

# Approaches to modelling climate migration

## A systematic intercomparison in the United States

Tom Bearpark<sup>\*†1</sup>, Jordan Rosenthal-Kay<sup>\*2</sup>, Nic Choquette-Levy<sup>3</sup>, Michael Oppenheimer<sup>4</sup>, and Tingyin Xiao<sup>4</sup>

<sup>1</sup>University of Exeter

<sup>2</sup>Federal Reserve Bank of San Francisco

<sup>3</sup>Pennsylvania State University

<sup>4</sup>Princeton University

June 24, 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. 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: cold regions gain population and hot regions lose it, by under one percent of the population over a 25-year horizon. The disagreement that remains is dominated by researchers' implementation choices, which contribute as much variance as the choice of approach itself and roughly twice as much as statistical uncertainty. Our framework provides a transparent, reproducible template, analogous to the intercomparison efforts that transformed climate and crop science, that the field can extend to new regions, hazards, and models.

---

<sup>\*</sup>These authors contributed equally to this work.

<sup>†</sup>Corresponding author: [t.bearpark@exeter.ac.uk](mailto: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, engineers, and demographers use *agent-based* models to simulate individual migration behaviour in response to climate stressors. Statisticians and demographers 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. However, despite their common goal of understanding climate-driven migration, these approaches have largely developed in isolation, making direct comparisons across approaches challenging.

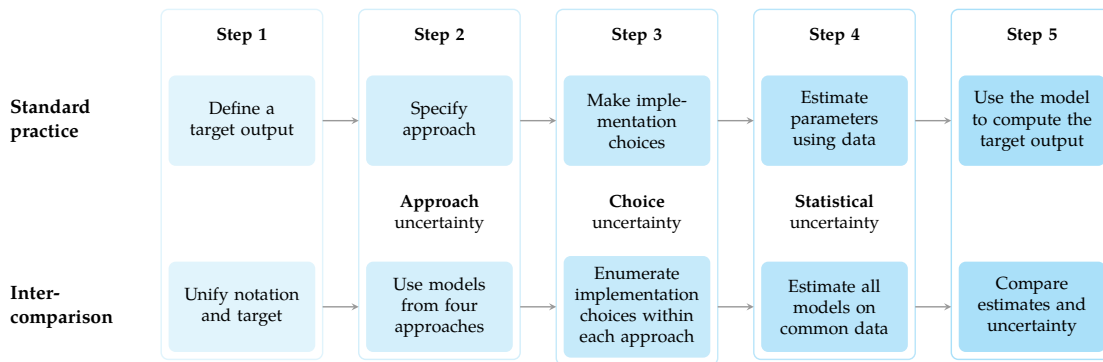
Existing meta-analyses (Cattaneo et al., 2019; Beine and Jeusette, 2021; Hoffmann et al., 2020, 2021; Šedová et al., 2021; Moore and Wesselbaum, 2022; Schewel et al., 2024) document substantial heterogeneity in estimates of climate-induced migration. The studies underlying these reviews differ widely 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 in both modelling approaches, as well as in the study contexts, datasets, definitions of migration and climate exposure, and underlying research objectives. Methodological choices are bundled within disciplinary traditions, making it challenging to disentangle the influence of modelling approach from other differences between studies. 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 mod-

els.<sup>1</sup> We consider four widely used *approaches* to modelling climate migration: causal inference, agent-based, gravity, and general equilibrium.<sup>2</sup> 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, we estimate more than 500 different models and compute our common target quantity thousands of times across two warming scenarios. 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 key part of our contribution is to 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) methods used for estimating model parameters, and (5) the way they present their results.



**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.

<sup>1</sup>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.

<sup>2</sup>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.

## Step 1: A common target quantity

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), and allows us to precisely define a common target parameter: the change in the spatial distribution of population caused by warming, measured at a fixed future horizon (Supplementary Fig. A1). Concretely, it is the 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. Instead, modelling assumptions are required in order to estimate it.

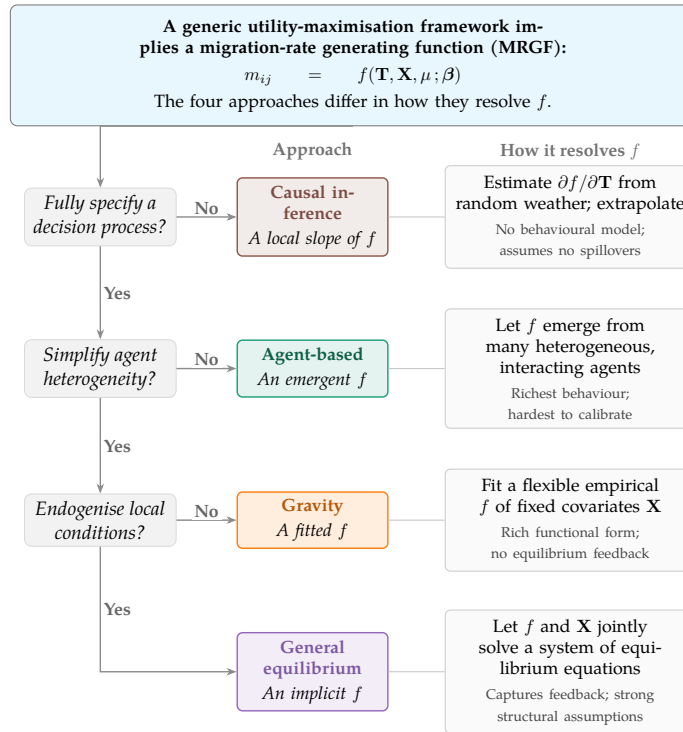
We adopt this quantity for two reasons. First, fixing the time horizon and warming scenario makes it a single, well-defined parameter, and all four approaches can be adapted to estimate it (see details in Fig. 2 and Supplementary Section B.2), so disagreement among them does not merely reflect differences in what is measured. Second, it matters for policy: where people end up shapes demands on housing, infrastructure, and public services.

## 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 bilateral barriers 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*.

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 it requires a sequence of modelling decisions (Fig. 2). These modelling decisions determine which approach is used (Fig. 1, Step 2).

The first decision is whether to model the underlying behavioural processes explicitly at all.



**Figure 2: The four approaches as different ways of resolving one common function.**

A generic utility-maximisation framework implies a *migration-rate generating function* (MRGF),  $m_{ij} = f(\mathbf{T}, \mathbf{X}, \mu; \beta)$ , that maps local conditions to migration rates (Methods). The four modelling traditions estimate this same object (Supplementary Section B.2) but differ in how they resolve  $f$ . Full specifications are given in Supplementary Sections B.5–B.8.

Causal-inference models do not. Instead, they exploit random year-to-year weather variation to identify the local slope of the migration-rate generating rate function (which directly encodes how migration rates respond to a marginal change in temperature) and then extrapolate these estimates to future climates. 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 locations do not affect one another (no spillovers), and that the response to short-run weather is informative about the response to long-run climate change.

Each of the other three approaches models the behavioural process in full. 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, including social-network effects, interactions between movers, path dependence, and the possibility of multiple equilibria. 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.

The remaining two approaches sacrifice some heterogeneity in exchange for tractability, work-

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

Gravity models abstract from such feedbacks. They specify a behavioural model that implies we can treat migration flows as a direct function of observable characteristics of origins and destinations only, and focus on obtaining detailed empirical representations of the migration-rate generating function. This permits the use of rich covariate sets, flexible functional forms, and machine-learning methods. However, because migration is assumed to not alter the environment in which future decisions are made, gravity models are 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, markets endogenously shape wages, housing costs, amenities, and other local conditions, which in turn influence migration decisions. Population distributions therefore 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 assumptions and parameter assumptions that cannot generally be identified from migration data alone. To remain tractable, it also typically restricts the set of factors influencing migration decisions.

### **Step 3: Implementation choices**

We define a model as a unique estimator of a given target parameter, 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 quantity, 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.<sup>3</sup>

#### **Step 4: Estimation**

Any estimate will require calibration of unknown model parameters to 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. Much of the existing literature instead focuses on rapid-onset climate shocks in lower-income, agriculturally dependent settings. 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. Migration flows are derived from IRS Statistics of Income data, which record changes in tax-filing locations and the associated number of tax exemptions, providing 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

---

<sup>3</sup>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, 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 an conservative lower bound on the disagreement a fuller exploration of the model space would reveal.

warming scenarios, one with moderate warming, and one with extreme warming (Supplementary Fig. A2). 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 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.

Because every approach requires unknown parameters to be estimated, each projection inherits statistical uncertainty from the data used to fit it. We quantify this uncertainty across approaches using a parametric bootstrap, re-estimating the target quantity 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 brings these estimates together 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 536 models in total — 128 causal-inference, 120 agent-based, 144 gravity and 144 general-equilibrium specifications (Fig. 6b). The preferred model within each approach additionally carries a 500-draw 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.

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). The *agent-based* model simulates 100,000 individuals who choose where

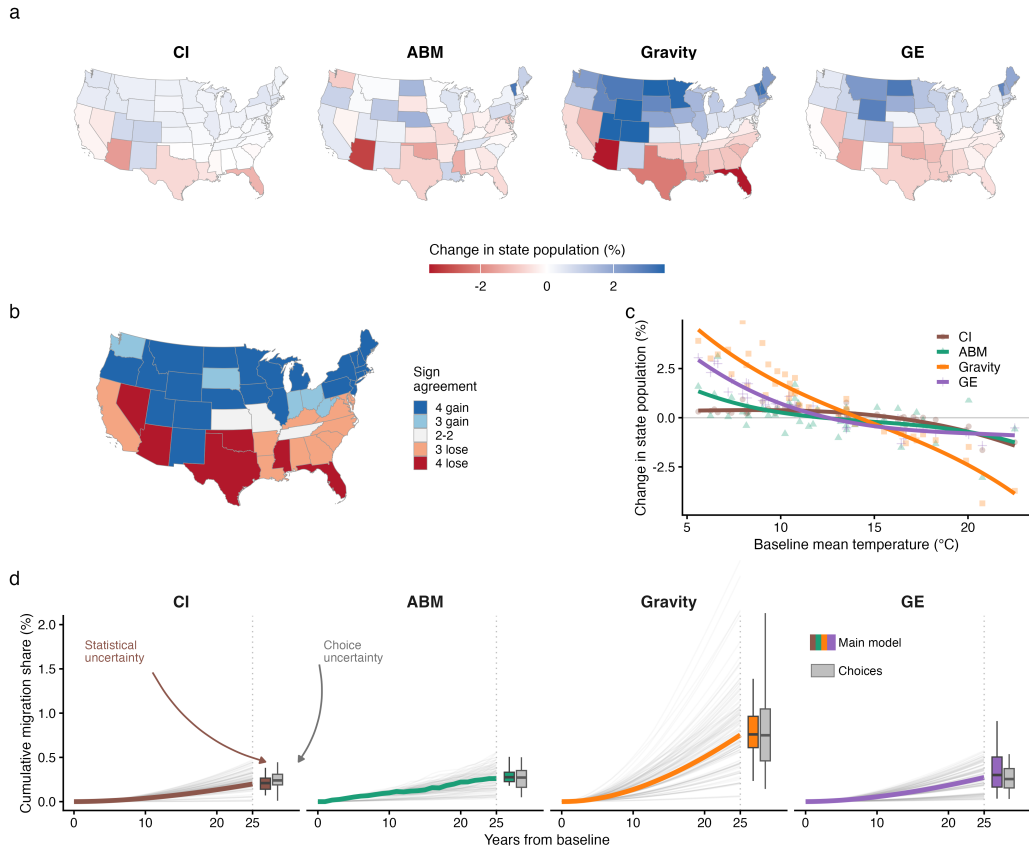
to live, pulled by a destination-temperature amenity and by the locations of their social ties, on a network anchored to Facebook’s Social Connectedness Index that evolves as agents move, connecting our approach to agent-based models that account for social networks when modelling climate migration (Choquette-Levy et al., 2021). In the agent-based model, we iterate over how damages and social preferences are estimated, and how agents’ social networks are constructed (e.g., whether anchored networks are random or small-world, á la Watts and Strogatz (1998)) and evolve over time (Supplementary Section B.6). The *gravity* model is a two-stage bilateral-flow regression: a Poisson (PPML) gravity equation augmented with a machine-learning layer that lets the temperature response interact flexibly with the other covariates (Supplementary Section B.7), similar to gravity models that capture complex interdependencies using frontier statistical tools (e.g., Xiao et al., 2022). Our implementation choices span how temperature, controls, and fixed effects are used in the first layer, and what machine learners are used in the second. The *general-equilibrium* 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 equilibrium. We allow for a variety of forces: interstate trade, agglomeration, local housing markets, and roundabout production, where the strength of these forces, and the 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 *magnitude*. For example, climate change is projected to cause North Dakota to gain between 0.3% (causal inference) and 4% (gravity) of its population, while a hot state, Arizona, is projected to lose between roughly 1.4% and 4%. This redistribution is concentrated. Across the choice grid, the few most-affected states account for the bulk of all population moved (Supplementary Fig. A6).

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: under the gravity approach, for instance, Florida’s projected population loss spans



### 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. A2). 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 (per-state uncertainty in Supplementary Fig. A7). (b) For each state, how many of the four approaches 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 share (the common target, Supplementary Fig. A1) 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.

roughly 1% to 7% once statistical error is accounted for, and widens further across implementation choices (Supplementary Fig. A7), while for some states the combined range runs from substantial loss to no significant change.

To assess the overall net-migration impact of climate change, we aggregate state-level estimates into a single migration share for each approach (Supplementary Fig. A1b), under our 1°C and 2°C warming scenarios (Fig. 3d). We find broad agreement on the rough magnitude: all four approaches project that warming relocates under one percent of the population over the period, with the causal-inference, agent-based and general-equilibrium models clustered near a quarter of a percent and the gravity model roughly threefold higher. The migration share scales approximately linearly with the warming anomaly, so each approach’s moderate-warming projection is close to half its extreme-warming one (Supplementary Fig. A5).

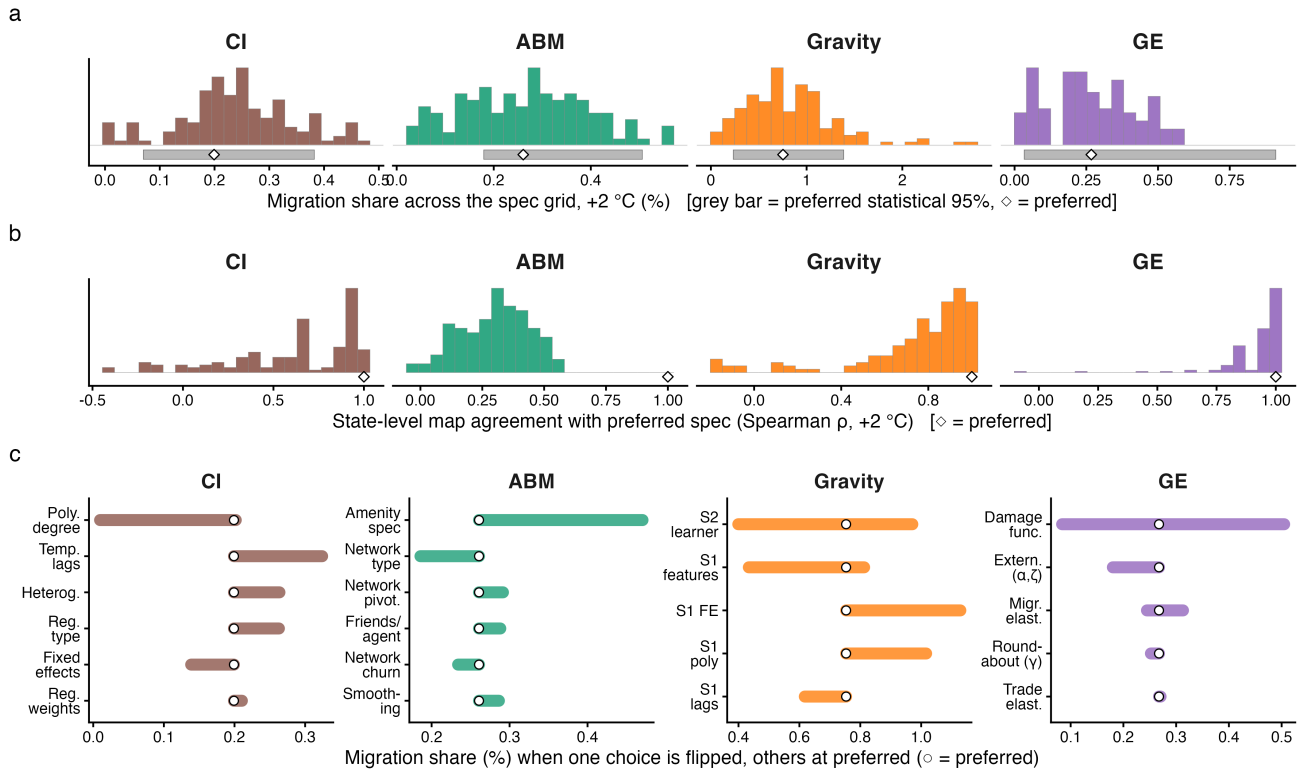
### Explaining variance in the results

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 nearly 3% of the population depending only on which specification is used (Fig. 4a; both warming scenarios are shown in Supplementary Fig. A4). For causal inference, the agent-based model, and gravity, the choice spread is as wide as, or wider than, the statistical 95% band of the preferred model (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 of winners and losers.

