layer fit on the Stage-1 residuals with a large suite of bilateral covariates including temperature. We accompany the preferred model with a 144-choice grid, and we propagate two sources of uncertainty: *statistical*, from a $K = 500$ origin-clustered Dirichlet bootstrap of the preferred two-stage model; and *choice*, from re-projection across the 144-choice grid.

B.7.1 Approach overview

Our gravity implementation is built around the migration-rate generating function f under multinomial choice probabilities so $m_{ijt}/m_{iit} \propto \exp\{V_{jt} - V_{it} - D(d_{ij};\delta)\}$, where V_{jt} collects destination-side covariates including temperature, V_{it} collects origin-side covariates, and $D(\cdot)$ proxies for bilateral moving costs (Section 4). That is, we again assume that agents' idiosyncratic taste shocks ϵ_{ajt} in the model (1d) to be i.i.d. Type-I extreme value (Gumbel), generating multinomial logit choice probabilities. Dividing the move probability m_{ijt} by the stay probability m_{iit} cancels the origin-specific logit denominator (the multilateral-resistance term), leaving exactly the exponentiated term. We therefore work throughout with the *ratio transform* m_{ijt}/m_{iit} (the bilateral migration flow relative to the origin's own stayers).

We estimate the model in two stages. We split the choice probabilities' argument into a parametric main-effects part and a residual interaction part:

$$\log \frac{m_{ijt}}{m_{iit}} = \underbrace{V_{jt} - V_{it} - D(d_{ij})}_{\text{Stage 1: PPML main-effects index}} + \underbrace{B\big(T_{it}, T_{jt}, \mathbf{X}_{ijt}; \boldsymbol{\theta}^{Gr}\big)}_{\text{Stage 2: regularised interaction residual}}, \tag{20}$$

where Stage 1 is the standard bilateral PPML linear index in temperature, controls, and fixed effects, and B is a cross-validated learner fit on the temperature$\times$covariate interactions that the additive Stage 1 omits ($\mathbf{X}_{ijt}$ collects the non-temperature covariates). The Stage 2 layer uses machine learning to predict the residual as a function of temperature and non-temperature covariates. Our approach differs from the 'standard' gravity literature by using this two-step learner embedded in the multinomial choice model (which obeys adding-up constraints, thus preserving population, and capturing multilateral resistance).

B.7.2 Preferred model

Our preferred specification has the following Stage-1 regression:

$$\log E\left(\frac{M_{ijt}}{M_{iit}}\right) = \sum_{p=1}^{2}\sum_{l=0}^{2}\Big[\beta_{p,l}^{Gr,o}\,T_{i,t-l}^{p} + \beta_{p,l}^{Gr,d}\,T_{j,t-l}^{p}\Big] + \delta_0^{Gr}\,d_{ij} + \gamma^{Gr,o}\,\mathrm{GSP}_{i,t-1} + \gamma^{Gr,d}\,\mathrm{GSP}_{j,t-1} + \alpha_j^{Gr} + \alpha_j^{Gr,\,\mathrm{self}}\cdot\mathbf{1}(i=j) + \delta_i^{Gr}\cdot t. \tag{21}$$

The regressand is the ratio transform m_{ij}/m_{ii} defined in the Approach overview, which Stage 1 takes directly as its dependent variable. We define the right-hand-side components in turn.

Temperature and continuous controls. Temperature, the object of interest, enters as a quadratic in both origin and destination temperature at 0-, 1-, and 2-year lags, alongside continuous controls for lagged origin and destination gross state product (GSP, per filer, lag 1) and bilateral distance (in thousands of miles). In projection the non-temperature controls are held at their baseline values, so only the temperature terms move. Thus Stage 1 resembles the specifications used in both the ABM and GE approaches.

Fixed effects and identification. The Stage 1 specification controls for a destination fixed effect α_j^{Gr} and a per-origin linear trend $\delta_i^{Gr}\cdot t$. We do not include an origin fixed effect as it is not implied by the model (dividing by m_{iit} eliminates the multilateral resistance part of the lefthandside). The per-origin linear trend instead absorbs persistent state-specific demographic drift while preserving the cross-state variation in long-run temperature that identifies the climate response.

Self-flow anchor. On the self-flow rows ($i=j$) the flow equals the stayer count, so the within-origin ratio is identically one; the self-flow-by-destination dummy block $\alpha_j^{Gr,\,\mathrm{self}}$ absorbs this anchor, so we do not try to fit it. Moreover, it supplies the per-state home-bias coefficient, and renders the resulting $\widehat{m}_{ij}$ diagonal *data-driven* rather than imposed. We fit Stage 1 on the full bilateral panel.

Stage-2 lasso residual layer. Stage 2 assumes the residual term $B()$ of Equation 20 as a cross-validated lasso. It regresses the bilateral count on the Stage-1 prediction (as an offset) plus a penalised interaction basis,

$$\log E(\mathrm{flow}_{ijt}) = \underbrace{\log \mathrm{stayers}_{it} + \hat{f}_{ijt}^{\mathrm{S1}}}_{\text{offset}} + \sum_{k=1}^{36}\theta_k\,Z_{ijt}^{(k)}, \qquad \hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}}\Big\{\mathrm{dev}_{\mathrm{Pois}}(\boldsymbol{\theta}) + \lambda\,\|\boldsymbol{\theta}\|_1\Big\}, \tag{22}$$

where $\hat{f}^{\mathrm{S1}}_{ijt}$ is the fitted Stage-1 linear predictor (Equation 21) and the $\{Z^{(k)}\}$ are the 36 temperature $\times$ non-temperature interactions: 12 temperature features ($\{T, T^2\}$ for origin and destination at lags $0, 1, 2$) times the 3 non-temperature features ($d_{ij}, \mathrm{GSP}_{i,t-1}, \mathrm{GSP}_{j,t-1}$). Using $\hat{f}^{\mathrm{S1}}_{ijt}$ as an offset in Stage 2 means that our learners fit only the residuals of the Stage 1 model. For the lasso, the penalty term λ is selected by cross-validation. For the lasso, the chosen penalty, the number of interactions retained, and the fit improvement are reported in the Estimation subsection below (Section B.7.4).

We fit all Stage-2 learners to the same training panel, with the same outcome and the same Stage-1 offset, and all share a single year-blocked partition of the sample into five folds: the lasso and the gradient-boosted trees select their tuning parameters by five-fold cross-validation over these folds, while the neural networks train on the first four blocks and use the final block for early stopping. We include gradient boosting as the frontier tree-based benchmark, as it can capture non-parametric interactions beyond those in the lasso's pre-specified set of interacted covariates.

Temperature responses. Figure A14 plots our preferred model's Stage 1 temperature responses, summing the contemporaneous and lagged temperature responses. On the "push" side (panel a), out-migration rates, relative to staying, are elevated for cold origins, and flatten through the middle of the observed baseline temperature support. Its confidence band widens sharply where the WS$\times$2 projection extrapolates beyond the in-sample support. On the "pull" side (panel b), state in-migration rates are depressed at cold destinations and peak for destinations near $18°$C. The two responses are close to mirror images (the summed linear terms are -0.089 and $+0.089$). In our gravity framework, bilateral rates respond to the destination-minus-origin temperature gap.

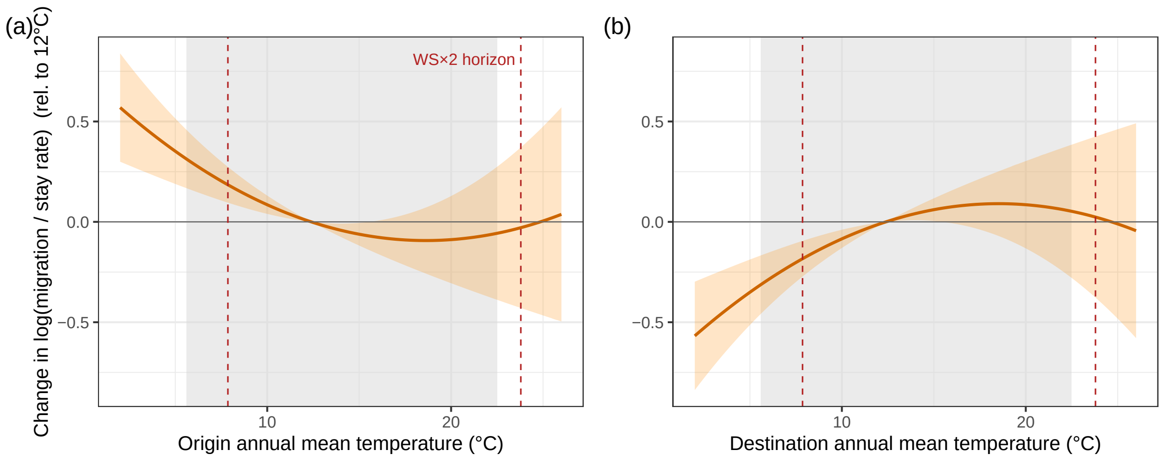

Figure A14: Preferred-model temperature responses of the gravity model: the change in the log bilateral migration rate relative to staying, $\log(m_{ij}/m_{ii})$, as **(a)** the origin's and **(b)** the destination's temperature shifts permanently (summing the contemporaneous and lag-1/lag-2 terms), with delta-method 95% bands and a shared vertical axis. The grey band marks the in-sample temperature support; the dashed lines bound the projected WS$\times$2 horizon temperatures across the 48 states.

B.7.3 Choice grid

We enumerate 144 choices in Table A6.

Table A6: Gravity: The 144-choice grid.

