Identification in Mover Designs

Vahid Moghani*

November 2026

Abstract

Mover designs regress movers' changes in outcome on the gap in mean outcomes between destination and origin. The coefficient is read as the slope that random assignment to locations would produce: how much better people do where residents do one unit better. I show that it equals this slope, whatever the place effects, if and only if movers' routes are aligned. Alignment means that, for every location, the gaps in mean outcomes to the locations its movers connect it to sum, up to a common factor, to what random assignment would give. Mover flows and mean outcomes suffice to check alignment. When alignment fails, the coefficient places no restriction on the slope, even with additive effects and exogenous mobility. Yet under nonnegative sorting between person and place effects, the coefficient bounds the slope. For childhood moves across U.S. commuting zones, the bounds for adult income are 0.25 and 0.88.

Keywords: mover designs, place effects, identification, sorting, two-way fixed effects, partial identification.

*Erasmus School of Economics, Tinbergen Institute, and Erasmus Center for Health Economics Rotterdam. I am especially grateful to Patrick Kline for insightful comments and observations. I also thank Hans van Kippersluis, Owen O'Donnell, Sonia Bhalotra, Olivier Marie, Jonathan Roth, Peter Hull, Jonathan Rossi, Lisa Voois, and Joaquim Vidiella-Martin for their helpful comments. I gratefully acknowledge financial support from the Erasmus Initiative Smarter Choices for Better Health. I declare that there are no conflicts of interest.

I Introduction

A move compares two locations, and the researcher does not choose which two. Imagine four regions formed by two factors: doctors who treat intensively or sparingly, and residents who seek care often or rarely. Intensive doctors add one unit to everyone's care use. Residents who seek care often use one unit more wherever they live. In this setup, half of the difference in average care across regions comes from the doctors, and half from the residents. Now follow a person who moves. Suppose she moves from a region with sparing doctors to one with intensive doctors, with similar residents. The destination's average use is one unit higher, and her own use rises by one unit. Her individual change matches the regional gap entirely. Suppose instead she moves between regions that differ only in their residents and not in their doctors. The destination's average use is still one unit higher, but her own use does not change. Her individual change matches none of the regional gap. A regression of the change in own use on the change in mean use gives one if all moves are of the first kind. It gives zero if all are of the second. It gives one half if the two kinds of moves are equally common. The regions, their doctors, and their residents are the same in every case. Only the routes movers take differ.

Mover designs answer a question about locations with a regression of this kind. Outcomes differ across regions, hospitals, and neighborhoods. Part of the difference is the place, and part is the people who sort into it. To separate place from people, the design relates the change in a mover's outcome to the gap in mean outcomes between her destination $d(i)$ and her origin $o(i)$,

$$\Delta Y_i = \text{constant} + \theta \left(\bar{Y}_{d(i)} - \bar{Y}_{o(i)} \right) + u_i. \tag{1}$$

I call the coefficient of this regression the *mover coefficient*. It is the share of the gap between destination and origin that movers take on, on average over the moves they make. Finkelstein et al. [2016] estimate it in an event study around the move, to ask why health care spending differs across regions. Chetty and Hendren [2018a] use children who move with their parents, to ask why adult incomes differ across the neighborhoods where children grow up. Their coefficient is the part of the gap that a child takes on with each year of childhood spent in the destination. The design is now standard in health, labor, public, and political economics. It has been used to study how physicians practice [Molitor, 2018], opioid use [Finkelstein et al., 2025], voting [Cantoni and Pons, 2022], carbon emissions [Lyubich, 2025], and intergenerational mobility [Chetty et al., 2020, Alesina et al., 2021].[1]

The question motivating the design concerns all locations. Consider an economy with many

[1] In the event-study form, applications include geographic variation in health care [Finkelstein et al., 2016, Godøy and Huitfeldt, 2020, Salm and Wübker, 2020, Zeltzer et al., 2021, Leal et al., 2026], physician and organizational behavior [Molitor, 2018, Huang and Ullrich, 2026, Doyle and Staiger, 2025, Agha et al., 2023, Encinosa and Dor, 2025, Badinski et al., 2023], health behaviors and outcomes [Hinnosaar and Liu, 2022, Hut, 2020, Finkelstein et al., 2025, Ding, 2026, Chyn and Shenhav, 2025], insurance and program choice [Cabral et al., 2025], employers and research institutions [Ahammer et al., 2024, Chandra and Xu, 2025], voting [Cantoni and Pons, 2022], and carbon emissions [Lyubich, 2025]. In the age-at-move form, they include Chetty and Hendren [2018a], Bell et al. [2019], Chetty et al. [2020], Deutscher [2020], Alesina et al. [2021], Laliberté [2021], and Chetty et al. [2026].

locations, each with its own residents. Define the *place effect* of a location as the average outcome individuals in the economy would have if they all lived there. Now consider an experiment that draws individuals at random from the economy and assigns them at random to locations. Ask how much higher an individual's outcome is, on average, where residents' mean outcome is one unit higher. Because assignment is random, the answer is how much higher the place effect is, on average, where residents' mean outcome is one unit higher. I call this the *cross-sectional coefficient*. Regression (1) would yield the same coefficient if the experiment moved individuals between two locations drawn independently at random. The probabilities with which the experiment assigns individuals to locations determine the weights on locations. These weights determine whose locations matter: those of all residents, of children, or of individuals similar to those who move.

Chetty and Hendren [2018a] and Chetty et al. [2026] define their target by the experiment, for children and their outcomes as adults. Finkelstein et al. [2016] define their target as the share of the gap in healthcare utilization between two regions that is due to place. A weighted average of that share over all pairs of regions is the cross-sectional coefficient, so the two targets are the same coefficient, for different outcomes. The mover coefficient is also read as the share of the variation across locations that is due to place [Finkelstein et al., 2016, Molitor, 2018, Ding, 2026]. This share is the cross-sectional coefficient when variation means the gaps between locations. It is the variance share of place effects plus half the share of the covariance of place effects with the rest of mean outcomes. If mean outcomes predict place effects exactly, the variance share of place effects is the square of the cross-sectional coefficient. A coefficient of 0.20 then goes with a variance share of four percent.

In this paper I ask what a mover design, run on a sample of movers, can identify about the cross-sectional coefficient of the locations. I start from the assumption that, across movers, what a mover gains beyond the effect of the place is uncorrelated with the gap in mean outcomes between destination and origin.[2] I call this assumption the *assumption of the design*. I call the mean outcomes of all locations the *map*. For each location, each move into or out of it creates a comparison of that location with another location. The sum of the gaps in mean outcomes over those comparisons, divided by the weight of the location, is the *comparison gap* of that location.

I show, to my knowledge for the first time, that the mover coefficient identifies the cross-sectional coefficient, without additional assumptions on the place effects, under one condition and only under it. The comparison gaps must be exactly proportional to the comparison gaps under moves between two locations drawn at random. I say that the routes are *aligned* with the map when this condition holds. When the routes are not aligned with the map, I show that the difference between the mover coefficient and the cross-sectional coefficient is unrestricted.

[2]The mover coefficient is a statement about place effects if and only if this assumption holds (Section II.1). Mover studies state the assumption, or a stronger one, as their identifying assumption. Their evidence on trends before the move, and on the symmetry of gains and losses, is evidence on this assumption. Under this assumption alone, a two-way fixed-effects regression need not recover the place effect of any location (Section VI). The main results, Theorems 1 to 3, hold under every stronger assumption as well. One such assumption is the additive model, in which a location has the same effect on everyone and moves are unrelated to shocks.

Alignment can be checked without estimating the place effect of any location. The check uses only the flows of movers between locations, mean outcomes, and the weights of the question. How close the routes are to aligned is measured by the correlation, across locations, between the comparison gap and the mean outcome. This correlation equals one when routes are aligned, and I call it the *alignment* of the routes. When each location is weighted by its moves, alignment can also be read off a plot. For each location, average the mean outcomes of the locations its movers connect it to. Plot this average against the mean outcome of the location itself. The routes are aligned if and only if the points lie on a line.

Routes that are not aligned for the population in the question may be aligned for another population of locations, or for none. Flows and mean outcomes show whether there is a population for which the routes are aligned. The sum of gaps of a location does not depend on the weights. This sum is positive when the mean outcome of the location is higher, on balance, than the mean outcomes of the locations its movers connect it to. The sum is negative when the mean outcome of the location is lower. Some weights align the routes if and only if every location with a positive sum has a higher mean outcome than every location with a negative sum.[3] When this condition fails, no weights align the routes. The mover coefficient then puts no limit on the cross-sectional coefficient of any population of locations.

When the routes are not aligned, I show that the difference between the two coefficients is a product of three numbers. The first number is how much the map misses about place effects: the part of place effects that mean outcomes do not predict. The second number is how far the routes are from aligned, and it is computed from flows and mean outcomes. The third number is the correlation between what the map misses and the part of the comparison gap that the map does not explain. A limit on the first number therefore bounds the difference. The data on movers give no such limit. Suppose the assumption of the design holds and no more. The outcomes of every mover before and after the move then identify the mover coefficient and nothing else about place effects.[4]

More assumptions narrow the answer. A sign for sorting gives bounds, and the assumption behind a two-way fixed-effects regression gives a number. Sorting is not negative when locations with larger place effects do not, on average, have residents who do worse wherever they live. Studies that estimate person and place effects report a positive correlation between the two for healthcare use and for voter turnout [Finkelstein et al., 2016, Cantoni and Pons, 2022]. The sign remains the researcher's judgment. Nonnegative sorting limits how much the map misses. The mover coefficient and the alignment then give bounds on the cross-sectional coefficient. No tighter bounds are valid under the assumption of the design and nonnegative sorting. For a mover coefficient of one half and an alignment of 0.98, the bounds are 0.40 and 0.60. The bounds are about what the data identify, and they remain with any number of movers.

[3] The statement is for the case in which no location has a sum of zero (Theorem 2).

[4] This statement, and the statement below that data on movers cannot test the stronger assumption, are for outcomes without bounds that are known in advance. Known bounds, as ranks have, add inequalities. These inequalities are weak when few movers have lived in any given location (Section V.1).

A two-way fixed-effects regression gives a number in place of the bounds, under a stronger assumption. At every location, the movers who arrive must gain as much beyond the place effect as the movers who leave. The regression then recovers the place effect of each location, and therefore the cross-sectional coefficient, for any routes that connect all locations. The additive model of Abowd et al. [1999], in which a location has the same effect on everyone, satisfies this assumption. Chetty and Hendren [2018b] describe the two requirements for their own estimates. The effect of each of 3,000 counties needs a condition for each county. The average exposure effect of the authors' first paper needs a single condition that holds on average. The outcomes of movers cannot test the stronger assumption.

For childhood moves between commuting zones in the United States, public data show routes that are not aligned with the map [Chetty and Hendren, 2018b].[5] The outcome is the income rank at age 26 of children whose parents are at the 25th percentile of income. For each ordered pair of zones with at least 25 children who made that move, the public files contain the number of children and an estimate. The estimate is what a year of childhood in the destination adds to the outcome, compared with a year in the origin. The mover coefficient computed from these estimates is 0.61. The alignment of the routes with the map is 0.77 for population weights, and no other weights align the routes. Draw a line anywhere between high and low mean outcomes. At least 106 of the 595 zones then lie above the line with a negative sum of gaps, or below the line with a positive sum.

With nonnegative sorting, the bounds on the cross-sectional coefficient are 0.25 and 0.88. The public files also contain two-way estimates of the place effect of each zone. A researcher who accepts the stronger assumption can take these estimates as the place effects. The cross-sectional coefficient is then 0.75. With the same place effects, the routes give a mover coefficient of 0.63. The two coefficients differ by 0.12, and with aligned routes they would be equal. A number in place of the bounds is what the stronger assumption buys.

The paper builds on the studies that use movers to separate place effects from sorting [Bronnenberg et al., 2012, Finkelstein et al., 2016, Chetty and Hendren, 2018a]. Several of these studies report shares computed from estimated place effects next to the mover coefficient [Finkelstein et al., 2016, Cantoni and Pons, 2022]. Lyubich [2025] notes that reading the mover coefficient as a share requires a condition on sorting. Card et al. [2024, 2025] show that industry and regional wage effects reflect the particular firms that movers leave and join. Hull [2018] shows that with more than two locations a regression with an effect for each location need not average causal effects. Under the assumption of the design alone, a two-way regression need not recover the place effects. I show that the mover coefficient identifies the cross-sectional coefficient under that assumption when the routes are aligned. When the routes are not aligned, the mover coefficient puts no limit on the cross-sectional coefficient, even in the additive model.

The work on firm wage effects separates the assumptions on what movers gain. Kline [2024]

[5] The public files list the pairs of zones with at least 25 children. These pairs carry three quarters of the moves, and all numbers are for these pairs. The counts of children pool all parental incomes, while the mean outcomes and the estimates are for parents at the 25th percentile. The numbers do not reproduce a published estimate (Section VII).

shows that the wage changes of workers who move between two firms identify an average effect for those workers. The average effects for pairs of firms fit together into an effect of each firm when workers do not select on their gains. Sorkin and Warwar [2026] compare the pay of two workers who make the same move between firms. In Brazilian and Italian data, the authors reject that workers do not select on their gains. Under the assumption of the design in this paper, movers may select on their gains, as long as these gains are unrelated to the gap in mean outcomes.

Closest to the condition on routes are results on how well observational estimates forecast causal effects. Abaluck et al. [2021] call the coefficient of causal effects on observational estimates a forecast coefficient. These authors show that an instrument identifies the forecast coefficient under a condition on the choices of the people whom the instrument moves. They call this condition the fallback condition. Kline [2024] notes that a regression of wages on worker effects and firm characteristics need not recover how firm effects vary with those characteristics. In a mover design the routes are the instrument. The counterpart of the fallback condition involves place effects, and place effects are not observed. This paper gives a condition on flows and mean outcomes only, under which the mover coefficient answers the question whatever the place effects are. The paper also gives the populations of locations, if any, for which the condition holds. When the condition fails, the paper gives bounds on the answer for a researcher who assumes nonnegative sorting.

Limited-mobility bias in two-way models is a separate matter [Andrews et al., 2008, Kline et al., 2020, Bonhomme et al., 2023]: the results of this paper hold with any number of movers. As with regressions whose weights the researcher did not choose [De Chaisemartin and d'Haultfoeuille, 2020, Sun and Abraham, 2021, Borusyak et al., 2024], the weights decide what is estimated. In a mover design the weights matter even when place effects are the same for everyone.

Section II states the setting, and Section III the question and the design. Sections IV and V contain the main results, Section VI the two-way answer, and Section VII the evidence. Proofs are in the Appendix.

II Setting

II.1 What a Move Shows

There are J locations: regions, hospitals, or neighborhoods. Individuals are observed in two periods, $t = 1, 2$. Individual i lives in location $j(i,t)$ and has outcome Y_{it}. Write $o(i) = j(i,1)$ for the origin and $d(i) = j(i,2)$ for the destination. A *mover* has $d(i) \neq o(i)$; everyone else is a stayer. A *route* is an ordered pair of locations, an origin and a destination. The mean outcome of the residents of location j is $\bar{Y}_j$. For an attribute f of locations, $\Delta f_i = f_{d(i)} - f_{o(i)}$ is its change along the move of i.

The *place effect* ψ_j of location j is the mean outcome that the movers would have if they lived in j. It is defined on the whole population of movers, whatever route each of them took. Only differences between place effects matter below. What the data show on a route is the mean change

in the outcome of those who took it,

$$\mathbb{E}[\Delta Y_i \mid o(i) = j,\, d(i) = k] = \psi_k - \psi_j + s_{jk}. \tag{2}$$

The first part is what the average mover would gain from this change of place. The second part is what the movers on this route gain beyond that, and I call s_{jk} their *extra gain*. It is positive when people move to locations that suit them better than they suit others. It also contains any change at the move that has nothing to do with the two places. Write $s_i = s_{o(i)d(i)}$ for the extra gain on the route of mover i.

Assumption 1 (What a move shows)**.** Among movers, the extra gain is uncorrelated with the gap in mean outcomes between destination and origin: $\mathrm{Cov}(s_i, \Delta\bar{Y}_i) = 0$.

The assumption lets people move to locations that suit them, and gain more from the move than the average mover would. It asks that these extra gains are unrelated to the gap in mean outcomes: the movers who gain more are not, on average, the movers who go to locations with higher mean outcomes. Take a patient who moves to a region with a specialized hospital, or a person who moves to be near family. Each may gain more from the move than the average mover would. The assumption allows both cases. It fails if movers like these two go, on average, to regions with higher mean healthcare use.

A mover coefficient needs exactly this assumption to be a statement about places. By (2), the coefficient of regression (1) is

$$\frac{\mathrm{Cov}(\Delta Y_i, \Delta\bar{Y}_i)}{\mathrm{Var}(\Delta\bar{Y}_i)} = \frac{\mathrm{Cov}(\Delta\psi_i, \Delta\bar{Y}_i)}{\mathrm{Var}(\Delta\bar{Y}_i)} + \frac{\mathrm{Cov}(s_i, \Delta\bar{Y}_i)}{\mathrm{Var}(\Delta\bar{Y}_i)}. \tag{3}$$

The first term involves only place effects and the routes of movers. The second term is zero if and only if Assumption 1 holds. Without the assumption, the coefficient mixes what places do with who gains from moving to them.

The leading special case is the two-way model of Abowd et al. [1999], in which the outcome of individual i in period t is

$$Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}. \tag{4}$$

Here α_i is the *person component*, and the place effect ψ_j is the same for everyone. Mobility is *exogenous* when the shocks ε_{it} have mean zero given all person components, all place effects, and where everyone lives in both periods. This rules out moves that respond to shocks, and it leaves sorting free: who lives where, and who moves where, may depend on the person component in any way. The mean change on the route from j to k is then $\psi_k - \psi_j$. The extra gain is zero on every route, so Assumption 1 holds in its strongest form. Finkelstein et al. [2016] derive the event-study design under this model, and the coefficient of their event study converges to that of the two-period regression studied here (Supplement H).