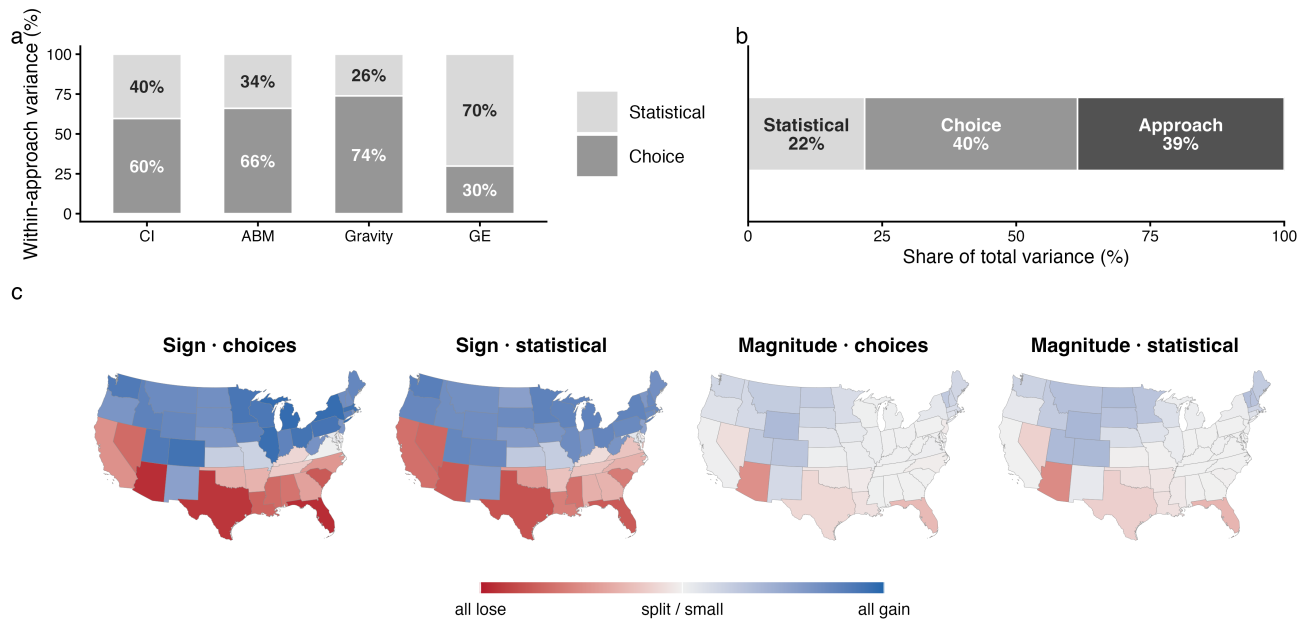
A few implementation choices tend to drive the majority of the variance in projections across choices within each approach (Fig. 4c; the full per-axis distributions are in Supplementary Fig. A3). For causal inference it is the polynomial degree of the temperature response; for the agent-based model, 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 result is driven almost entirely by the choice of damage-function estimator, while the structural elasticities most debated in that literature — the trade and migration elasticities — barely move it.

Decomposing the variance confirms that implementation choice is a first-order source of dis-



**Figure 4: Which implementation choices drive the disagreement.**

All panels use the extreme-warming scenario (Supplementary Fig. A2) and the headline *migration share* — the percentage of the population that relocates because of warming, relative to the no-climate baseline (Supplementary Fig. A1); each approach is 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 ( $\diamond$ ) its point estimate. (b) The same choices also reshape the *geography*: for each specification, the rank correlation (Spearman  $\rho$ ) of its state-level changes with those of the preferred model ( $\diamond$ ,  $\rho = 1$  by construction). Values near 1 mean the map of winners and losers 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 ( $\circ$  = preferred), ordered by range.



**Figure 5: What the approaches agree on, and where their disagreement comes from.**

All panels use the extreme-warming scenario (Supplementary Fig. A2). (a) Within each approach, the share of the variance in the migration share attributable to statistical estimation versus implementation choice. (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. (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 statistical draw.

agreement. Within approaches (Fig. 5a), choice uncertainty dominates statistical uncertainty for causal inference (60% vs 40%), the agent-based model (66% vs 34%) and gravity (74% vs 26%); general equilibrium is the mirror image (30% vs 70%), the only approach for which statistical estimation outweighs implementation choice. Aggregating across all approaches (Fig. 5b), implementation choice and the choice of approach contribute almost equally — 40% and 39% of the total variance — while statistical estimation accounts for just 22%. In other words, the modelling choices a researcher makes *within* an approach explain as much of the overall disagreement as the choice of approach itself, and roughly twice as much as the sampling error that is conventionally the only uncertainty reported.

Beneath this disagreement over magnitudes lies 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; only the agreement on its *magnitude* washes out (Fig. 5c). That is, the spatial pattern of uncertainty generated by implementation choices and statistical error is similar. The geography of where each source of uncertainty dominates is mapped in Supplementary Fig. A7.

## 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).

Firstly, we find that there is strong consensus across models from all approaches that 1°C or 2°C warming alone is unlikely to drive extreme population redistribution in the United States. The low likelihood assigned across all models underscores the importance of considering non-climatic drivers when projecting future migration flows. Second, there is agreement that climate change will lead to a net increase in the population of colder regions and a decline in hotter regions.

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 three — even though such outcomes remain unlikely. Finally, the balance of these two uncertainties is itself approach-dependent. While the general-equilibrium model diverges primarily through statistical estimation, implementation choice is comparable to or larger than statistical uncertainty for the causal-inference

and agent-based approaches, and is the dominant source of divergence for gravity.

## 4 Discussion

We have brought together four families of model that the climate-migration literature has until now deployed in isolation: causal inference, agent-based, gravity, and general equilibrium. By recasting each as a different way of making inferences on a common migration-rate generating function, we made traditions estimate the same quantity 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; 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.

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 can range from 1 to 7% of baseline population, while cold states grow by anywhere between 1 to 4% across approach and implementation choices, and the full uncertainty range admits movements two to three times the central estimate.

More broadly, our framework aims to provide the climate-migration field the infrastructure that reshaped climate and crop science: a common target, a shared dataset, 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.<sup>4</sup> 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.

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.

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, and we exclude extreme cases that would reverse or eliminate the temperature response. The true spread of defensible models is wider than the choice uncertainty we report, which should be read as a lower bound.

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 it is neither the data nor the method, but the 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.

---

<sup>4</sup><https://environmenthalfcentury.princeton.edu/research/2024/researchers-around-world-attend-climate-migration-workshop-princeton-0>

# 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  $i \in \{1, \dots, N_i\}$  and in each period  $t \in \{1, \dots, N_t\}$  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

$$j_{ait}^* = \operatorname{argmax}_j U_{ajt}(\mathbf{T}_{jt}, \mathbf{X}_{jt}, \mu_{ijt}, \epsilon_{ajt}) \quad (1a)$$

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

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

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

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  $\mu_{ijt}$  capture the tangible and intangible costs of moving  $i \rightarrow j$ , and an idiosyncratic preference  $\epsilon_{ajt} \sim g(\cdot)$  captures factors specific to the agent. 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  $\operatorname{argmax}$  (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  $\epsilon_{ajt}$  is random, the probability that an agent moves  $i \rightarrow j$  in period  $t$  can be written as

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

where  $f$  is the *migration-rate generating function* (MRGF),  $\mu := \{\mu_{ijt}\}$  collects migration costs, and  $\beta$  is a vector of unknown parameters. We write  $m_{ijt}(\mathbf{T}) = f(\mathbf{T}, \mathbf{X}, \mu; \beta)$  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  $\mathbf{m}_t(\mathbf{T}) = [m_{ijt}(\mathbf{T})]_{ji}$  and letting  $\mathbf{L}_t(\mathbf{T})$  be the population vector under climate  $\mathbf{T}$  with  $\mathbf{L}_\tau(\mathbf{T}_0) = \mathbf{L}_0$  observed at baseline, the population distribution evolves as  $\mathbf{L}_\tau(\mathbf{T}) = \mathbf{m}_\tau(\mathbf{T}) \mathbf{L}_{\tau-1}(\mathbf{T})$ . Our common *target* is the warming-induced change in the population distribution at a fixed horizon  $\bar{t}$ ,

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

the difference between where people are projected to live under warming and where they would have lived under a baseline climate (see a conceptual illustration in Fig. A1).

## 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 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,  $\approx +1^\circ\text{C}$ , and extreme,  $\approx +2^\circ\text{C}$ , at the horizon; Supplementary Fig. A2 and Section B.3); we project over a 25-year horizon. 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  $M_{ijt} = L_{ijt} / \sum_k L_{kit-1}$ . 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  $\beta$  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

The four approaches share the same endpoints but differ in the stages in between and in where the implementation choices enter. They also differ in what they can separately identify. The reduced-form approaches (CI, gravity) recover only the composite temperature response, whereas the structural models (ABM, GE) separate a location’s underlying climate-relevant value from the mobility that turns it into flows, and in whether there are interactions between the warming effects of each state. Supplementary Section B.4 makes these distinctions precise.

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.

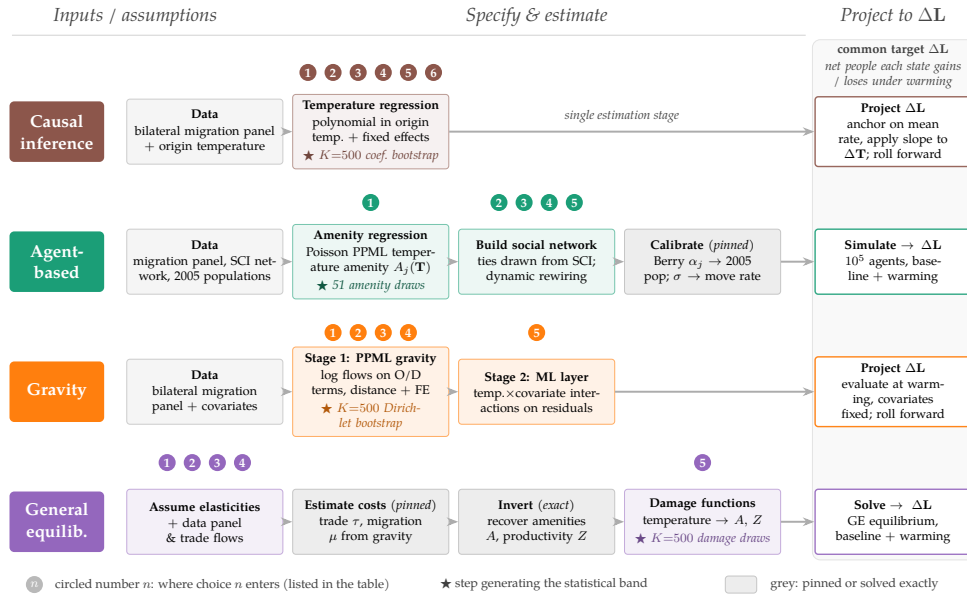
### Causal inference

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

$$\log E(M_{ijt}) = h(\mathbf{T}_{it}; \beta^{CI}) + \gamma \cdot \mathbf{W}_{it}, \quad \forall i \neq j, \quad (4)$$

where  $h$  is a polynomial in (lagged) origin temperature and  $\mathbf{W}_{it}$  are fixed effects and trends that, together with the always-included origin-destination pair effect, render the remaining temperature variation as good as randomly assigned. Identification is local: the recovered  $d \log f / d\mathbf{T}$  combines the direct effect of temperature with adaptation through  $\mathbf{X}$ , and turning it into a projection requires the assumption that one origin’s out-migration is unaffected by other locations’ temperatures (a no-interference, or stable-unit-treatment-value, assumption) and that the in-sample slope applies at the projected climate. We project by anchoring on the in-sample average rate and applying a first-order temperature adjustment,  $\hat{m}_{ijt} = \bar{\mathbf{m}}_{ij} [1 + h(\mathbf{T}_{it}; \hat{\beta}^{CI}) - h(\bar{\mathbf{T}}_i; \hat{\beta}^{CI})]$ , with the diagonal fixed by the column-sum constraint (full form in Section B.5, with the estimated response in Fig. A8). Six choices make up the CI grid (Table A4): regression type (OLS or Poisson), the fixed-effect/trend block, regression weights, the polynomial degree, the temperature lag structure, and origin-destination interactions (choices 1–6 in Fig. 6b).

**a**



**b**

**Causal inference** (= 128 models)

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

**Gravity** (= 144 models)

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

**Agent-based** (= 120 models)

No.	Choice	What it controls	Options ( <i>preferred bold</i> )
1	Amenity specification	how the climate signal (temperature amenity) is estimated	<b>pop.-weighted Poisson amenity</b> unweighted Poisson pop.-weighted + destination trend
2	Network type	structure of social ties that guide moves	<b>small-world (Facebook SCI)</b> random
3	Friends per agent	density of each agent's social network	10 25
4	Within-state smoothing	how strongly friend locations are smoothed	<b>0.01</b> 0.10
5	Network-rewiring regime	whether ties evolve as agents move (amplifies the signal)	static dynamic ( $\sigma_0, \kappa$ ), ( <b>5,1,5</b> ) preferred

**General equilibrium** (= 144 models)

No.	Choice	What it controls	Options ( <i>preferred bold</i> )
1	Migration elasticity ( $\epsilon$ )	how readily people relocate for a given gain	1 3 5
2	Trade elasticity ( $\sigma$ )	substitutability of goods across states	2 5 8
3	Externalities ( $\alpha, \zeta$ )	agglomeration / congestion feedbacks	<b>off</b> (0,0) on (0.3, 0.05)
4	Roundabout share ( $\gamma$ )	input-output linkages that amplify shocks	<b>0</b> 0.5
5	Damage-function spec.	how temperature maps to productivity & amenity damage	<b>quadratic, levels-with-lag</b> 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 (the net number of people each of the 48 contiguous US states gains or loses under warming, where  $L$  is the vector of state populations). Circled numbers mark the step at which each discretionary modelling choice enters;  $\star$  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).