Choice	Values	Justification	$\lvert\cdot\rvert$
Stage-1 fixed effects	{full, drop origin, drop origin + trends}	origin-FE treatment of long-run drift	3
Stage-1 controls	{baseline, + inequality/precip., + SCI/log-pop.}	controls beyond temperature	3
Temperature polynomial	{quadratic, cubic}	polynomial degree on T	2
Temperature lags	{lag 0 only, lags 0–2}	lag structure	2
Stage-2 learner	{none, lasso, neural net, gradient boosting}	residual learner	4
Cartesian product			144

The three fixed-effect choices are central choices. The *full* specification is a nearly-saturated bilateral PPML with origin, destination, and self-flow-by-destination effects. The *drop-origin* specification removes the origin fixed effect entirely, allowing the secular trend in origin temperature to identify the climate response, as the macro climate-migration literature (Cai et al., 2016) implicitly does when it estimates panels with year fixed effects only. The *drop-origin-plus-trends* specification (our preferred specification) adds an origin-specific linear trend on top, absorbing changes in population due to “demographic drift” while preserving the long-run temperature variation that identifies the response. All three retain the self-flow-by-destination fixed effects (i.e., state-specific "home bias" controls) and all three preserve the textbook negative correlation $\mathrm{corr}(T_0, \Delta L_{25}) < 0$ at WS$\times 2$.

The three control sets are: *baseline*, the continuous controls $\{d_{ij}, \mathrm{GSP}_{i,t-1}, \mathrm{GSP}_{j,t-1}\}$ only; + *inequality/precipitation*, which adds origin and destination income inequality (the relative mean deviation) and lagged precipitation (origin and destination); and + *SCI/log-population*, which adds the bilateral log Social Connectedness Index and origin/destination log-population.

The two polynomial choices span the degrees commonly used in the literature: the quadratic retains degrees $\{1, 2\}$ in temperature, the cubic degrees $\{1, 2, 3\}$. The two lag choices control dynamic effects: one builds temperature terms at lag 0 only, the other at lags $\{0, 1, 2\}$. Our preferred model uses the latter.

The four Stage-2 learners are: *none* (Stage-1 only); *lasso* (the preferred Stage 2, a cross-validated lasso on temperature$\times$non-temperature interactions); a *neural network* ensemble (five 64–32–1 multilayer perceptrons); and *gradient-boosted* Poisson trees. The neural network implementations are deliberately a sensitivity row rather than a candidate for the headline: Stage 1 extrapolates beyond the in-sample temperature support for every state, and at the extrapolated WS$\times 2$ temperature, neural network ensembles that fit indistinguishably in-sample produce divergent extrapolations (the well-known interpolation–extrapolation gap). The lasso is preferred because it lives within the same generalised-linear-model class as the Stage-1 PPML and so extrapolates predictably, and because its cross-validated fit is consistent across implementations without overfitting (the selection counts and deviance reduction are reported in the Estimation subsection).

Pinned choices. We cluster standard errors at the origin level (using a Dirchlet bootstrap, see Estimation section below). The projection baseline is the historical reference cross-section ($t = 0$, the first year of the projection horizon). All models are estimated on the same bilateral panel covering training years 1992–2010: although the migration data begin in 1990, the lag-1 GSP control is undefined in 1990 and 1991, so those rows drop out of every fit. Self-flows are retained in training in every model.

B.7.4 Estimation

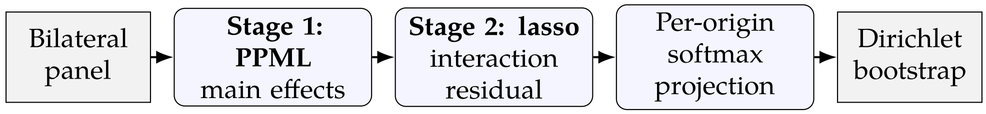

Figure A15: The two-stage gravity pipeline: a Stage-1 PPML for the main effects; a Stage-2 cross-validated lasso for the temperature interactions (a Frisch–Waugh–Lovell residual layer); a per-origin softmax projection under warming; and an origin-clustered Dirichlet bootstrap for the statistical band.

Each Stage-1 specification is fit by Poisson pseudo-MLE with origin-clustered standard errors. The 36 unique Stage-1 specifications (the product of the three fixed-effect, three control-set, two polynomial, and two lag choices, $3 \times 3 \times 2 \times 2 = 36$ — distinct from the 36 candidate interactions of the Stage-2 lasso in Equation 22) are fit once and reused; the four Stage-2 learners are then attached to each, giving 144 choices in total.

We implement the following Stage 2 learners. The *lasso* (the preferred Stage 2) uses a (negative) Poisson log-likelihood as the loss function with an ℓ_1 penalty scaled by λ (standardised inputs, no intercept). The penalty λ is chosen by 5-fold year-blocked cross-validation (each fold splits by calendar year, which preserves within-state serial dependence in errors) to minimise the mean Poisson deviance. For the preferred model, the selected penalty is $\lambda = 1.018$ and the lasso selects 24 of the 36 candidate interactions. The two non-preferred learners — gradient-boosted Poisson trees (depth 5, with early stopping) and a five-member ensemble of 64–32–1 multilayer neural network (i.e., perceptrons with ReLU activations and dropout, trained on the 15 continuous features with early stopping) — both treat the Stage 1 prediction as an offset when minimising the Poisson deviance.

In 15 of the 144 models — mostly cubic-temperature models with extra controls — the selected lasso penalty sits at the floor of the candidate grid, suggesting that a wider penalty would shrink them further. We retain these choice nonetheless, with that caveat. The preferred model is unaffected.

B.7.5 Projection

For each model across our set of choices, we form the projected migration matrix by evaluating Stage 1 (and Stage 2, if present) at the SSP3-7.0 projected origin temperature in each year, holding all non-temperature features (distance, GSP, SCI, log-population, lagged precipitation) at their baseline values. The Poisson first-order conditions guarantee the Stage 1 fit exactly replicates the baseline population distribution, but our Stage 2 modification means that this match is not exact. Under counterfactual warming, only the temperature features change. The per-state population trajectory is then forward-rolled through the migration matrices in the warming and no warming scenarios:

$$L^{(c)}_{i,t+1} = \sum_j m^{(c)}_{jit} \cdot L^{(c)}_{j,t}, \qquad t \in \{0, \dots, 24\}, \quad c \in \{0, 1\},$$

where $c = 0$ is the no-warming trajectory and $c = 1$ the warming trajectory (computed at both WS$\times 1$ and WS$\times 2$; the headline uses WS$\times 2$). The migration-share metric at WS$\times 2$ is $100/2 \cdot \sum_i |L_{1,i,25} - L_{0,i,25}| / \sum_i L_{i,0}$, identical to the metric used by all four approaches. Total population is conserved over the horizon. The same per-state L_0, L_1 pair is the input to the Statistical and Choice bands.

B.7.6 Uncertainty quantification

We report three uncertainty types in parallel with the other three approaches:

None. The preferred model point estimate. The other 143 choices contribute their own point estimates, which enter only the Choice band.

Statistical. An origin-clustered Dirichlet bootstrap of the preferred two-stage model. We draw $K = 500$ origin-level weight vectors $w^{(k)} \sim \mathrm{Dirichlet}(\mathbf{1})$, and on each draw re-fit *both* stages: Stage-1 PPML with weights $w^{(k)}$, and Stage-2 lasso with the same weights, holding the penalty fixed at the preferred model's selected value. Each draw therefore moves the Stage-1 coefficients, the Stage-2 lasso coefficients, and the selected-feature set (the weights change which interactions are zeroed), but not the penalty level itself. The reported band is the $\{q_{2.5}, q_{25}, q_{75}, q_{97.5}\}$ quantiles of the resulting $K = 500$ migration-share values.

Choice. The same four quantiles across the 144-choice grid at each model's point estimate, capturing choice (researcher-degree-of-freedom) variance.

The Dirichlet primitive is origin-clustered: the same weight is applied to every (i, j, t) observation sharing origin i, so the bootstrap mimics resampling *origins* with replacement, which is the natural cluster level given that origin-level serial dependence is the main standard-error concern. The Dirichlet (rather than nonparametric) bootstrap is used since it puts continuous weights on each observation, meaning we can estimate all fixed effects in each bootstrap iteration. The per-draw weight vector is rescaled to sum to the number of training observations, which keeps the penalty's scale relative to the likelihood constant across draws.

B.7.7 Limitations and robustness

We highlight three caveats.

First, the projection assumes the Stage-1 PPML semi-elasticity at the in-sample mean applies at the out-of-sample WS$\times 2$ temperature; every one of the 48 states has its horizon WS$\times 2$ temperature outside its 1992–2010 training range — on average $+0.8\,^\circ$C beyond the hottest year observed in sample, and $+1.3\,^\circ$C beyond it at most (distinct from the full WS$\times 2$ anomaly relative to the average year, $\approx +2\,^\circ$C, reported in Table A3) — so every projection is mechanically an extrapolation. The temperature-polynomial choices (quadratic vs. cubic) exposes the sensitivity; cubic specifications produce wider Statistical bands because cubic polynomials extrapolate more aggressively.

Second, the Stage-2 lasso is fit on the same panel as Stage 1, so its degrees of freedom enter the cross-validation deviance reduction but not a separately quantified standard error. The bootstrap implicitly handles this: at each Dirichlet draw, both stages are re-fit, so the resulting band absorbs the second-stage selection variance (across the 500 draws all 36 candidate interactions are selected on at least one draw, and none is retained on all 500).

Third, the gravity approach is *partial-equilibrium*: equilibrium feedback through prices, wages, and externalities is not modelled, which is the role of the GE approach (Appendix B.8). The gravity model's response is a behavioural reduced-form that conflates utility-side and equilibrium-side responses; under our preferred specification the implied migration share is the largest of the four point estimates — roughly threefold the causal-inference, agent-based, and general equilibrium cluster — consistent with a reduced form that translates the temperature signal into movement without the equilibrium feedback that tempers the GE response.