Scope. The results on the mover coefficient hold under every assumption that implies Assump-

tion 1, the additive model with exogenous mobility included. They say when the coefficient answers the question (Theorems 1 and 2) and what it bounds when it does not (Theorem 3). Each compares two functions of place effects, and the comparison does not depend on why the regression recovers the first. A difficulty that arises in the additive model is therefore not caused by differences in effects across people. The statement that the outcomes of movers identify nothing beyond the coefficient is of another kind (Proposition 4). It holds at Assumption 1, and it fails under the stronger assumptions of Section VI, which identify more.

Mover studies state Assumption 1, or a stronger condition, as their identifying assumption. Cantoni and Pons [2022] require that changes in what drives the turnout of a mover are unrelated to the difference in average turnout between her destination and her origin. Finkelstein et al. [2016] rule out trends that vary with the origin and the destination. In the age-at-move form of the design the outcome is observed once, and the condition is on selection in levels: Chetty and Hendren [2018a] require that selection does not vary with the age at the move. Chetty and Hendren [2018b] describe this requirement as a single condition that must hold on average: exposure to places where permanent residents do better is unrelated to the other inputs into a child's outcome. The same authors describe the condition for the effect of each county as a stronger version of that single condition.

The evidence that the studies offer is about the relation between the changes of movers and the gap. Before the move, the outcomes of movers show no trend that is related to the gap [Finkelstein et al., 2016, Cantoni and Pons, 2022, Badinski et al., 2023]. Changes are symmetric for moves to locations with higher and with lower mean outcomes [Finkelstein et al., 2016, Chetty and Hendren, 2018a]. Whether the assumption is reasonable depends on why people move. It is plausible where the outcome is not what they move for, and less so where people move to where their own skills pay best [Dahl, 2002, Kennan and Walker, 2011]. Supplement B goes through the assumptions of the studies.

II.2 Movers, Locations, and the Map

The data are about movers, and the question is about everyone. Let ψ_j^{pop} be the mean outcome that all individuals, movers and stayers, would have if they lived in j.

Assumption 2 (Representative movers)**.** A change of location does on average the same for movers as for everyone: $\psi_k - \psi_j = \psi_k^{\text{pop}} - \psi_j^{\text{pop}}$ for all locations j and k.

The assumption concerns movers as a whole. It says nothing about the movers on a given route, or about those who arrive in a location or leave it, who may gain more or less than others. That is left to the extra gains. Nor does it say that movers and stayers have the same outcomes: movers may be younger or healthier, and their outcomes may follow other trends. No result below needs the assumption as a statement about movers. It turns each result into a statement about the whole population, and I maintain it for that reading. Finkelstein et al. [2016] state it directly: their

approach relies on place effects being the same for movers and for non-movers. In the additive model it holds by construction.

Locations are fixed, and individuals are drawn at random. The number of locations J and the place effects $\psi_1, \ldots, \psi_J$ are parameters, the same in every sample. Each individual comes with a history: the locations where she lives in the two periods and her outcomes there. Histories are independent draws from one population distribution, and the sizes of locations, their mean outcomes, and the flows between them are features of it. The number of individuals grows and J stays fixed, so "the coefficient" of a regression means its probability limit. The results are statements about this limit, for any number of locations and any pattern of routes. The coefficient, the alignment, and the bounds use no estimate of a place effect. Limited-mobility bias, the noise in estimated place effects when movers are few, is therefore a separate matter [Andrews et al., 2008, Kline et al., 2020, Bonhomme et al., 2023]. Few movers on a route make the estimated flows imprecise, and Supplement J has a split-sample estimate of the alignment for that case.

Assumption 3 (Stationary locations)**.** Each location has the same size and the same share of movers in both periods, and its residents have the same mean outcome in both periods.

Locations are in a steady state: people come and go, and each location keeps its size and its mix of people. The assumption restricts totals, location by location, and not who moves where. Origins and destinations may be paired in any way.

The steady state makes the mean outcome of a location an attribute of the location, and not of a period. The regressor of the design is one number for each location, so the design already treats the mean outcome as an attribute of the location. Finkelstein et al. [2016] average the mean outcome of each region over the years of their sample. The authors report a correlation of 0.9 between the rank of a region by healthcare use in the first half of the sample and in the second half. For permanent residents with parents at the 25th percentile of income, Chetty and Hendren [2018a] report a correlation of 0.95 between the mean outcomes of adjacent birth cohorts.

In a steady state each location also receives as many movers as it sends. The gap in mean outcomes is then zero on average among movers, and the direction of a move does not matter for the mover coefficient (Proposition 1). The data show whether flows balance, and when they do not, the results hold with one more term for the net flows (Supplement I). In the data on U.S. commuting zones of Section VII, net flows add up to a quarter of all moves. Allowing for these net flows moves the mover coefficient from 0.609 to 0.606.

Under Assumption 3 the mean outcome $\bar{Y}_j = \mathbb{E}[Y_{it} \mid j(i,t) = j]$ is the same in both periods. I call the vector $(\bar{Y}_1, \ldots, \bar{Y}_J)$ the *map*. It is what the data show before anyone moves: healthcare spending by region, adult income by childhood neighborhood. Some studies take the mean over all residents [Finkelstein et al., 2016], and some over permanent residents, who do not move [Chetty and Hendren, 2018a]. The results hold for either, with Assumptions 1 and 3 read for the mean that

is used. Write $\bar{\alpha}_j := \bar{Y}_j - \psi_j$ for the part of the mean that is not the place effect, so that

$$\bar{Y}_j = \psi_j + \bar{\alpha}_j. \tag{5}$$

I call $\bar{\alpha}_j$ the *composition* of the location. It is what the residents of j add to what the place does for the average mover. In the additive model it is their mean person component, $\bar{\alpha}_j = \mathbb{E}[\alpha_i \mid j(i,t) = j]$. The map mixes what locations do with who lives in them, and the question is how much of it is place.

III The Question and the Coefficient

III.1 The Question

The question behind a mover design is about locations, and the studies state it in two ways that are one target. Chetty and Hendren [2018a] and Chetty et al. [2026] define a slope by a hypothetical experiment. Children are assigned at random to grow up in different locations, and their outcomes are regressed on the mean outcome of the residents of the location they were assigned to. A child who moves at birth would "pick up about 80%" of the difference in the outcomes of residents [Chetty and Hendren, 2018a]. Daysal et al. [2026] study a setting close to this experiment. Refugee children were assigned to Danish municipalities quasi-randomly. The authors regress their income rank in adulthood on the rank that children of low-income permanent residents reach in the municipality of assignment. Finkelstein et al. [2016] define the share of the gap in mean outcomes between two areas that is due to place. Cantoni and Pons [2022], Chyn and Shenhav [2025] and Cabral et al. [2025] report such shares for voting, birth weight, and insurance choice. The slope is the average of these shares over all pairs of locations (Proposition 2).

To state the slope, fix weights $\pi_j > 0$ that sum to one. The weights say whom the question is about. A policy that reaches everyone weights locations by their residents. A question about children weights them by children. A question about people like those who move weights them by how often movers visit them. The weights belong to the question, and the researcher chooses them. Moments across locations with these weights carry the subscript π.

Definition 1 (Cross-sectional coefficient)**.**

$$\theta_\pi := \frac{\mathrm{Cov}_\pi(\psi_j, \bar{Y}_j)}{\mathrm{Var}_\pi(\bar{Y}_j)}. \tag{6}$$

The *cross-sectional coefficient* θ_π is the slope of place effects on the map: how much better a location is, on average, when its residents do one unit better. It is the coefficient of the hypothetical experiment in which each person is assigned to location j with probability π_j. The persons are movers, on whom the place effects are defined. Under Assumption 2 they can be anyone. The slope is also the return to a policy that tilts this assignment toward the locations whose residents do better than average. For each unit by which the tilt raises the mean outcome of the residents of a

person's location, it raises her own outcome by θ_π.

The studies also report shares of the variance of the map. The slope is the variance of place effects plus their covariance with composition, over the variance of the map:

$$\theta_\pi = \frac{\text{Var}_\pi(\psi_j) + \text{Cov}_\pi(\psi_j, \bar{\alpha}_j)}{\text{Var}_\pi(\bar{Y}_j)}. \tag{7}$$

It is the variance share of place effects, $\text{Var}_\pi(\psi_j)/\text{Var}_\pi(\bar{Y}_j)$, only if place effects and composition are uncorrelated across locations, a condition that Chetty et al. [2026, footnote 32] state for their slope. Their estimate from a housing experiment is 0.68. The variance share is never below the square of the slope. It is at most the slope if the two are not negatively related. A slope of 0.68 then fits any variance share between 0.46 and 0.68.

I state the results for the slope, the reading that a mover coefficient most often receives. One may also ask what share of the gap between two groups of locations is place. The groups can be the top and the bottom quartile of the map, or the locations above and below the median. Supplement D restates the results for these shares and for shares of the variance.[6]

III.2 The Design

Among movers, the researcher regresses the change in the outcome on the change in the map,

$$\Delta Y_i = \text{constant} + \theta\, \Delta \bar{Y}_i + u_i. \tag{8}$$

The two sides of the regression change for different reasons. By (2), the mean change in the outcome on a route is the change in the place effect plus the extra gain. By (5), $\Delta \bar{Y}_i = \Delta\psi_i + \Delta\bar{\alpha}_i$: the regressor changes because the mover gets another place effect and other neighbors. Under Assumption 1 the second term of (3) is zero, and the coefficient is

$$\theta_D = \frac{\text{Cov}(\Delta\psi_i, \Delta\bar{Y}_i)}{\text{Var}(\Delta\bar{Y}_i)}, \tag{9}$$

with moments taken over the population of movers. It is large when moves change place effects a lot and neighbors little. Throughout, some movers travel between locations with different mean outcomes, so both coefficients are defined.

Only pairs of locations enter this ratio. Let F_{jk} be the probability that a mover goes from j to k, and let

$$a_{jk} := F_{jk} + F_{kj}$$

be the probability that a mover travels between j and k, in either direction. I call the a_{jk} *comparison*

[6] Comparisons between groups are common. Finkelstein et al. [2016] report the share for regions above and below the median and for the top and bottom quartiles. In studies of firm wage effects, Card et al. [2013] plot the earnings of workers who move between quartiles of mean coworker wages. The authors compare the gains of those who move up with the losses of those who move down [see Kline, 2024].

weights, and a set D of such weights a *design*. For two attributes f and h of locations, let

$$\mathcal{L}_D(f,h) := \sum_{j<k} a_{jk}\,(f_j - f_k)(h_j - h_k). \tag{10}$$

Each pair of locations contributes the product of its difference in f and its difference in h, with the comparison weight of the pair.

Proposition 1 (The mover coefficient)**.** *Under Assumptions 1 and 3, the coefficient of regression* (8) *is*

$$\theta_D = \frac{\mathcal{L}_D(\psi,\bar{Y})}{\mathcal{L}_D(\bar{Y},\bar{Y})}, \tag{11}$$

with comparison weights $a_{jk} = F_{jk} + F_{kj}$. It depends on the mobility process only through the comparison weights.

The direction of moves does not matter. Take two economies in which the same pairs of locations exchange the same numbers of movers. They have the same coefficient, even if every move in one is the reverse of a move in the other. The converse also holds. Any place effects, any map, and any comparison weights are generated by some additive economy with exogenous mobility that satisfies Assumption 3 (Lemma 2 in the Appendix). The place effects, the map, and the comparison weights therefore describe every case that the theory covers. No model of who moves where is needed.

The design that compares children who move at different ages gives a coefficient of the same form. The outcome is observed once, in adulthood, and children who leave the same origin at the same age are compared across their destinations [Chetty and Hendren, 2018a]. The coefficient at age m is

$$b_m = \omega_m\,\theta_{D_m} + \xi_m. \tag{12}$$

Here ω_m is the share of childhood that remains at age m, and ξ_m is the selection of families. The coefficient θ_{D_m} is (11) with comparison weights on pairs of destinations, $a_{jk} \propto \sum_o P_m(o)\,p_m(j \mid o)\,p_m(k \mid o)$. Here $P_m(o)$ is the share of age-m movers who leave o, and $p_m(\cdot \mid o)$ is the distribution of their destinations (Proposition 11 in the Supplement). A design is therefore any set of comparison weights, and the results below apply to both regressions.

The cross-sectional coefficient is a moment of levels: it has one term for each location, with the weights of the question. The mover coefficient is a moment of differences: it has one term for each pair of locations, with the weights of the routes.

III.3 An Average over Pairs

Both coefficients are averages of the same shares, one for each pair of locations, and they differ in the pairs that they weight. For two locations with different mean outcomes, let

$$S_{jk} := \frac{\psi_j - \psi_k}{\bar{Y}_j - \bar{Y}_k}$$

be the share of the gap between them that is place. This is the pair share of Finkelstein et al. [2016].

Proposition 2 (Pair shares)**.**

(i) *Both coefficients are weighted averages of pair shares, over the pairs with different mean outcomes:*

$$\theta_D = \frac{\sum_{j<k} a_{jk} \, (\bar{Y}_j - \bar{Y}_k)^2 \, S_{jk}}{\sum_{j<k} a_{jk} \, (\bar{Y}_j - \bar{Y}_k)^2}, \qquad \theta_\pi = \frac{\sum_{j<k} \pi_j \pi_k \, (\bar{Y}_j - \bar{Y}_k)^2 \, S_{jk}}{\sum_{j<k} \pi_j \pi_k \, (\bar{Y}_j - \bar{Y}_k)^2}.$$

(ii) *Fix the place effects and the map. As the comparison weights vary, θ_D takes every value between the smallest and the largest pair share, and no other value.*

The mover coefficient is the weighted average that Finkelstein et al. [2016] describe, and its weights are the routes. The cross-sectional coefficient weights each pair by the sizes of the two locations. These are the pairs that two independent draws from π would produce. I call comparison weights of this form, $a_{jk} \propto \pi_j \pi_k$, *independent pairing*. In the four regions of the introduction, the pair share is one for two regions that differ in their doctors, and zero for two regions that differ in their residents. With moves of these two kinds only, the mover coefficient is the share of the first kind. The cross-sectional coefficient of four equally large regions is one half. The weights would not matter if all pair shares were equal. They are equal only when the place effect is an exact linear function of the map. Otherwise the weights matter, and by part (ii) they can matter a great deal.

IV When the Coefficient Answers the Question

IV.1 A Slope Across Locations

To compare the two coefficients, write the mover coefficient in levels. Let $a_j := \sum_{k \neq j} a_{jk}$ be the comparison weight of location j, and let

$$\bar{Y}_j^P := \frac{1}{a_j} \sum_{k \neq j} a_{jk} \, \bar{Y}_k \tag{13}$$

be the mean outcome of its *partners*, the locations it is compared with. For each location let

$$g_j := \frac{1}{\pi_j} \sum_{k \neq j} a_{jk} \, (\bar{Y}_j - \bar{Y}_k) = \frac{a_j}{\pi_j} \, (\bar{Y}_j - \bar{Y}_j^P). \tag{14}$$

I call g_j the *comparison gap* of location j. Sum the gaps between the mean outcome of a location and the mean outcomes of the locations its movers connect it to. The comparison gap is this sum divided by the weight of the location. It is positive for a location whose mean outcome is above that of its partners, and negative for a location whose mean outcome is below.

Proposition 3 (The coefficient in levels)**.** *For every attribute f of locations, $\mathcal{L}_D(f, \bar{Y}) = \mathrm{Cov}_\pi(f_j, g_j)$. Hence*

$$\theta_D = \frac{\mathrm{Cov}_\pi(\psi_j, g_j)}{\mathrm{Cov}_\pi(\bar{Y}_j, g_j)}, \qquad \textit{while} \qquad \theta_\pi = \frac{\mathrm{Cov}_\pi(\psi_j, \bar{Y}_j)}{\mathrm{Var}_\pi(\bar{Y}_j)}. \tag{15}$$

Both coefficients are slopes of place effects on the map, across locations. The cross-sectional coefficient is the least-squares slope. The mover coefficient is the instrumental-variables slope of the same regression, with the comparison gap as the instrument. The researcher chooses the regression. The movers choose the instrument.

The two slopes differ when the instrument is related to what the regressor leaves out. Write the least-squares regression as

$$\psi_j = \text{constant} + \theta_\pi \, \bar{Y}_j + e_j. \tag{16}$$

The residual e_j is the part of the place effect that the map does not predict. It has mean zero and is uncorrelated with the map across locations. Chetty and Hendren [2018b] estimate the variance of its counterpart for commuting zones and counties in the United States. Their counterpart is the part of an area's effect that the outcomes of its permanent residents do not capture. The size of the residual relative to the map is $\eta := \mathrm{sd}_\pi(e_j) / \mathrm{sd}_\pi(\bar{Y}_j)$. This number says how much the map misses about place effects. The variance share of place effects is $\theta_\pi^2 + \eta^2$. If the map misses nothing, it is the square of the slope.

IV.2 Alignment

Theorem 1 (Identification)**.** *Let D be a design with $\mathcal{L}_D(\bar{Y}, \bar{Y}) > 0$.*

(i) The difference. $\theta_D - \theta_\pi = \dfrac{\mathrm{Cov}_\pi(e_j, g_j)}{\mathrm{Cov}_\pi(\bar{Y}_j, g_j)}$.