## 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 sharpens the experiment on a single behavioural channel and distinguishes the ABM from the GE model, which also damages productivity. Each agent maximises

$$U_{aijt} = \gamma \log G_{aj,t-1} + A_{jt}(\beta_j, T_{jt}) + \alpha_j - \delta(d_{ij}) + \sigma \epsilon_{aijt}, \quad \epsilon_{aijt} \sim \text{Gumbel}(0, 1), \quad (5)$$

where  $G_{aj,t-1}$  is the number of the agent’s social ties resident in  $j$ ,  $A_{jt}$  a temperature amenity,  $\alpha_j$  a location effect and  $\delta(d_{ij})$  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  $\alpha_j$  are then re-fit by Berry contraction so the model reproduces observed 2005 populations, and the logit scale  $\sigma$  is calibrated to the observed year-1 move rate (Section B.6.4, Fig. A10). Networks are dynamic: ties are re-sampled as agents move, with a rich-get-richer rule. We then simulate the population forward under baseline and warming temperatures. 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).

## Gravity

Our gravity model treats migration flows as a direct function of observable origin and destination characteristics, with no equilibrium feedback, and pairs the standard estimator with a machine-learning layer. A Type-I extreme-value assumption on  $\epsilon_{ajt}$  delivers a multinomial-logit migration probability whose ratio to the stay probability gives the estimating equation

$$\frac{E(M_{ijt})}{E(M_{iit})} = \exp\left(\gamma_2 \mathbf{X}_{jt} + h(\mathbf{T}_{jt}; \beta^{Gr}) - (\gamma_1 \mathbf{X}_{it} + h(\mathbf{T}_{it}; \beta^{Gr})) - D(d_{ij}; \delta)\right), \quad (6)$$

where  $D(d_{ij}; \delta)$  proxies moving costs. We fit this in two stages: Stage 1 is a Poisson pseudo-maximum-likelihood (PPML) regression with origin/destination terms, distance and fixed effects; Stage 2 fits a cross-validated regularised learner on the Stage-1 residuals, allowing the temperature response to interact flexibly with the other features without sacrificing the interpretability of the Stage-1 coefficients. Projection plugs the fitted model into counterfactual temperatures, holding non-climate covariates fixed (full specification in Section B.7, with the pipeline in Fig. A13). Five choices make up the gravity grid (Table A6): the Stage-1 fixed-effect block, the Stage-1 feature set, the 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 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 idiosyncratic shocks; goods are produced from local productivity  $Z_{jt}$  and traded subject to costs  $\tau_{ij}$ ; 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. Five structural parameters are not identified from migration data and must be assumed: the migration elasticity  $\varepsilon$ , the trade elasticity  $\sigma$ , the congestion ( $\alpha$ ) and agglomeration ( $\zeta$ ) externalities, and the roundabout intermediate-input share  $\gamma$ . Conditional on these, we recover migration and trade costs from structural gravity regressions, 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. A14 and the estimated damage functions in Fig. A15). Five choices make up the GE grid (Table A7):  $\varepsilon$ ,  $\sigma$ , the externality regime ( $\alpha, \zeta$ ) (varied jointly, off or on), the roundabout share  $\gamma$ , and the damage-function specification (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 536 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. 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 nested kinds of uncertainty, defined identically across the four so they are comparable (Section B.3). NONE is the preferred model’s point estimate. STATIS-

TICAL uncertainty reflects sampling error in the fitted parameters: we draw repeatedly from the sampling distribution of the *preferred model only* ( $K = 500$  draws for causal inference, gravity and general equilibrium; 51 for the ABM, where simulation cost limits the count) — by 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 — 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.

# Supplementary Materials

## Contents

<b>A</b>	<b>Supplementary Tables and Figures</b>	<b>25</b>
<b>B</b>	<b>Supplementary Methods</b>	<b>34</b>
B.1	Data and forcing construction . . . . .	35
B.2	Linking observed migration rates to migration probability . . . . .	38
B.3	Common structure of the four implementations . . . . .	40
B.4	What each approach identifies . . . . .	41
B.5	Causal inference . . . . .	43
B.6	Agent-based . . . . .	51
B.7	Gravity . . . . .	60
B.8	General equilibrium . . . . .	68

# A Supplementary Tables and Figures

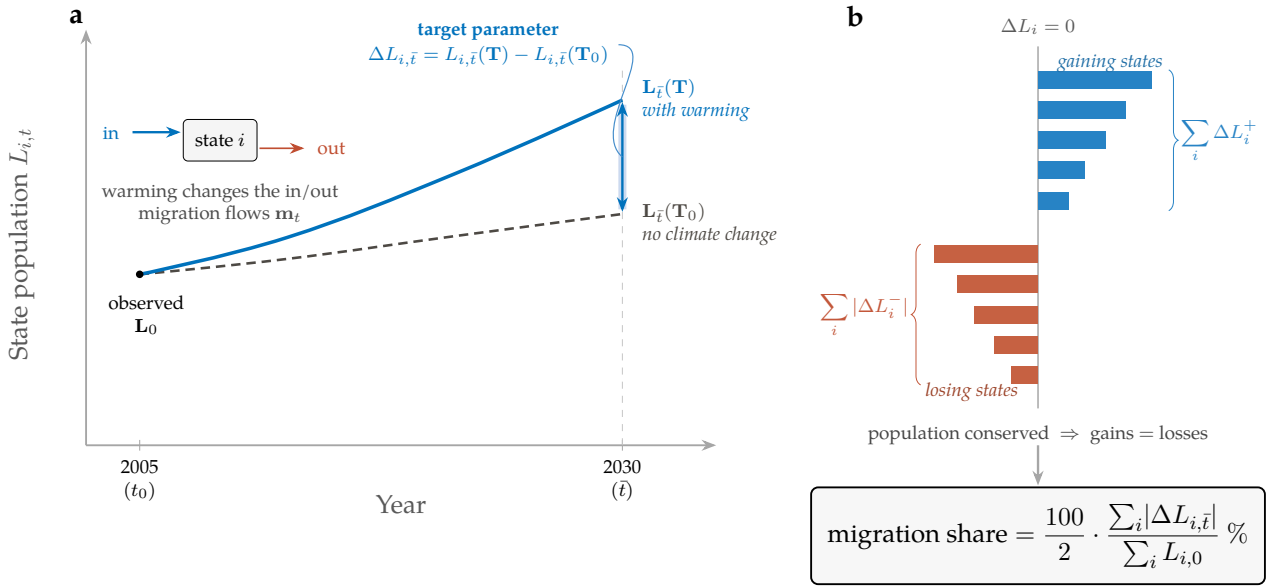
Dataset	Source / provider	Variable(s) and role	Used by			
			CI	ABM	Grav.	GE
IRS Statistics of Income migration	US IRS (via Hauer and Byars, 2019 replication)	Bilateral state-to-state migration flows; tax-exemption counts as the population/agent base (1990–2010)	•	•	•	•
ERA5 / ERA5-Land re-analysis	ECMWF / Copernicus (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 and state population			•	•
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 → bilateral great-circle distance (moving-cost proxy)		•	•	•
Commodity Flow Survey 2012 (PUMF)	US Census Bureau	Interstate trade flows → trade costs $\tau_{ij}$				•
State inequality / income shares	Frank et al. (2015)	Inequality covariates (relative mean deviation, Gini, top-1% share)			•	

**Table A1: Datasets used in the analysis and the approaches that use each.**

	CI	ABM	Gravity	GE	Ensemble
<i>What the models imply about climate-driven migration</i>					
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	High
<i>Where the disagreement comes from</i>					
Implementation choices contribute as much uncertainty as statistical estimation	Medium	Medium	High	Low	Medium

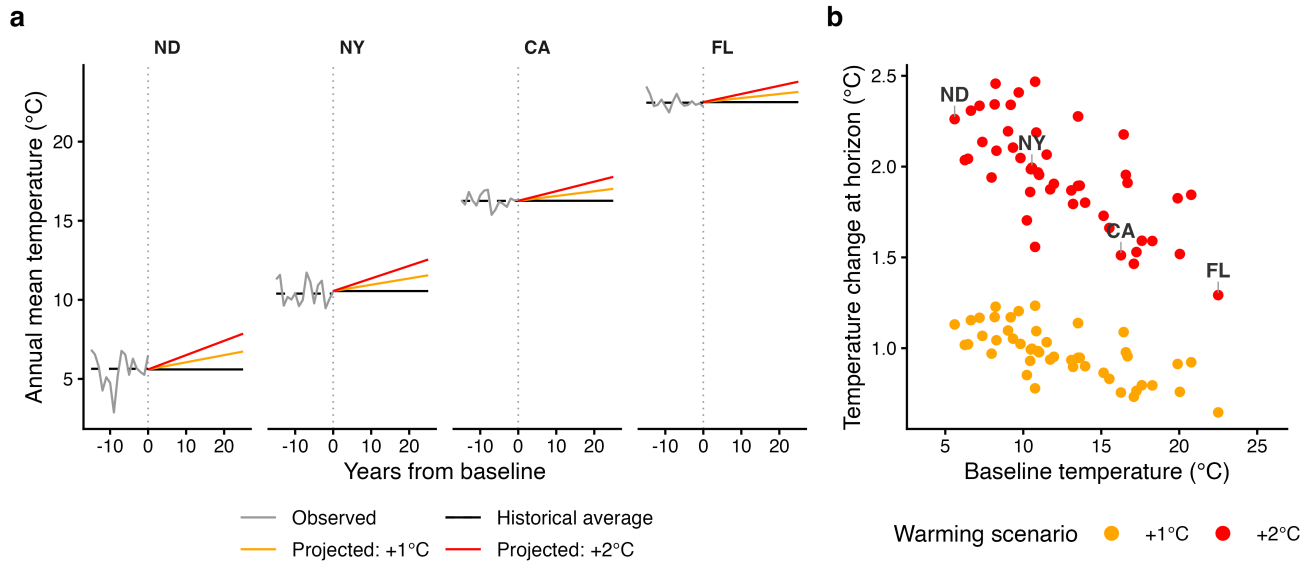
**Table A2: Synthesis: strength of evidence for policy-relevant statements, by approach and pooled.**

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 at the extreme-warming scenario ( $\approx 2^\circ\text{C}$  at the 25-year horizon; Fig. A2); the Ensemble column pools across approaches. 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. Across the full ensemble, implementation choice and the choice of approach each contribute more variance than statistical estimation (Fig. 5b).




**Figure A1: Defining the target parameter and the national migration share.**

A stylised illustration of the quantity every model in the intercomparison estimates and how the state-level estimates are combined. (a) For a single state, two projected population trajectories share the observed baseline  $L_0$  and then diverge: one holds the climate at its baseline ( $L_{\bar{t}}(\mathbf{T}_0)$ , dashed) and one applies a warming scenario ( $L_{\bar{t}}(\mathbf{T})$ , solid). Because warming alters the in- and out-migration flows  $m_t$ , the two trajectories separate over time; the *target parameter* is the gap between them at the horizon  $\bar{t}$ ,  $\Delta L_{i,\bar{t}} = L_{i,\bar{t}}(\mathbf{T}) - L_{i,\bar{t}}(\mathbf{T}_0)$  (Equation 3). The illustration shows a state that gains population; a state that loses population has the opposite sign. (b) The state-level changes  $\Delta L_i$  partition into gaining states (blue) and losing states (red). Because population is conserved, total gains equal total losses, so summing absolute changes counts each move twice. The aggregate *migration share* — the headline reported in Fig. 3d and used throughout Figs. 4 and 5 — is therefore half the summed absolute change, expressed as a percentage of the baseline national population.




**Figure A2: 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  $\times 1$  ( $\approx +1^\circ\text{C}$ ) and  $\times 2$  ( $\approx +2^\circ\text{C}$ ) by the projection horizon, shown against the observed series and the historical baseline mean. (b) Baseline annual mean temperature against the projected temperature change at the horizon for all 48 states, under both warming scenarios.



**Figure A3: Marginal effect of each implementation choice on the migration share.**

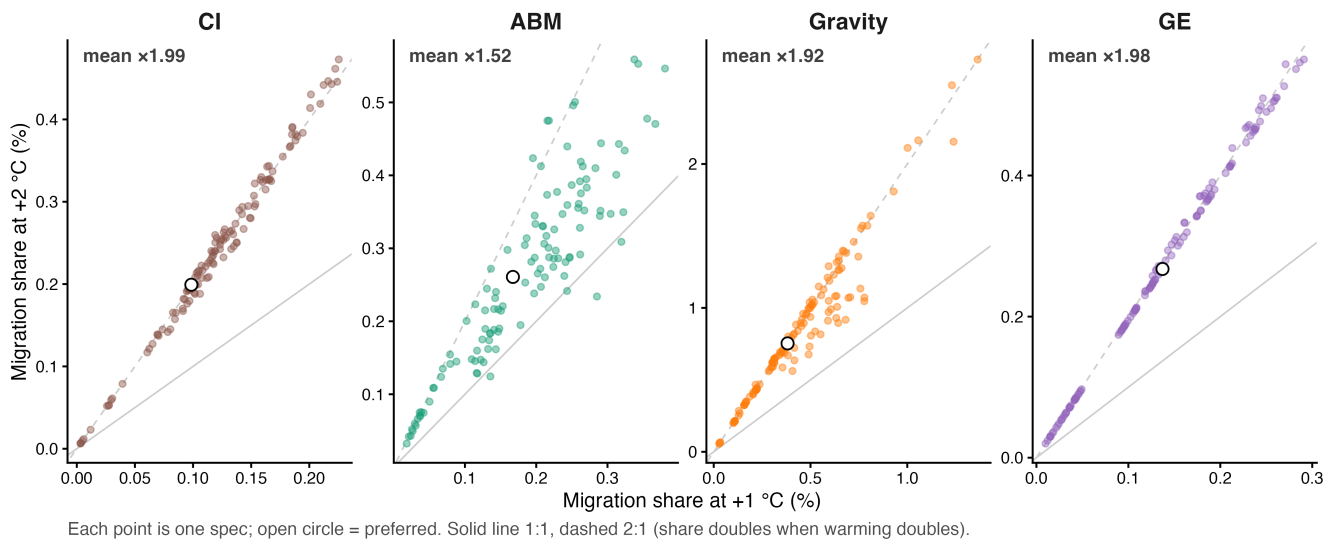
One row per approach (causal inference, agent-based, gravity, general equilibrium), for the extreme-warming scenario ( $\approx +2^\circ\text{C}$  at the 25-year horizon). Within each panel, a box plot shows the distribution of the aggregate migration share across every specification in the choice grid that shares a given level of that axis. Axes whose boxes are well separated are the choices that move the headline most; axes whose boxes overlap are immaterial. This is the full per-axis view underlying the tornado summary in Fig. 4c.



Histogram: migration share across the spec grid (choice uncertainty), height-normalised per panel. Box below: preferred spec's statistical band (box = IQR, whiskers = 2.5–97.5%, line = median); diamond = preferred point estimate. Shared x-axis across approaches.