B.8 General equilibrium

This appendix gives the full specification of our general equilibrium (GE) implementation. The preferred model is a dynamic spatial general equilibrium model with Fréchet migration and constant-elasticity-of-substitution-based (CES) goods trade across the 48 contiguous states, evaluated at central migration and trade elasticities ($\varepsilon^{GE} = 3$, $\sigma^{GE} = 5$), and a levels-with-lag quadratic damage function. We accompany the preferred model with a 128-choice grid — 32 structural choices, including two different cost calibration regimes, each fit under 4 damage-function specifications, and we propagate two distinct sources of uncertainty: *statistical*, from $K = 500$ multivariate-normal draws of the preferred model's damage-function coefficients; and *choice*, from re-projection across the full 128-choice grid.

Notation: each structural parameter carries its ^{GE} superscript where it is first defined; because every parameter in this appendix is specific to the general equilibrium model, the superscript is omitted thereafter for brevity (σ, ε, α, ζ, γ denote σ^{GE}, ε^{GE}, α^{GE}, ζ^{GE}, γ^{GE} in this appendix section).

B.8.1 Approach overview

Our GE implementation specifies the migration-rate generating function f structurally: the bilateral migration rates m_{ijt} are the Fréchet choice probabilities of a dynamic spatial general equilibrium model of the 48 contiguous US states. Unlike the other three approaches, climate enters through temperature-driven damages to *both* fundamental productivity $\bar{Z}_{it}$ and fundamental amenities $\bar{A}_{it}$, and the projected migration response is an equilibrium outcome: agents re-sort across locations while wages, goods prices, and the population distribution adjust simultaneously. The defining feature relative to the gravity and CI approaches is this equilibrium feedback — a state with climate-damaged productivity can adapt via trade (residents consume more of their neighbours' goods) and by congestion relief (out-migration raises the real wage of those who remain) — so the projected redistribution accounts for these general-equilibrium adjustments, rather than the "partial equilibrium" elasticities measured in the data.

Model environment The economy is composed of N locations indexed by i. Time is discrete and indexed by t. At the start of each period t, there is a distribution of population across space $L_{i,t-1}$ where $\sum_i L_{i,t-1} = \bar{L}$. Agents choose where to live, subject to migration costs, and how much to consume. Agents inelastically supply labour to a representative goods firm in destination earning a wage w_i. Each location produces a unique traded variety consumed by agents.

Agents' preferences In period t, an agent a residing in i in period $t-1$ faces the following utility maximisation problem for location choice,

$$j^*_{ait} = \text{argmax}_j \ (V_{jt}/\mu_{ij})\ \epsilon_{aijt}, \tag{23}$$

where ϵ_{aijt} is Fréchet distributed with location parameter 1 and shape parameter ε^{GE}. Note that the log of a Fréchet random variable is Type-1 Extreme Value distributed. Agents' rewards come from the solution to the following utility maximisation problem,

$$V_{jt} = \max_{\{c_k\}} U_{aijt} = \underbrace{A_{jt}}_{\substack{\text{destination}\\ \text{amenities}}} \times \underbrace{\left(\sum_{k=1}^{N} c_k^{\frac{\sigma^{GE}-1}{\sigma^{GE}}}\right)^{\frac{\sigma^{GE}}{\sigma^{GE}-1}}}_{\text{consumption utility}} \quad \text{subject to} \quad \sum_k p_{kjt}c_k \leq w_{jt}. \tag{24}$$

The maximised value V_{jt} is the indirect utility of residing in j; it depends on the destination and period only (through local prices and the wage), not on the agent or origin, and is the object that enters the location choice (23).

Goods production Each location has a representative firm that produces a location-unique variety with technology, $y_j = Z_j L_j$. Goods are shipped with iceberg trade costs τ_{ij}, so that τ_{ij} units of an i good must be shipped to j for one unit to arrive.

Agglomeration and congestion forces Following a long tradition in urban and spatial economics, we allow for local externalities in the form of local increasing returns,

$$A_{it} = \bar{A}_{it}\big(L_{it}\big)^{-\alpha^{GE}}, \quad Z_{it} = \bar{Z}_{it}\big(L_{it}\big)^{\zeta^{GE}}. \tag{25}$$

Here, α^{GE} captures local congestion externalities, which are negative externalities from co-location – increased traffic, pollution, and so on, that deteriorate local fundamental amenities $\bar{A}_{it}$ They may also capture the effect of rising prices of local services, like housing. ζ^{GE} captures the extent of local positive externalities from large regions, like knowledge spillovers, increased product variety, and more fluid local labour markets, all of which amplify local fundamental productivities $\bar{Z}_{it}$.

Definition (a period-t equilibrium) Given trade and migration cost matrices, $\{\mu_{ij}\}$, $\{\tau_{ij}\}$, amenities, A_{it}, productivities Z_{it}, parameters $\{\sigma, \varepsilon, \alpha, \zeta\}$, and an initial vector of population $L_{i,t-1}$, a period-t equilibrium consists of,

1. all agents, taking goods prices, wages, and amenities as given, maximise (24) by optimally picking j and $\{c_k\}$;

2. producers, taking goods prices, wages, and productivities as given, maximise profits, by

optimally choosing labour demand L_{it},

3. the value of amenities and productivities endogenously update according to (25);
4. and goods markets clear, $y_{it} = \sum_j \tau_{ij} c_{jt} L_{jt}$.

These conditions ensure that all agents behave optimally, every agent lives somewhere, and all goods produced are consumed or 'melt' in transport, and that workers earn their marginal product, $w_j = p_j Z_j$.

A dynamic equilibrium of the model is a sequence of period-t equilibria interlocked through the dependence of a period-t equilibrium on $L_{i,t-1}$.

A period-t equilibrium of the model can then be described as a system of $2N$ nonlinear equations in the $2N$ unknowns $\{w_{it}, L_{it}\}$,

$$\begin{aligned} L_{jt} &= \sum_i m_{ijt} L_{i,t-1} \\ w_{it} L_{it} &= \sum_j p_{ijt} c_{ijt} L_{jt}, \end{aligned} \tag{26}$$

where the second equation writes goods market clearing in terms of *values*, not quantities. The variables c_{ij} and m_{ijt} are endogenous and come from the functional form assumptions. In particular,

$$m_{ijt} = \frac{\mu_{ij}^{-\varepsilon} \left(V_{jt}\right)^{\varepsilon}}{\sum_k \mu_{ik}^{-\varepsilon} \left(V_{kt}\right)^{\varepsilon}}, \quad V_{jt} = \frac{A_{jt} w_{jt}}{P_{jt}}$$

where the term $P_{jt} = \left(\sum_i p_{ijt}^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$, and is referred to as the *price index*. Consequently V_{jt} is the amenity-adjusted real wage. Prices reflect marginal costs gross of transportation costs, $p_{ij} = \tau_{ij} \frac{w_i}{Z_i}$. Additionally, goods demand is, $c_{ijt} = \left(\frac{p_{ijt}}{P_{jt}}\right)^{-\sigma} \frac{w_{jt}}{P_{jt}}$. The solution to the system of equations is homogeneous of degree 0 in prices, which reflects the fact that only *relative* prices matter in determining agents' utility. This fact demands that we enumerate prices in units of a good, rather than dollars, by picking a numeraire and setting one price equal to 1. The choice of numeraire has no impact on migration rates.

We assume that both fundamental amenities $\bar{A}_{it}$ and fundamental productivities $\bar{Z}_{it}$ are functions of location temperature, T_{it}. Thus, given $L_{i,t-1}$ and T_{it}, updating L_{it} requires solving the nonlinear system (26). This is a very rich nonlinear structure and allows for many different types of dependencies. We now discuss several examples.

First, suppose a state i is negatively affected by climate, but nearby states are not. Residents of i can adapt to climate change not only by migrating, but also through trade, by consuming more of their neighbours' goods instead. However, if damages from climate change are spatially correlated, a negative shock to one region may be amplified by the fact that its neighbours are also worse off. Second, this model introduces some local congestion. Suppose one location is relatively attractive following a change in temperature. In equilibrium, not all agents would want to move there. The reason is that the more people move to a location, the more output that location produces. By making that location's good relatively less scarce, its price falls – workers in other states might consequently reap a higher wage.

Algorithm for computation. Given $L_{i,t-1}$, and current-period temperature T_{it}, the following algorithm computes $\{w_{it}, L_{it}\}$. The algorithm is iterative, and begins with a guess of $w_{it}^{(0)}$ and $L_{i,t}^{(0)}$, where the superscript denotes the iteration step. A good guess for wages and population are, $L_{i,t}^{(0)} = L_{i,t-1}$ and $w_{it}^{(0)} = w_{i,t-1}$, as population movements are generally slow.