(ii) Alignment. $\theta_D = \theta_\pi$ *for every vector of place effects if and only if the comparison gap is proportional to the map: $g_j = \kappa \, (\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}])$ for some $\kappa > 0$ and all j.*

(iii) Anything goes. *Suppose the comparison gap is not proportional to the map, and take any two numbers t and t'. There is an additive economy with exogenous mobility that satisfies Assumption 3 and has the same map, the same weights π, and the same comparison weights. In it, $\theta_D = t$ and $\theta_\pi = t'$.*

(iv) Independent pairing. *The comparison gap is proportional to the map for every map if and only if $a_{jk} \propto \pi_j \pi_k$.*

I say that routes are *aligned* with the map when the condition in (ii) holds. Multiply both sides by the weight of the location. The left side is then the sum of the gaps between the mean outcome

of the location and the mean outcomes of the locations its movers connect it to. The right side is proportional to the same sum under independent pairing. Aligned routes give every location the sum that independent pairing would give it, up to a common factor.

The theorem leaves no intermediate case. If routes are aligned with the map, the mover coefficient is the cross-sectional coefficient, whatever the place effects are. If they are not, the mover coefficient carries no information about the cross-sectional coefficient: any value of one is compatible with any value of the other. This holds in the additive model with exogenous mobility, the case most favorable to the design. Alignment is a condition on flows and location means, which are observed.

Example 1 (Two kinds of moves)**.** The outcome is the use of health care. Four equally large regions are indexed by (p, z), where p and z are each -1 or 1. The place effect is $\psi = p$: in regions with $p = 1$ providers treat more intensively, and anyone who lives there uses more care. Composition is $\bar{\alpha} = cz$: regions with $z = 1$ have residents who seek care often, and they use more care wherever they live. The two are unrelated across regions, so $\theta_\pi = 1/(1 + c^2)$. A share ρ of moves are *across practice styles*: p changes and z does not. The other moves are *across patient populations*: z changes and p does not. The map is $p + cz$, and the comparison gap of region (p, z) is $4\{\rho\, p + (1 - \rho)\, cz\}$. The routes are therefore aligned with the map only at $\rho = \frac{1}{2}$, and

$$\theta_D = \frac{\rho}{\rho + c^2(1 - \rho)}$$

equals θ_π only there. As ρ goes from zero to one, the coefficient goes from zero to one, while the regions, their providers, and their residents stay the same. Every region also sends and receives the same number of movers for every ρ. It matters which locations movers compare, and not only which locations they visit. With $c = 1$ this is the economy of the introduction, apart from the unit of use. With $c = \frac{1}{2}$ the cross-sectional coefficient is 0.8. The pair shares are one across practice styles, zero across patient populations, and 0.67 and 2 along the two diagonals.

Part (iii) can be seen in the example. Take $c = \frac{1}{2}$ and $\rho = 0.9$. With the place effect p and composition $z/2$, the coefficients are $\theta_D = 0.97$ and $\theta_\pi = 0.80$. Now let the place effect be $1.05\, p - 0.9\, z$ and composition be $-0.05\, p + 1.4\, z$. The map is the same, so are the routes, and so is the mover coefficient. The cross-sectional coefficient is 0.48. A reader who has the map, the routes, and the mover coefficient cannot tell the two economies apart.

In the example, those who move from one region to another gain what those who move the other way lose, for every ρ. Symmetry compares the two directions of a move, and the coefficient does not depend on the direction (Proposition 1). Symmetry is evidence on the assumption of the design. Alignment is about which locations movers compare.[7]

[7] Finkelstein et al. [2016] plot the change in the utilization of movers against the gap. The relation is symmetric and linear, and the authors take this as support for their additive model. A linear relation says that pair shares do not vary on average with the size of the gap along the routes. It can hold for any alignment. In Example 1 with $c = 1$, the mean change of movers is ρ times the gap for every ρ, and the routes are aligned only at $\rho = \frac{1}{2}$. More generally, the relation is linear whenever the changes in place effects and in composition are jointly normal across moves.

Part (i) describes one economy. The two coefficients agree in it if and only if what the map misses about place effects is unrelated, across locations, to the comparison gap. This is the exclusion restriction of the instrument, and it cannot be examined without place effects. The condition has a counterpart in the literature on value added. That literature asks how well observational estimates forecast causal effects, with quasi-experiments that change a person's teacher, school, or health plan [Chetty et al., 2014, Angrist et al., 2016]. Abaluck et al. [2021] call the coefficient of causal effects on observational estimates the forecast coefficient. Their instrument is the termination of a person's health plan. It identifies the forecast coefficient under a condition on the plans people then choose. The instrument must be unrelated to what observational mortality misses about those plans. They call this the fallback condition. In a mover design the routes are the instrument, and part (i) is the counterpart of that condition.

For firm wage effects, Kline [2024, Section 5.1] notes that a regression of wages on worker effects and firm characteristics need not recover the projection of firm effects on those characteristics. The part of firm effects that the characteristics miss is related to who works where. With the mean outcome as the characteristic, that regression is the mover regression, and the projection is the cross-sectional coefficient. Parts (ii) to (iv) go from a condition that involves place effects to a condition on flows and location means.

IV.3 Which Routes Are Aligned?

Routes are aligned with every map only under independent pairing, by part (iv). Independent pairing describes movers whose origin and destination are two independent draws from the population of locations. The hypothetical experiment of Section III.1 produces it: random assignment makes the destination independent of the origin. Random assignment also makes what else changes for a person unrelated to where she goes. Independent pairing does not, so the mover regression still needs Assumption 1. People who choose where to move do not produce independent pairing. Distance, housing costs, and family ties connect the destination to the origin, so some pairs of locations are compared often, and most are never compared. Exogenous mobility does not deliver independent pairing either. It restricts the shocks of movers and leaves their routes free.

Alignment with the map at hand asks for less: $J-2$ restrictions, where independent pairing imposes $J(J-1)/2-1$. It can be read off a plot. Let the weights be the shares of moves that involve each location, $\pi_j = a_j / \sum_k a_k$. Then $a_j/\pi_j = 2$ for every location, and by (14) the comparison gap is $g_j = 2\,(\bar{Y}_j - \bar{Y}_j^P)$. Alignment is therefore $\bar{Y}_j^P = (1-\kappa/2)\,\bar{Y}_j + (\kappa/2)\,\mathbb{E}_\pi[\bar{Y}]$: the mean outcome of the partners of a location is a linear function of the mean outcome of the location itself. Locations may exchange movers mostly with similar locations. Alignment asks that this similarity be the same everywhere on the map. For other weights, plot the comparison gap of each location against its mean outcome. The routes are aligned if and only if the points lie on a line. With population weights, the first plot serves as well when every location has the same share of movers.

To measure the distance from alignment, split the comparison gap into the part that the map

explains and the rest,

$$g_j = \kappa\,(\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}]) + g_j^\perp, \qquad \mathrm{Cov}_\pi(g_j^\perp, \bar{Y}_j) = 0, \tag{17}$$

and let

$$r_D := \mathrm{Corr}_\pi(g_j, \bar{Y}_j) \in (0, 1]. \tag{18}$$

Routes are aligned with the map if and only if $g^\perp = 0$, that is, $r_D = 1$. I call r_D the *alignment* of the design. It is computed from flows and location means.

IV.4 For Which Weights?

Alignment refers to the weights π of the question. Routes that are not aligned for these weights may be aligned for other weights. The mover coefficient is then the cross-sectional coefficient of another population of locations. Flows and location means show whether such a population exists. By (14), the comparison gap of a location is positive if $\bar{Y}_j > \bar{Y}_j^P$ and negative if $\bar{Y}_j < \bar{Y}_j^P$, whatever the weights.

Theorem 2 (For which weights)**.** *Let every location be compared with some other location, and let no location have exactly the mean outcome of its partners, so that no comparison gap is zero.*

(i) Existence. *There are weights π for which the routes are aligned with the map if and only if every location with a negative comparison gap has a lower mean outcome than every location with a positive comparison gap.*

(ii) The weights. *In that case, let m be a level strictly between the mean outcomes of the two kinds of locations. The routes are aligned for the weights*

$$\pi_j \;\propto\; a_j \,\frac{\bar{Y}_j - \bar{Y}_j^P}{\bar{Y}_j - m}, \tag{19}$$

and for no weights that are not of this form. Under these weights m is the mean of the map, and $\theta_D = \theta_\pi$ for every vector of place effects.

(iii) No weights. *Otherwise the routes are aligned for no weights, and part (iii) of Theorem 1 holds for every choice of π.*

Weights cannot change the sign of a comparison gap. A location with a positive comparison gap is compared, on balance, with locations of lower mean outcomes, whatever the weights. In a cross-section, a location is high when its mean outcome is above the mean of the map. Alignment needs the two to agree for every location: the locations with positive comparison gaps must be the upper part of the map. Part (i) says that this is also enough.

The condition fails where routes are local and the map is not. Take a prosperous town in a poor region and a poor town in a rich region, and let people move mostly within regions. The first town

has a positive comparison gap, and the second a negative one. If the poor town in the rich region has the higher mean outcome of the two, there are no weights. Movers compare each town with its neighbors, and no way of weighting locations turns these comparisons into those of a cross-section.

Example 1 with $c = \frac{1}{2}$ has both cases. The mean use of the four regions is 1.5, 0.5, -0.5, and -1.5. If more than a third of the moves are across practice styles, the two regions with intensive providers have positive comparison gaps and the other two negative ones, so weights exist. At $\rho = 0.9$ and $m = 0$ they put 36 percent on each of the two regions in the middle of the map and 14 percent on each of the other two. For this population the cross-sectional coefficient is 0.97, which is the mover coefficient. For four equally large regions it is 0.80. If fewer than a third of the moves are across practice styles, the region with intensive providers and residents who seek care rarely has a negative comparison gap. The region with sparing providers and residents who seek care often has a positive one, and its mean use is the lower of the two. There are then no weights.

When weights exist, the mover coefficient answers the question for a population of locations, though not for the one the researcher had in mind. By (19), this population gives more weight to a location the further it is from its partners and the closer it is to the middle of the map. When no weights exist, the two coefficients agree, for any population of locations, only for particular place effects. The mover coefficient remains an average of pair shares with the weights of the routes. For the routes between U.S. commuting zones in Section VII there are no weights.

V When Routes Are Not Aligned

V.1 The Outcomes of Movers

A researcher has more than the coefficient, the map, and the flows. She has the outcome of every mover before and after the move. Under Assumption 1, these outcomes add nothing about place effects.

Proposition 4 (What the outcomes of movers identify)**.** *Let $J \geq 3$, let the routes with movers connect all locations, and let outcomes have no bounds that are known in advance. Under Assumptions 1 and 3, a vector of place effects is consistent with the joint distribution of origins, destinations, and outcomes in the two periods if and only if*

$$\mathcal{L}_D(\psi, \bar{Y}) = \theta_D \, \mathcal{L}_D(\bar{Y}, \bar{Y}). \tag{20}$$

The same holds when Assumption 2 is added. The vector ψ then also gives the differences between the place effects of the whole population.

The data restrict place effects through this equation and through nothing else. Its right side is known: the mover coefficient θ_D is a regression coefficient, and $\mathcal{L}_D(\bar{Y}, \bar{Y})$ is computed from flows and location means. Its left side is a weighted sum of differences in place effects. With three or more locations, the equation determines neither the place effect of a location nor, in general, the difference between two. A place effect is the mean outcome that all movers would have in

a location, and most movers never lived in that location. Assumption 1 ties these means to the data through a covariance, and the mover coefficient is that covariance. If outcomes have known bounds, as ranks do, the bounds add inequalities. These are weak when few movers have lived in any given location.

With Theorem 1, the proposition settles what the data on movers say when routes are not aligned. The place effects that the data allow are those with the observed mover coefficient. By part (iii) of the theorem, they include place effects with every value of the cross-sectional coefficient. The data on movers therefore leave θ_π free, for a researcher with the full panel as for a reader with the coefficient. The bounds of Theorem 3 are sharp for both.

Granting that movers are representative does not help. Assumption 2 adds no equation, because no stayer is observed in a location where she does not live. It makes the proposition a statement about the whole population. The data identify a single combination of place effects because the movers on each route are not representative: who takes which route is left free, up to a covariance.

For firm wage effects, Kline [2024] establishes the cases on either side of this result. With no restriction on who gains from a move, the mean wage change of the workers who move between two firms is an average effect for those workers. Such effects need not fit together into an effect of each firm. If workers do not select on their gains, they do, and every difference between firms is identified. Proposition 4 is the case between the two. It has a restriction on who gains along the map, which is what a mover coefficient needs, and a single identified combination of place effects. Under the stronger assumptions of Section VI, the outcomes of movers identify every place effect.

V.2 The Difference between the Two Coefficients

When $\eta > 0$ and $r_D < 1$, part (i) of Theorem 1 is a product,

$$\theta_D - \theta_\pi = \underbrace{\eta}_{\text{what the map misses}} \times \underbrace{\frac{\sqrt{1 - r_D^2}}{r_D}}_{\text{how far routes are from aligned}} \times \underbrace{\text{Corr}_\pi\left(e_j, g_j^\perp\right)}_{\text{how the two are related}} . \tag{21}$$

The first factor says how much the map misses about place effects. The second says how far the routes are from aligned with the map, and it is computed from flows and location means. The third says whether the two go together. It is far from zero when the locations where the routes are not aligned are also the locations whose place effects the map mispredicts. The two coefficients agree when any of the factors is zero. Neither independent pairing nor a map that predicts place effects exactly is necessary.

The third factor is a correlation, so the first two bound the difference. The first involves place effects, and Assumption 1 gives no way to estimate it. A sign for sorting limits it.

V.3 Nonnegative Sorting and the Bounds

Assumption 4 (Nonnegative sorting)**.** $\mathrm{Cov}_\pi(\psi_j, \bar{\alpha}_j) \geq 0$.

Locations with larger place effects do not have lower composition, on average. In the additive model, their residents do not have lower person components. This is the situation that motivates mover designs. The concern is that good locations look better than they are because of who lives there. The assumption restricts the sign of sorting and not its size. It caps what the map misses: under it, $\eta \leq \sqrt{\theta_\pi(1-\theta_\pi)} \leq \frac{1}{2}$.

The sign of sorting is the researcher's judgment. Under Assumption 1, the data on movers cannot establish the sign, by Proposition 4. Without the sign, or another limit on what the map misses, there are no bounds, by Theorem 1. A researcher who assumes more has a number in place of the bounds. Under the condition of Section VI, a two-way regression estimates every place effect, and the sign of sorting is then estimated and not assumed. Finkelstein et al. [2016] report a positive correlation between patient and place effects for healthcare use across U.S. regions. Cantoni and Pons [2022] report a positive correlation between voter and state effects for turnout (Supplement L.2).

Theorem 3 (Bounds)**.** *Let* $r_D < 1$.

(i) Given what the map misses. *For all place effects,*

$$|\theta_D - \theta_\pi| \;\leq\; \eta\, \frac{\sqrt{1-r_D^2}}{r_D}. \tag{22}$$

Equality holds if and only if e_j *is proportional to* $g_j^\perp$*. For every map, every design, and every value of* θ_π *and* η*, there are place effects that attain the bound.*

(ii) Given the sign of sorting. *Let Assumption 4 hold. The values of* θ_π *that are consistent with the mover coefficient, the map, and the comparison weights are those with* $\theta_- \leq \theta_\pi \leq \theta_+$*, where the bounds are*

$$\theta_\pm = \tfrac{1}{2}\Big[\, 1 - r_D^2\,(1-2\theta_D) \;\pm\; \sqrt{1-r_D^2}\,\sqrt{1-r_D^2\,(1-2\theta_D)^2}\,\Big]. \tag{23}$$

Such values exist if and only if $|1-2\theta_D| \leq 1/r_D$*, and each of them is attained.*

The bounds depend on the data through the coefficient and the alignment. With aligned routes both are the coefficient itself. Less alignment moves them apart in a known way. For a coefficient of one half, θ_π lies between $\frac{1}{2}(1-\sqrt{1-r_D^2})$ and $\frac{1}{2}(1+\sqrt{1-r_D^2})$. That is from 0.40 to 0.60 at $r_D = 0.98$, and from 0.28 to 0.72 at $r_D = 0.90$. An alignment that looks high leaves wide bounds. The bounds are a statement about identification. They are not a confidence interval, and they would remain with any number of movers. In a simulated economy with 256 locations, only the routes change. Across four patterns of routes the mover coefficient ranges from 0.25 to 0.61, around a cross-sectional coefficient of 0.49, and the bounds contain that coefficient throughout (Supplement K).

The assumption can be rejected. With aligned routes, the coefficient lies between zero and one. With routes that are not aligned, it can leave that range by at most $\frac{1}{2}(1/r_D - 1)$. A larger departure contradicts nonnegative sorting. Other assumptions on what the map misses work in the same way: a researcher who is willing to bound η has the bounds of part (i). The logic is that of Manski [2003]: the conclusion is as strong as the assumption on what the map misses. The closest relatives are bounds for instrumental-variables estimates with imperfect instruments [Conley et al., 2012, Nevo and Rosen, 2012]. For firm wage effects, Kline [2024, Section 3.3.2] suggests a different route to partial identification, a bound on how much movers select on their gains. The bounds here keep Assumption 1 and restrict the sign of sorting.

VI The Two-Way Answer

A two-way fixed-effects regression gives a number for the cross-sectional coefficient, whatever the routes are. It rests on a condition at every location. The extra gains of those who arrive in a location must equal the extra gains of those who leave it:

$$\sum_k F_{kj}\, s_{kj} = \sum_k F_{jk}\, s_{jk} \qquad \text{for every location } j. \tag{24}$$

No location then looks better than it is because of who moves in or out. I call the mean change in the outcome of those who move from j to k the *route contrast*. By (2) it is $\psi_k - \psi_j + s_{jk}$. The *two-way regression* is the least-squares regression of the route contrasts on indicators of destination and origin, weighted by flows. Under (24), and if the routes with movers connect all locations, it returns every difference between two place effects. With the map, these give θ_π for any weights. No alignment is needed: the flows affect the precision of the estimates and not what they estimate. The additive model with exogenous mobility implies (24), because it makes the extra gain zero on every route.