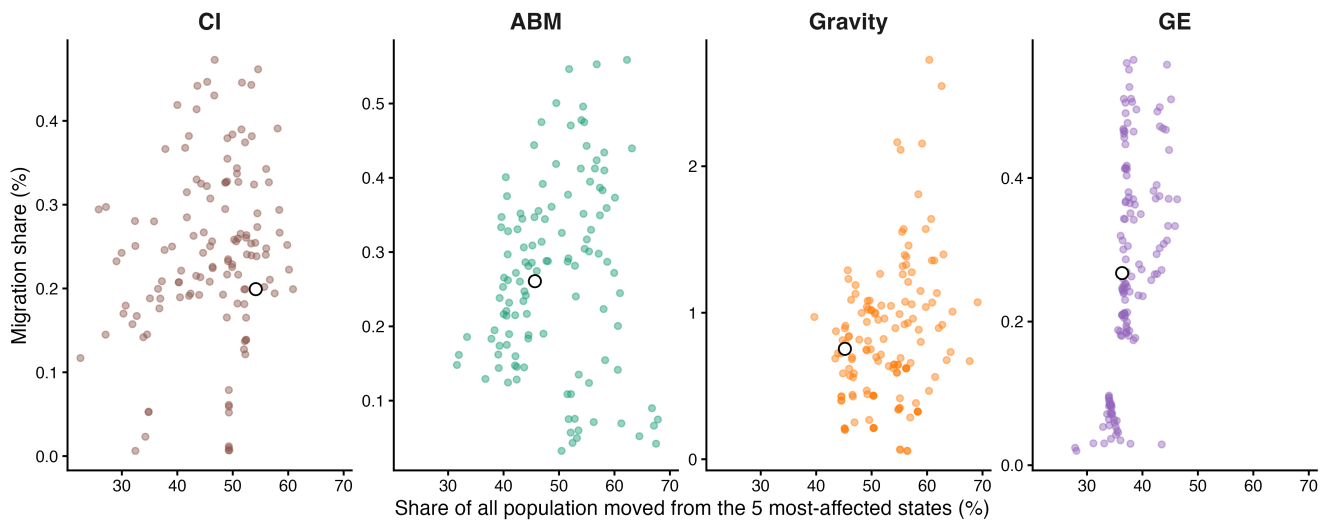
**Figure A4: Distribution of the migration share across the choice grid, with the preferred model's statistical band.**

Rows are approaches; columns are the moderate (+1 °C) and extreme (+2 °C) warming scenarios. In each panel the histogram is the migration share across the full choice grid (choice uncertainty, height-normalised), and the box plot directly beneath is the preferred model's statistical uncertainty (parametric bootstrap: box = interquartile range, whiskers = 2.5–97.5%, line = median; diamond = preferred point estimate). The choice spread is comparable to or wider than the statistical band for every approach except general equilibrium, in both scenarios. This is the two-scenario companion to the extreme-warming panel in Fig. 4a.



**Figure A5: Linearity of the migration response in warming intensity.**

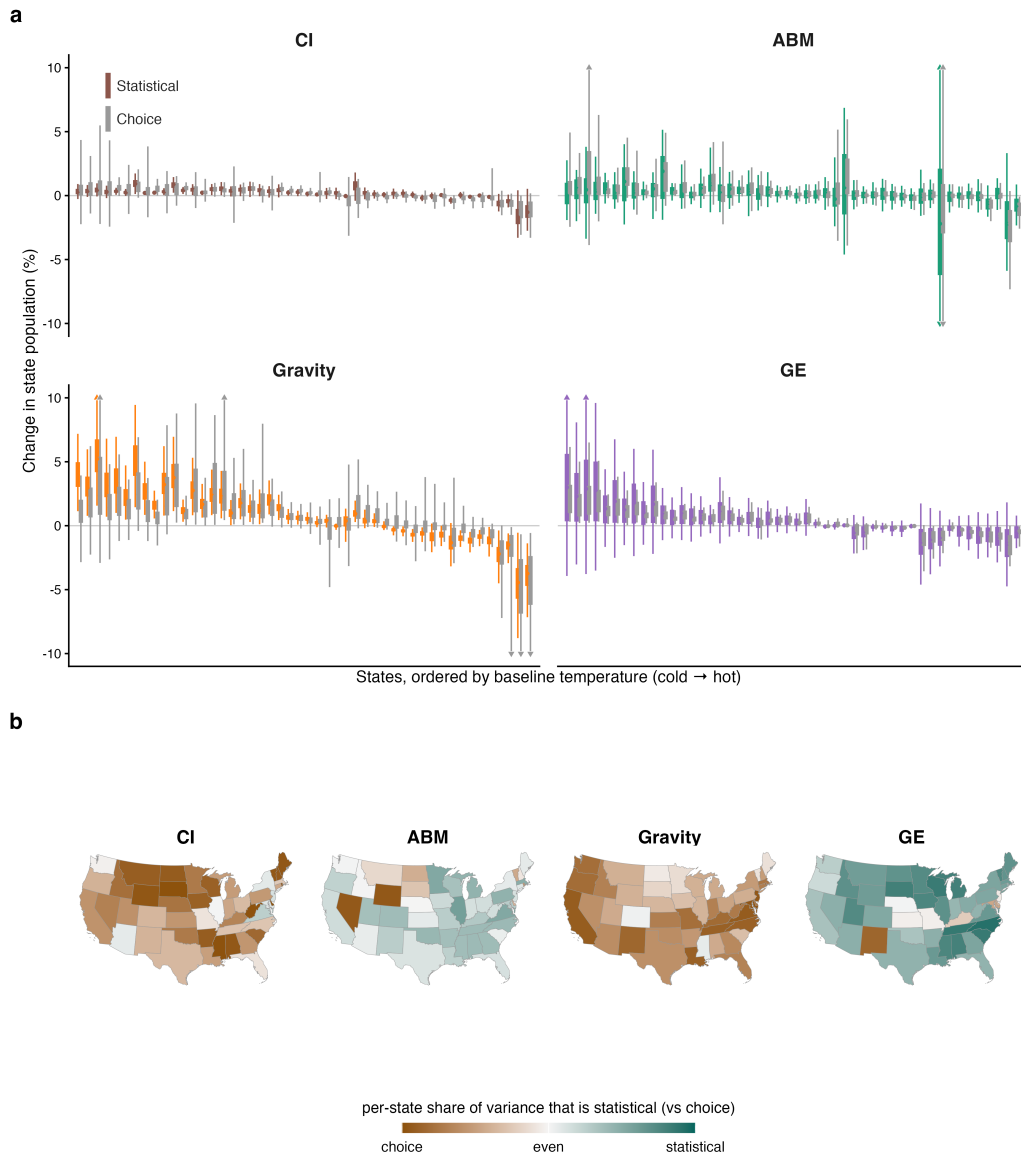
Each point is one specification: its migration share under moderate warming (+1 °C, horizontal) against extreme warming (+2 °C, vertical); open circle = preferred model. The solid line is 1:1 and the dashed line 2:1, so a cloud lying along the 2:1 line means doubling the warming anomaly doubles the migration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



Each point is one spec; open circle = preferred. Concentration uses absolute people moved, so small states with large percentage swings do not dominate. An even split across 48 states would put ~10% in any 5 states.

**Figure A6: 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  $\approx 10\%$  in any five. The redistribution is robustly concentrated in a handful of states across the entire choice grid.



**Figure A7: State-level statistical versus choice uncertainty in the projected population change, by approach.**

Both panels use the extreme-warming scenario ( $\approx +2^\circ\text{C}$  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) from the preferred model's parametric bootstrap, 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  $\pm 10\%$ ; bars extending beyond are flagged with a triangle (chiefly small-population states under the agent-based model, where a swing of a few simulated agents is a large percentage). **(b)** For each approach and state, the share of the per-state variance in the population change that is statistical (preferred-model bootstrap) rather than implementation choice,  $\text{var}_{\text{stat}} / (\text{var}_{\text{stat}} + \text{var}_{\text{choice}})$ .

## B Supplementary Methods

This Supplementary Section sets out the full details of 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](#) sets out the reporting structure common to 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 implementations respectively.

## B.1 Data and forcing construction

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

**Migration.** Bilateral migration flows are the IRS county-to-county migration series, measured in exemptions (a person-equivalent), collapsed to the 48 states. The flow  $L_{ijt}$  from origin  $i$  to destination  $j$  in year  $t$  is the summed exemptions over the corresponding county pairs; the state  $\times$  state  $\times$  year grid is balanced with zeros for unobserved pairs. 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 that maps to no state; this unattributable mass is dropped when the county data are collapsed to states, rather than reassigned to the observed cells. 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). Self-flows ( $i = j$ ) are retained, since they supply the “stayer” counts used by the gravity and ABM normalisations. The usable span is 1991–2010 (a rate in year  $t$  requires the  $t-1 \rightarrow t$  flow). These flows are the county-to-county compilation of Hauer and Byars (2019).

**Temperature (in sample).** The key treatment variable is each state’s population-weighted annual mean near-surface (2 m) air temperature, in °C. To calculate this, gridded daily temperature from reanalysis is weighted by gridded population to a daily state series and averaged within the calendar year. The polynomial specifications use powers  $T_{it}, T_{it}^2, \dots$  of this annual mean, and one- and two-year lags are constructed for the lag axes. The reanalysis product is ERA5-Land (Muñoz-Sabater et al., 2021). The gridded population used to form the weights is from LandScan Global 2000 (Dobson et al., 2000).

**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 path for each state and convert it to °C. The total warming signal (the projected mid-century anomaly relative to the historical reference period) is spread linearly over the horizon, giving a per-state annual increment that attains the projected level after 25 years. 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 — supplies both the population weights and the baseline population vector  $L_0$  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  $48 \times 48$  state grid. Each bilateral quantity is one variable shown over its two subsets — the self-pairs ( $i = j$ , the diagonal: stayers for migration, within-state links for the SCI, intra-state trade, and zero distance) and the off-diagonal pairs ( $i \neq j$ ). The panel covers 1991–2010, and the self-flow migration rate can exceed 100% because its denominator is the prior-year 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, $M_{ij}$ (%)	$i=j$ , $48 \times 20$ yr	97.429	1.391	87.363	104.161	960
	$i \neq j$ , $48 \times 47 \times 20$ yr	0.029	0.075	0.000	2.456	45,120
Bilateral migration flow (exemptions)	$i=j$ , $48 \times 20$ yr	4,460,810	4,680,796	374,212	28,681,489	960
	$i \neq j$ , $48 \times 47 \times 20$ yr	1,291	3,612	0	84,614	45,120
Bilateral distance (miles)	$i=j$ , 48 pairs	0	0	0	0	48
	$i \neq j$ , $48 \times 47$ 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
	$i \neq j$ , $48 \times 47$ 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
	$i \neq j$ , $48 \times 47$ pairs	20.2	49.1	0.0	1,043.4	2,256
Origin temperature, $T_{it}$ (°C)	48 states $\times$ 20 yr	12.4	4.2	2.9	23.0	960
Origin annual total precipitation	48 states $\times$ 20 yr	0.49	0.16	0.06	0.90	960
Real GSP per capita (USD)	48 states $\times$ 20 yr	41,282	8,677	23,269	71,155	960
Origin population (IRS filers)	48 states $\times$ 20 yr	4,548,116	4,720,806	393,084	28,832,510	960
Income inequality (rel. mean dev.)	48 states $\times$ 20 yr	0.814	0.047	0.721	0.978	960
Projected warming, WS $\times$ 1 (°C)	48 states	0.98	0.14	0.65	1.23	48
Projected warming, WS $\times$ 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  $\epsilon_{aijt}$  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:

$$m_{aijt} = P(U_{aijt} \geq \max_k U_{aikt}). \quad (7)$$

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,

$$m_{ijt} := E(m_{aijt}) = \frac{1}{N_{\mathbf{a}_{it}^*}} \sum_{\mathbf{a}_{it}^*} m_{aijt} \quad (8)$$

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

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

Then, the observed migration rate is,

$$M_{ijt} = \frac{L_{ijt}}{\sum_k L_{kit-1}} \quad (9)$$

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  $\{m_{aijt}\}$ , leaving only the agents' period- $t$  destination draws random. Then

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

Averaging over realised populations gives the unconditional  $E(M_{ijt}) = E(m_{ijt})$ . When the resi-

dents 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; the three uncertainty types defined below follow Section 4.

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

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

where the warming-scaling factor is  $\text{WS} = 1$  (the native SSP3-7.0 anomaly,  $\approx +1^\circ\text{C}$  at the horizon) or  $\text{WS} = 2$  (twice the anomaly,  $\approx +2^\circ\text{C}$ ); the no-warming counterfactual holds temperature at its baseline ( $\text{WS} = 0$ ). The projection runs 25 years forward from a baseline cross-section; throughout we index time by years since that baseline,  $t = 0$  (baseline) to  $t = 25$  (the horizon). The baseline is calibrated to a historical reference period — the 2005 cross-section in our data — but nothing in the method ties it to that calendar year, and the horizon is simply the baseline plus 25 years, not the calendar year 2030. In scenario terms,  $\text{WS}\times 1$  is the native SSP3-7.0 median projection at the horizon (a US-average warming of  $\approx +1^\circ\text{C}$  from baseline, at the scenario’s  $\approx 0.4^\circ\text{C}/\text{decade}$  rate), whereas  $\text{WS}\times 2$  ( $\approx +2^\circ\text{C}$ ) is the warming the same projection reaches around mid-century (2040–2059).  $\text{WS}\times 2$  is therefore best read as the migration response to a larger warming *amount* — a plausible mid-century level — rather than as a faster-than-scenario forecast for the horizon year. The same forcing enters every approach, as origin and/or destination temperature as the model requires. Throughout,  $m_{ij}$  denotes the probability that a resident of  $i$  relocates to  $j$  (so  $\sum_j m_{ij} = 1$ , the diagonal  $m_{ii}$  being the stay probability) and  $M_{ij}$  its observed counterpart.

Each approach designates a single *preferred model* — the specification we judge most defensible — 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).

Notation is local to each appendix unless stated otherwise: symbols such as  $\sigma$  and  $T$  carry approach-specific meanings, defined where they appear.

## B.4 What each approach identifies

Although the four approaches estimate  $f$  differently, climate enters the migration-rate generating function (2) through a single scalar per location, which clarifies what each can and cannot separate. Let  $v_j^c(\mathbf{T})$  denote the *climate-relevant value* of location  $j$  — the part of utility that warming moves, common across agents (the superscript  $c$  distinguishes it from the real wage  $v_{jt}$  of the GE inversion, Section B.8). The structural approaches make it explicit: the ABM carries the temperature amenity only,  $v_j^c = A_j$ , while the GE model adds real income,  $v_j^c = \log A_j + \log(w_j/P_j)$ . Because climate acts *only* through  $v^c$ , the rate is a composition  $m_{ij}(\mathbf{T}) = f(v^c(\mathbf{T}), \mathbf{X}, \mu; \beta)$ , and, holding the other locations' values fixed, its own-temperature gradient factors by the chain rule into a behavioural and a preference piece,

$$\frac{\partial m_{ij}}{\partial T_j} = \underbrace{\frac{\partial f}{\partial v_j^c}}_{\text{mobility}} \underbrace{\frac{\partial v_j^c}{\partial T_j}}_{\text{marginal value}}, \quad (10)$$

where the mobility term  $\partial f/\partial v_j^c$  — how strongly flows respond to a unit of value — is governed by the dispersion of the shocks  $g$ . 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 (10) 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. Taking  $f$  to be increasing in value (so  $\partial f/\partial v_j^c > 0$ ), a useful consequence follows: the *location* of any peak is robust, since  $\partial m_{ij}/\partial T_j = 0 \Leftrightarrow \partial v_j^c/\partial T_j = 0$ , so 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 elasticity that the GE model supplies by assumption. This own, ceteris-paribus effect is also distinct from the response that warming realises. Equation (10) holds the rest of the field  $\mathbf{T}_{-j}$  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 re-solving over the whole field. Warming shifts that field together: with the warming-scaling factor WS of

the shared forcing above (so  $\mathbf{T} = \mathbf{T}(\text{WS})$ ,  $\text{WS} = 0$  at baseline and  $\dot{T}_j \equiv dT_j^{\text{proj}}/d\text{WS} = T_j^{\text{raw}} - \bar{T}_j$ ), the realised response is the total derivative

$$\frac{dm_{ij}}{d\text{WS}} = \underbrace{\frac{\partial m_{ij}}{\partial T_j} \dot{T}_j}_{\text{own (reported)}} + \underbrace{\sum_{k \neq j} \frac{\partial m_{ij}}{\partial T_k} \dot{T}_k}_{\text{rivals warm too}}. \quad (11)$$

Because 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 (10) 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-cell 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-cell 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

$$\log E(M_{ijt}) = h(\mathbf{T}_{it}; \boldsymbol{\beta}^{CI}) + \gamma \cdot \mathbf{W}_{it}, \quad \forall i \neq j, \quad (12)$$

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 recovers that derivative, write the migration rate as  $m_{ijt} = f(\mathbf{T}_{it})$  (suppressing the non-temperature arguments) and take a first-order Taylor expansion of  $\log f$  in temperature about the in-sample origin mean  $\bar{\mathbf{T}}_i$ ,

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

where the omitted higher-order terms are carried by the polynomial  $h$ . The leading term  $\log f(\bar{\mathbf{T}}_i)$  is origin-specific and time-invariant, so it is absorbed by the fixed-effect and trend block  $\mathbf{W}_{it}$ ; the second term is the local semi-elasticity. Provided temperature is conditionally exogenous given  $\mathbf{W}_{it}$  — that is,  $E[\epsilon_{ijt} \mid \mathbf{T}, \mathbf{W}] = 0$  once fixed effects and trends are partialled out, 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 recovers  $d \log f / d\mathbf{T}$  at the in-sample mean.