1. Using $L_{i,t}^{(n)}$ and $w_{it}^{(n)}$ construct A_{it}, Z_{it}, and P_{it},

$$A_{it} = \bar{A}_{it}(T_{it}) \left(L_{it}^{(n)}\right)^{-\alpha}, \quad Z_{it} = \bar{Z}_{it}(T_{it}) \left(L_{it}^{(n)}\right)^{\zeta}, \quad \text{and} \quad P_{jt} = \left(\sum_i \left(\tau_{ij} \frac{w_{it}^{(n)}}{Z_{it}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

2. Construct,

$$V_{jt} = \frac{A_{jt} w_{jt}^{(n)}}{P_{jt}}, \quad \text{and} \quad m_{ijt} = \frac{\mu_{ij}^{-\varepsilon} (V_{jt})^{\varepsilon}}{\sum_k \mu_{ik}^{-\varepsilon} (V_{kt})^{\varepsilon}}$$

3. Update population,

$$L_{jt}^{(n+1)} = \sum_i m_{ijt} L_{i,t-1}$$

4. Form a guess for the wage,

$$\tilde{w}_{jt}^{(n+1)} = \left(L_{jt}^{(n)}\right)^{-1/\sigma} Z_{jt}^{1-1/\sigma} \left(\sum_i (\tau_{ji})^{1-\sigma} P_{it}^{\sigma-1} w_{it}^{(n)} L_{it}^{(n)} \right)^{1/\sigma}$$

5. Compute the error,

$$\text{error} = \max_j \mid \tilde{w}_{jt}^{(n+1)} - w_{jt}^{(n)} \mid$$

   If the error falls below a pre-specified tolerance (in the implementation, 10^{-15}), convergence is achieved and stop. Otherwise, update the wage,

$$w_{it}^{(n+1)} = w_{it}^{(n)} + \theta \left(\tilde{w}_{it}^{(n+1)} - w_{it}^{(n)} \right),$$

   where $\theta \in (0, 1)$ is a relaxation parameter, and repeat steps 1–5 until convergence. We fix $\theta = 0.8$; after each wage update the numeraire is re-imposed by normalising w by its first entry.

Roundabout (intermediate-input) production. We also use a *roundabout* variant of the production block as a choice, in which a share γ^{GE} of gross output is itself used as an intermediate input. In these models, the production function for each location combines labour and an intermediate bundle x_i (the same CES aggregate of state goods that consumers buy, which is priced at the same local price index P_i) via a Cobb–Douglas technology $Q_i = Z_i \ell_i^{1-\gamma} x_i^{\gamma}$. Under cost minimisation, the factory-gate price is

$$p_i = c_\gamma \, w_i^{1-\gamma} \, P_i^{\gamma} / Z_i, \qquad c_\gamma = (1-\gamma)^{-(1-\gamma)} \cdot \gamma^{-\gamma}.$$

Now, the local price index now enters the cost of production: a state's goods are an input to every other state's output; a local productivity shock propagates across states through the input–output network. Because labour remains the only primary factor, GDP per capita still equals the wage, and goods market clearing — $w_i L_i = p_i^{1-\sigma} \sum_j \tau_{ij}^{1-\sigma} \, w_j L_j / P_j^{1-\sigma}$ — is nearly identical to the $\gamma = 0$ case (the factor $1/(1-\gamma)$ cancels from both sides). The roundabout share therefore lives entirely inside the definition of p, the key equilibrium price in the model (i.e., the solution to high dimensional system of nonlinear equations). Two consequences follow. (i) In the "forward solve" of the model, p and P must be solved jointly at fixed (w, L) as a nested inner $p \rightleftarrows P$ fixed point, after which the wage update uses the modified market-clearing exponent $\xi = \sigma - \gamma(\sigma - 1)$,

$$w_i = \left[c_\gamma^{1-\sigma} \, L_i^{-1} \, P_i^{\gamma(1-\sigma)} \, \tilde{Z}_i^{\sigma-1} \, \Phi_i \right]^{1/\xi}, \qquad \tilde{Z}_i = Z_i L_i^{\zeta},$$

with $\Phi_i = \sum_j \tau_{ij}^{1-\sigma} w_j L_j / P_j^{1-\sigma}$. The migration block remains unchanged. (ii) In the inversion under $\gamma \neq 0$, the (p, P) market-clearing fixed point is γ-invariant; γ enters only the closed-form $Z_i = c_\gamma \, w_i^{1-\gamma} \, P_i^{\gamma} / p_i$, so the recovered Z_i change with γ even though the iteration does not (as Z_i change, the damage function must be re-estimated). At $\gamma = 0$, $\xi \to \sigma$, and both the forward solve and the inversion collapse exactly to the economy of the baseline model (Equation 26); $p_i = w_i / Z_i$, recovering the algorithm above.

B.8.2 Preferred model

Our preferred model is $\varepsilon = 3, \ \sigma = 5, \ \alpha = 0, \ \zeta = 0, \ \gamma = 0$, a gravity-based estimate of the migration and trade costs (projecting μ and τ onto distance), with the levels-with-lag quadratic damage form (damage specification m_2); in other words, the literature-standard migration and trade elasticities, both the congestion/agglomeration externalities and the roundabout share switched off, smoothed trade and migration costs from a gravity regression, and the modal damage specification. This implementation is the headline specification; the other 127 models contribute their own point estimates, which enter the Choice band.

B.8.3 Choice grid

The grid is not a single six-way Cartesian sweep over independent scalars. It is a 32-choice *structural* grid — 'the product of two migration-elasticity values, two trade-elasticity values, two externality regimes, two roundabout shares, and two sets of μ and τ from different calibration procedures ($2^5 = 32$)' — each of which is then fit under each of 4 damage-function specifications, giving 128 implementation choices (Table A7). The four structural parameters cannot be identified from the within-model gravity step and must be assumed; the trade and migration cost calibration approach navigates a bias-variance tradeoff (Dingel and Tintelnot, 2026); the damage-function form is a modelling choice in the temperature-response estimation. The role of the grid is to render the resulting choice variance explicit.

Table A7: GE: The 128-choice grid (32 structural choices $\times$ 4 damage functions).

Choice	Values	Justification	$\lvert\cdot\rvert$
ε	$\{1, 3\}$	migration elasticity (Fréchet shape)	2
$\sigma - 1$	$\{4, 8\}$	trade elasticity	2
(α, ζ)	$\{(0,0), (0.3,0.05)\}$	congestion + agglomeration externalities, off/on (varied jointly)	2
γ	$\{0, 0.5\}$	roundabout intermediate-input share	2
$\{\mu\}, \{\tau\}$	{ Distance projected, Exact (Head-Ries) }	Two different calibration procedures	2
Structural choices (Cartesian product of the above)			32
damage spec.	$\{m_1, m_2, m_3, m_4\}$	quadratic-FD / levels+lag quadratic (pref.) / levels+lag cubic / pop.-weighted	4
Total (32×4)			128

The migration elasticity ε governs the strength of agents' response to differences in destination utility. We bracket the consensus range with $\varepsilon \in \{1, 3\}$: $\varepsilon = 1$ matches roughly Suárez Serrato and Zidar (2016) (they find an estimate of $\varepsilon = 1.25$) and the lower end of Caliendo et al. (2019); $\varepsilon = 3$ is a central estimate consistent with Monte et al. (2018); Burzyński et al. (2022).

The trade elasticity (here $\sigma - 1$) governs the substitutability of state-level goods (σ is the elasticity of substitution across state varieties). Published estimates again span a wide range. We bracket the range $\sigma \in \{5, 9\}$, anchored against Broda and Weinstein (2006); Simonovska and Waugh (2014) ($\sigma - 1 \approx 4$–6, near our central value) and the higher end of Eaton and Kortum (2002) (for countries, they find a trade elasticity $\sigma - 1$, of around 8).

The externality parameters α and ζ govern congestion (negative amenity externality from co-location, capturing effects like local price pressure on housing) and agglomeration (positive productivity externality from co-location, capturing effects like knowledge spillovers). α acts only on utility, while ζ affects productivity and therefore other locations through trade. Estimates of α can be large (generally, it is required that $\alpha > \zeta$), while estimates of ζ tend to be near zero. We vary the pair *jointly* (i.e., in lockstep) between an *off* regime $(\alpha, \zeta) = (0, 0)$ and an *on* regime $(\alpha, \zeta) = (0.3, 0.05)$, so the externality implementation choices contribute a single binary (off/on) choice rather than an independent 2×2 grid. The case $\alpha = \zeta = 0$ shuts down both externalities and reduces the model to a pure 'Armington' gravity model in goods plus frictional Fréchet migration; the case $(0.3, 0.05)$ activates both at the upper credible bound.

The roundabout share $\gamma \in \{0, 0.5\}$ governs the fraction of gross output used as an intermediate input in production. $\gamma = 0$ recovers the baseline economy with linear-in-labor production, while $\gamma = 0.5$ introduces production linkages that amplify how a local climate shock propagates through prices and wages locally and across space, calibrated to a value close the U.S. intermediate input share. As with the elasticities and externalities, γ is not identified within the model and is varied to expose its contribution to our implementation choices.

The choice of the migration and trade cost (μ and τ) calibration method navigates two extremes of a bias-variance tradeoff as explained in Dingel and Tintelnot (2026). We model the costs as either a two parameter function of distance (slope and intercept), estimated within a "structural" gravity regression, or as $N(N-1)/2$ parameters that *exactly* fit the observed flows (implicitly assuming no measurement error, the so-called Head-Ries index).

The damage-function specification choices enumerate four (non-nested) strategies for estimating the temperature response of the inverted (A_{it}, Z_{it}) series. All four share the same tight fixed-effect structure (state-by-outcome and year-by-outcome effects) and state-clustered standard errors, and differ only in functional form, left-hand-side transform, and weighting (detailed under Estimation). They are: m_1 a quadratic in first differences (the exact dimensional match to the projection, used as a conservative anchor); m_2 a quadratic in levels with a one-period lag of the outcome as a control (*preferred*); m_3 a cubic in levels with the same lag control; and m_4 a population-weighted variant of m_2. These four trace a roughly continuous spectrum of implied redistribution and bracket the modelling choices that produce qualitatively different damage estimates in this small panel.

Pinned choices. We do not enumerate over (i) the initialisation (we always initialise from the inverted baseline data, not from a forced steady state, because the model does not converge on the projection horizon if forced; we verified that the alternative initialisation produces qualitatively identical ΔL time series but a different initial-period level), or (ii) the GE numeraire (the model is homogeneous of degree zero in prices; the numeraire choice does not affect migration rates).