When condition (24) fails, the normal equations of least squares say what the two-way regression returns. For firm wage effects, Kline [2024, Section 3.1.2] writes them as a system in the mean wage changes on each route, with the Laplacian of the mobility network as its matrix. In the notation here, the estimates m_j solve

$$\sum_k a_{jk}\,(m_j - m_k) = \sum_k F_{kj}\, G_{kj} - \sum_k F_{jk}\, G_{jk} \qquad \text{for every location } j, \tag{25}$$

where G_{jk} is the route contrast. With $G_{jk} = \psi_k - \psi_j + s_{jk}$, the estimates are $m_j = \psi_j + \chi_j$ up to a constant, where

$$\sum_k a_{jk}\,(\chi_j - \chi_k) = \nu_j, \qquad \nu_j := \sum_k F_{kj}\, s_{kj} - \sum_k F_{jk}\, s_{jk}. \tag{26}$$

I call ν_j the *net extra gain* of location j: the extra gain of those who arrive minus the extra gain of those who leave. Condition (24) is $\nu_j = 0$ for every j. It holds if and only if $\mathrm{Cov}(s_i, \Delta x_i) = 0$ among movers for every attribute x of locations. Assumption 1 is the same statement for the map alone:

$\sum_j \bar{Y}_j \nu_j = 0$, which is $\mathcal{L}_D(\chi, \bar{Y}) = 0$.

Assumption 1 asks that extra gains be unrelated to the change in the mean outcome. Condition (24) asks this of the change in every attribute of locations: climate, size, a specialized hospital, a school, family nearby. People move for many reasons, and they may gain from what they moved for. When they do, the two-way regression returns the place effect plus χ_j, an attribute of locations along which movers gain more than the average mover would. Under Assumption 1 this attribute is unrelated to the map along the routes, so the mover coefficient is unaffected. The estimated effect of each location is not.

Proposition 5 (Two-way estimates under Assumption 1)**.** *Let Assumption 3 hold, and let the routes with movers connect all locations.*

(i) The two-way answer. *The cross-sectional coefficient computed from the two-way estimates is* $\theta_\pi + \mathrm{Cov}_\pi(\chi_j, \bar{Y}_j) / \mathrm{Var}_\pi(\bar{Y}_j)$. *Under Assumption 1, the absolute value of the second term is at most*

$$\frac{\mathrm{sd}_\pi(\chi_j)}{\mathrm{sd}_\pi(\bar{Y}_j)} \sqrt{1 - r_D^2},$$

and this value is attained. The second term is zero for every χ *that Assumption 1 allows if and only if routes are aligned.*

(ii) No test. *Under the conditions of Proposition 4, any route contrasts are reproduced by some place effects and extra gains that satisfy* (24).

Estimating place effects is therefore not a way around alignment. Under Assumption 1 alone, the cross-sectional coefficient computed from two-way estimates is right when routes are aligned, and it can be wrong when they are not. Its error is at most the variation of χ across locations relative to that of the map, times $\sqrt{1 - r_D^2}$. With an alignment of 0.9 and a χ that varies a quarter as much as the map, the error can reach 0.11. With an alignment of 0.77, a χ that varies a sixth as much as the map is enough for the same error.

No pattern in the route contrasts speaks for or against condition (24). By part (ii), every set of them is consistent with it.[8] The reason is in (25): the residuals of the two-way regression satisfy the condition by construction, so the regression imposes it and cannot check it. The evidence that mover studies report is evidence along the map. Outcomes before the move show no trend that is related to the gap, and changes are symmetric for moves to locations with higher and with lower mean outcomes. Both concern how the changes of movers vary with the gap in the map, which is what Assumption 1 restricts.

Stronger assumptions can be tested. The additive model with exogenous mobility makes the extra gain zero on every route, so that route contrasts add up around cycles of moves. Kline [2024] shows that these sums are all that the model restricts in the mean changes of movers. He finds that wage changes between Italian firms come close to satisfying them, without doing so exactly. In the

[8] For outcomes with known bounds, the inequalities of Section V.1 apply to the two-way estimates as well.

additive model, two workers who make the same move also keep the gap between them, apart from transitory changes. Sorkin and Warwar [2026] reject this for firms in Brazil and Italy, and they reject the weaker hypothesis that workers do not move on gains of their own. Such comparisons show how extra gains differ between two movers on a route. Condition (24) is about average extra gains, which cancel in the comparison. It asks only that they net out at every location.

Example 2 (Locations that suit some people)**.** Locations differ in an attribute that raises the outcome of half of the people by t and lowers that of the other half by t. A region may organize care around specialists, or a neighborhood may have schools with a particular program. On average the attribute does nothing, so it is no part of the place effect, and it is unrelated to everything else about a location. A mover chooses the attribute of her destination with an eye on her own gain. Let one region in four have the attribute, and let a quarter of all movers arrive in such a region. Let selection be strong enough that the attribute suits 83 percent of them, against half of all movers.

Extra gains are independent of the gap in the map, so Assumption 1 holds, and the mover coefficient is what it would be without any selection. Condition (24) fails. The two-way regression credits such a region with an effect of $0.22\,t$ that the average mover does not experience (Supplement C). The attribute is unrelated to the map, so these estimates still give the right cross-sectional coefficient here. They give the wrong effect for each such region.

The literature draws the same line between a condition along the map and a condition for every location. Chetty and Hendren [2018b] require an orthogonality condition for each of 3,000 counties, where the average exposure effect of their first paper requires a single condition that holds on average. Deryugina and Molitor [2021] make the same point for health. A destination chosen because relatives live there does no harm if the presence of relatives is unrelated to the characteristic of destinations that is studied. For firm wage effects, Kline [2024] separates the effect of a move on the workers who make it from effects that rank all firms on one scale. The second requires that workers do not select on their gains.

Other assumptions give more than bounds and less than every place effect. Split the movers into a few groups that travel different kinds of routes, and let Assumption 1 hold in each group as well as among all movers. Each group then identifies a coefficient of its own. If the map is a combination of the comparison gaps of the groups, the cross-sectional coefficient is a weighted sum of the group coefficients. The weights are computed from flows and location means. Otherwise the groups tighten the bounds. In Example 1, the movers across practice styles and the movers across patient populations are two such groups. Their coefficients are one and zero, and together they give θ_π for every ρ (Proposition 7 in the Supplement). Under (24), a mover regression can also be aimed at θ_π, with an instrument built from flows and the map (Supplement E).

VII Childhood Moves between U.S. Commuting Zones

For childhood moves between commuting zones in the United States, public files contain every piece of the argument [Chetty and Hendren, 2018b]. They give the map, the routes, an estimate

for each pair of zones, and a two-way estimate of the effect of each zone. The map is the mean income rank at age 26 of permanent residents whose parents are at the 25th percentile. The routes are those of the children in the authors' sample of movers. For each ordered pair of zones with at least 25 such children, the files give the number of children and an estimate for the pair. The estimate is how the outcome varies with the age at the move, which measures a year of childhood in the destination against a year in the origin.[9]

These are 1.4 million children on 7,078 pairs of zones, between 595 zones. The listed pairs account for 75 percent of the moves across zones in the authors' sample, as they report. The coefficient, the alignment, and the two-way estimates below all use the listed pairs, so they describe the same routes. Effects are in units of twenty years of childhood.

I use two sets of comparisons. One is origin against destination. The other is two destinations reached from the same origin, which is what an age-at-move regression with an intercept for each origin compares. Locations are weighted in two ways. Population weights use the population of each zone. Mover weights use how often the comparisons involve each zone.

The alignment is computed from these counts and the map. Resampling the children across the listed pairs, with the map held fixed, gives it a standard error below 0.001. The sampling error of the bounds therefore comes from the coefficient. Table 1 reports, below each pair of bounds, a confidence interval that covers the cross-sectional coefficient with a probability of 95 percent [Imbens and Manski, 2004]. Supplement J has a split-sample estimate of the alignment for samples with few movers, and a test of alignment.

The coefficient. The estimates for pairs give the mover coefficient without any place effects. Across all zones it is 0.61 for origin against destination, with a standard error of 0.07, and 0.47 for two destinations, with a standard error of 0.10. Under Assumption 1, each is an average of pair shares over the pairs of zones that the children compare.

The flows. The routes are not aligned with the map. The alignment is between 0.77 and 0.96 in Table 1, and it is near one in a single row: the large zones, comparisons between destinations, and mover weights. No other weights on zones align the routes either. For origin against destination across all zones, at least 106 of the 595 zones have a comparison gap of the wrong sign at every level of the map. For two destinations among the 211 large zones, at least 20 do (Theorem 2). For every population of zones, the two coefficients therefore agree only for particular place effects.

A reader who assumes that sorting is not negative has the bounds of Theorem 3. Across all zones with population weights, the bounds are 0.25 and 0.88 for origin against destination, and 0.19 and 0.78 for two destinations. With the sampling error of the coefficient, the confidence intervals run from 0.17 to 0.94 and from 0.09 to 0.88. The tightest of the eight are 0.29 and 0.57, with a confidence interval from 0.11 to 0.75. An alignment of 0.84 looks high, and it leaves bounds of 0.25

[9]The map and the two-way estimates are in online data table 3 of Chetty and Hendren [2018b], at `https://opportunityinsights.org/paper/neighborhoodsii/`. The routes and the estimates for pairs are in online data table 5 of the replication files of the two papers. The authors publish the estimates for pairs to facilitate research with other models of neighborhood effects. Supplement L gives the construction.

Table 1: Childhood Moves between U.S. Commuting Zones

	Pair estimates, flows, and the map			With two-way estimates of place effects		
Comparisons, weights on locations	r_D	θ_D	Bounds	θ_π	θ_D	Difference
A. All 595 commuting zones						
Origin and destination, population	0.77	0.609	[0.25, 0.88]	0.747	0.625	−0.121
		(0.074)	(0.17, 0.94)	(0.057)	(0.065)	(0.046)
Origin and destination, mover weights	0.81	0.609	[0.28, 0.86]	0.707	0.625	−0.081
		(0.074)	(0.19, 0.93)	(0.053)	(0.065)	(0.039)
Two destinations, population	0.81	0.474	[0.19, 0.78]	0.740	0.602	−0.138
		(0.096)	(0.09, 0.88)	(0.057)	(0.059)	(0.039)
Two destinations, mover weights	0.91	0.474	[0.27, 0.68]	0.674	0.602	−0.072
		(0.096)	(0.14, 0.82)	(0.054)	(0.059)	(0.024)
B. The 211 commuting zones with at least 250,000 people						
Origin and destination, population	0.84	0.532	[0.25, 0.79]	0.470	0.560	0.091
		(0.094)	(0.14, 0.90)	(0.071)	(0.078)	(0.042)
Origin and destination, mover weights	0.88	0.532	[0.29, 0.76]	0.504	0.560	0.056
		(0.094)	(0.17, 0.88)	(0.068)	(0.078)	(0.036)
Two destinations, population	0.85	0.419	[0.18, 0.70]	0.470	0.521	0.051
		(0.119)	(0.05, 0.85)	(0.071)	(0.072)	(0.040)
Two destinations, mover weights	0.96	0.419	[0.29, 0.57]	0.529	0.521	−0.008
		(0.119)	(0.11, 0.75)	(0.069)	(0.072)	(0.020)

Notes. The map is the mean income rank at age 26 of permanent residents with parents at the 25th percentile. Routes are the children in the sample of movers of Chetty and Hendren [2018b], for ordered pairs of commuting zones with at least 25 such children. "Origin and destination" uses the comparison weights of Proposition 1, and "two destinations" those of the age-at-move design in Section III.2. Eight zones are compared with no other destination and drop from the latter in panel A. Population weights are the population in 2000. Mover weights are how often the comparisons involve each zone. In the left block, θ_D is computed from the published estimates for pairs of zones, without place effects. For origin and destination it is the regression of the pair estimate on the gap in the map between destination and origin, weighted by the number of children. For two destinations it is the regression of the pair estimate on the map of the destination with an intercept for each origin. Its standard error treats the estimates for pairs as independent. The bounds in brackets are those of Theorem 3 at this coefficient and r_D, and they assume nonnegative sorting. Below it, in parentheses, is a 95 percent confidence interval for θ_π [Imbens and Manski, 2004]. It adds the sampling error of the coefficient; that of the alignment is below 0.001. The right block uses the published two-way estimates of the effect of each zone. Its θ_D is the coefficient that the same comparisons give with these estimates, and the difference is this coefficient minus θ_π. Its standard errors reflect the sampling error of the published estimates only, and treat the estimates of different zones as independent. These standard errors are computed from the standard errors that the files label raw. With the bootstrapped standard errors of the same files they are up to 13 percent larger. Effects are in units of twenty years of childhood.

and 0.79.

The pairs that are not listed. A quarter of the moves are on pairs of zones with fewer than 25 children, which the files do not list. The results apply to the listed pairs as they stand, because the coefficient and the alignment come from the same routes. The alignment of all moves may be higher. Suppose the missing moves were paired independently, in proportion to the weights of origin and destination. Let them fall on the pairs that are not listed, with at most 24 children on each. The alignment for origin against destination with population weights would then be 0.92 in place of 0.77, and between 0.92 and 0.99 in the other rows of the table. Independent pairing is aligned with every map, so this case is favorable to alignment. In every row some zones keep a comparison gap of the wrong sign, from 2 of 211 to 55 of 595. No weights align the routes in this case either. At the coefficient of the listed pairs, the bounds of the first row would be 0.40 and 0.78.

The two-way assumption. The published effects of each zone come from a two-way regression of

the estimates for pairs, and they are estimates of place effects under a condition of the form (24). They put the cross-sectional coefficient at 0.75 across all zones and at 0.47 across the large zones, with population weights.[10] Taking them as the place effects, the routes between all zones give a mover coefficient of 0.63. The difference of 0.12 between the two coefficients has a standard error of 0.05. The condition thus replaces bounds of 0.25 and 0.88 by a number. In every row of the table, the two-way answer lies within the bounds. Across the rows, the two coefficients computed from the place effects differ by between -0.14 and 0.09.

These numbers hold under that condition, and the data cannot say whether it holds. By Proposition 5, the two-way answer is exposed to attributes along which movers gain. At an alignment of 0.77, one that varies a sixth as much across zones as the map moves the answer by up to 0.11. Trends before the move and the symmetry of gains and losses speak to selection along the map, and not to such an attribute. For two destinations, the place effects give a mover coefficient of 0.60 across all zones, where the estimates for pairs give 0.47. In the additive model the two estimate the same number.

The sign of sorting. The bounds assume that sorting is not negative. A reader who accepts the two-way condition does not need the bounds, and for this reader the sign is estimated. Sampling error is about two thirds of the variance of the published effects across zones, and the files give two standard errors for each zone. Across all zones, the slope of composition on place effects is -0.21 with the standard errors that the files label raw. The standard error of this slope is 0.16. With the bootstrapped standard errors the slope is positive. Across the 211 large zones the slope is positive with both, 0.05 and 0.33. The bounds therefore assume a sign that these estimates do not settle across all zones. A reader who does not accept the two-way condition has no estimate of the sign, and the bounds rest on her judgment. Table 5 in the Supplement reports what the map misses and its relation to the routes with both standard errors.

The numbers of this section do not reproduce a published estimate. The counts of movers pool all parental incomes, while the map and the estimates are for parents at the 25th percentile. The standard errors treat estimates that come from one regression as independent, so I do not single out any difference as distinguishable from zero. Supplement L adds the migration flows of all tax filers, which give a similar alignment. It also has results for parents at the 75th percentile and for groups of comparisons.

VIII Conclusion

A mover design answers a question about a map of locations with a regression on the people who travel across it. The question is a moment of levels, with one term for each location. The coefficient is a moment of differences, with one term for each pair of locations that movers connect, and the routes of movers are its weights.

[10] The weights matter. With weights equal to the precision of the estimates, measured by their bootstrapped standard errors, the slope across all zones is 0.032 per year of childhood, which is the estimate that Chetty and Hendren [2018b] report. With population weights it is 0.037 per year, or 0.75 over twenty years.

The coefficient answers the question, whatever the place effects are, exactly when routes are aligned with the map. Flows and location means are enough to check this, and to say whether the routes are aligned for any other population of locations. The routes of an experiment that assigns locations at random are aligned with every map, for the weights of the experiment. The routes that people choose need not be aligned with the map at hand, for the weights of the question or for any other.

What a reader can claim depends on what she has and on what she assumes. With the coefficient and the assumption of the design, she has the share of place in the gaps between the locations that movers travel between. With the flows and the map, she knows whether this share answers the question, and for which population of locations. When this share does not answer the question, a sign for sorting gives bounds. For adult income and childhood moves between U.S. commuting zones, the alignment is 0.77, and the bounds on the answer are 0.25 and 0.88. With the assumption that extra gains cancel at every location, she has a number, 0.75 in the same data. That assumption asks that extra gains be unrelated to every attribute of locations. Data on movers cannot test it, and its failure matters more the further the routes are from aligned. A mover coefficient can therefore be reported with the alignment of its routes and the bounds that the two imply, so that a reader sees what each assumption adds.