### B.5.2 Preferred model

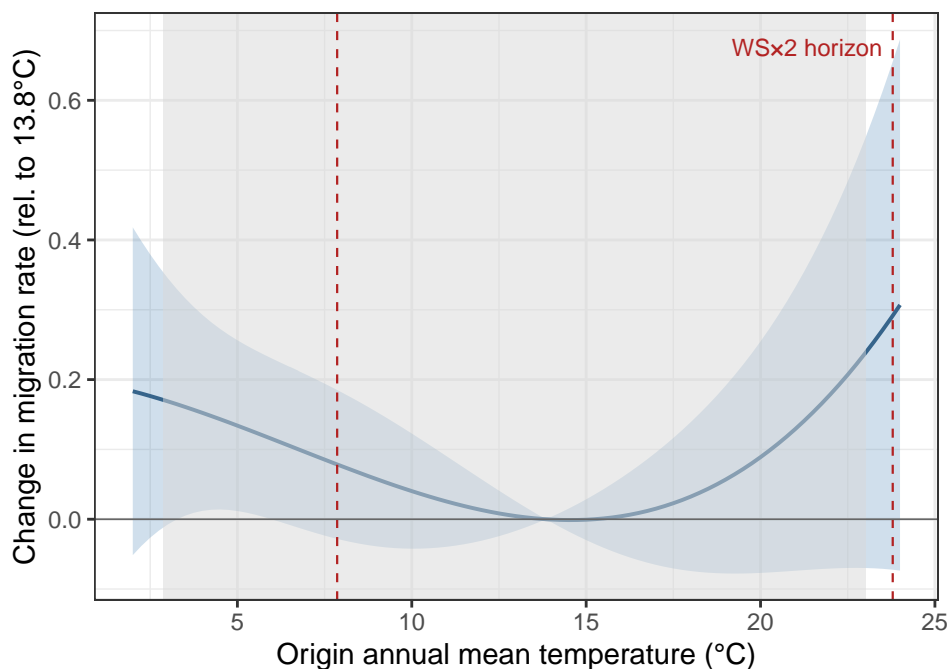
Our preferred CI model is a Poisson PPML regression with population weighting, full bilateral fixed effects, and a cubic polynomial in contemporaneous origin temperature:

$$\log E(M_{ijt}) = \sum_{p=1}^3 \beta_p T_{it}^p + \sum_{q=1}^2 \eta_q R_{it}^q + \alpha_{ij} + \alpha_t + \delta_{1,i} \cdot t + \delta_{2,j} \cdot t, \quad \forall i \neq j, \quad (13)$$

The fixed-effect block  $\mathbf{W}_{it} = \{\alpha_{ij}, \alpha_t, \delta_{1,i} \cdot t, \delta_{2,j} \cdot t\}$  comprises bilateral pair effects  $\alpha_{ij}$  (absorbing all time-invariant features of origin-destination pairs, including distance and historical network density), year effects  $\alpha_t$  (absorbing nationwide trends in migration intensity), and *origin- and destination-specific linear trends* ( $\delta_{1,i} \cdot t$  and  $\delta_{2,j} \cdot t$ ). 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  $\eta_q$ , shown explicitly in Equation 13) is included in every cell as a continuous control but is not the object of interest. The estimator is Poisson pseudo-MLE with population weights equal to the time-average of origin-state tax filers; this gives parameters with the standard semi-elasticity interpretation and accommodates zero migration flows without ad-hoc transformations (Wooldridge, 2010). Self-flows ( $i = j$ ) are dropped from the estimation sample, and observations with missing migration rate are excluded.

The polynomial degree  $P = 3$  is selected 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 the quadratic  $P = 2$ ; the grid below enumerates  $P \in \{1, 2, 3, 4\}$  so the sensitivity is reported explicitly.

A theoretical point governs how this fit is used. The high-dimensional bilateral fixed effects  $\alpha_{ij}$  are incidental parameters: with a panel of bounded length 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 the projection on the observed in-sample mean rate  $\bar{m}_{ij}$  and add only the temperature-driven *deviation* implied by the consistently-estimated slope of  $h$  (Equation 15); the inconsistent fixed effects cancel in this difference. The object we project is thus a *deviation* from an observed baseline, not a predicted level of migration — which is both why the baseline is the observed  $\bar{m}_{ij}$  and why the projection is insulated from the incidental-parameter problem in  $\{\alpha_{ij}\}$ .



**Figure A8:** Preferred-model temperature response function. The estimated response  $h(T; \hat{\beta}^{CI})$  of the preferred cubic-PPML specification, plotted as the proportional change in the migration rate relative to the in-sample mean origin temperature (13.8 °C), with delta-method 95% confidence bands. The grey band marks the in-sample temperature support; the dashed lines bound the projected WS×2 temperatures at the 25-year horizon across the 48 states. The response is U-shaped — migration rises at both cold and hot extremes — and the bands widen sharply where the projection extends beyond the sampled range.

### B.5.3 Choice grid

We enumerate 128 specifications corresponding to the Cartesian product of six axes (Table A4). Four additional choices (treatment variable, outcome variable, cluster level, base year of the projection) are fixed across the grid because they are dictated by the intercomparison framework and the shared dataset. The six varying axes are the ones over which the CI literature differs most heavily.

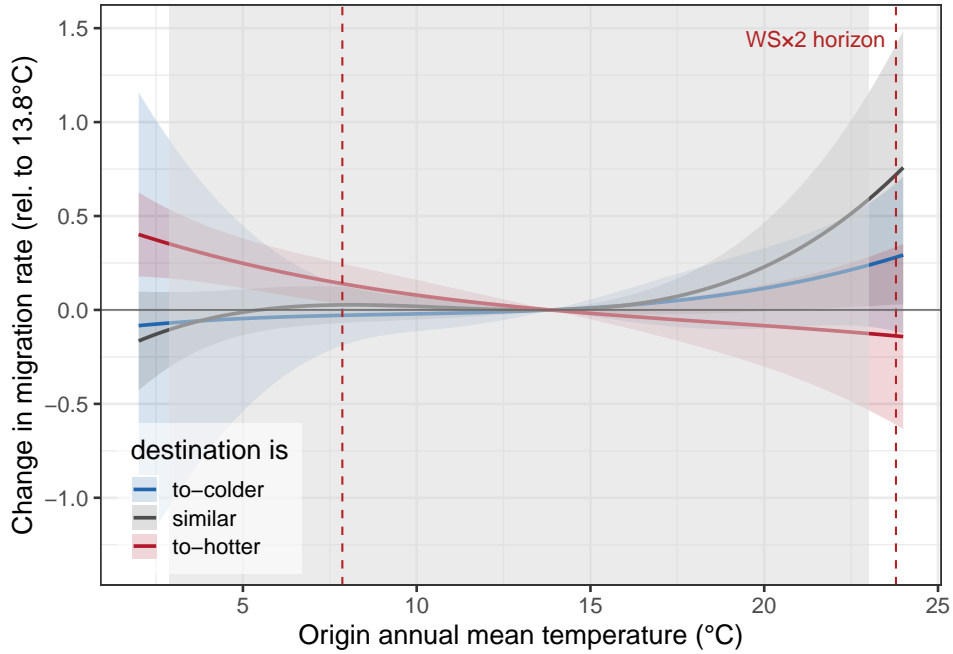
**Table A4:** The 128-cell CI choice grid.

Axis	Values	Justification	·
Estimator	{OLS, PPML}	PPML is the field standard; OLS retained as sensitivity	2
Fixed effects	{full, year-only}	full = $\alpha_{ij} + \alpha_t + \delta_{1i}t + \delta_{2j}t$ ; year-only diagnostic	2
Temp. poly degree	{1, 2, 3, 4}	degree of the polynomial in $T_{it}$ ; preferred 3	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 values were chosen as follows. Polynomial degrees up to  $P = 4$  are conventional in the temperature-response literature (Carleton et al., 2022); we include  $P = 1$  as a sensitivity row even though linear-only specifications are known to underestimate nonlinear responses. The full block always includes the bilateral-pair fixed effect  $\alpha_{ij}$ ; the diagnostic year-only specification drops the state-by-time trend block, 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 weight 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 as a sensitivity row. The OLS cells fit the same right-hand side on the level of the migration rate  $M_{ijt}$  rather than on its log-expectation, so zero flows enter directly as  $M_{ijt} = 0$  with no logarithm or ad-hoc offset (the PPML cells instead accommodate zeros through the Poisson log-link); the OLS cells constitute exactly half of the 128-cell grid and feed the CHOICE band on the same footing as the PPML cells. The lag axis is  $\{0\}$  vs  $\{0, 1\}$ : in pre-tests reported in Leduc and Wilson (2024) and replicated on our panel, the joint significance of leads at  $l \in \{1, 2\}$  is below 5% while contemporaneous and lag-1 effects are jointly significant at well below 1%.

The directional-heterogeneity axis stratifies the polynomial in temperature into three buckets according to the relative climate of origin and destination. Each origin and destination state is assigned its panel-average annual temperature, and every bilateral pair is classified by the destination–origin temperature gap: *to-colder* ( $C_{ij} = 1$ ) if the destination is more than  $2^\circ\text{C}$  cooler than the origin, *similar* ( $C_{ij} = 2$ ) if within  $\pm 2^\circ\text{C}$ , and *to-hotter* ( $C_{ij} = 3$ ) if more than  $2^\circ\text{C}$  warmer. The cutoff  $2^\circ\text{C}$  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 held fixed through the projection: warming rescales the within-bucket temperature response but never reclassifies a pair across the  $2^\circ\text{C}$  boundary. The interacted specification estimates a separate polynomial in  $T_{it}$  for each bucket and is motivated by recent findings that climate-induced movers select destinations with climates similar to their origin (Obolensky et al., 2024; Leduc and Wilson, 2024):

$$\log E(M_{ijt}) = \sum_{c=1}^3 \sum_{p=1}^P \mathbf{1}(C_{ij} = c) \cdot \beta_{c,p} T_{it}^p + \mathbf{W}_{it}. \quad (14)$$



**Figure A9:** Direction-interacted temperature response functions. Bucket-specific responses from the direction-interacted specification (Equation 14), estimated separately for pairs whose destination is colder than (*to-colder*), similar to (*similar*), or hotter than (*to-hotter*) the origin. Plotted relative to 13.8 °C with delta-method 95% bands; shading and dashed lines as in Figure A8.

**Pinned axes.** The following choices are held fixed across all 128 cells. 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  $\bar{m}_{ij}$  (the within-pair time mean of the *observed* rate over 1991–2010; the migration rate in year  $t$  is built from the  $t-1 \rightarrow t$  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 cell is fit by Poisson pseudo-maximum-likelihood (PPML, with iterative absorption of the high-dimensional fixed effects) or, in the OLS cells, by ordinary least squares. Cluster-robust standard errors are computed at the origin-state level via the standard sandwich estimator. For each cell we retain the estimated variance–covariance matrix of the temperature coefficients  $\hat{\beta}^{CI}$

(denoted  $\hat{\Sigma}$ ), which is the input to the statistical-uncertainty bootstrap.

For the STATISTICAL band we draw  $K = 500$  coefficient vectors  $\beta^{(k)} \sim \text{MVN}(\hat{\beta}^{CI}, \hat{\Sigma})$  and propagate each through the projection. This bootstrap is intentionally a *coefficient resample*, not a data refit: it captures sampling uncertainty in the estimated coefficients, conditional on the model specification, and is inexpensive because no re-estimation is required per draw. Following the cross-approach convention, the headline propagates the  $K = 500$  draws through the preferred cell alone, while every cell is evaluated at its coefficient mean to populate the CHOICE band.

### B.5.5 Projection

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

$$\widehat{m_{ijt}^{CI(k)}}(\mathbf{T}) = \bar{m}_{ij} \left[ 1 + h(\mathbf{T}_{it}; \hat{\beta}^{(k)}) - h(\bar{\mathbf{T}}_i; \hat{\beta}^{(k)}) \right], \quad (15)$$

which respects the semi-elasticity interpretation of the PPML coefficients. For OLS cells the update is instead *additive*,  $\widehat{m_{ijt}^{CI(k)}} = \bar{m}_{ij} + [h(\mathbf{T}_{it}; \hat{\beta}^{(k)}) - h(\bar{\mathbf{T}}_i; \hat{\beta}^{(k)})]$ , reflecting that the OLS coefficients are derivatives of the level of  $M_{ijt}$  rather than of its log. Here  $\bar{m}_{ij}$  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  $t \in \{0, \dots, 25\}$  (years since the baseline) at warming scaling factor  $WS \in \{1, 2\}$  (the projection is anchored at the base value and scaled,  $\mathbf{T}_{it}^{\text{proj}} = \bar{\mathbf{T}}_i + (\mathbf{T}_{it}^{\text{raw}} - \bar{\mathbf{T}}_i) \cdot WS$ ). For the preferred cell ( $P = 3$ , no interaction, no lag),  $h(\mathbf{T}; \beta) = \sum_{p=1}^3 \beta_p T^p$ ; for directional-heterogeneity cells the projection applies the bucket-specific coefficients to each bilateral pair according to its  $C_{ij}$  label, and for one-lag cells the temperature difference is summed over the contemporaneous and one-year-lagged terms. The diagonal  $\widehat{m_{iit}^{CI(k)}}$  is recovered by the column-sum constraint  $\sum_j m_{ijt} = 1$  (Equation 1c). Negative predicted rates are clipped at 0 before the constraint is applied (there is no upper clip). 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 (cell, draw) pair. The aggregated target output (migration share at  $WS \times 2$ ) is computed from the resulting  $L_0$  (no-warming) and  $L_1$  (with-warming) trajectories via the

migration-share metric used elsewhere in this paper (Fig. A1). 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, so the denominator time is immaterial; we use the baseline population for definiteness.

### B.5.6 Uncertainty quantification

We report three types of uncertainty for each warming scenario:

**NONE.** The point estimate from the preferred cell at the coefficient mean  $\hat{\beta}^{CI}$ . The other 127 cells 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 cell, capturing sampling uncertainty in  $\hat{\beta}^{CI}$  at the chosen specification.

**CHOICE.** The same four quantiles of the 128 projected migration-share values across the full grid, evaluated at each cell’s coefficient mean. This captures choice (researcher-degree-of-freedom) variance.