B.8.4 Estimation

Estimating the GE model We estimate the GE model in three steps. First, we use the model structure and data on trade and migration flows across states to recover τ_{ij} and μ_{ij}. Second, we use data on population, GDP, and our estimated trade and migration costs to recover Z_{it} and A_{it} by solving the nonlinear system 26, treating these parameters as unknowns. Third, we estimate the relationship between temperature T_{it}, amenities A_{it}, and productivities Z_{it}.

Once we have conducted the three steps above, we can use the model to simulate counterfactual migration rates under a different climate. Figure A16 summarises the pipeline.

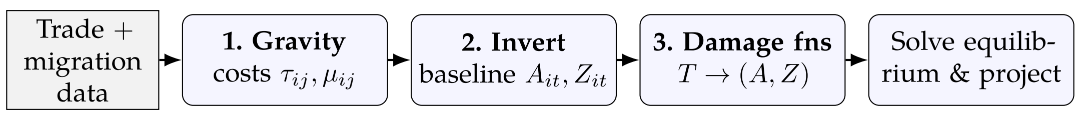

Figure A16: The GE estimation pipeline: recover bilateral trade and migration costs from gravity regressions; invert the baseline equilibrium for fundamental amenities A_{it} and productivities Z_{it}; estimate their temperature damage functions; then solve the equilibrium forward under warming.

Estimate $\{\tau_{ij}\}$ and $\{\mu_{ij}\}$ The model gives log-linear relationships between migration and trade flows, and variables governing the costs of moving goods and agents across space.

First, we note that the model delivers a log-linear relationship between migration rates and their determinants,

$$\log \frac{L_{ijt}}{L_{i,t-1}} = -\varepsilon \log \mu_{ij} + \varepsilon \left(V_{jt}\right) - \log \sum_k \mu_{ik}^{-\varepsilon} (V_{kt})^{\varepsilon}.$$

This equation gives us two approaches to measuring the costs μ_{ij} and τ_{ij}. First, we can assume the costs are a function of distance plus a mean-one error term, and estimate with a regression model. That is,

$$\log m_{ijt} = -\varepsilon\delta \log d_{ij} + \xi_{it} + \xi_{jt} + e^{\mu}_{ijt}$$

where $\mu_{ijt} = d_{ij}^{-\delta} e^{\mu}_{ij}$, and $E(e^{\mu}_{ij} \mid d_{ij}, \xi_{it}, \xi_{jt}) = 1$, then we can estimate this relationship with a Poisson PMLE by fitting,

$$\log E(m_{ijt}) = -\varepsilon\delta \log d_{ij} + \xi_{it} + \xi_{jt}.$$

Subsequently we estimate $-\varepsilon \log \mu_{ij}$ with $-\widehat{\varepsilon\delta} \log d_{ij}$. In practice, we also include a dummy equal to one when $i = j$ and normalize costs by this measure so on average μ_{ii} is equal to one. Note that while ε is unidentified in this regression, μ_{ij} only enters the model raised to this power. The migration-cost regression is fit on the multi-year bilateral migration panel by Poisson pseudo-MLE, with origin-by-year and destination-by-year fixed effects (ξ_{it} = origin×year, ξ_{jt} =

destination×year) and a self-flow dummy; the fitted $\widehat{\mu}_{ij}$ is then collapsed to a single bilateral matrix and normalised by the mean of its self-flow ($i = j$) diagonal.

Likewise, using the demand equation for traded goods in values,

$$\log p_{ij} y_{ij} = -(\sigma - 1)\kappa \log d_{ij} + \varsigma_{it} + \varsigma_{jt} + e^{\tau}_{ij}$$

where again we assume the error is such that $E(e^{\tau}_{ij} \mid d_{ij}, \varsigma_{it}, \varsigma_{jt}) = 1$. Consequently, $-(\sigma - 1)\log \tau_{ij}$ can be estimated with $-\widehat{(\sigma - 1)}\kappa \log d_{ij}$, again using a Poisson PMLE estimator. The regression which calibrates trade costs is estimated on distinct dataset that no other approach uses, the 2012 Commodity Flow Survey (CFS) (a single cross-section which we aggregate to state-to-state shipment values), which uses origin-state and destination-state fixed effects with no year dimension (in contrast to the multi-year migration regression's origin×year and destination×year effects). As with μ, the recovered $\widehat{\tau}^{1-\sigma}_{ij}$ matrix is normalised by the mean of its self-flow diagonal so that $\tau^{1-\sigma}_{ii}$ is on the order of unity. Both regressions are fit by Poisson pseudo-MLE; the cost point estimates, not their uncertainty, are what feed the GE model.

Our alternative approach assumes no idiosyncratic error and directly recovers the migration costs by rearranging the structural gravity equations. For example, assuming $\mu_{ij} = \mu_{ji}$,

$$\sqrt{\frac{L_{ij} L_{ji}}{L_{ii} L_{jj}}} = \mu^{-\varepsilon}_{ij}, \qquad \sqrt{\frac{X_{ij} X_{ji}}{X_{ii} X_{jj}}} = \tau^{1-\sigma}_{ij}$$

where $X_{ij} = p_{ij} y_{ij}$; these are the so-called Head-Ries indices (Head and Ries, 2001). They exactly fit the data, but can mistake idiosyncratic noise or censoring in the data for implausibly large or small migration barriers. The advantage of using them is they can recover migration or trade corridors that do not follow a distance-based gravity law.

The cost estimates implied by the estimated gravity equations are reported in Table A8 (origin- and destination-by-year fixed effects for migration, origin and destination fixed effects for trade; origin-clustered standard errors). The distance coefficient is precisely determined in both equations and the model fit is good — in contrast to the damage-function estimates, which are statistically imprecise (see below).

Dependent Variable:	Migration Flow	Trade Flow
Model:	(1)	(2)
Variables		
log distance	-0.8190***	-0.9893***
	(0.0860)	(0.0535)
self-flow	2.668***	0.5180***
	(0.4296)	(0.1696)
Fixed-effects		
origin×year	Yes	—
destination×year	Yes	—
origin	—	Yes
destination	—	Yes
Fit statistics		
Observations	46,080	2,304
Pseudo R^2	0.99788	0.91213

Origin-clustered standard errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A8: Gravity (cost) estimates: our preferred calibration. The migration column carries origin×year and destination×year fixed effects; the trade column, estimated on a single cross-section, carries origin and destination fixed effects. Both use origin-clustered standard errors. The migration column is the multi-year bilateral migration panel; the trade column is the 2012 Commodity Flow Survey state-to-state cross-section. In the structural inversion the trade costs are recovered from this single cross-section, so the trade regression has no year dimension.

It is important to note that the migration elasticity ε and the trade elasticity σ are not identified in this model: the gravity regressions recover $\widehat{\varepsilon\delta}$ and $\widehat{(\sigma-1)\kappa}$ as composites, and the cost matrices $\mu_{ij}^{-\varepsilon}$ and $\tau_{ij}^{1-\sigma}$ enter the model only through these composites. We therefore take ε and σ as assumptions, calibrated from the external literature, and varying ε and σ across models shows the sensitivity of our results to alternative assumptions across the consensus range $\varepsilon \in \{1, 3\}$, $\sigma \in \{5, 9\}$.

Recover $\{A_{it}\}$ and $\{Z_{it}\}$ in sample For a given set of estimates for τ_{ij} and μ_{ij}, we can invert the system described by (26) to solve for the state-level amenities and productivity. Specifically, we observe for each year t the population L_{it} and the population in the previous period $L_{i,t-1}$. We calculate the wage by dividing the state's GDP by its population (so w_{it} is GDP per capita, which equals value added per worker because labour is the only primary factor). Given these inputs, we invert the system in two passes. *First*, holding the wage w_{it} and population fixed, the goods-market-clearing (p, P) system is iterated to a fixed point in p, so that $p_{it} = c_\gamma\, w_{it}^{1-\gamma} P_{it}^{\gamma}/Z_{it}$, i.e. $Z_{it} = c_\gamma\, w_{it}^{1-\gamma} P_{it}^{\gamma}/p_{it}$ (which collapses to $Z_{it} = w_{it}/p_{it}$ at $\gamma = 0$). *Second*, with (w, P) now fixed, we write $v_{jt} = w_{jt}/P_{jt}$ so that $V_{jt} = A_{jt}\, v_{jt}$; the amenity is recovered from inverting the migration flow equations, $L_{jt} = \sum_i m_{ijt} L_{i,t-1}$. These "Fréchet" shares give the fixed-point relation $A_{jt} = L_{jt}^{1/\varepsilon}\big(\sum_i m^0_{ijt}\, L_{i,t-1}\big)^{-1/\varepsilon}$, where m^0 is evaluated at the current amenity guess and iterated to convergence. Amenities are inverted *after* productivity, because it enters only through $V_{jt} = A_{jt}w_{jt}/P_{jt}$ and therefore requires that we first construct the price index P_{jt}; the externality-free fundamentals $\bar{A}_{it}, \bar{Z}_{it}$ are obtained simply by dividing out the congestion and agglomeration terms of Equation 25. Because the (p, P) system is homogeneous of degree zero, the numeraire $p_1 = 1$ is fixed before the read-off; any constant level shift in $\log Z$ or $\log A$ is later absorbed by the fixed effects of the damage regression. The choice of numeraire has no bearing on the results. This inversion is performed conditional on a particular choice of ε, σ, the externality regime, and the roundabout share γ (which, as noted, shifts the recovered Z_{it} even though the (p, P) iteration itself is γ-invariant).