A Proofs

Sampling frame. The frame is written for the additive model (4). In the general setting of Section II.1, a history is the two locations and the two outcomes of an individual, and the same definitions apply. Stack the $2N$ observations as

$$\mathbf{Y} = \mathsf{U}_N \boldsymbol{\alpha} + \mathsf{Z}_N \boldsymbol{\psi} + \boldsymbol{\varepsilon},$$

where U_N and Z_N are the incidence matrices of individuals and of locations. Let $\mathcal{F}_N$ be the information in $(\boldsymbol{\alpha}, \mathsf{Z}_N, \boldsymbol{\psi})$. Given $\mathcal{F}_N$, this is the two-way model with fixed effects, and exogenous mobility is $\mathbb{E}[\varepsilon_{it} \mid \mathcal{F}_N] = 0$. The vector $\boldsymbol{\psi}$ is a fixed parameter. The histories $(\alpha_i, j(i,1), j(i,2), \varepsilon_{i1}, \varepsilon_{i2})$ are independent and identically distributed across individuals, and outcomes have bounded second moments. The two shocks of an individual may be correlated.

Population objects are defined under the distribution of one history. The size of location j in period t is $\Pr\{j(i,t) = j\}$, and its share of movers is $\Pr\{d(i) \neq o(i) \mid j(i,t) = j\}$. Assumption 3 says that neither depends on t, and that $\mathbb{E}[\alpha_i \mid o(i) = j,\, d(i) \neq j] = \mathbb{E}[\alpha_i \mid d(i) = j,\, o(i) \neq j]$ for every location with movers. The flows are $F_{jk} = \Pr\{o(i) = j,\, d(i) = k \mid d(i) \neq o(i)\}$. Composition is $\bar{\alpha}_j = \mathbb{E}[\alpha_i \mid j(i,t) = j]$, and the map is $\bar{Y}_j = \mathbb{E}[Y_{it} \mid j(i,t) = j]$. The coefficient θ_D is the probability limit of the least-squares coefficient of (8) as $N \to \infty$ with J fixed.

Notation. Attributes of locations are vectors in $\mathbb{R}^J$. Let L_D be the symmetric matrix with off-diagonal entries $-a_{jk}$ and diagonal entries $\sum_k a_{jk}$. Let $\Pi = \operatorname{diag}(\pi)$ and $M_\pi = \Pi - \pi\pi'$. Then

$\mathcal{L}_D(f,h) = f'L_D h$ and $\mathrm{Cov}_\pi(f_j, h_j) = f'M_\pi h$. Both matrices are positive semidefinite and annihilate constant vectors. The comparison gap is $g = \Pi^{-1}L_D\bar{Y}$. Moments with the subscript mov are taken over movers.

Lemma 1 (Covariances in pairs)**.** *For attributes f and h of locations,* $\mathrm{Cov}_\pi(f_j, h_j) = \sum_{j<k} \pi_j\pi_k\,(f_j - f_k)(h_j - h_k)$.

Proof. $\sum_{j<k} \pi_j\pi_k(f_j - f_k)(h_j - h_k) = \frac{1}{2}\sum_{j,k} \pi_j\pi_k(f_j - f_k)(h_j - h_k) = \mathbb{E}_\pi[f_j h_j] - \mathbb{E}_\pi[f_j]\,\mathbb{E}_\pi[h_j]$. □

Proof of Proposition 1. By (2), $\mathbb{E}[\Delta Y_i \mid o(i), d(i)] = \Delta\psi_i + s_i$ among movers. Under Assumption 3 the map is the same in both periods, so the regressor $\Delta\bar{Y}_i = \Delta\psi_i + \Delta\bar{\alpha}_i$ is a function of the two locations of i. Hence $\mathrm{Cov}_{\mathrm{mov}}(\Delta Y_i, \Delta\bar{Y}_i) = \mathrm{Cov}_{\mathrm{mov}}(\Delta\psi_i, \Delta\bar{Y}_i) + \mathrm{Cov}_{\mathrm{mov}}(s_i, \Delta\bar{Y}_i)$. The last term is zero by Assumption 1, and the coefficient is (9). In the additive model with exogenous mobility, $\Delta Y_i = \Delta\psi_i + \Delta\varepsilon_i$ and $s_i = 0$. For attributes f and h,

$$\mathbb{E}_{\mathrm{mov}}[\Delta f_i\,\Delta h_i] = \sum_{j\neq k} F_{jk}(f_k - f_j)(h_k - h_j) = \sum_{j<k}(F_{jk} + F_{kj})(f_j - f_k)(h_j - h_k) = \mathcal{L}_D(f,h).$$

Each location has the same size and the same share of movers in both periods, so it receives as many movers as it loses, and $\mathbb{E}_{\mathrm{mov}}[\Delta f_i] = \sum_j f_j\big(\sum_k F_{kj} - \sum_k F_{jk}\big) = 0$. Hence $\mathrm{Cov}_{\mathrm{mov}}(\Delta f_i, \Delta h_i) = \mathcal{L}_D(f,h)$, which gives (11). Only the weights a_{jk} enter. □

Proof of Proposition 2. (i) Pairs with $\bar{Y}_j = \bar{Y}_k$ contribute zero to both sums in (11). For the others, $(\psi_j - \psi_k)(\bar{Y}_j - \bar{Y}_k) = (\bar{Y}_j - \bar{Y}_k)^2 S_{jk}$. For θ_π, apply Lemma 1 to the numerator and the denominator of (6). (ii) By (i), θ_D is an average of pair shares with nonnegative weights, so it lies between the smallest and the largest. A design with all its weight on one pair has that pair's share as its coefficient. A design that divides its weight between the pairs with the smallest and the largest share attains every value in between. □

Proof of Proposition 3. By symmetry of the weights,

$$\mathcal{L}_D(f, \bar{Y}) = \frac{1}{2}\sum_{j\neq k} a_{jk}(f_j - f_k)(\bar{Y}_j - \bar{Y}_k) = \sum_j f_j \sum_{k\neq j} a_{jk}(\bar{Y}_j - \bar{Y}_k) = \mathbb{E}_\pi[f_j g_j].$$

Taking f constant gives $\mathbb{E}_\pi[g_j] = 0$, so $\mathbb{E}_\pi[f_j g_j] = \mathrm{Cov}_\pi(f_j, g_j)$. Apply this to the numerator and the denominator of (11). □

Proof of Theorem 1. (i) By (16) and Proposition 3, $\mathrm{Cov}_\pi(\psi_j, g_j) = \theta_\pi\,\mathrm{Cov}_\pi(\bar{Y}_j, g_j) + \mathrm{Cov}_\pi(e_j, g_j)$. Divide by $\mathrm{Cov}_\pi(\bar{Y}_j, g_j) = \mathcal{L}_D(\bar{Y}, \bar{Y}) > 0$.

(ii) and (iii). Use the decomposition (17). Both g and $\bar{Y} - \mathbb{E}_\pi[\bar{Y}]$ have mean zero under π, so $g^\perp$ does too, and $\kappa = \mathcal{L}_D(\bar{Y}, \bar{Y})\,/\,\mathrm{Var}_\pi(\bar{Y}_j) > 0$. Since e is uncorrelated with the map, $\mathrm{Cov}_\pi(e_j, g_j) =$

$\mathrm{Cov}_\pi(e_j, g_j^\perp)$. If $g^\perp = 0$, the two coefficients coincide by (i). If $g^\perp \neq 0$, let

$$\psi = t'\,\bar{Y} + b\,g^\perp, \qquad b = \frac{(t - t')\ \mathrm{Cov}_\pi(\bar{Y}_j, g_j)}{\mathrm{Var}_\pi(g_j^\perp)}.$$

Then $\theta_\pi = t'$, because $g^\perp$ is uncorrelated with the map. The residual is $e = b\,g^\perp$, so by (i) $\theta_D = t' + b\,\mathrm{Var}_\pi(g_j^\perp)/\,\mathrm{Cov}_\pi(\bar{Y}_j, g_j) = t$. By Lemma 2, these place effects, the map, and the comparison weights are generated by an additive economy with exogenous mobility that satisfies Assumption 3.

(iv) If $a_{jk} = \kappa\pi_j\pi_k$, then $g_j = \kappa\sum_k \pi_k(\bar{Y}_j - \bar{Y}_k) = \kappa(\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}])$ for every map. Conversely, alignment with the map y means $L_D y = \kappa(y) M_\pi y$. With two locations there is one pair and nothing to prove. Otherwise take y and y' such that y, y' and the constant vector are linearly independent. Then

$$\kappa(y)M_\pi y + \kappa(y')M_\pi y' = L_D(y + y') = \kappa(y + y')\,M_\pi(y + y').$$

The kernel of M_π is the constants, so $M_\pi y$ and $M_\pi y'$ are linearly independent, and $\kappa(y) = \kappa(y + y') = \kappa(y')$. A map $y'' = \lambda y + \mu\mathbf{1}$ with $\lambda \neq 0$ has $L_D y'' = \lambda L_D y = \kappa(y) M_\pi y''$, so $\kappa(y'') = \kappa(y)$ as well. Hence κ is the same for all maps, $L_D = \kappa M_\pi$, and the off-diagonal entries give $a_{jk} = \kappa\pi_j\pi_k$. □

Proof of Theorem 2. Let $v_j := \sum_k a_{jk}(\bar{Y}_j - \bar{Y}_k) = a_j(\bar{Y}_j - \bar{Y}_j^P)$, so that $g_j = v_j/\pi_j$. Because $a_{jk} = a_{kj}$, $\sum_j v_j = 0$. By assumption $v_j \neq 0$ for every j, so some v_j are positive and some are negative, and $\mathcal{L}_D(\bar{Y}, \bar{Y}) > 0$.

(i) and (ii). By Theorem 1(ii), the routes are aligned for the weights π if and only if $v_j = \kappa\,\pi_j\,(\bar{Y}_j - m_\pi)$ for all j, with $\kappa > 0$ and $m_\pi = \mathbb{E}_\pi[\bar{Y}]$. Suppose this holds. Since $\pi_j > 0$ and $v_j \neq 0$, the difference $\bar{Y}_j - m_\pi$ has the sign of v_j. Every location with a negative comparison gap therefore has a mean outcome below m_π, and every location with a positive one has a mean outcome above it. So the condition in (i) holds, m_π lies strictly between the two kinds of locations, and $\pi_j = v_j/\{\kappa(\bar{Y}_j - m_\pi)\}$, which is (19). Conversely, let the condition in (i) hold, and take m strictly between the two kinds of locations. Then $v_j/(\bar{Y}_j - m) > 0$ for every j. Let $\kappa := \sum_j v_j/(\bar{Y}_j - m)$ and $\pi_j := v_j/\{\kappa(\bar{Y}_j - m)\}$. These weights are positive and sum to one. They give $\sum_j \pi_j(\bar{Y}_j - m) = \sum_j v_j/\kappa = 0$, so $m = \mathbb{E}_\pi[\bar{Y}]$, and $g_j = v_j/\pi_j = \kappa\,(\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}])$. The routes are aligned, and $\theta_D = \theta_\pi$ by Theorem 1(ii).

(iii) If the condition in (i) fails, then by the first part of the argument the comparison gap is proportional to the map under no weights, and Theorem 1(iii) applies to each choice of π. □

Proof of Proposition 4. Necessity is Proposition 1. For sufficiency, let G_{jk} be the mean change in the outcome on the route from j to k, let ψ satisfy the equation, and put $s_{jk} = G_{jk} - (\psi_k - \psi_j)$. Under Assumption 3 each location receives as many movers as it sends, so $\mathbb{E}_{\mathrm{mov}}[\Delta\bar{Y}_i] = 0$ and the coefficient of regression (8) is $\theta_D = \sum_{j\neq k} F_{jk}G_{jk}(\bar{Y}_k - \bar{Y}_j)/\mathcal{L}_D(\bar{Y}, \bar{Y})$. Hence $\mathrm{Cov}_{\mathrm{mov}}(s_i, \Delta\bar{Y}_i) = \sum_{j\neq k} F_{jk}\,s_{jk}\,(\bar{Y}_k - \bar{Y}_j) = \theta_D\,\mathcal{L}_D(\bar{Y}, \bar{Y}) - \mathcal{L}_D(\psi, \bar{Y}) = 0$, and Assumption 1 holds. It remains to give every mover an outcome in each location where she did not live, such that ψ is the vector of place

effects up to a constant. Let μ^1_{jk} and μ^2_{jk} be the mean outcomes of the movers on the route from j to k before and after the move. Let $a_l = \sum_k a_{lk}$ be the share of movers who lived in location l in one of the two periods, and let $o_l := \sum_k F_{lk}\,\mu^1_{lk} + \sum_k F_{kl}\,\mu^2_{kl}$. The data show both. If $a_l = a_{l'} = 1$ for two locations, every move is between these two, and the routes cannot connect $J \geq 3$ locations. So $a_l < 1$ for every location except at most one, say h. Let $c := o_h - \psi_h$ if there is such a location, and $c := 0$ otherwise. A mover keeps her observed outcomes in the two locations where she lived. In every other location l, let her outcome be $(\psi_l + c - o_l)/(1 - a_l)$. No bound that is known in advance excludes this value. The mean outcome that all movers would have in l is then

$$o_l + (1 - a_l)\,\frac{\psi_l + c - o_l}{1 - a_l} = \psi_l + c,$$

and in h it is $o_h = \psi_h + c$. The place effects are therefore ψ up to a constant. The distribution of origins, destinations, and observed outcomes is unchanged. So the extra gains are the s_{jk} above, and Assumption 3 holds because it holds in the data. The differences between place effects are therefore restricted by the equation and by nothing else.

Assumption 2 adds no restriction. It requires that the mean outcome that stayers would have in location j, minus ψ_j, be the same for all j. The data contain the outcome of a stayer in her own location, and in no other. Suppose first that stayers live in two or more locations. Then in every location j a positive share of stayers does not live. Their outcomes in j are not observed and have no known bounds, so the mean outcome of all stayers in j can take any value. Suppose instead that all stayers live in one location. Their mean outcome there is observed and fixes the constant, and their mean outcome in every other location can take any value. In both cases the assumption can be met for every ψ that satisfies the equation. □

Proof of Theorem 3. (i) By Theorem 1(i) and the Cauchy–Schwarz inequality,

$$|\theta_D - \theta_\pi| = \frac{|\operatorname{Cov}_\pi(e_j, g_j^\perp)|}{\operatorname{Cov}_\pi(\bar{Y}_j, g_j)} \leq \frac{\operatorname{sd}_\pi(e_j)\,\operatorname{sd}_\pi(g_j^\perp)}{\operatorname{Cov}_\pi(\bar{Y}_j, g_j)},$$

with equality if and only if e is proportional to $g^\perp$. Now $\operatorname{sd}_\pi(g_j^\perp) = \operatorname{sd}_\pi(g_j)\sqrt{1 - r_D^2}$ and $\operatorname{Cov}_\pi(\bar{Y}_j, g_j) = r_D\,\operatorname{sd}_\pi(g_j)\,\operatorname{sd}_\pi(\bar{Y}_j)$, which gives (22). The bound is attained by $\psi = \theta_\pi \bar{Y} \pm \eta\,\operatorname{sd}_\pi(\bar{Y}_j)\,g^\perp / \operatorname{sd}_\pi(g_j^\perp)$, and Lemma 2 supplies an economy with these place effects.

(ii) Write $\theta = \theta_\pi$ and $\Lambda = \sqrt{1 - r_D^2}/r_D$, so that $1 + \Lambda^2 = 1/r_D^2$. Let $d = g^\perp / \operatorname{Cov}_\pi(\bar{Y}_j, g_j)$, so that $\theta_D - \theta = \operatorname{Cov}_\pi(e_j, d_j)$ and $\operatorname{sd}_\pi(\bar{Y}_j)\,\operatorname{sd}_\pi(d_j) = \Lambda$. By (16), $\operatorname{Var}_\pi(\psi_j) = (\theta^2 + \eta^2)\operatorname{Var}_\pi(\bar{Y}_j)$ and $\operatorname{Cov}_\pi(\psi_j, \bar{Y}_j) = \theta\,\operatorname{Var}_\pi(\bar{Y}_j)$. Assumption 4 is $\operatorname{Cov}_\pi(\psi_j, \bar{Y}_j - \psi_j) \geq 0$, that is, $\eta^2 \leq \theta(1 - \theta)$.

Step 1. Fix θ. A residual e is admissible if $\mathbb{E}_\pi[e_j] = 0$, $\operatorname{Cov}_\pi(e_j, \bar{Y}_j) = 0$, and $\operatorname{Cov}_\pi(e_j, d_j) = \theta_D - \theta$. The first two conditions define a space of dimension $J - 2$ that contains $d \neq 0$. The admissible residual of smallest variance is a multiple of d, and it has $\eta = |\theta_D - \theta|/\Lambda$. If $J \geq 4$, adding a vector of that space that is uncorrelated with d keeps the residual admissible and raises η to any larger value. If $J = 3$, the space is spanned by d and $\eta = |\theta_D - \theta|/\Lambda$.

Step 2. Let $q(\theta) = \theta^2 + (\theta_D - \theta)^2/\Lambda^2$. By Step 1, θ is consistent with the mover coefficient, the map, the comparison weights, and the assumption if and only if $q(\theta) \leq \theta$, that is,

$$(1+\Lambda^2)\,\theta^2 - (2\theta_D + \Lambda^2)\,\theta + \theta_D^2 \leq 0.$$

The roots are

$$\frac{2\theta_D + \Lambda^2 \pm \Lambda\sqrt{\Lambda^2 + 4\theta_D(1-\theta_D)}}{2\,(1+\Lambda^2)}.$$

Multiplying the numerator and the denominator by r_D^2 and using $r_D^2\Lambda^2 = 1 - r_D^2$ gives (23). The roots are real if and only if $1 - r_D^2(1-2\theta_D)^2 \geq 0$. Each θ between the roots is attained by the residual of smallest variance in Step 1, and Lemma 2 supplies the economy. At $\theta_D = 1 + \frac{1}{2}(1/r_D - 1)$ and at $\theta_D = -\frac{1}{2}(1/r_D - 1)$ the condition $|1 - 2\theta_D| \leq 1/r_D$ holds with equality, which gives the largest departure of the coefficient from the range between zero and one. □

Proof of Proposition 5. The normal equations. Write $G_{jk} = \psi_k - \psi_j + s_{jk}$ for the route contrast. The two-way regression chooses $m \in \mathbb{R}^J$ to minimize $\sum_{j\neq k} F_{jk}\,(G_{jk} - m_k + m_j)^2$. The first-order condition for m_j is $\sum_k F_{kj}\,(G_{kj} - m_j + m_k) = \sum_k F_{jk}\,(G_{jk} - m_k + m_j)$, which is (25), that is, $L_D m = b$ with $b_j = \sum_k F_{kj}G_{kj} - \sum_k F_{jk}G_{jk}$. The objective is convex, and L_D has rank $J-1$ when the routes connect the locations, so the solution is unique up to a constant. With the route contrasts above, $b = L_D\psi + \nu$, so $m = \psi + \chi$ with $L_D\chi = \nu$, which is (26). Under (24), $\nu = 0$ and the solutions are ψ plus a constant.