### B.5.7 Limitations and robustness

Four caveats deserve emphasis. *First*, the linear extrapolation in Equation 15 assumes that the local derivative  $d \log f / dT$  estimated in-sample applies at the projected WS×2 temperature for every state. Every one of the 48 states has its horizon WS×2 temperature outside its 1991–2010 sampled range — on average  $+0.8^\circ\text{C}$  beyond the hottest year observed in sample, and  $+1.3^\circ\text{C}$  beyond it at most (distinct from the full WS×2 anomaly relative to the average year,  $\approx +2^\circ\text{C}$ , reported in Table A3) — so the projection is mechanically out-of-sample at the state level; only the polynomial *degree* is identified, not its behaviour at the extrapolated temperature. The temperature-polynomial-degree axis is designed to expose this sensitivity: as  $P$  increases, the variance of the projected migration share grows because higher-order polynomials extrapolate more aggressively. This is most acute for the cells that combine the quartic polynomial ( $P = 4$ ) with directional heterogeneity, which are retained in the grid but whose uncertainty bands are unstable and 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 secular component is absorbed by the trends); 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  $\nu = 0.01$ , and dynamic rewiring with churn  $\sigma_0 = 0.5$  and pivotality  $\kappa = 1.5$ . We accompany the preferred ABM with a 120-cell 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-cell 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. Common random numbers are used across choice-grid cells, so that differences between cells reflect their parameters rather than simulation noise. Aggregate statistics are then read off these run-averaged trajectories.

### 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,

$$U_{ajt} = \underbrace{\beta_{\text{sci}} \log \tilde{G}_{aj,t-1}}_{\text{social-network pull}} + \underbrace{A_{jt}(\boldsymbol{\beta}, T_{jt})}_{\text{destination temperature amenity}} + \underbrace{\alpha_j}_{\text{location fixed effect}} - \underbrace{\delta(d_{ij})}_{\text{distance cost}} + \sigma_{EV} \varepsilon_{ajt}, \quad \varepsilon_{ajt} \stackrel{\text{iid}}{\sim} \text{Gumbel}(0, 1). \quad (16)$$

We define each term in turn.

**Social-network pull**,  $\beta_{\text{sci}} \log \tilde{G}_{aj,t-1}$ . 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,

$$\tilde{G}_{aj,t-1} = (1 - \nu) \frac{g_{aj,t-1}}{N_a} + \nu \bar{S}_j, \quad (17)$$

where  $\bar{S}_j$  is the (sum-normalised) average SCI weight to destination  $j$  and  $\nu$  is the smoothing parameter ( $= 0.01$  in the preferred cell). The smoothing floor  $\nu \bar{S}_j$  keeps  $\tilde{G}_{aj,t-1} > 0$  for every destination, so the logarithm is always defined and each state carries a small information weight even when  $a$  has no friends there. The coefficient  $\beta_{\text{sci}}$  is the coefficient on log SCI from the upstream 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  $\tilde{G}_{aj,t-1}$  is specific to the individual agent  $a$ : two agents resident in the same origin state generally have friends in different destinations, so their social-pull terms differ. The network thus enters utility as genuine agent-level heterogeneity with no sufficient statistic at the state level, and there is no representative-agent reduction that preserves it. This is exactly why the ABM must be *simulated agent by agent*, and why (unlike the gravity and GE approaches) it cannot be collapsed into a closed-form bilateral gravity equation or solved as an aggregate equilibrium fixed point: doing so would average away the individual information sets that drive the dynamics.

**Destination temperature amenity**,  $A_{jt}(\beta, T_{jt})$ . Climate enters utility only through this term. The amenity is a quadratic in the destination's own temperature,  $A_{jt} = \beta_1 T_{jt} + \beta_2 T_{jt}^2$ , with coefficients  $\beta = (\beta_1, \beta_2)$  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); under warming each state's  $A_{jt}$  moves along this quadratic as its projected temperature changes.

**Location fixed effect**,  $\alpha_j$ . A time-invariant index of destination attractiveness, absorbing everything about  $j$  — size, wages, housing, unmodelled amenities — not already carried by temperature, network, or distance. The  $\alpha_j$  are *not* read off the regression: they are re-fit by Berry contraction (Section B.6.4) so that the ABM's choice rule reproduces the observed baseline population *shares*. This is critical, because the bilateral-flow PPML returns destination effects de-

fixed *conditional* on its own fixed-effect block, whereas the choice rule needs effects that reproduce population shares — a different object.

**Distance cost,  $\delta(d_{ij})$ .** A bilateral cost of moving, recovered from the distance terms of the gravity regression (Section B.6.4): positive off-diagonal and increasing in distance, and a home-bias bonus on the diagonal. Because it depends only on the origin–destination pair, it carries no agent-level heterogeneity, in contrast to the network term.

**Idiosyncratic taste shock,  $\sigma_{EV}\varepsilon_{aijt}$ .** Each agent–destination pair draws an i.i.d. Gumbel taste shock, which makes the choice a multinomial logit. The scale  $\sigma_{EV}$  sets the weight of idiosyncratic noise relative to deterministic utility, and hence the overall propensity to move; it is calibrated to the observed move rate (Section B.6.4). Each agent relocates to  $\arg \max_j U_{aijt}$ , where the choice set includes staying ( $j = i$ ).

**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 A11a). It is constructed in three stages from the SCI matrix  $G$  (with a fixed random seed shared with the calibration step, so the inverted  $\alpha_j$  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  $\beta_i = 1 - G_{ii} / \sum_j G_{ij}$ , the SCI-implied share of  $i$ 's connectedness that points outside the state; the  $\beta_i$  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  $\beta_i$ , 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  $5 \times 10^5$  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  $\sigma_- = \sigma_0 / (1 + d_a / \bar{d})$  (so high-degree agents are stickier), and gains new friends in its destination state in number  $\sim \text{Poisson}(\sigma_+ \cdot n_a^{\text{lost}})$  with the pivotality multiplier

$\sigma_+ = \max(1 + \kappa(d_a/\bar{d} - 1), 0)$ , where  $\bar{d}$  is the reference (mean) degree. The multiplier  $\sigma_+$  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  $\sigma_0 = 0.5$  (the network “churn” parameter) and  $\kappa = 1.5$  (the “pivotality” parameter). This rewiring regime is the dominant uncertainty axis in the ABM (Section 3); the alternative static regime ( $\sigma_0 = \kappa = 0$ ) sits well below the median of the 120-cell migration-share distribution.

### B.6.3 Choice grid

Our ABM choice grid enumerates 120 cells across the five axes that the literature has identified as the most consequential discretionary modelling choices for migration ABMs (Table A5).

**Table A5:** The 120-cell ABM choice grid.

Axis	Values	Justification	·
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 $\nu$	2
Rewiring regime	5 levels	static + 2 × 2 (churn × 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* ( $\sigma_0 = 0, \kappa = 0$ ); *low-churn / low-pivotality* ( $\sigma_0 = 0.3, \kappa = 0.5$ ); *high-churn / low-pivotality* ( $\sigma_0 = 0.5, \kappa = 0.5$ ); *low-churn / high-pivotality* ( $\sigma_0 = 0.3, \kappa = 1.5$ ); *high-churn / high-pivotality* ( $\sigma_0 = 0.5, \kappa = 1.5$ , the preferred cell). The 2 × 2 design over churn and pivotality is motivated by the distinct roles these parameters play:  $\sigma_0$  scales the magnitude of per-mover friendship turnover,  $\kappa$  governs how aggressively high-degree agents accumulate ties relative to low-degree agents. Holding the two axes separate prevents confounding the magnitude and the heterogeneity of network rewiring. When  $\sigma_0 = 0$  the rewiring machinery is disabled and  $\kappa$  is set to 0, so the static regime is a single cell rather than two.


The two mean-degree values (10 vs 25) represent two assumptions on neighbourhood size. The two smoothing values ( $\nu = 0.01$  vs  $\nu = 0.10$ ) bracket the smoothing strength applied to the

friend share in Equation 17; smoothing also keeps the log-share term defined for agents with no friends in a destination. The network-topology axis contrasts the SCI-anchored small-world topology with 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); the random graph serves as a null in which the local clustering structure of friendships is removed.

**Pinned axes.** Population is 100,000 agents in all cells (a balance between simulation noise and tractability; for cells with  $N_{\text{agents}} = 10^6$  in pilot runs the central tendency was unchanged). The simulation runs over the 25-year horizon, indexed  $t \in \{0, \dots, 25\}$  (years since the baseline). The number of independent shock realisations per cell is  $N_{\text{sim}} = 10$ . The Gumbel scale  $\sigma_{EV}$  is bisection-calibrated to match the empirical year-1 move rate at 1.5%, separately for each (amenity regression, network topology, mean degree, smoothing) combination. The destination-fixed effects  $\alpha_j$  are Berry-contracted to match the baseline population vector, again per combination. We do not enumerate over  $\sigma_{EV}$  or  $\alpha_j$ : both are pinned by external moment conditions.

#### B.6.4 Estimation and calibration

The ABM is calibrated in three steps, all of which sit upstream of the simulation and run *sequentially*; the two moment conditions are matched one at a time, not jointly. Figure A10 summarises the pipeline.



**Figure A10:** The ABM calibration pipeline. A Poisson gravity regression recovers the structural parameters; a Berry contraction inverts the destination effects to match observed baseline population shares; a one-dimensional bisection sets the logit scale to match the observed move rate. Each step targets the moment that identifies it, and none requires searching the joint parameter space.

**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 16) the probability that an agent in  $i$  selects  $j$  is multinomial logit, proportional to  $\exp(V_{aij}/\sigma_{EV})$  in the deterministic utility  $V_{aij}$ . Replacing the agent-specific friend share with its origin-level average — using the aggregate SCI matrix  $\text{SCI}_{ij}$  in place of the realised  $\tilde{G}_{aj}$ , 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

$$\log E(m_{ijt}) = \beta_{\text{sci}} \log \text{SCI}_{ij} + \gamma_1 \log d_{ij} + \gamma_2 (\log d_{ij})^2 + \beta_1 (T_{jt} - T_{it}) + \beta_2 (T_{jt} - T_{it})^2 + \xi_{it} + \alpha_j, \quad i \neq j, \quad (18)$$

with origin-by-year effects  $\xi_{it}$  and destination effects  $\alpha_j$  (an origin-specific self-flow term carries the diagonal). It is estimated by population-weighted Poisson pseudo-maximum-likelihood, clustered on origin. Because  $\xi_{it}$  absorbs the origin component of the temperature gap, the amenity coefficients  $(\beta_1, \beta_2)$  are identified from *destination* temperatures, which is what licenses reading off the destination amenity  $A_{jt} = \beta_1 T_{jt} + \beta_2 T_{jt}^2$  (concave, peaking near 13°C; Figure A11c); the distance terms give the cost  $\delta(d_{ij})$  and the log-SCI coefficient gives  $\beta_{\text{sci}}$ . The simulation then *reintroduces* the agent-level network heterogeneity that this aggregation averaged away.

*Step 1 (gravity calibration)*. The structural coefficients  $(\beta^{Gr}, \beta_{\text{sci}}, \delta, \alpha_j^{Gr})$  are estimated by the population weighted PPML fit of Equation 18 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  $\delta(d_{ij})$ , so  $\delta$  is positive off-diagonal (a cost of moving) and negative on the diagonal (a home-bias bonus).

*Step 2 (Berry contraction)*. Holding the gravity coefficients fixed at their PPML values, the regression destination effect  $\alpha_j^{Gr}$  is iteratively refined to a new vector  $\alpha_j^{ABM}$  such that, under the ABM's choice rule (Equation 16), the simulated baseline population vector reproduces observed shares. The contraction (performed at unit scale,  $\sigma_{EV} = 1$ ) is

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


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

*Step 3 ( $\sigma_{EV}$ -calibration)*. The Gumbel scale is calibrated *separately*, after Step 2, to match the empirical year-1 cross-state move rate of 1.5% (tolerance 0.001). Because the move rate is monotone increasing in  $\sigma_{EV}$ , calibration is a grid-bracket over  $\sigma_{EV} \in \{0.1, 0.3, 0.5, 0.7, 1.0\}$  followed by bisection within the bracketing interval; each candidate  $\sigma_{EV}$  is evaluated by a short three-year ( $t = 0$  to  $t = 2$ ) ABM run with the Berry-inverted  $\alpha_j$ , reading the year-1 move rate of the no-warming baseline. Calibrated scales are  $\sigma_{EV}^* \approx 0.55\text{--}0.87$  across cells (0.74 at the preferred cell). Both the inverted  $\alpha_j$  and the calibrated  $\sigma_{EV}$  are computed once per (amenity regression, network topology, mean degree, smoothing) combination and shared across the five rewiring

values, since the rewiring axis affects neither the baseline population nor the year-1 move rate (calibration is performed on a static, year-1 network).

Within each shock realisation the Gumbel draws are held fixed across choice-grid cells, so that differences between cells reflect their parameters rather than simulation noise (common random numbers across cells).

**Why targeted moment-matching, not a parameter search.** A natural alternative would be to fix the parameters by brute force — Approximate Bayesian Computation, a simulated method of moments, or a grid search that tries many parameter combinations and keeps those whose simulated moments match the data. That is infeasible here, and avoiding it is one of the methodological contributions of this approach. The model must reproduce 48 baseline population shares and the aggregate move rate, but the free parameters that do so — the 48 destination effects  $\alpha_j$  and the scale  $\sigma_{EV}$  — live in a high-dimensional space, and *every* trial parameter vector requires a full stochastic simulation (averaged over  $N_{\text{sim}}$  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 so that each parameter is matched by the cheapest sufficient device: the gravity PPML recovers the bilateral and amenity parameters in a single convex estimation; the Berry contraction is a guaranteed fixed-point map that inverts the 48 destination effects *exactly* — to a maximum share error of  $3 \times 10^{-9}$  in a few hundred softmax evaluations, with no simulation in the loop; and the logit scale is then a single monotone one-dimensional root, found by bisection in roughly eight short simulations. Figure A11b shows the outcome: the calibrated baseline shares sit on the  $45^\circ$  line, whereas the uncalibrated gravity destination effects do not. The whole calibration runs in seconds per cell, rather than the simulation-days a comparable parameter search would demand.



**Figure A11:** 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  $45^\circ$  line; maximum error  $3 \times 10^{-9}$ ), whereas the raw, uncalibrated gravity destination fixed effects do not reproduce them. (c) The preferred-model destination temperature amenity  $A(T) = \beta_1 T + \beta_2 T^2$  is concave, peaking near  $13^\circ\text{C}$ ; the grey band is the in-sample range of state baseline temperatures.

### B.6.5 Projection

For each cell, 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  $t \in \{0, \dots, 25\}$  (years since the baseline) and shock realisation  $s \in \{1, \dots, 10\}$ . 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-cell migration share (Fig. A1), 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-cell 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 cell, 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  $(\beta_1, \beta_2)$ . The other structural parameters —  $\beta_{sci}$ , the distance cost, the destination effects  $\alpha_j$ , and

$\sigma_{EV}$  — 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  $N_{\text{sim}}$  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-(draw, 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 cells 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  $\approx 100$  agents by the 25-year horizon; a swing of a few agents produces  $\pm 10\%$  per-state  $\Delta L$ , which is small- $N$  simulation noise rather than meaningful climate response.

*Thirdly,* and most importantly, our ABM sits at the simple end of the migration agent-based-modelling literature. It has a single agent type, one behavioural channel (a temperature amenity), and a parsimonious friendship network, and it deliberately omits features that richer migration ABMs incorporate: bounded rationality and heuristic or satisficing decision rules in place of exact utility maximisation; agent heterogeneity by age, income, education, or housing tenure; liquidity, credit, and housing-market frictions; learning, adaptive expectations, and imperfect information beyond the network; multiplex or geographically embedded network structures that evolve on more than the single rewiring margin we allow; and endogenous local feedbacks such as congestion, wages, and prices. This minimalism isolates the social-network channel and keeps the ABM comparable to the other three deliberately stylised approaches — but it means the model should be read as a controlled experiment in one mechanism rather than a full behavioural simulation of US migration. 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 as additional axes.