Specifying a damage function We match the recovered amenities and productivity above to annual average temperatures and estimate the temperature response by OLS. The two outcomes are stacked into a single regression: each observation is a state-year-by-outcome triple indexed by an outcome label $v \in \{A, Z\}$, and every right-hand-side term is interacted with v, with state-by-outcome and year-by-outcome fixed effects and standard errors clustered by state. The outcome interaction makes every coefficient outcome-specific, so the temperature responses and the lagged-outcome control are estimated separately for A and Z rather than sharing a common slope. For the preferred (levels-with-lag quadratic) specification this is

$$\begin{aligned}
\log A_{it} &= \beta_1^A T_{it} + \beta_2^A T_{it}^2 + \rho^A \log A_{i,t-1} + \gamma_i^A + \gamma_t^A + \epsilon_{it}^A, \\
\log Z_{it} &= \beta_1^Z T_{it} + \beta_2^Z T_{it}^2 + \rho^Z \log Z_{i,t-1} + \gamma_i^Z + \gamma_t^Z + \epsilon_{it}^Z,
\end{aligned} \tag{27}$$

where the lagged-outcome term enters as a control whose coefficient ρ^m is *not* restricted to be common across A and Z, and the fixed effects γ_i^m, γ_t^m are the outcome-interacted state and year effects (state-by-outcome and year-by-outcome). The four damage specifications differ as follows: m_1 replaces both sides with first differences and drops the lagged-outcome control ($d \log Y \sim dT + dT^2$); m_2 is the levels-with-lag quadratic above (preferred); m_3 adds a cubic term $\beta_3^m T_{it}^3$; and m_4 is m_2 estimated with population weights. All four use the identical state-by-outcome and year-by-outcome fixed effects and state-clustered standard errors. The four specifications are non-nested. These modelling choices produce qualitatively different damage estimates in this small panel.

For interpretation and as a sense-check, the estimated quadratic temperature response can be re-expressed in a Gaussian "optimal temperature" form. Imposing that there is some optimal temperature for amenities and productivities, μ_m, and that deviations from it affect the outcome according to

$$D_m(T) = \exp\left\{ -\,(T - \mu_m)^2 \,\frac{1}{2 s_m} \right\}, \quad m \in \{A, Z\}, \tag{28}$$

the quadratic coefficients map to the location and shape parameters via $\beta_1^m = \mu_m/s_m$ and $\beta_2^m = -\frac{1}{2s_m}$.[14] We stress that this Gaussian form is *not* the object that drives the projection (see Projection, below): it is used only to reparameterise the estimates for interpretation, and as a sense-check that the first-differences projection reproduces the structural $D_m(T)$ implied by m_1. The projection itself applies the regression temperature coefficients directly in first differences.

The recovered damage functions, and their dependence on the assumed elasticities ε and σ, exhibit several features (Figure A17). First, the estimates are very statistically imprecise: at low elasticities the 95% confidence interval for the amenity-damage location parameter μ_A can span roughly -7 to 19°C, so the point estimates should be interpreted with caution and the sensitivity checks of the choice grid are essential.[15] Second, some estimates do not correspond to priors on their magnitude: influential studies such as Burke et al. (2015) find larger and more convex impacts. Third, there are systematic patterns across the assumed elasticities: higher migration elasticities ε are associated with higher optimal temperatures and less steeply declining amenity damages but leave the productivity-damage parameters essentially unchanged, whereas the assumed trade elasticity σ mainly affects the productivity-damage function (higher σ implies a higher optimal productivity temperature with a shallower slope).

[14] Since $\hat{D}(T) = \exp\left(-\left(T-\mu_m\right)^2/2s_m\right)$ implies $\log D = -\frac{1}{2s_m}T^2 + T\mu_m/s_m - \mu_m^2/2s_m$, matching coefficients against the quadratic in Equation (27) gives $\beta_1^m = \mu_m/s_m$ and $\beta_2^m = -1/(2s_m)$.

[15] Alternative specifications of Equation (27) — for example, estimating in first differences — can lead to qualitatively different damage-function estimates, further highlighting that the damage function estimated in this sample is noisy and imprecise. This is precisely why the damage-estimation strategy is one of the choices.

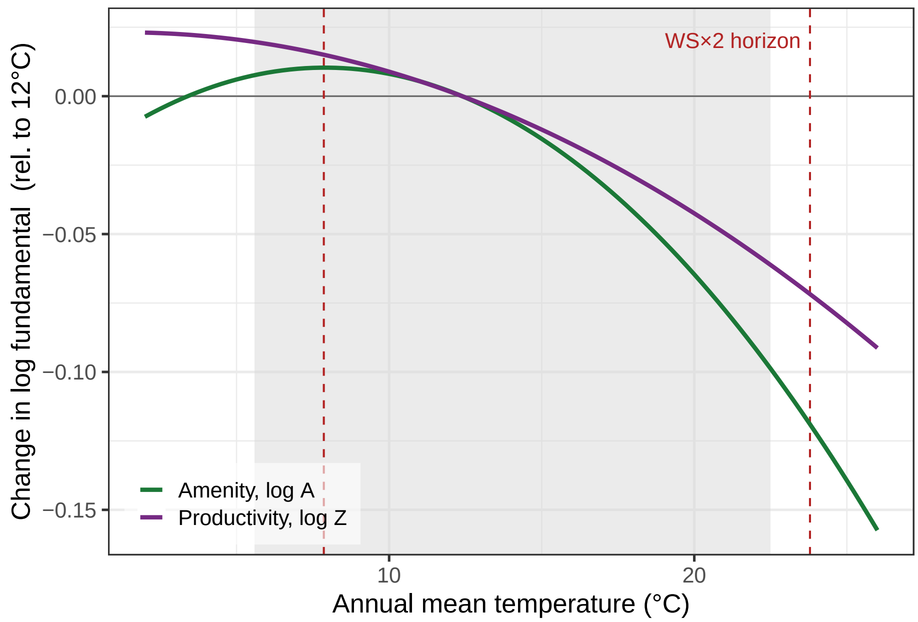

Figure A17: Preferred-model (m_2) temperature damage functions: the change in log fundamental amenity A and log fundamental productivity Z as temperature shifts, relative to a $12\,°\text{C}$ reference, over the in-sample support (grey band) with the WS×2 horizon range marked. Both fundamentals decline with warming across most of the sampled range, however the estimates are statistically imprecise.

B.8.5 Projection

For each implementation and damage-coefficient draw k, the projection produces a 25-year trajectory $(t = 0, \ldots, 25)$ of $(L_{it}^{(k)}, w_{it}^{(k)})$ under three scenarios: no warming, WS$\times$1, and WS$\times$2. The mechanism is a *first-differences* update of the damaged fundamentals: for each outcome the log fundamental is rolled forward year-on-year by

$$\Delta \log Y_{it} = \sum_j \beta_j^Y \, \Delta\big(T_{it}^j\big)\,, \qquad \Delta\big(T_{it}^j\big) = \Big(T_{i,t-1} + \tilde{T}_{it}\Big)^j - T_{i,t-1}^j,$$

where $\{\beta_j^Y\}$ are the regression temperature coefficients for outcome $Y \in \{A, Z\}$ and $\tilde{T}_{it}$ is the (warming-scaled) projected temperature increment; the sum runs to the polynomial order of the chosen specification (so the cubic spec m_3 uses $j \in \{1, 2, 3\}$). The first-differences form is the exact dimensional match to the quadratic-FD spec m_1; for the levels-with-lag specs (m_2–m_4) the lagged-outcome term is treated purely as an estimation control and the projection extracts only the temperature coefficients. The Gaussian form of Equation (28) does *not* enter this update; it appears only as the sense-check that compares this projection against the structural $D_m(T)$ for m_1. The damaged $(A_{it}^{(k)}, Z_{it}^{(k)})$ trajectory is then the input to the equilibrium solve.

The damage coefficients are themselves resampled: we draw $K = 500$ vectors of the damage-function temperature coefficients (β_j^A, β_j^Z) from their joint multivariate-normal posterior implied by the regression covariance. To keep the computation tractable the full $K = 500$ propagation is carried out for the preferred structural choices only; the other 31 structural choices are propagated at their coefficient mean, which suffices for the Choice band. For each implementation, the equilibrium (Equation 26) is solved at every year $t \in \{0, \ldots, 25\}$ under the no-warming baseline and under SSP3-7.0 warming at WS$\times 1$ and WS$\times 2$, using the algorithm above ($\gamma = 0$ choices through the single-loop system, $\gamma > 0$ choices through the nested price system). If an iteratation produces non-finite values, we halve the relaxation parameter θ and re-try the solve a bounded number of times. For roundabout implementations ($\gamma > 0$) that fail to converge from an uninformative initial guess, the equilibrium is traced by homotopy: we start from $\gamma = 0$ and increase γ to the desired value (0.5) in small steps, initialising each step at the previous solution.

The migration share at WS$\times 2$ is computed by the same migration-share metric used throughout the paper (Equation 4).

A non-trivial implementation choice is whether to initialise the GE state at the inverted baseline distribution (which is not in steady state) or at a steady state forced under baseline climate. We adopt the former: empirically, populations were shifting at the baseline is indeed not a steady state. Forcing a steady state at $t = 0$ is therefore, in some sense, counterfactual. The population vector L_0 is initialised at the IRS baseline share. The cost of the non-steady-state initialisation is that in the no-warming trajectory, the population drifts over the projection horizon as the GE state slowly converges to its (counterfactually unchanging) steady-state; this drift cancels in the WS$\times 2$ minus no-warming differential that the migration-share metric isolates. This makes the GE approach to computing the target output identical to the ABM and gravity approaches.