Each location receives as many movers as it sends, so $\mathbb{E}_{\text{mov}}[\Delta x_i] = 0$ for every attribute x of locations, and $\text{Cov}_{\text{mov}}(s_i, \Delta x_i) = \sum_{j\neq k} F_{jk}\,s_{jk}\,(x_k - x_j) = \sum_j x_j\,\nu_j$. The sum is zero for every x if and only if $\nu = 0$. With $x = \bar{Y}$, Assumption 1 is $\bar{Y}'\nu = 0$, and $\bar{Y}'\nu = \bar{Y}'L_D\chi = \mathcal{L}_D(\chi, \bar{Y})$.

(i) The two-way estimates are $\psi + \chi$, which gives the first statement. By Proposition 3, $\mathcal{L}_D(\chi, \bar{Y}) = \text{Cov}_\pi(\chi_j, g_j)$, so under Assumption 1 χ is uncorrelated with the comparison gap. Write $\bar{Y}_j = \text{constant} + b\,g_j + \bar{Y}_j^\perp$ with $\text{Cov}_\pi(\bar{Y}_j^\perp, g_j) = 0$. Then $\text{Var}_\pi(\bar{Y}_j^\perp) = (1 - r_D^2)\,\text{Var}_\pi(\bar{Y}_j)$ and $\text{Cov}_\pi(\chi_j, \bar{Y}_j) = \text{Cov}_\pi(\chi_j, \bar{Y}_j^\perp)$, and the bound is the Cauchy–Schwarz inequality. If routes are aligned, $\bar{Y}^\perp = 0$ and the second term is zero. If they are not, take $\chi = c\,\bar{Y}^\perp$ for a number c. It is uncorrelated with g, so $\mathcal{L}_D(\chi, \bar{Y}) = 0$. By Lemma 3, an economy that satisfies Assumptions 1, 2 and 3 has the extra gains $s_{jk} = \chi_k - \chi_j$, for which $L_D\chi = \nu$. The second term is then $c\,(1 - r_D^2)$, and $\text{sd}_\pi(\chi_j)/\,\text{sd}_\pi(\bar{Y}_j) = |c|\sqrt{1 - r_D^2}$, so the bound holds with equality for every value of this ratio.

(ii) Let m solve (25) for the given contrasts G_{jk}, and put $\psi = m$ and $s_{jk} = G_{jk} - (m_k - m_j)$. Then (25) is (24), which implies Assumption 1. By the argument in the proof of Proposition 4, m is a possible vector of place effects. □

A.1 Representation Lemmas

The results take the place effects, the map, and the comparison weights as given. The first lemma shows that every such configuration comes from an economy that satisfies the assumptions of the paper, so nothing is lost by this. The second adds extra gains.

Lemma 2 (Representation)**.** *Take J locations with any place effects ψ_j, any composition $\bar{\alpha}_j$, any sizes $n_j > 0$, and any shares of movers $w_j \in [0,1)$ that are not all zero. Let $m_j := w_j n_j / \sum_l w_l n_l$ be the share of location j among all movers. Take any flows $F_{jk} \geq 0$ with $F_{jj} = 0$ and $\sum_k F_{jk} = \sum_k F_{kj} = m_j$ for every j. There is an additive economy with exogenous mobility that satisfies Assumption 3 and has these place effects, this composition, these sizes and shares of movers, and these flows. Its comparison weights are $a_{jk} = F_{jk} + F_{kj}$. Its map $\psi_j + \bar{\alpha}_j$ can be the mean outcome of all residents or that of permanent residents.*

The condition on the flows restricts nothing in the results that use the lemma. In a steady state a location sends and receives the same number of movers. Given comparison weights a_{jk} and any sizes, take the flows $F_{jk} = a_{jk}/2$, and shares of movers that are proportional to $\sum_k a_{jk}/n_j$ and below one. A location without movers has $w_j = 0$.

Proof. Let $M = \sum_l w_l n_l$ be the number of movers. In the first period location j has n_j residents, of whom $w_j n_j$ move. If $w_j > 0$, a mover from j goes to k with probability F_{jk}/m_j, and these probabilities sum to one. The flow from j to k, as a share of all movers, is then $m_j F_{jk}/m_j = F_{jk}$. Location j receives $M \sum_k F_{kj} = w_j n_j$ movers, as many as it sends. Its size and its share of movers are therefore the same in both periods.

Every mover draws her person component from one distribution with mean zero, whatever her origin and her destination. The movers who arrive in a location then have the same mean person component as the movers who leave it. If the map is the mean outcome of all residents, the stayers of location j draw theirs from a distribution with mean $\bar{\alpha}_j/(1 - w_j)$. The mean person component of the residents of j is then $\bar{\alpha}_j$ in both periods. If the map is the mean outcome of permanent residents, the stayers of j draw theirs from a distribution with mean $\bar{\alpha}_j$. Outcomes are $Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}$, with shocks that have mean zero and are independent of person components and locations. Mobility is exogenous, the mean outcome that defines the map is $\psi_j + \bar{\alpha}_j$ in both periods, and Assumption 3 holds. □

Lemma 3 (Representation with extra gains)**.** *Take place effects, a composition, sizes, shares of movers, and flows as in Lemma 2, with routes that connect $J \geq 3$ locations. Let χ be an attribute of locations with $\mathcal{L}_D(\chi, \bar{Y}) = 0$, where $\bar{Y}_j = \psi_j + \bar{\alpha}_j$. There is an economy that satisfies Assumptions 1, 2 and 3, has these place effects, this map, these sizes, shares of movers, and flows, and has the extra gain $s_{jk} = \chi_k - \chi_j$ on every route.*

Proof. Sizes, shares of movers, and flows are as in the proof of Lemma 2. Let $a_l = \sum_k a_{lk}$ be the share of movers who live in location l in one of the two periods. As in the proof of Proposition 4, $a_l < 1$ except for at most one location. If there is one, add a constant to χ so that $\chi_l = 0$ there, which changes no s_{jk}. Every mover has a person component α_i with mean zero, whatever her route. In a location l where she lives in one of the two periods, her outcome is $\alpha_i + \psi_l + \chi_l + \varepsilon_{it}$. In a location l where she never lives, it would be $\alpha_i + \psi_l - \chi_l a_l/(1 - a_l) + \varepsilon_{it}$. Shocks have mean zero and are independent of everything else. The mean outcome that all movers would have in l is $\psi_l + a_l\chi_l - (1 - a_l)\chi_l a_l/(1 - a_l) = \psi_l$, so ψ is the vector of place effects. The mean change

on the route from j to k is $\psi_k + \chi_k - \psi_j - \chi_j$, so $s_{jk} = \chi_k - \chi_j$. Flows balance at every location, so $\text{Cov}_{\text{mov}}(s_i, \Delta\bar{Y}_i) = \mathcal{L}_D(\chi, \bar{Y}) = 0$, which is Assumption 1.

A stayer has the outcome $\alpha_i + \psi_l + \varepsilon_{it}$ in every location l. Place effects are then the same for stayers and for movers up to a constant, which is Assumption 2. Let the stayers of location j have mean person component $(\bar{\alpha}_j - w_j\chi_j)/(1 - w_j)$. The movers who leave j and the movers who arrive in it both have mean outcome $\psi_j + \chi_j$ there. The mean outcome of the residents of j is therefore $(1 - w_j)\psi_j + \bar{\alpha}_j - w_j\chi_j + w_j(\psi_j + \chi_j) = \psi_j + \bar{\alpha}_j$ in both periods, and Assumption 3 holds. If the map is the mean outcome of permanent residents, let the stayers of j have mean person component $\bar{\alpha}_j$ instead. □

References

Jason Abaluck, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc. Mortality effects and choice across private health insurance plans. *The Quarterly Journal of Economics*, 136(3):1557–1610, 2021.

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999. ISSN 1468-0262.

Leila Agha, Keith Marzilli Ericson, and Xiaoxi Zhao. The impact of organizational boundaries on health care coordination and utilization. *American Economic Journal: Economic Policy*, 15(3): 184–214, 2023.

Alexander Ahammer, Analisa Packham, and Jonathan Smith. Firms and worker health, 2024. National Bureau of Economic Research, Working Paper No. w32011.

Alberto Alesina, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. Intergenerational mobility in Africa. *Econometrica*, 89(1):1–35, 2021.

M. J. Andrews, L. Gill, T. Schank, and R. Upward. High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society Series A: Statistics in Society*, 171(3):673–697, 2008. ISSN 0964-1998.

Joshua Angrist, Peter Hull, Parag Pathak, and Christopher Walters. Interpreting tests of school VAM validity. *American Economic Review*, 106(5):388–392, 2016.

Ivan Badinski, Amy Finkelstein, Matthew Gentzkow, and Peter Hull. Geographic variation in healthcare utilization: The role of physicians. Technical report, National Bureau of Economic Research, 2023.

Alex Bell, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. Who becomes an inventor in America? The importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713, 2019.

Stéphane Bonhomme, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler. How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics*, 41(2):291–322, 2023. ISSN 0734-306X.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285, 2024.

Bart J Bronnenberg, Jean-Pierre H Dubé, and Matthew Gentzkow. The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, 102(6):2472–2508, 2012.

Marika Cabral, Colleen Carey, and Jinyeong Son. Partial outsourcing of public programs: Evidence on determinants of choice in medicare. *Review of Economics and Statistics*, 2025. Forthcoming.

Enrico Cantoni and Vincent Pons. Does context outweigh individual characteristics in driving voting behavior? evidence from relocations within the united states. *American Economic Review*, 112(4):1226–1272, 2022.

David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013. ISSN 0033-5533.

David Card, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70, 2018. ISSN 0734-306X.

David Card, Jesse Rothstein, and Moises Yi. Industry wage differentials: A firm-based approach. *Journal of Labor Economics*, 42(S1):S11–S59, 2024.

David Card, Jesse Rothstein, and Moises Yi. Location, location, location. *American Economic Journal: Applied Economics*, 17(1):297–336, 2025.

Amitabh Chandra and Connie Xu. Person and place effects in scientific discovery. Working Paper 33996, National Bureau of Economic Research, 2025.

Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects. *The Quarterly Journal of Economics*, 133(3):1107–1162, 2018a.

Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility II: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228, 2018b.

Raj Chetty, John N. Friedman, and Jonah E. Rockoff. Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9):2593–2632, 2014.

Raj Chetty, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter. Race and economic opportunity in the United States: An intergenerational perspective. *The Quarterly Journal of Economics*, 135(2):711–783, 2020.

Raj Chetty, John N. Friedman, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter. The opportunity atlas: Mapping the childhood roots of social mobility. *American Economic Review*, 116(1):1–51, 2026.

Eric Chyn and Na'ama Shenhav. Place effects and geographic inequality in health at birth. *American Economic Journal: Economic Policy*, 17(4):260–291, 2025.

Timothy G. Conley, Christian B. Hansen, and Peter E. Rossi. Plausibly exogenous. *The Review of Economics and Statistics*, 94(1):260–272, 2012.

Gordon B. Dahl. Mobility and the return to education: Testing a Roy model with multiple markets. *Econometrica*, 70(6):2367–2420, 2002.

N. Meltem Daysal, Joseph J. Doyle, Jr., Naeim S. Samandari, Emil A. L. Simonsen, and Mircea Trandafir. Randomized to opportunity: Measuring place effects for refugee children. Working Paper 35853, National Bureau of Economic Research, 2026.

Clément De Chaisemartin and Xavier d'Haultfoeuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American economic review*, 110(9):2964–2996, 2020.

Tatyana Deryugina and David Molitor. The causal effects of place on health and longevity. *Journal of Economic Perspectives*, 35(4):147–170, 2021.

Nathan Deutscher. Place, peers, and the teenage years: Long-run neighborhood effects in Australia. *American Economic Journal: Applied Economics*, 12(2):220–249, 2020.

Hui Ding. Geographic variation in mental health treatment utilization: Evidence from migration. *American Economic Journal: Economic Policy*, 18(1):35–68, 2026.

Joseph J. Doyle and Becky Staiger. Physician group influences on treatment intensity and health: Evidence from physician switchers. *American Economic Journal: Economic Policy*, 17(2):471–505, 2025.

William Encinosa and Avi Dor. Does physician-hospital vertical integration signal care-coordination? evidence from mover-stayer analysis of commercially insured enrollees. *Journal of Health Economics*, 101, 2025.

Amy Finkelstein, Matthew Gentzkow, and Heidi Williams. Sources of geographic variation in health care: Evidence from patient migration. *The Quarterly Journal of Economics*, 131(4):1681–1726, 2016. ISSN 0033-5533.

Amy Finkelstein, Matthew Gentzkow, and Heidi Williams. Place-based drivers of mortality: Evidence from migration. *American Economic Review*, 111(8):2697–2735, 2021. ISSN 0002-8282.

Amy Finkelstein, Matthew Gentzkow, and Dean Li. What drives risky prescription opioid use? evidence from migration. *The Quarterly Journal of Economics*, 140(4), 2025.

Anna Godøy and Ingrid Huitfeldt. Regional variation in health care utilization and mortality. *Journal of Health Economics*, 71:102254, 2020.

Marit Hinnosaar and Elaine M. Liu. Malleability of alcohol consumption: Evidence from migrants. *Journal of Health Economics*, 85, 2022.

Shan Huang and Hannes Ullrich. Provider effects in antibiotic prescribing: Evidence from physician exits. *Journal of Human Resources*, 61(4):1159–1191, 2026.

Peter Hull. Estimating treatment effects in mover designs. *arXiv preprint arXiv:1804.06721*, 2018.

Stefan Hut. Determinants of dietary choice in the US: Evidence from consumer migration. *Journal of Health Economics*, 2020.

Guido W. Imbens and Charles F. Manski. Confidence intervals for partially identified parameters. *Econometrica*, 72(6):1845–1857, 2004.

John Kennan and James R. Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251, 2011.

Patrick Kline. Firm wage effects. In *Handbook of Labor Economics*, volume 5, pages 115–181. Elsevier, 2024.

Patrick Kline, Raffaele Saggio, and Mikkel Sølvsten. Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898, 2020. ISSN 1468-0262.

Patrick M Kline. Branching fixed effects: A proposal for communicating uncertainty. Technical report, National Bureau of Economic Research, 2025.

Jean-William Laliberté. Long-term contextual effects in education: Schools and neighborhoods. *American Economic Journal: Economic Policy*, 13(2):336–377, 2021.

Roxana Leal, William P. Luan, Jonathan Skinner, and John Zhou. Place-based variation in health care: Evidence from mandatory movers in the U.S. military health system. *Journal of Health Economics*, 110:103192, 2026.

Eva Lyubich. The role of people versus places in individual carbon emissions. *American Economic Review*, 115(5):1439–1484, 2025.

Charles F. Manski. *Partial Identification of Probability Distributions*. Springer, New York, 2003.

David Molitor. The evolution of physician practice styles: Evidence from cardiologist migration. *American Economic Journal: Economic Policy*, 10(1):326–356, 2018.

Aviv Nevo and Adam M. Rosen. Identification with imperfect instruments. *The Review of Economics and Statistics*, 94(3):659–671, 2012.

Martin Salm and Ansgar Wübker. Sources of regional variation in healthcare utilization in Germany. *Journal of Health Economics*, 69, 2020.

Isaac Sorkin and Lucas Warwar. Paired movers. Working paper, Stanford University, 2026.

Jörg Stoye. More on confidence intervals for partially identified parameters. *Econometrica*, 77(4): 1299–1315, 2009.

Liyang Sun and Sarah Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199, 2021.

Dan Zeltzer, Liran Einav, Avichai Chasid, and Ran D. Balicer. Supply-side variation in the use of emergency departments. *Journal of Health Economics*, 2021.

Supplement to "Identification in Mover Designs"

B The Assumptions of the Design in the Literature

B.1 What the Studies Assume

Assumption 1 is the identifying assumption that mover studies state, without the model that usually surrounds it. Cantoni and Pons [2022] require that changes in what drives the turnout of a mover are unrelated to the difference in average turnout between her destination and her origin. Molitor [2018] requires that nothing that changes with the move, other than the environment, is related to the change in the environment. Badinski et al. [2023] rule out shocks at the time of the move that are related to utilization in the origin and the destination. Finkelstein et al. [2016] rule out trends that vary with the two. Finkelstein et al. [2025] require that the selection of movers does not change at the move. In the age-at-move form of the design, Chetty and Hendren [2018a] require that selection does not vary with the age at the move. Each is one condition on what changes at a move, in relation to the gap between destination and origin.