## B.7 Gravity

This appendix gives the full specification of our gravity implementation. The preferred gravity model is a two-stage estimator: a Poisson pseudo-maximum-likelihood (PPML) bilateral flow regression with a destination fixed effect, a data-driven self-flow-by-destination anchor, and per-origin linear time trends (Stage 1), augmented by a cross-validated lasso interaction layer fit on the Stage-1 residuals (Stage 2). We accompany the preferred model with a 144-cell 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-cell grid.

### B.7.1 Approach overview

Our gravity implementation is built around the migration-rate generating function  $f$  in its bilateral form:  $m_{ijt} = E(M_{ijt})$  is modelled as the exponentiated linear-index function  $\exp\{V_{jt} - V_{it} - D(d_{ij}; \delta)\}$ , 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). This bilateral form is *implied by* the location-choice model, under a particular assumption about the structure of individual preferences. Taking the idiosyncratic preference  $\epsilon_{ajt}$  of that model (1d) to be i.i.d. Type-I extreme value (Gumbel) makes the destination-choice probabilities multinomial logit (McFadden, 1974), and 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 linear index above. We therefore work throughout with the *ratio transform*  $m_{ijt}/m_{iit}$  — the bilateral flow relative to the origin’s own stayers — whose logarithm is the linear index  $V_{jt} - V_{it} - D(d_{ij})$ .

We estimate this index in two stages, splitting it into a parametric main-effects part and a residual interaction part:

$$\log \frac{m_{ijt}}{m_{iit}} = \underbrace{V_{jt} - V_{it} - D(d_{ij})}_{\text{Stage 1: PPML main-effects index}} + \underbrace{g(T_{it}, T_{jt}, \mathbf{X}_{ijt}; \boldsymbol{\theta})}_{\text{Stage 2: regularised interaction residual}}, \quad (19)$$

where Stage 1 is the standard bilateral PPML linear index in temperature, controls, and fixed effects, and  $g$  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). Holding  $g$  on the residual scale — a Frisch–Waugh–Lovell construction — lets the temperature–migration relationship interact with the panel’s other features (lagged temperature and gross state product, distance, social connectedness) while preserving the closed-form interpretability of the Stage-1 coefficients. This Stage-2 machine-learning layer, together with an explicit enumeration of 144 specification

choices is what distinguishes our implementation from the standard gravity literature.

### B.7.2 Preferred model

Our preferred specification has the following Stage-1 regression:

$$\log E\left(\frac{M_{ijt}}{M_{iit}}\right) = \sum_{p=1}^2 \sum_{l=0}^2 [\beta_{p,l}^o T_{i,t-l}^p + \beta_{p,l}^d T_{j,t-l}^p] + \gamma^o d_{ij} + \gamma_1^d \text{GSP}_{i,t-1} + \gamma_2^d \text{GSP}_{j,t-1} + \alpha_j + \alpha_j^{\text{self}} \cdot \mathbf{1}(i = j) + \delta_i \cdot t. \quad (20)$$

The regressand is the ratio transform  $m_{ij}/m_{ii}$  defined in the Approach overview, which Stage 1 takes directly as its PPML 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.

**Fixed effects and identification.** The specification controls for a destination fixed effect  $\alpha_j$  and a per-origin linear trend  $\delta_i \cdot t$ . We do not include an origin fixed effect as we do not want to absorb all origin-level variation. 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. The destination fixed effect is retained because it absorbs equilibrium-side variation that is not the target of identification (e.g., destination wages, amenities).

**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  $\alpha_j^{\text{self}}$  absorbs this anchor, supplies the per-state home-bias coefficient, and renders the diagonal *data-driven* rather than imposed. Stage 1 is fit on the full bilateral panel with self-flows included; excluding them materially shifts the diagonal anchor and changes the per-state response curves.

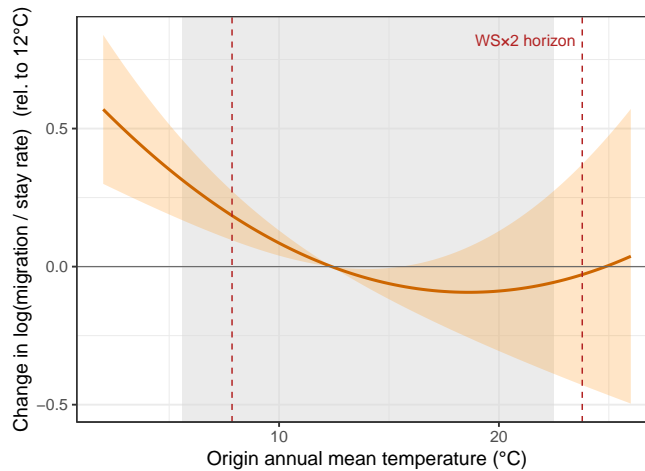
**Stage-2 lasso residual layer.** Stage 2 instantiates the residual term  $g$  of Equation 19 as a cross-validated lasso. It regresses the bilateral count on the Stage-1 prediction (as an offset) plus a

penalised interaction basis,

$$\log E(\text{flow}_{ijt}) = \underbrace{\log \text{stayers}_{it}}_{\text{offset}} + \hat{f}_{ijt}^{\text{S1}} + \sum_{k=1}^{36} \theta_k Z_{ijt}^{(k)}, \quad \hat{\theta} = \arg \min_{\theta} \left\{ \text{dev}_{\text{Pois}}(\theta) + \lambda \|\theta\|_1 \right\}, \quad (21)$$

where  $\hat{f}_{ijt}^{\text{S1}}$  is the fitted Stage-1 linear predictor (Equation 20) 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 ( $d_{ij}, \text{GSP}_{i,t-1}, \text{GSP}_{j,t-1}$ ). The  $\log(\text{stayers}_{it})$  offset carries the ratio transform onto the count scale, and entering  $\hat{f}_{ijt}^{\text{S1}}$  as a fixed offset makes the lasso a *pure residual layer* that fits only the interaction structure the additive Stage-1 main effects leave behind. The penalty  $\lambda$  is selected by cross-validation; the chosen penalty, the number of interactions retained, and the fit improvement are reported in the Estimation subsection below.

**Origin-temperature response.** Figure A12 plots the preferred cell’s origin-temperature response, summing the contemporaneous and lagged origin terms: out-migration relative to staying is elevated for cold origins, flattens through the middle of the sampled range, and its confidence band widens sharply where the WS $\times$ 2 projection extrapolates beyond the in-sample support.



**Figure A12:** Preferred-model origin-temperature response of the gravity model: the change in the log bilateral migration rate relative to staying,  $\log(m_{ij}/m_{ii})$ , as a state’s own temperature shifts permanently (summing the contemporaneous and lag-1/lag-2 origin terms), with delta-method 95% bands. 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 cells from the Cartesian product of five axes (Table A6).

**Table A6:** The 144-cell gravity choice grid.

Axis	Values	Justification	·
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	{1, 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 the central axis. The *full* specification is the standard 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) adds an origin-specific linear trend on top, absorbing demographic drift while preserving the long-run temperature variation that identifies the response (Preferred model). All three retain the self-flow-by-destination anchor and all three preserve the textbook negative correlation  $\text{corr}(T_0, \Delta L_{25}) < 0$  at  $\text{WS} \times 2$ .

The three control sets are: *baseline*, the continuous controls  $\{d_{ij}, \text{GSP}_{i,t-1}, \text{GSP}_{j,t-1}\}$  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.

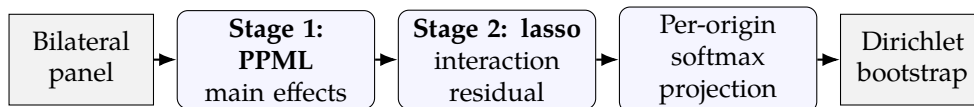
The two polynomial choices bracket the degrees commonly used in the literature: the quadratic retains degrees  $\{1, 2\}$  in temperature, the cubic degrees  $\{1, 2, 3\}$ . The two lag choices bracket the lag structures: the one-lag setting builds temperature terms at lag 0 only, while the two-lag setting builds lags  $\{0, 1, 2\}$ . The two-lag specification is preferred; one lag is included as a sensitivity row.

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 cells are deliberately a sensitivity row rather than a candidate for the headline: Stage 1 extrapolates be-

yond 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 cells without overfitting (the selection counts and deviance reduction are reported in the Estimation subsection).

**Pinned axes.** The treatment is annual average origin temperature. Standard errors are clustered at the origin level (using a Dirichlet bootstrap, see Estimation section below). The projection baseline is the historical reference cross-section ( $t = 0$ , the first year of the projection horizon). All cells 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 cell (Preferred model).

#### B.7.4 Estimation



**Figure A13:** 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 21) are fit once and reused; the four Stage-2 learners are then attached to each, giving 144 cells in total.

The Stage-2 learners are configured as follows. The *lasso* (the preferred Stage 2) uses a Poisson likelihood with an  $\ell_1$  penalty (standardised inputs, no intercept); the penalty is chosen by 5-fold year-blocked cross-validation (folds split by calendar year, preserving within-state serial dependence within each fold) to minimise mean Poisson deviance, with the candidate grid extended well below the unpenalised solution. At the preferred cell the lasso selects 24 of the 36 candidate interactions, with a cross-validated deviance reduction over Stage-1 alone of 3.64% — modest but consistent across cells, indicating genuine non-additivity rather than overfitting. The two non-preferred learners — gradient-boosted Poisson trees (depth 5, with early stopping)

and a five-member ensemble of 64–32–1 multilayer perceptrons (ReLU activations and dropout, trained on the 15 continuous features with early stopping) — both take the Stage-1 prediction as an offset and serve only as sensitivity rows; their full hyperparameter settings are recorded with the code.

In 15 of the 144 cells—mostly cubic-temperature cells with extra controls—the selected lasso penalty sits at the floor of the candidate grid, indicating that a wider penalty path would shrink them further; these cells still enter the CHOICE band, with that caveat. The preferred cell is unaffected.

### B.7.5 Projection

For each cell 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. Within-sample, the implied baseline migration rates exactly equal the observed rates (by construction of the PPML estimator’s first-order condition); under counterfactual warming, only the temperature features change. The per-origin migration matrix is recovered from the projected log-rate index by a per-origin softmax, which simultaneously normalises each origin’s outflows and recovers the self-flow diagonal. The per-state population trajectory is then forward-rolled through the migration matrix:

$$L_{i,t+1}^{(c)} = \sum_j m_{jit}^{(c)} \cdot L_{j,t}^{(c)}, \quad t \in \{0, \dots, 24\}, \quad c \in \{0, 1\},$$

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, so the choice of denominator time is immaterial. 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-cell point estimate. The other 143 cells 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  $w^{(k)} \sim \text{Dirichlet}(\mathbf{1})$ , 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 cell’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-cell grid at each cell’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  $+0.8^\circ\text{C}$  beyond the hottest year observed in sample, and  $+1.3^\circ\text{C}$  beyond it at most (distinct from the full  $WS \times 2$  anomaly relative to the average year,  $\approx +2^\circ\text{C}$ , reported in Table A3) — so every projection cell is mechanically an extrapolation. The temperature-polynomial axis (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.

The Stage-2 neural-network cells deserve a final note: across the 144-cell grid, 36 cells use the neural network as the Stage-2 learner, and they consistently produce wider CHOICE dispersion than the lasso or the gradient-boosted trees. We interpret this as evidence that the network's flexibility leads to ensemble-to-ensemble divergence under extrapolation, not as evidence of richer non-additivity.

## 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 CES goods trade across the 48 contiguous states, evaluated at central migration and trade elasticities ( $\varepsilon = 3, \sigma = 5$ ), and a levels-with-lag quadratic damage function. We accompany the preferred model with a 144-cell grid — 36 structural cells, 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 cell’s damage-function coefficients; and *choice*, from re-projection across the full 144-cell grid.

### 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 climate-damaged state is partly cushioned by 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 reflects general-equilibrium adjustment rather than a partial-equilibrium elasticity.

**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,

$$j_{ait}^* = \operatorname{argmax}_j (V_{jt}/\mu_{ij}) \epsilon_{aijt}, \quad (22)$$

where  $\epsilon_{aijt}$  is Fréchet distributed with location parameter 1 and shape parameter  $\varepsilon$ . 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,

$$V_{jt} = \max_{\{c_k\}} U_{ajt} = \underbrace{A_{jt}}_{\text{destination amenities}} \times \underbrace{\left( \sum_{k=1}^N c_k^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}}_{\text{consumption utility}} \quad \text{subject to} \quad \sum_k p_{kjt} c_k \leq w_{jt}. \quad (23)$$

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 (22).

**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  $\tau_{ij}$ , so that  $\tau_{ij}$  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,

$$A_{it} = \bar{A}_{it} (L_{it})^{-\alpha}, \quad Z_{it} = \bar{Z}_{it} (L_{it})^{\zeta}. \quad (24)$$

Here,  $\alpha$  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.  $\zeta$  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 (23) 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 (24);
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  $\{w_{it}, L_{it}\}$ ,

$$\begin{aligned} L_{jt} &= \sum_i m_{ijt} L_{i,t-1} \\ w_{it} L_{it} &= \sum_j p_{ijt} c_{ijt} L_{jt}, \end{aligned} \tag{25}$$

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,

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

where the term  $P_{jt} = (\sum_i p_{ijt}^{1-\sigma})^{\frac{1}{1-\sigma}}$ , 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 (25). 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,

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

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.** The grid also enumerates a *roundabout* variant of the production block, in which the share  $\gamma$  of gross output is itself used as an intermediate input. Instead of producing with labour alone, each location combines labour and an intermediate bundle (the same CES aggregate of state goods that consumers buy, priced at the local

index  $P_i$ ) via a Cobb–Douglas technology  $Q_i = Z_i \ell_i^{1-\gamma} m_i^\gamma$ . Cost minimisation makes the mill (factory-gate) price

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

so the local price index now enters the cost of production: a state’s goods are an input to every other state’s output, and a local damage propagates through the input–output network. Because labour is the only primary factor, GDP per capita still equals the wage, and goods market clearing written in mill prices —  $w_i L_i = p_i^{1-\sigma} \sum_j \tau_{ij}^{1-\sigma} w_j L_j / P_j^{1-\sigma}$  — is *identical in form* to the  $\gamma = 0$  case (the factor  $1/(1-\gamma)$  cancels from both sides). The roundabout share therefore lives entirely inside the definition of  $p$ . Two consequences follow. (i) In the forward solve the mill price is no longer explicit:  $p$  and  $P$  must be solved jointly at fixed  $(w, L)$  as a nested inner  $p \rightleftharpoons P$  fixed point, after which the wage update uses the modified market-clearing exponent  $\xi = \sigma - \gamma(\sigma - 1)$ ,

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

with  $\Phi_i = \sum_j \tau_{ij}^{1-\sigma} w_j L_j / P_j^{1-\sigma}$  and the migration block unchanged. (ii) In the inversion the  $(p, P)$  market-clearing fixed point is  $\gamma$ -invariant;  $\gamma$  enters only the closed-form read-off  $Z_i = c_\gamma w_i^{1-\gamma} P_i^\gamma / p_i$ , so the recovered  $Z_i$  shift with  $\gamma$  even though the iteration does not. At  $\gamma = 0$ ,  $c_\gamma \rightarrow 1$ ,  $P_i^\gamma \rightarrow 1$ ,  $\xi \rightarrow \sigma$ , and both the forward solve and the inversion collapse exactly to the value-added economy of the baseline model (Equation 25);  $p_i = w_i / Z_i$ , recovering the algorithm above.

## B.8.2 Preferred model

Our preferred cell is  $\varepsilon = 3$ ,  $\sigma = 5$ ,  $\alpha = 0$ ,  $\zeta = 0$ ,  $\gamma = 0$  with the levels-with-lag quadratic damage form (damage specification  $m_2$ ): the central migration and trade elasticities, both the congestion/agglomeration externalities and the roundabout share switched off, and the modal damage specification. This cell is the headline specification; the other 143 model×structural cells contribute their own point estimates, which enter the CHOICE band.