B.8.6 Uncertainty quantification

We report three uncertainty types in parallel with the other three approaches:

None. The preferred model point estimate at the coefficient mean. The other 127 model$\times$structural choices contribute their own point estimates, which feed the Choice band.

Statistical. The $\{q_{2.5}, q_{25}, q_{75}, q_{97.5}\}$ quantiles of the $K = 500$ migration-share values from the preferred model, one per damage-coefficient draw. This captures sampling uncertainty in the damage-function regression, conditional on the elasticity assumptions, the externality regime, the roundabout share, and the preferred damage-function form.

Choice. The same four quantiles across the 128 model$\times$structural choices at the coefficient mean, capturing the elasticity, externality, roundabout, and damage-form uncertainty that the structural model cannot resolve internally.

B.8.7 Limitations and robustness

We have three caveats for the our general equilibrium approach.

First, the damage-function estimation is noisy in our small panel: the 95% confidence interval for the amenity-damage location parameter can span roughly $[-7, 19]$°C at low elasticities, and qualitatively different specifications (notably the first-difference form) yield different damage functions. The $K = 500$ MVN draws from the damage-function regression reflect this, and the resulting Statistical uncertainty is correspondingly wide. However, this statistical uncertainty only reflects coefficient sampling within the preferred damage function estimation strategy. The damage estimation strategy is itself one of our implementation choices and therefore part of Choice uncertainty.

Second, the elasticities ε and σ are not identified internally; the choice grid therefore recognizes that any GE projection is conditional on assumed elasticities, and our Choice uncertainty quantifies this.

Third, the externality choices (α, ζ) and the roundabout share γ are similarly assumed: the externalities are varied jointly between an *off* regime ($\alpha = \zeta = 0$) and an *on* regime ($\alpha = 0.3$, $\zeta = 0.05$), and γ between 0 (no roundabout) and 0.5; none of these choices are data-driven in our setting.

The GE implementation produces qualitatively similarresults to the gravity implementation — migration-share medians both fall in the 0.2–0.9% range under SSP3-7.0 at the 25-year horizon — but with different cross-state response patterns driven by different model mechanisms: for example, the GE approach allows hot states to partially adapt via trade, while the gravity model treats each state's response in near-isolation (the only "interaction" in the Gravity approach is through multilateral resistance).

B.8.8 Extensions

Several extensions are possible, and our implementation here is necessarily narrow. Extensions fall into two categories: first, modifying our baseline model to include additional economic interactions; and second, modifying the model to include richer dynamics and forward-looking behaviour. See Desmet and Rossi-Hansberg (2024) for a full review of this literature.

To modify the baseline model, the resolution of the model could be reduced to a finer geographic scale to include interregional commuting behaviour, local labour markets, and housing markets, in the style of Monte et al. (2018). One could introduce nonhomothetic preferences and sectoral heterogeneity, in the style of Conte (2022) or Nath (2025). Finally, one could include additional climate damages, not only to amenities and productivities, but also to land availability, which erodes as global warming floods nations' coastlines (Desmet et al., 2021).

To add dynamics, one could add forward-looking worker dynamics, so agents account for future damages when making their migration decision, like in Balboni (2019) or Cruz (2024), who extend the model of Caliendo et al. (2019) to account for climate damages. Additionally, one could add dynamics that account for the forward-looking behaviour of agents who accumulate capital, and include damages on the capital stock, as in Bilal and Rossi-Hansberg (2023).

References

Abel, Guy J, Michael Brottrager, Jesus Crespo Cuaresma, and Raya Muttarak, "Climate, conflict and forced migration," *Global environmental change*, 2019, *54*, 239–249. Publisher: Elsevier.

AgMIP, Agricultural Model Intercomparison and Improvement Project, "Agricultural Model Intercomparison and Improvement Project (AgMIP)," November 2023. Publisher: AgMIP.

Allen, Treb and Dave Donaldson, "Persistence and path dependence in the spatial economy," Technical Report, National Bureau of Economic Research 2020.

Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong, "Social Connectedness: Measurement, Determinants, and Effects," *Journal of Economic Perspectives*, August 2018, *32* (3), 259–280.

Balboni, Clare Alexandra, "In harm's way? infrastructure investments and the persistence of coastal cities." phd, London School of Economics and Political Science 2019.

Baylis, Patrick, Prashant Bharadwaj, Jamie T. Mullins, and Nick Obradovich, "Climate and migration in the United States," *Journal of Public Economics*, 2025, *249.*

Beine, Michel and Lionel Jeusette, "A Meta-analysis of the Literature on Climate Change and Migration," *JODE - Journal of Demographic Economics*, September 2021, *87* (3), 293–344. Publisher: Cambridge University Press.

Bell, Andrew Reid, Carlos Calvo-Hernandez, and Michael Oppenheimer, "Migration, Intensification, and Diversification as Adaptive Strategies," *Socio-Environmental Systems Modelling*, January 2019, *1*, 16102–16102.

Benveniste, Hélène, Peter Huybers, and Jonathan Proctor, "Global Climate Migration Is a Story of Who and Not Just How Many," *Nature Communications*, September 2025, *16* (1), 7752.

Benveniste, Hélène, Michael Oppenheimer, and Marc Fleurbaey, "Climate change increases resource-constrained international immobility," *Nature Climate Change*, July 2022, *12* (7), 634–641. Number: 7 Publisher: Nature Publishing Group.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Bertoli, Simone and Jesús Fernández-Huertas Moraga, "Multilateral resistance to migration," *Journal of Development Economics*, May 2013, *102*, 79–100.

Beyer, Robert M., Jacob Schewe, and Guy J. Abel, "Modeling climate migration: dead ends and new avenues," *Frontiers in Climate*, 2023, *5.* Publisher: Frontiers.

Bilal, Adrien and Esteban Rossi-Hansberg, "Anticipating climate change across the united states," Technical Report, National Bureau of Economic Research 2023.

Black, Richard, W. Neil Adger, Nigel W. Arnell, Stefan Dercon, Andrew Geddes, and David Thomas, "The effect of environmental change on human migration," *Global Environmental Change*, December 2011, *21*, S3–S11.

Broda, Christian and David E Weinstein, "Globalization and the gains from variety," *The Quarterly journal of economics*, 2006, *121* (2), 541–585. Publisher: MIT Press.

Bryan, Gharad and Melanie Morten, "The aggregate productivity effects of internal migration: Evidence from Indonesia," *Journal of Political Economy*, 2019, *127* (5), 2229–2268. Publisher: The University of Chicago Press Chicago, IL.

Burke, Marshall, Solomon Hsiang, and Edward Miguel, "Global non-linear effect of temperature on economic production," *Nature*, November 2015, *527* (7577), 235–239.

Burzyński, Michał, Christoph Deuster, Frédéric Docquier, and Jaime de Melo, "Climate Change, Inequality, and Human Migration," *Journal of the European Economic Association*, June 2022, *20* (3), 1145–1197.

Cai, Ruohong, Shuaizhang Feng, Michael Oppenheimer, and Mariola Pytlikova, "Climate variability and international migration: The importance of the agricultural linkage," *Journal of Environmental Economics and Management*, September 2016, *79*, 135–151.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro, "Trade and labor market dynamics: General equilibrium analysis of the china trade shock," *Econometrica : journal of the Econometric Society*, 2019, *87* (3), 741–835. Publisher: Wiley Online Library.

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*, April 2022, p. qjac020.

Cattaneo, Cristina, 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*, July 2019, *13* (2), 189–206. Publisher: The University of Chicago Press.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *The Quarterly Journal of Economics*, 2024, *139* (2), 891–936. Publisher: President and Fellows of Harvard College.

Choquette-Levy, Nicolas, Matthias Wildemeersch, Michael Oppenheimer, and Simon A. Levin, “Risk transfer policies and climate-induced immobility among smallholder farmers,” *Nature Climate Change*, December 2021, *11* (12), 1046–1054. Number: 12.

Clement, Viviane, Kanta Kumari Rigaud, Alex de Sherbinin, Bryan Jones, Susana Adamo, Jacob Schewe, Nian Sadiq, and Elham Shabahat, “Groundswell Part 2: Acting on Internal Climate Migration,” Technical Report, World Bank, Washington, DC September 2021. Accepted: 2021-09-09T16:29:51Z.

Climate Impact Lab, “Impact Map: United States Climate Projections,” `https://impactlab.org/map` 2023. State-level CMIP6 projections; data updated 10 November 2022.

Conte, Bruno, “Climate change and migration: the case of africa,” 2022. Publisher: CESifo working paper.

Cruz, José-Luis, “Global warming and labor market reallocation,” *Available at SSRN 4946752*, 2024.

— **and Esteban Rossi-Hansberg**, “The Economic Geography of Global Warming,” *Review of Economic Studies*, 2024, *91* (2), 899–939. Publisher: Oxford University Press US.

Desmet, Klaus and Esteban Rossi-Hansberg, “Climate change economics over time and space,” *Annual Review of Economics*, 2024, *16.* Publisher: Annual Reviews.

— **, 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*, April 2021, *13* (2), 444–486.

DeWaard, Jack, Mathew Hauer, Elizabeth Fussell, Katherine J. Curtis, Stephan D. Whitaker, Kathryn McConnell, Kobie Price, David Egan-Robertson, Michael Soto, and Catalina Anampa Castro, “User Beware: Concerning Findings from the Post 2011–2012 U.S. Internal Revenue Service Migration Data,” *Population Research and Policy Review*, April 2022, *41* (2), 437–448.

Dingel, Jonathan I and Felix Tintelnot, “Spatial economics for granular settings,” *Econometrica*, 2026, *94* (2), 407–464.