Most of these conditions are stated in the additive model. What changes at a move, apart from the place, is then a shock. Assumption 1 is the same condition without that model. It lets the change contain a gain from the place that differs across people. Where the studies let effects differ, they add that the differences are unrelated to the moves that people make. Finkelstein et al. [2016] take differences in the share of place to be unrelated to the other variables of their model. Chetty and Hendren [2018b] read their estimates as mean effects if differences across children are unrelated to where families move. Hull [2018] shows that with more than two places a mover regression has a causal reading only under restrictions on how effects differ across people. Every one of these studies implies Assumption 1 for the movers it compares, in the form of Section III.2 where the outcome is observed once. None of them assumes less, and a mover coefficient that is read as a statement about places cannot assume less. What the studies add buys more than a coefficient, and Section VI says what.

Whether the assumption is reasonable depends on why people move. It is plausible where the outcome is not what they move for. A person who moves to be near family does not choose a region for how its physicians practice. She may still gain more or less than others from the move. The assumption asks only that this gain is unrelated to the gap in mean outcomes between the two regions. It is less plausible where people move for the outcome itself. Models of migration have people move where their own skills pay best [Dahl, 2002], and move on when a location turns out to suit them badly [Kennan and Walker, 2011]. Where such choices are related to the map, Assumption 1 fails, and the results of this paper do not apply.

The studies also offer evidence, and it is evidence along the map. Before the move, the outcomes of movers show no trend that is related to the gap [Finkelstein et al., 2016, Cantoni and Pons, 2022, Badinski et al., 2023]. Changes are symmetric for moves to locations with higher and with lower mean outcomes [Finkelstein et al., 2016, Chetty and Hendren, 2018a], and the coefficient is similar

for families displaced by shocks [Chetty and Hendren, 2018a]. The counterpart for firms is the symmetry test of Card et al. [2013]. This evidence is about how the changes of movers vary with the gap in the map, which is what Assumption 1 restricts. It does not establish the assumption. An extra gain that is proportional to the gap leaves the changes symmetric and only changes the slope.

Finkelstein et al. [2016] state Assumption 2 directly: their approach relies on place effects being the same for movers and for non-movers. Like Cantoni and Pons [2022], they let movers differ from non-movers in the level of the outcome. Chetty and Hendren [2018a] note that their design does not require movers to be like permanent residents, and report that the two groups have similar characteristics. Their target is defined for children who are assigned to neighborhoods at random, which is a statement about everyone.

B.2 What the Studies Report and What They Add

The design is narrower than the question. Nearly every study reports one coefficient: the change in a mover's outcome per unit of the difference in mean outcomes between her destination and her origin. Statements of what this coefficient measures refer to movers. Finkelstein et al. [2016] describe it as a weighted average, across movers, of the share of the gap between origin and destination that is due to place. Finkelstein et al. [2025] describe it as the average share of that gap that the move closes.

Some studies also estimate a place effect for every area, compute the shares of Section III.1 from these estimates, and report that the coefficient is in line with them [Finkelstein et al., 2016, Cantoni and Pons, 2022]. Finkelstein et al. [2016] add that the two are not the same exercise: the decomposition compares two groups of regions, and the event study averages the shares of pairs across movers. Others read the coefficient itself as the share of the variation across locations that is due to place [Molitor, 2018, Ding, 2026]. In the studies of childhood exposure it is typically the only estimate. Lyubich [2025] notes that this reading needs an additional assumption on sorting.

The studies that use movers assume more than Assumption 1, and what they add buys the lower rows of Table 2. A common effect, as in the additive model with moves that are unrelated to shocks, makes the extra gain zero on every route. It buys the effect of every location. Outcomes that are observed only after the move take a restriction on selection in levels, because there is no change in the outcome to compare. Chetty and Hendren [2018a] assume that selection does not vary with the age at the move. Finkelstein et al. [2021] assume that unobserved and observed health are related to the destination in the same proportion as to the origin. Structure across locations or over time buys the dynamics and the components of place. Examples are a gap that closes partly at the move and partly with the years since [Bronnenberg et al., 2012], the transition to addiction and the availability of opioids [Finkelstein et al., 2025], and physicians as a separate part of place [Badinski et al., 2023].

C Stronger Assumptions

Assumption 1 asks that extra gains be unrelated to the gap in the map, and by Proposition 4 the outcomes of movers then identify the mover coefficient and no other feature of place effects. This section collects what stronger assumptions add: a condition at every location or on every route, the assumption of the design in groups of movers, and the economy of Example 2. It ends with other economies that generate a given map and given comparison weights.

C.1 At Every Location and on Every Route

Section VI states the condition at every location, equation (24). A stronger version holds on every route.

On every route. The extra gain is zero on every route with movers, $s_{jk} = 0$. Each route contrast is then the difference of two place effects. This is the same as requiring (24) for any weights on the routes.

The additive model with exogenous mobility implies both versions. It adds that the place effect is the same for people who do not move, and that composition is the mean of a component that residents carry with them.

Proposition 6 (What stronger assumptions add)**.** *Let Assumption 3 hold, and let the routes with movers connect all locations.*

(i) Place effects. *Under* (24)*, the least-squares regression of the route contrasts on indicators of destination and origin, weighted by flows, returns every difference $\psi_k - \psi_j$. With the map, these determine θ_π, η, every Q_c of Section D, the sign of sorting, and the difference $\theta_D - \theta_\pi$, whatever the comparison weights are.*

(ii) One sum. *The mover coefficient uses one weighted sum of the route contrasts, $\sum_{j \neq k} F_{jk} \left(\bar{Y}_k - \bar{Y}_j\right) \left(\psi_k - \psi_j + s_{jk}\right)$. Under Assumption 1 it equals $\mathcal{L}_D(\psi, \bar{Y})$.*

Part (i) is the two-way estimator. If each route contrast were a difference of two place effects, contrasts would add up along any path of routes. Condition (24) makes them add up on average at each location, and that is enough. Place effects can then be rebuilt from any set of routes that connects the locations. The estimator returns to levels, and there it can weight locations as the question requires. The flows affect its precision and not what it estimates. It needs no alignment. For firm effects, Kline [2024, Section 3.1.2 and footnote 6] gives the algebra behind this. Least-squares estimates are unchanged when the wage changes of movers are shifted around a cycle of moves, in inverse proportion to the number of movers on each edge of the cycle. Condition (24) says that the extra gains are shifts of this kind.

The mover regression stays in differences. It weights each route by its flow and by its gap in the map. One weighted sum answers one question, and the routes choose which. Under the stronger versions, a mover regression can also deliver the cross-sectional coefficient. The first

Table 2: What Each Version of the Assumption Delivers

	Assumption 1	At every location	On every route	Additive model
$\theta_D = \theta_\pi$ when routes are aligned (Theorem 1)	Yes	Yes	Yes	Yes
The bound (Theorem 3(i))	Yes	Yes	Yes	Yes
The bounds, if sorting is not negative (Theorem 3(ii))	Yes	Yes	Yes	Yes
Place effects, θ_π, and the difference $\theta_D - \theta_\pi$	No	Yes	Yes	Yes
The sign of sorting	No	Yes	Yes	Yes
θ_π from reweighted routes	No	No	Yes	Yes
The split of the variance between places and people	No	No	No	Yes

Notes. Assumption 1: among movers, extra gains are uncorrelated with the gap in the map. At every location: equation (24). On every route: $s_{jk} = 0$. Each column implies the columns to its left. All statements are about movers. Under Assumption 2 they are also about the whole population. In the additive model that assumption holds by construction. Last row: the variance share of place effects is computed from place effects and the map in the second and third columns as well. Reading the rest of the variance as what people account for needs composition to be a component that residents carry with them.

version takes an instrument that is built from flows and the map, and the second takes reweighted routes (Section E).

Table 2 collects what each version delivers.

A reader of a published study has the coefficient, and can have the map and the flows. Whatever version holds, that reader is in the first column. The third version has testable content in the route contrasts. Under it, the contrast from j to k is minus the contrast from k to j, and contrasts add up along paths. The first of these is the symmetry examined by Card et al. [2013]. Kline [2024] shows that such sums around cycles of moves are all that the additive model restricts in the mean changes of movers. He finds that wage changes between Italian firms come close to satisfying them, without doing so exactly. One coefficient cannot reveal either. The second version has no testable content (Proposition 5). When even Assumption 1 fails, a causal reading of a mover regression takes other restrictions. Hull [2018] replaces the additive model by parallel trends. With more than two locations, he shows, the regression then needs restrictions on how effects differ across locations and over time.

Proof of Proposition 6. (i) Write $G_{jk} = \psi_k - \psi_j + s_{jk}$ for the route contrast. The regression chooses $m \in \mathbb{R}^J$ to minimize $\sum_{j \neq k} F_{jk} \left(G_{jk} - m_k + m_j\right)^2$. The first-order condition for m_j is $\sum_k F_{kj} \left(G_{kj} - m_j + m_k\right) = \sum_k F_{jk} \left(G_{jk} - m_k + m_j\right)$. At $m = \psi$ it reads $\sum_k F_{kj} s_{kj} = \sum_k F_{jk} s_{jk}$, which is (24). The objective is convex, and its Hessian is $2L_D$, which has rank $J - 1$ when the routes connect the locations. The minimizers are therefore ψ plus a constant. The map and the differences of place effects determine θ_π, η, every Q_c, $\mathrm{Cov}_\pi(\psi_j, \bar{\alpha}_j)$, and θ_D.

(ii) The part of the sum in ψ is $\sum_{j \neq k} F_{jk} (\bar{Y}_k - \bar{Y}_j)(\psi_k - \psi_j) = \sum_{j<k} (F_{jk} + F_{kj})(\bar{Y}_j - \bar{Y}_k)(\psi_j - \psi_k) = \mathcal{L}_D(\psi, \bar{Y})$. The part in s is $\mathbb{E}_{\mathrm{mov}}[s_i \, \Delta \bar{Y}_i]$. Under Assumption 3 each location receives as many movers as it sends, so $\mathbb{E}_{\mathrm{mov}}[\Delta \bar{Y}_i] = 0$, and this part is $\mathrm{Cov}_{\mathrm{mov}}(s_i, \Delta \bar{Y}_i)$. It is zero under Assumption 1. □

C.2 *Groups of Movers*

Split the movers into K groups, and let the assumption of the design hold in each group:

$$\text{Cov}\left(s_i, \Delta\bar{Y}_i \mid i \text{ in group } k\right) = 0 \qquad \text{for } k = 1, \ldots, K. \tag{27}$$

The extra gain s_i is defined with the place effects of all movers. If the movers of a group share routes with other movers, s_i is the extra gain of her group on her route: the mean change in the outcome of the movers of group k on the route from j to l is $\psi_l - \psi_j + s^k_{jl}$, and $s_i = s^k_{o(i)d(i)}$. A group can be the movers on a set of routes: short moves and long moves, or routes that many movers take and routes that few take. In the age-at-move design it can be the children who leave a set of origins.

Each group is a design of its own. Let θ_{D_k} be the coefficient of regression (8) among the movers of group k. Let g^k be the comparison gap of the group, defined as in (14) with the comparison weights of the group.[11] Let $\bar{Y}^G_j$ and ψ^G_j be the fitted values of the least-squares regressions, with weights π, of the map and of the place effects on the K comparison gaps. Let

$$R^2_D := \frac{\text{Var}_\pi(\bar{Y}^G_j)}{\text{Var}_\pi(\bar{Y}_j)}, \qquad t_G := \frac{\text{Cov}_\pi(\psi^G_j, \bar{Y}_j)}{\text{Var}_\pi(\bar{Y}_j)}, \qquad v_G := \frac{\text{Var}_\pi(\psi^G_j)}{\text{Var}_\pi(\bar{Y}_j)}.$$

I call R_D the alignment of the groups together. It is the multiple correlation of the map with the comparison gaps of the groups. With one group it is the alignment r_D.

Proposition 7 (Groups of movers)**.** *Let the conditions of Proposition 4 hold, with (27) in the place of Assumption 1, and let the gap in the map vary among the movers of every group.*

(i) K numbers. *A vector of place effects is consistent with the data if and only if* $\text{Cov}_\pi(\psi_j, g^k_j) = \theta_{D_k}\text{Cov}_\pi(\bar{Y}_j, g^k_j)$ *for every group k. These equations determine ψ^G, and with it t_G and v_G.*

(ii) Identification. *If $R_D = 1$, then $\theta_\pi = t_G$ for every vector of place effects. It is a weighted sum of the group coefficients, $\theta_\pi = \sum_k w_k\,\theta_{D_k}$. The weights sum to one and are computed from flows and location means. If $R_D < 1$, the data do not restrict θ_π.*

(iii) Bounds. *Let Assumption 4 hold and $R_D < 1$. The values of θ_π that are consistent with the data are those on or between the bounds*

$$t_G + \tfrac{1}{2}\left(1 - R^2_D\right) \;\pm\; \sqrt{\left(1 - R^2_D\right)\left\{\tfrac{1}{4}\left(1 - R^2_D\right) + t_G - v_G\right\}}. \tag{28}$$

Such values exist if and only if $\frac{1}{4}(1 - R^2_D) + t_G - v_G \geq 0$. With one group, $t_G = r^2_D\,\theta_D$ and $v_G = r^2_D\,\theta^2_D$, and the bounds are those of Theorem 3.

[11] If the movers of a group do not arrive at each location and leave it in equal numbers, the comparison gap of the group has one more term. Let a^k_{jl} be the comparison weights of the group, n^k_j its net inflow into location j, both as shares of all movers, and μ_k the mean of $\Delta\bar{Y}_i$ in the group. Then $g^k_j = \pi_j^{-1}\{\sum_l a^k_{jl}(\bar{Y}_j - \bar{Y}_l) - n^k_j\,\mu_k\}$. The proposition holds with this gap, as its proof shows.

No group needs aligned routes. It is enough that the map is a combination of the comparison gaps of the groups. Take Example 1 with $c = \frac{1}{2}$. Let the moves across practice styles be one group, and the moves across patient populations the other. Their coefficients are one and zero. Neither group is aligned with the map: the alignments are 0.89 and 0.45. But the map is a combination of the two comparison gaps, so $R_D = 1$. Part (ii) gives $\theta_\pi = 0.8 \times 1 + 0.2 \times 0 = 0.8$, for every share ρ of moves across practice styles. The coefficient of all movers together is 0.8 only at $\rho = \frac{1}{2}$.

The groups must differ in their routes, or their comparison gaps are nearly the same and a second group adds little. One choice that works is by how common a comparison is. Independent pairing would give the pair of locations j and k the comparison weight $\pi_j \pi_k$. Sort the pairs by the ratio $a_{jk}/(\pi_j \pi_k)$, and form groups of pairs with similar ratios. Between the large U.S. commuting zones, five such groups raise an alignment of 0.85 to 0.98 (Section L). A group is there a set of comparisons between two destinations, and not a set of children.

Condition (27) asks more than Assumption 1. Under Assumption 1, a positive relation between extra gains and the gap in the map among long moves can offset a negative one among short moves. Under (27) each group must satisfy the condition on its own.[12] It remains a condition about the map, and it does not restrict how extra gains are related to other attributes of locations.

With one group the proposition is Proposition 4 with Theorems 1 and 3. With $J-1$ groups whose comparison gaps are linearly independent, the equations of part (i) determine every difference between two place effects. Each added group adds a coefficient to estimate, and the weights w_k can be large and of both signs. A few groups with different routes are the useful case.

Proof of Proposition 7. Write $\langle f, h \rangle := \mathrm{Cov}_\pi(f_j, h_j)$ and $x_j := \bar{Y}_j - \mathbb{E}_\pi[\bar{Y}]$. For group k, let F^k_{jl} be its flows as shares of all movers, $a^k_{jl} = F^k_{jl} + F^k_{lj}$, $n^k_j = \sum_l (F^k_{lj} - F^k_{jl})$, and let μ_k be the mean of $\Delta\bar{Y}_i$ in the group. For an attribute f of locations, the covariance of Δf_i and $\Delta\bar{Y}_i$ in the group, times the share of the group among movers, is

$$\ell_k(f) := \sum_{j \neq l} F^k_{jl} \, (f_l - f_j) \, (\bar{Y}_l - \bar{Y}_j - \mu_k) = \sum_j f_j \Big\{ \sum_l a^k_{jl} (\bar{Y}_j - \bar{Y}_l) - n^k_j \, \mu_k \Big\}.$$

The term in braces sums to zero over j, so $\ell_k(f) = \langle f, g^k \rangle$ with the gap g^k of the footnote. If the movers of the group arrive at each location and leave it in equal numbers, $n^k = 0$ and g^k is (14) with the comparison weights of the group.

(i) Let G^k_{jl} be the mean change in the outcome on the route from j to l in group k. The coefficient of regression (8) in the group is $\theta_{D_k} = \sum_{j \neq l} F^k_{jl} \, G^k_{jl} \, (\bar{Y}_l - \bar{Y}_j - \mu_k) / \ell_k(\bar{Y})$. With $s^k_{jl} = G^k_{jl} - (\psi_l - \psi_j)$, condition (27) for group k is $\theta_{D_k} \, \ell_k(\bar{Y}) = \ell_k(\psi)$, which is the equation in (i). Conversely, if ψ satisfies the K equations, these extra gains satisfy (27). By the argument in the proof of Proposition 4, ψ is

[12] When all groups have the same mean extra gain, or the same mean gap in the map, (27) implies Assumption 1. Otherwise (27) need not imply Assumption 1, and Assumption 1 does not imply (27). If Assumption 1 is maintained together with (27), add the comparison gap of all movers to the K gaps. The proposition holds with these $K+1$ gaps, and the alignment of the groups together is then at least r_D. Without it, R_D can be smaller than r_D when the movers of a group do not arrive at each location and leave it in equal numbers.

then a possible vector of place effects. The coefficients of the regression of ψ on the gaps depend on ψ only through the $\langle \psi, g^k \rangle$, which gives the last statement.