## B.8.3 Choice grid

The grid is not a single five-way Cartesian sweep over independent scalars. It is a 36-cell *structural* grid — the product of three migration-elasticity values, three trade-elasticity values, two externality regimes, and two roundabout shares ( $3 \times 3 \times 2 \times 2 = 36$ ) — each of which is then fit under each of 4 damage-function specifications, giving 144 model×structural cells (Table A7). The four structural parameters cannot be identified from the within-model gravity step and

must be assumed; 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:** The 144-cell GE choice grid (36 structural cells  $\times$  4 damage forms).

Axis	Values	Justification	·
$\varepsilon$	{1, 3, 5}	migration elasticity (Fréchet shape)	3
$\sigma$	{2, 5, 8}	trade elasticity ( $\sigma - 1$ )	3
$(\alpha, \zeta)$	{(0,0), (0.3,0.05)}	congestion + agglomeration externalities, off/on (varied jointly)	2
$\gamma$	{0, 0.5}	roundabout intermediate-input share	2
Structural cells (Cartesian product of the above)			36
damage spec.	{ $m_1, m_2, m_3, m_4$ }	quadratic-FD / levels+lag quadratic (pref.) / levels+lag cubic / pop.-weighted	4
Total (36 $\times$ 4)			144

The migration elasticity  $\varepsilon$  governs the strength of agents' response to differences in destination utility; published estimates span an order of magnitude. We bracket the consensus range with  $\varepsilon \in \{1, 3, 5\}$ :  $\varepsilon = 1$  matches Suárez Serrato and Zidar (2016) and the lower end of Caliendo et al. (2019);  $\varepsilon = 5$  matches the upper end of Monte et al. (2018);  $\varepsilon = 3$  is a central estimate consistent with Burzyński et al. (2022). The two extreme values serve as informative bounds rather than independently credible point estimates.

The trade elasticity (here  $\sigma - 1$ , with  $\sigma$  the elasticity of substitution across state varieties) governs the substitutability of state-level goods; published estimates again span an order of magnitude. We bracket the range  $\sigma \in \{2, 5, 8\}$ , anchored against Broda and Weinstein (2006) ( $\sigma \approx 4-6$ , near our central value) and the higher end of Eaton and Kortum (2002); the low value  $\sigma = 2$  is an informative lower bound.

The externality parameters  $\alpha$  and  $\zeta$  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). Each is conventionally near zero in spatial-economics estimates. We vary the pair *jointly* (in lockstep) between an *off* regime  $(\alpha, \zeta) = (0, 0)$  and an *on* regime  $(\alpha, \zeta) = (0.3, 0.05)$ , so the externality axis contributes a single binary (off/on) factor rather than an independent  $2 \times 2$  grid. The case  $\alpha = \zeta = 0$  shuts down both externalities and reduces the model to a pure Eaton–Kortum gravity in goods plus 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 (the input–output “roundabout” structure of the preceding paragraph);  $\gamma = 0$  recovers the baseline value-added economy, while  $\gamma = 0.5$  introduces production link-

ages that amplify how a local climate shock propagates through prices and wages. As with the elasticities and externalities,  $\gamma$  is not identified within the model and is varied to expose its contribution to the CHOICE band.

The damage-function specification axis enumerates 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 axes.** 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  $\Delta L$  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  $\tau_{ij}$  and  $\mu_{ij}$ . 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 25, 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 A14 summarises the pipeline.



**Figure A14:** 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,

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

Therefore,

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

where  $\mu_{ijt} = d_{ij}^{-\delta} e_{ijt}^\mu$ , and  $E(e_{ijt}^\mu \mid d_{ij}, \xi_{it}, \xi_{jt}) = 1$ , 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  $-\widehat{\varepsilon} \delta \log d_{ij}$ . Note that while  $\varepsilon$  is unidentified in this regression,  $\mu_{ij}$  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 ( $\xi_{it} = \text{origin} \times \text{year}$ ,  $\xi_{jt} = \text{destination} \times \text{year}$ ) 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,

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

where again we assume the error is such that  $E(e_{ij}^\tau \mid d_{ij}, \varsigma_{it}, \varsigma_{jt}) = 1$ . Consequently,  $-(\sigma - 1) \log \tau_{ij}$  can be estimated with  $-\widehat{(\sigma - 1)\kappa} \log d_{ij}$ , again using a Poisson PMLE estimator. The trade-cost regression is estimated on a distinct dataset — a single 2012 Commodity Flow Survey (CFS) cross-section, aggregated to state-to-state shipment values — and therefore carries *origin-state and destination-state* fixed effects with no year dimension (in contrast to the multi-year migration regression's *origin*  $\times$  *year* and *destination*  $\times$  *year* effects). As with  $\mu$ , the recovered  $\widehat{\tau}_{ij}^{1-\sigma}$  matrix is normalised by the mean of its self-flow diagonal so that  $\tau_{ii}^{1-\sigma}$  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.

The cost estimates implied by these two gravity equations are reported in Table A8 (*origin*  $\times$  *year* and *destination*  $\times$  *year* fixed effects, 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: Model:	Migration Flow (1)	Trade Flow (2)
<i>Variables</i>		
log distance	-0.8190*** (0.0860)	-0.9893*** (0.0535)
self-flow	2.668*** (0.4296)	0.5180*** (0.1696)
<i>Fixed-effects</i>		
origin×year	Yes	Yes
destination×year	Yes	Yes
<i>Fit statistics</i>		
Observations	46,080	2,304
Pseudo R <sup>2</sup>	0.99788	0.91213

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

**Table A8:** Gravity (cost) estimates: origin×year and destination×year fixed effects, 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  $\varepsilon$  and the trade elasticity  $\sigma$  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  $\varepsilon$  and  $\sigma$  as assumptions, calibrated from the external literature, and the  $\varepsilon$  and  $\sigma$  axes of the choice grid (above) show the sensitivity of the model results to alternative assumptions across the consensus range  $\varepsilon \in \{1, 3, 5\}$ ,  $\sigma \in \{2, 5, 8\}$ .

**Recover  $\{A_{it}\}$  and  $\{Z_{it}\}$  in sample** Now that we have estimates for  $\tau_{ij}$  and  $\mu_{ij}$ , we can invert the system described by (25) 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 mill-price coordinates, and productivity is read off from the mill-price identity  $p_{it} = c_\gamma w_{it}^{1-\gamma} P_{it}^\gamma / Z_{it}$ , 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, write  $v_{jt} = w_{jt} / P_{jt}$  so that  $V_{jt} = A_{jt} v_{jt}$ ; the amenity is recovered from the migration accounting  $L_{jt} = \sum_i m_{ijt} L_{i,t-1}$ , the Fréchet shares giving the fixed-point relation  $A_{jt} = L_{jt}^{1/\varepsilon} (\sum_i m_{ijt}^0 L_{i,t-1})^{-1/\varepsilon}$ , where  $m^0$  is evaluated at the current amenity guess and iterated to convergence. The amenity is recovered *after* productivity because it enters only through  $V_{jt} = A_{jt} w_{jt} / P_{jt}$  and so requires the price

index  $P_{jt}$  produced by the first pass; the externality-free fundamentals  $\bar{A}_{it}, \bar{Z}_{it}$  then follow by dividing out the congestion/agglomeration terms of Equation 24. 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. This inversion is performed conditional on a particular choice of  $\varepsilon, \sigma$ , the externality regime, and the roundabout share  $\gamma$  (which, as noted, shifts the recovered  $Z_{it}$  even though the  $(p, P)$  iteration itself is  $\gamma$ -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

$$\begin{aligned}\log A_{it} &= \beta_1^A T_{it} + \beta_2^A T_{it}^2 + \rho^A \log A_{i,t-1} + \gamma_i^A + \gamma_t^A + \epsilon_{it}^A, \\ \log Z_{it} &= \beta_1^Z T_{it} + \beta_2^Z T_{it}^2 + \rho^Z \log Z_{i,t-1} + \gamma_i^Z + \gamma_t^Z + \epsilon_{it}^Z,\end{aligned}\tag{26}$$

where the lagged-outcome term enters as a control whose coefficient  $\rho^m$  is *not* restricted to be common across  $A$  and  $Z$ , and the fixed effects  $\gamma_i^m, \gamma_t^m$  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 and bracket the modelling choices that produce qualitatively different damage estimates in this small panel.

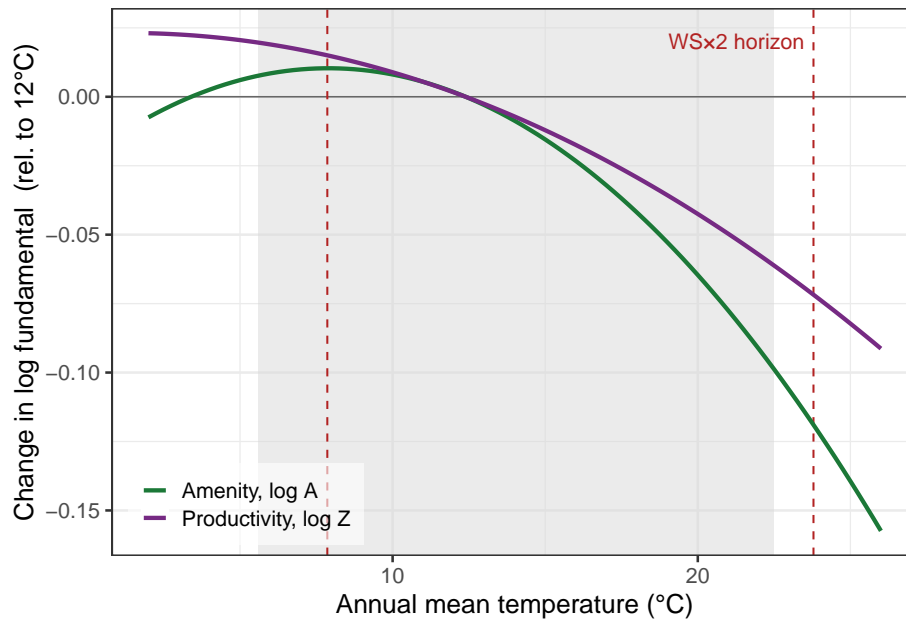
For interpretation and as a sense-check, the estimated quadratic temperature response can be re-expressed in a Gaussian “optimal-temperature” form. Imposing that there is some optimal temperature for amenities and productivities,  $\mu_m$ , and that deviations from it affect the outcome according to

$$D_m(T) = \exp \left\{ - (T - \mu_m)^2 \frac{1}{2s_m} \right\}, \quad m \in \{A, Z\},\tag{27}$$

the quadratic coefficients map to the location and shape parameters via  $\beta_1^m = \mu_m/s_m$  and

$\beta_2^m = -\frac{1}{2s_m}$ .<sup>5</sup> 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  $\varepsilon$  and  $\sigma$ , exhibit several features (Figure A15). First, the estimates are very statistically imprecise: at low elasticities the 95% confidence interval for the amenity-damage location parameter  $\mu_A$  can span roughly  $-7$  to  $19^\circ\text{C}$ , so the point estimates should be interpreted with caution and the sensitivity checks of the choice grid are essential.<sup>6</sup> 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  $\varepsilon$  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  $\sigma$  mainly affects the productivity-damage function (higher  $\sigma$  implies a higher optimal productivity temperature with a shallower slope).



**Figure A15:** 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  $12^\circ\text{C}$  reference, over the in-sample support (grey band) with the  $\text{WS}\times 2$  horizon range marked. Both fundamentals decline with warming across most of the sampled range, however the estimates are statistically imprecise.

<sup>5</sup>Since  $\hat{D}(T) = \exp\left(-\frac{(T - \mu_m)^2}{2s_m}\right)$  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 (26) gives  $\beta_1^m = \mu_m/s_m$  and  $\beta_2^m = -1/(2s_m)$ .

<sup>6</sup>Alternative specifications of Equation (26) — 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 choice axes.

### B.8.5 Projection

For each cell and damage-coefficient draw  $k$ , the projection produces a 25-year trajectory ( $t = 0, \dots, 25$ ) 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

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

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 (27) 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  $(\beta_j^A, \beta_j^Z)$  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 cell only; the other 35 structural cells are propagated at their coefficient mean, which suffices for the CHOICE band. For each cell and draw the equilibrium (Equation 25) is solved at every year  $t \in \{0, \dots, 25\}$  under the no-warming baseline and under SSP3-7.0 warming at  $WS \times 1$  and  $WS \times 2$ , using the Gauss–Seidel iteration of the algorithm above ( $\gamma = 0$  cells through the single-loop system,  $\gamma > 0$  cells through the nested mill-price system). If an iterate produces non-finite values the relaxation parameter  $\theta$  is halved and the solve retried, a bounded number of times; for roundabout cells ( $\gamma > 0$ ) that still fail to converge from an uninformative initial guess, the equilibrium is traced by homotopy from  $\gamma = 0$  to the target 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 (Fig. A1).

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, and forcing a steady state at  $t = 0$  would override observed initial conditions with a model artefact. The population vector  $L_0$  is initialised at the IRS baseline share; the wage vector is initialised at  $w_0^{(0)} = \mathbf{1}$

and Gauss–Seidel-iterated to fixed point at each year. The cost of the non-steady initialisation is that the no-warming trajectory drifts over the projection horizon as the GE state slowly converges to its (counterfactually unchanging) steady state; this drift cancels in the WS×2 minus no-warming differential that the migration-share metric isolates.

### B.8.6 Uncertainty quantification

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

**NONE.** The preferred-cell point estimate at the coefficient mean. The other 143 model×structural cells 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 cell, 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 144 model×structural cells 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

Four caveats.

*First*, the damage-function estimation is noisy in our small panel: the 95% confidence interval for the amenity-damage location parameter can span roughly  $[-7, 19]^{\circ}\text{C}$  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 band is correspondingly wide. The damage-estimation strategy is itself one of the CHOICE axes, so this specification uncertainty enters the CHOICE band; the STATISTICAL band reflects only coefficient sampling within the preferred strategy.

*Second*, the elasticities  $\varepsilon$  and  $\sigma$  are not identified internally; the choice grid is an explicit acknowledgement that the GE projection is conditional on assumed elasticities, and the CHOICE band quantifies the consequences.

*Third*, the externality regime  $(\alpha, \zeta)$  and the roundabout share  $\gamma$  are similarly assumed: the externalities are varied jointly between an *off* regime ( $\alpha = \zeta = 0$ ) and an *on* regime ( $\alpha = 0.3, \zeta = 0.05$ ) that brackets the literature, and  $\gamma$  between 0 (no roundabout) and 0.5; none is data-driven for our panel.

*Fourth*, the GE solver occasionally produces non-finite iterates in the high- $\sigma$ /high- $\zeta$  corner and for  $\gamma > 0$  cells where the inner price block contracts slowly; the bounded  $\theta$ -halving and the  $\gamma$ -homotopy fallback recover the equilibrium, so all 144 cells contribute valid points to the Statistical and CHOICE bands.

The GE implementation produces qualitatively similar headline numbers to the gravity implementation — migration-share medians both fall in the 0.2–0.8% range under SSP3-7.0 at the 25-year horizon — but with different cross-state response patterns: the GE allows hot states to partially adapt via trade, while the gravity model treats each state’s response in isolation.

### **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 (n.d.). 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.
- 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.
- Benveniste, H el ene, 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.
- Bertoli, Simone and Jes us Fern andez-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.
- 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.
- 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.
- 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.
- FAO**, "Migration, agriculture and climate change: Reducing vulnerabilities and enhancing resilience," Technical Report September 2018.
- 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.
- Hauer, Mathew and James Byars**, “IRS county-to-county migration data, 1990–2010,” *Demographic Research*, May 2019, 40, 1153–1166.
- 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.
- 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**, “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*, 2021, 13 (9), 4349–4383.
- 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*.
- 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.
- 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.
- 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.

- 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.