Dobson, Jerome E., Edward A. Bright, Phillip R. Coleman, Richard C. Durfee, and Brian A. Worley, “LandScan: A global population database for estimating populations at risk,” *Photogrammetric Engineering and Remote Sensing*, 2000, *66* (7), 849–857.

Eaton, Jonathan and Samuel Kortum, “Technology, geography, and trade,” *Econometrica : journal of the Econometric Society*, 2002, *70* (5), 1741–1779. Publisher: Wiley Online Library.

Entwisle, Barbara, Nathalie E. Williams, Ashton M. Verdery, Ronald R. Rindfuss, Stephen J. Walsh, George P. Malanson, Peter J. Mucha, Brian G. Frizzelle, Philip M. McDaniel, Xiaozheng Yao, Benjamin W. Heumann, Pramote Prasartkul, Yothin Sawangdee, and Aree Jampaklay, "Climate shocks and migration: an agent-based modeling approach," *Population and Environment*, September 2016, *38* (1), 47–71. Number: 1.

Eyring, Veronika, Sandrine Bony, Gerald A. Meehl, Catherine A. Senior, Bjorn Stevens, Ronald J. Stouffer, and Karl E. Taylor, "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization," *Geoscientific Model Development*, May 2016, *9* (5), 1937–1958. Publisher: Copernicus GmbH.

Fan, Qin, Karen Fisher-Vanden, and H. Allen Klaiber, "Climate Change, Migration, and Regional Economic Impacts in the United States," *Journal of the Association of Environmental and Resource Economists*, July 2018, *5* (3), 643–671.

FAO, "Migration, agriculture and climate change: Reducing vulnerabilities and enhancing resilience," Technical Report September 2018.

Feng, Shuaizhang, Michael Oppenheimer, and Wolfram Schlenker, "Climate Change, Crop Yields, and Internal Migration in the United States," Technical Report w17734, National Bureau of Economic Research, Cambridge, MA January 2012.

Ford, James D., E. C. H. Keskitalo, Tanya Smith, Tristan Pearce, Lea Berrang-Ford, Frank Duerden, and Barry Smit, "Case study and analogue methodologies in climate change vulnerability research," *WIREs Climate Change*, 2010, *1* (3), 374–392.

Frank, Mark, Mark Price, Emmanuel Saez, and Estelle Sommeiller, "Frank-Sommeiller-Price Series for Top Income Shares by US States since 1917," Technical Notes 201507, World Inequality Lab July 2015.

Frank, Mark W., "Inequality and Growth in the United States: Evidence from a New State-Level Panel of Income Inequality Measures," *Economic Inquiry*, 2009, *47* (1), 55–68.

Groeneveld, Jürgen, Birgit Müller, Carsten M. Buchmann, Gunnar Dressler, Cheng Guo, Niklas Hase, Falk Hoffmann, Florian John, Christian Klassert, Thomas Lauf, Veronika Liebelt, Henning Nolzen, Nadine Pannicke, Jule Schulze, Hanna Weise, and Nina Schwarz, "Theoretical foundations of human decision-making in agent-based land use models – a review," *Environmental Modelling & Software*, 2017, *87*, 39–48.

Hauer, Mathew and James Byars, "IRS county-to-county migration data, 1990–2010," *Demographic Research*, May 2019, *40*, 1153–1166.

Head, Keith and John Ries, "Increasing returns versus national product differentiation as an explanation for the pattern of US–Canada trade," *American Economic Review*, 2001, *91* (4), 858–876.

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 et al., "The ERA5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, 2020, *146* (730), 1999–2049.

Hoffmann, Roman, Anna Dimitrova, Raya Muttarak, Jesus Crespo Cuaresma, and Jonas Peisker, "A meta-analysis of country-level studies on environmental change and migration," *Nature Climate Change*, October 2020, *10* (10), 904–912. Number: 10 Publisher: Nature Publishing Group.

—, Barbora Šedová, and Kira Vinke, "Improving the evidence base: A methodological review of the quantitative climate migration literature," *Global Environmental Change*, November 2021, *71*, 102367.

IPCC, "Summary for Policymakers: Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change," 2022.

Kniveton, Dominic R., Christopher D. Smith, and Richard Black, "Emerging migration flows in a changing climate in dryland Africa," *Nature Climate Change*, June 2012, *2* (6), 444–447. Number: 6.

Kolstad, Charles D. and Frances C. Moore, "Estimating the Economic Impacts of Climate Change Using Weather Observations," *Review of Environmental Economics and Policy*, January 2020, *14* (1), 1–24. Publisher: The University of Chicago Press.

Leduc, Sylvain and Daniel J Wilson, "Climate Change and the Geography of the US Economy," Technical Report, Federal Reserve Bank of San Francisco 2023.

— and Daniel J. Wilson, "Snow Belt to Sun Belt Migration: End of an Era?," *Federal Reserve Bank of San Francisco, Working Paper Series*, July 2024, *2024* (21), 01–23.

McFadden, Daniel, "Conditional logit analysis of qualitative choice behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, Academic Press, 1974, pp. 105–142.

Mohaddes, Kamiar, Ryan N C Ng, M Hashem Pesaran, Mehdi Raissi, and Jui-Chung Yang, "Climate change and economic activity: evidence from US states," *Oxford Open Economics*, February 2023, *2*, odac010.

Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg, "Commuting, migration, and local employment elasticities," *American Economic Review*, 2018, *108* (12), 3855–3890. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Moore, Maya and Dennis Wesselbaum, "Climatic factors as drivers of migration: a review," *Environment, Development and Sustainability*, February 2022.

Muñoz-Sabater, Joaquín, Emanuel Dutra, Anna Agustí-Panareda, Clément Albergel, Gabriele Arduini, Gianpaolo Balsamo, Souhail Boussetta, Margarita Choulga, Shaun Harrigan, Hans Hersbach, Brecht Martens, Diego G. Miralles, María Piles, Nemesio J. Rodríguez-Fernández, Ervin Zsoter, Carlo Buontempo, and Jean-Noël Thépaut, "ERA5-Land: a state-of-the-art global reanalysis dataset for land applications," *Earth System Science Data*, September 2021, *13* (9), 4349–4383. Publisher: Copernicus GmbH.

Nath, Ishan, "Climate Change, the Food Problem, and the Challenge of Adaptation through Sectoral Reallocation," *Journal of Political Economy*, 2025, *133* (6), 1705–1756.

Obolensky, Marguerite, Marco Tabellini, and Charles Taylor, "Homeward Bound: How Migrants Seek Out Familiar Climates," January 2024.

Ou-Salah, Loubna and Lore Van Praag, "Gendered dynamics in Moroccan transnational migrant networks in Tenerife: environmental concerns, engagement, and remittances," *Population and Environment*, January 2025, *47* (1), 4.

Redding, Stephen J. and Esteban Rossi-Hansberg, "Quantitative Spatial Economics," *Annual Review of Economics*, August 2017, *9* (Volume 9, 2017), 21–58.

Rigaud, Kanta Kumari, Alex de Sherbinin, Bryan Jones, Jonas Bergmann, Viviane Clement, Kayly Ober, Jacob Schewe, Susana Adamo, Brent McCusker, Silke Heuser, and Amelia Midgley, "Groundswell: Preparing for Internal Climate Migration," Technical Report, World Bank, Washington, DC March 2018. Accepted: 2018-03-15T20:16:21Z.

Schewel, Kerilyn, Sarah Dickerson, B. Madson, and Gabriela Nagle Alverio, "How well can we predict climate migration? A review of forecasting models," *Frontiers in Climate*, January 2024, *5.* Publisher: Frontiers.

Schwarz, Nina, Gunnar Dressler, Karin Frank, Wander Jager, Marco Janssen, Birgit Müller, Maja Schlüter, Nanda Wijermans, and Jürgen Groeneveld, "Formalising theories of human decision-making for agent-based modelling of social-ecological systems: practical lessons learned and ways forward," *Socio-Environmental Systems Modelling*, 2020, *2.*

Serrato, Juan Carlos Suárez and Owen Zidar, "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms," *American Economic Review*, 2016, *106* (9), 2582–2624. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Simonovska, Ina and Michael E Waugh, "The elasticity of trade: Estimates and evidence," *Journal of international Economics*, 2014, *92* (1), 34–50.

Sinha, Paramita, Martha L. Caulkins, and Maureen L. Cropper, "Household Location Decisions and the Value of Climate Amenities," *Journal of Environmental Economics and Management*, November 2018, *92*, 608–637.

Tebboth, M.G.L., D. Conway, and W.N. Adger, "Mobility endowment and entitlements mediate resilience in rural livelihood systems," *Global Environmental Change*, January 2019, *54*, 172–183.

Thober, Jule, Nina Schwarz, and Kathleen Hermans, "Agent-based modeling of environment-migration linkages: a review," *Ecology and Society*, 2018, *23* (2), 41.

U.S. Census Bureau and Bureau of Transportation Statistics, "2012 Commodity Flow Survey," U.S. Department of Transportation, Washington, DC 2015. Public-use microdata.

Watts, Duncan J. and Steven H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, June 1998, *393* (6684), 440–442. Publisher: Nature Publishing Group.

Wooldridge, Jeffrey M, *Econometric analysis of cross section and panel data*, MIT press, 2010.

Xiao, Tingyin, Michael Oppenheimer, Xiaogang He, and Marina Mastrorillo, "Complex climate and network effects on internal migration in South Africa revealed by a network model," *Population and Environment*, March 2022, *43* (3), 289–318.

Šedová, Barbora, Lucia Čizmaziová, and Athene Cook, "A meta-analysis of climate migration literature," *CEPA Discussion Papers*, March 2021. Number: 29 Publisher: Center for Economic Policy Analysis.