(ii) and (iii) Let P be the projection, in the inner product $\langle \cdot, \cdot \rangle$ on centered vectors, on the span of the gaps. Then $\bar{Y}^G = Px$ and $\psi^G = P\psi$, up to constants. Put $x^\perp = x - Px$ and $\psi^\perp = \psi - P\psi$. By (i), ψ^G is known, and $\psi^\perp$ can be any centered vector that is orthogonal to the gaps. Since $\langle \psi, x \rangle = \langle \psi^G, x \rangle + \langle \psi^\perp, x^\perp \rangle$,

$$\theta_\pi = t_G + \frac{\langle \psi^\perp, x^\perp \rangle}{\langle x, x \rangle}.$$

If $R_D = 1$, then $x^\perp = 0$ and $\theta_\pi = t_G$. Write $x = \sum_k \lambda_k g^k$. Then $t_G \langle x, x \rangle = \langle \psi, x \rangle = \sum_k \lambda_k \theta_{D_k} \langle x, g^k \rangle$, so $w_k = \lambda_k \langle x, g^k \rangle / \langle x, x \rangle$, and $\sum_k w_k = \langle x, x \rangle / \langle x, x \rangle = 1$. If $R_D < 1$, then $x^\perp \neq 0$, and $\psi^\perp = c\, x^\perp$ gives every value of θ_π as c varies.

Assumption 4 is $\langle \psi, x - \psi \rangle \geq 0$, that is, $\langle \psi^\perp, \psi^\perp \rangle \leq (\theta_\pi - v_G) \langle x, x \rangle$. By the Cauchy–Schwarz inequality,

$$(\theta_\pi - t_G)^2 \langle x, x \rangle^2 = \langle \psi^\perp, x^\perp \rangle^2 \leq \langle \psi^\perp, \psi^\perp \rangle \, \langle x^\perp, x^\perp \rangle \leq (\theta_\pi - v_G) \, (1 - R_D^2) \, \langle x, x \rangle^2.$$

Hence $(\theta_\pi - t_G)^2 \leq (1 - R_D^2)(\theta_\pi - v_G)$, and the solutions of this inequality are the values on or between the bounds (28). Each of its points is attained by $\psi^\perp = c\, x^\perp$ with $c = (\theta_\pi - t_G)/(1 - R_D^2)$, for which both inequalities above hold. With one group, $\psi^G = \{ \langle \psi, g \rangle / \langle g, g \rangle \}\, g$. This gives $t_G = r_D^2\, \theta_D$ and $v_G = r_D^2\, \theta_D^2$, and (28) becomes (23). □

C.3 The Economy of Example 2

The example of Section VI is the following economy.

Locations differ in an attribute x_j that suits some people and not others. A region may organize care around specialists, or a neighborhood may have schools with a particular program. The outcome of person i in location j is $\alpha_i + \psi_j + \tau_i\, x_j$. For half of the people the attribute raises the outcome, $\tau_i = t > 0$. For the other half it lowers it, $\tau_i = -t$. On average it does nothing, so ψ_j is the place effect. The attribute is unrelated to everything else about a location.

People live for two periods and may move once, between the two. Where a person first lives is unrelated to her τ_i, and so is whether she moves. Assumption 2 therefore holds. A mover chooses the attribute of her destination with an eye on her own gain. She chooses the value x with a probability that is proportional to $\omega(x) \exp(\lambda \tau_i x)$. Here $\lambda \geq 0$ is the strength of the selection, and the weights $\omega(x) > 0$ are such that each value of the attribute is as common among destinations as among origins. She chooses the rest of her destination independently of τ_i, of the attribute she chose, and of the attribute of her origin. Every value of the attribute occurs in at least two locations. A new generation is born in every period, and each location receives as many movers as it sends. The economy is therefore in a steady state, and Assumption 3 holds. The map is the mean outcome of permanent residents. They never chose a location, so the map does not depend on the attribute.

By Lemma 4 below, the extra gain on the route from j to k is $T(x_k)\,(x_k - x_j)$. Here $T(x)$ is the mean of τ_i among the movers whose destination has attribute x, and it rises with x when $\lambda > 0$. Extra gains are independent of the gap in the map among movers, so Assumption 1 holds for every λ. Condition (24) fails. Those who arrive in a location with an unusual attribute are mostly people whom it suits, and their extra gain is large. Those who leave have the same mean extra gain at every location. The two-way estimator adds half of the mean extra gain of the arrivals to each place effect.

For numbers, let one region in four be special, $x = 1$, and the others standard, $x = 0$. Let all regions send the same number of movers, and let $\lambda t = 1$. In the steady state a quarter of all movers then arrive in a special region, which takes $\omega(1) = 0.21$ and $\omega(0) = 0.79$. The attribute suits half of all movers, and 83 percent of those who arrive in a special region. A mover from a standard to a special region gains $0.65\,t$ more than the average mover would. A mover from a special to a standard region gains $0.22\,t$ more, because she is more likely a person whom the special region did not suit. The mover coefficient is what it would be without any selection. The two-way estimator credits a special region with an effect of $0.22\,t$ that the average mover does not experience.

Assumption 1 holds exactly in the example because the attribute is unrelated to the rest of a location, and because the map is that of permanent residents. If the attribute were more common where mean outcomes are high, or if it had an effect on average, the extra gains would be related to the gap in the map. A map that is computed from all residents contains the arrivals, and with them their own gains. The assumption then holds up to a term that is proportional to the share of arrivals among residents. Take the regions of Example 1 with $c = \rho = \frac{1}{2}$, and let $t = 1$. If a tenth of all residents are arrivals, the mover coefficient is 0.802 where it is 0.800 with the map of permanent residents.

Lemma 4 (The economy of Example 2)**.** *In the economy of Example 2 with $\lambda > 0$, let $T(x) := t\,\tanh\{\lambda t\,(x - \tilde{x})\}$, where $\tilde{x}$ solves $\sum_x \omega(x)\,e^{\lambda t(x-\tilde{x})} = \sum_x \omega(x)\,e^{-\lambda t(x-\tilde{x})}$. With $\lambda = 0$, let $T(x) := 0$. Let $\bar{x}$ be the mean attribute of the origins of movers, let $g(x) := T(x)\,(x - \bar{x})$, and let $\bar{g}$ be the mean of g over the destinations of movers.*

(i) *The mean of τ_i among the movers whose destination has attribute x is $T(x)$, and the extra gain on the route from j to k is $s_{jk} = T(x_k)\,(x_k - x_j)$. It is zero on every route if and only if $\lambda = 0$ or all locations have the same attribute.*

(ii) *Among movers, s_i and $\Delta\bar{Y}_i$ are independent, so Assumption 1 holds for every λ.*

(iii) *Those who arrive in location j have a mean extra gain of $g(x_j)$, and those who leave it of $\bar{g}$. Hence $v_j = q_j\,\{g(x_j) - \bar{g}\}$, where q_j is the outflow of location j, and the two-way regression of Proposition 6(i) returns $\psi_j + \frac{1}{2}g(x_j)$, up to a constant.*

Proof of Lemma 4. (i) Let $Z_\pm = \sum_x \omega(x)e^{\pm\lambda t x}$. A mover with $\tau_i = \pm t$ chooses the attribute value x with probability $\omega(x)e^{\pm\lambda t x}/Z_\pm$, whatever her origin. The two values of τ_i are equally likely among

movers, so by Bayes' rule the mean of τ_i among movers whose destination has attribute x is

$$t\,\frac{e^{\lambda t x}/Z_+ - e^{-\lambda t x}/Z_-}{e^{\lambda t x}/Z_+ + e^{-\lambda t x}/Z_-} = t\,\tanh\{\lambda t\,(x-\tilde{x})\}, \qquad e^{2\lambda t \tilde{x}} = Z_+/Z_-,$$

and the last equation is the one that defines $\tilde{x}$. The origin and the choice of the rest of the destination are unrelated to τ_i, so $\mathbb{E}[\tau_i \mid o(i)=j, d(i)=k] = T(x_k)$. A mover has not chosen her origin, so the mean change on the route is $\psi_k - \psi_j + T(x_k)(x_k - x_j)$. If $\lambda > 0$ and two values $x < x'$ occur, then $T(x) < T(x')$, and the routes between locations with these two attributes have extra gains $T(x')(x'-x)$ and $T(x)(x-x')$, which are not both zero.

(ii) The extra gain of a mover is a function of the attributes of her origin and her destination. The gap in the map is a function of the rest of her origin and of her destination, because the map is the mean outcome of permanent residents and their τ_i is unrelated to where they live. The attribute is independent of the rest across locations, who moves is unrelated to the attribute, and the two parts of the destination are chosen independently. Hence s_i and $\Delta\bar{Y}_i$ are independent among movers.

(iii) The attribute of the origin is independent of the destination. Hence the mean of s_i among those who arrive in j is $T(x_j)(x_j - \bar{x}) = g(x_j)$. The mean of $T(x_{d(i)})$ over movers is the mean of τ_i over movers, which is zero. Hence the mean of s_i among those who leave j is $\mathbb{E}[T(x_{d(i)})\,x_{d(i)}] - x_j\,\mathbb{E}[T(x_{d(i)})] = \mathbb{E}[T(x_{d(i)})(x_{d(i)} - \bar{x})] = \bar{g}$. Inflow equals outflow, which gives ν_j. Put $\chi_j = \frac{1}{2}g(x_j)$. Each location receives as many movers as it sends, so the attributes of origins and of destinations have the same distribution among movers, and χ has mean $\frac{1}{2}\bar{g}$ under both. Hence $\sum_k (F_{jk}+F_{kj})(\chi_j - \chi_k) = q_j\,\{2\chi_j - \frac{1}{2}\bar{g} - \frac{1}{2}\bar{g}\} = \nu_j$, which is (26).

Steady state. In every period each location contains a young generation, the old who stayed, and the old who arrived, in the same numbers and with the same mean outcomes as in every other period. The mean outcome of permanent residents is the same in both periods of their lives. □

C.4 Other Economies behind Lemma 2

The economy in the proof of Lemma 2 gives every mover a person component with the same mean, and lets the stayers carry the composition of each location. This is enough for the results. The mover coefficient, the cross-sectional coefficient, and the alignment depend on an economy only through its place effects, its map, and its comparison weights. By Proposition 1, all flows with the same comparison weights give the same coefficient.

Movers can also be sorted. Give a mover who travels between j and k a person component with mean $\varphi_j + \varphi_k$, where φ solves $(I+W)\varphi = \bar{\alpha}$ and W is the matrix of comparison weights with each row scaled to sum to one. With flows that are the same in both directions, the movers who leave location j and the movers who arrive in it then both have mean person component $\bar{\alpha}_j$, and no stayers are needed. The system has a solution unless a connected set of locations splits into two classes with every move between the classes and none within them. Three locations that exchange movers with one another rule this out. In this economy a mover matches the mean composition of

both of her locations, and she carries one person component across the two. The simulations of Section K use it.

D Other Questions

The main text states the results for the slope θ_π. This section defines the other targets that mover studies report, relates them to the slope, and restates the results for them.

D.1 Shares of the Variance

The studies report two shares of the variance of the map. One is how much of it would disappear if place effects were equalized across areas. The other is how much would disappear if the characteristics of residents were equalized. Let

$$\theta_{\text{var}} := \frac{\text{Var}_\pi(\psi_j)}{\text{Var}_\pi(\bar{Y}_j)}, \qquad S_{\text{eq}} := 1 - \frac{\text{Var}_\pi(\bar{\alpha}_j)}{\text{Var}_\pi(\bar{Y}_j)}. \tag{29}$$

The *variance share* θ_{var} is the part of the variance of the map that would remain if residents were the same everywhere. The *equalization share* S_{eq} is the part that equalizing place effects would remove.

Lemma 5 (Three shares)**.** $\theta_\pi = \frac{1}{2}(\theta_{\text{var}} + S_{\text{eq}})$. *If* $\text{Cov}_\pi(\psi_j, \bar{\alpha}_j) \geq 0$, *then* $\theta_\pi^2 \leq \theta_{\text{var}} \leq \theta_\pi \leq S_{\text{eq}} \leq 2\theta_\pi - \theta_\pi^2$.

The slope lies halfway between the two shares. The variance share gives none of the covariance between place effects and composition to place, and the equalization share gives all of it. The slope is therefore a share of the variance only when place effects and composition are uncorrelated across locations, as in (7). A study should say which of the three it reports. Reading the two shares as what places and what people account for needs that composition is something residents carry with them, as in the additive model.

Proof of Lemma 5. Let V_ψ, V_α and C be the variances and the covariance of ψ_j and $\bar{\alpha}_j$ under π, and $V_Y = V_\psi + 2C + V_\alpha$. Then $\theta_\pi = (V_\psi + C)/V_Y$, $\theta_{\text{var}} = V_\psi/V_Y$, and $S_{\text{eq}} = (V_\psi + 2C)/V_Y$, which gives the identity. If $C \geq 0$, then $\theta_{\text{var}} \leq \theta_\pi \leq S_{\text{eq}}$. By the Cauchy–Schwarz inequality, $\text{Cov}_\pi(\psi_j, \bar{Y}_j)^2 \leq V_\psi V_Y$, which is $\theta_\pi^2 \leq \theta_{\text{var}}$. Then $S_{\text{eq}} = 2\theta_\pi - \theta_{\text{var}} \leq 2\theta_\pi - \theta_\pi^2$. □

Corollary 1 (Bounds on the shares of the variance)**.** *Let Assumption 4 hold, let* $r_D < 1$, *and let* $0 \leq \theta_D \leq 1$. *The values of the variance share and of the equalization share that are consistent with the mover coefficient, the map, and the comparison weights are*

$$\theta_{\text{var}} \in \left[\, r_D^2\, \theta_D^2,\ \theta_+ \right], \qquad S_{\text{eq}} \in \left[\, \theta_-,\ 1 - r_D^2\,(1-\theta_D)^2 \right], \tag{30}$$

with $\theta_\pm$ *as in* (23). *Every value on or between the bounds is attained. With aligned routes,* $r_D = 1$, *the bounds reduce to those of Lemma 5.*

Proof. Use the notation of the proof of Theorem 3 in the Appendix of the paper: $\theta = \theta_\pi$, $\Lambda = \sqrt{1 - r_D^2}/r_D$, and $q(\theta) = \theta^2 + (\theta_D - \theta)^2/\Lambda^2$. Let $v = \theta_{\text{var}}$. By definition $v = \theta^2 + \eta^2$ and $S_{\text{eq}} = 2\theta - v$, and Assumption 4 is $v \leq \theta$. By Step 1 of that proof, for a given θ the smallest attainable v is $q(\theta)$. If $J \geq 4$, the attainable pairs are $\{(\theta, v) : q(\theta) \leq v \leq \theta\}$. The function q is convex, with minimum $q(\theta^*) = \theta_D^2/(1 + \Lambda^2) = r_D^2\theta_D^2$ at $\theta^* = \theta_D/(1 + \Lambda^2)$, and $q(\theta^*) \leq \theta^*$ if and only if $0 \leq \theta_D \leq 1$. The largest v is θ_+. For the equalization share, $2\theta - v$ ranges from θ to $2\theta - q(\theta)$. The smallest value is θ_-. The function $2\theta - q(\theta)$ is concave, with maximum $1 - r_D^2(1 - \theta_D)^2$ at $\theta^{**} = (\theta_D + \Lambda^2)/(1 + \Lambda^2)$. The problem is symmetric under $(\theta, \theta_D) \mapsto (1 - \theta, 1 - \theta_D)$, so θ^{**} is feasible if and only if $0 \leq \theta_D \leq 1$. Intermediate values are attained by continuity. If $J = 3$, $v = q(\theta)$ on $[\theta_-, \theta_+]$. Since $q(\theta_\pm) = \theta_\pm$ and both θ^* and θ^{**} lie between θ_- and θ_+, the ranges of v and of $2\theta - v$ have the same bounds. □

D.2 Shares of a Gap and Other Contrasts

The share of a gap is an object of the same kind as the slope. Let c_j be any variable that contrasts locations, such as an indicator for the locations above the median, and put it in the place of the map in both covariances:

$$Q_c := \frac{\text{Cov}_\pi(\psi_j, c_j)}{\text{Cov}_\pi(\bar{Y}_j, c_j)}. \tag{31}$$

For the indicator, Q_c is the share of the gap between the two groups that is place. The slope is the case $c = \bar{Y}$. Each Q_c has the policy reading of the slope, for a policy that tilts the assignment toward locations with a high c_j.

The results of Sections IV and V of the paper compare the mover coefficient with the slope θ_π. They hold for every question Q_c.

Corollary 2 (Other questions)**.** *Let c be a variable that contrasts locations, with $\text{Cov}_\pi(\bar{Y}_j, c_j) > 0$.*

(i) Two gaps.

$$\theta_D - Q_c = \underbrace{(\theta_D - \theta_\pi)}_{\textit{routes}} + \underbrace{(\theta_\pi - Q_c)}_{\textit{question}}, \qquad \textit{with} \qquad \theta_\pi - Q_c = -\frac{\text{Cov}_\pi(e_j, c_j)}{\text{Cov}_\pi(\bar{Y}_j, c_j)}. \tag{32}$$

(ii) Identification. *$\theta_D = Q_c$ for every vector of place effects if and only if the comparison gap is proportional to $c_j - \mathbb{E}_\pi[c]$. Otherwise anything goes, as in Theorem 1*(iii).

(iii) Bound. *$|\theta_D - Q_c| \leq \eta \ \text{sd}_\pi(\bar{Y}_j) \ \text{sd}_\pi(w_j)$, where $w_j = g_j/\text{Cov}_\pi(\bar{Y}_j, g_j) - c_j/\text{Cov}_\pi(\bar{Y}_j, c_j)$. Under Assumption 4, $\eta \leq \frac{1}{2}$.*

Part (i) separates a gap that is about routes from a gap that is about the question. The first gap is about routes. It compares the movers with the experiment that defines the slope: the same question, other comparisons. Sections IV and V are about this gap. The second gap has nothing to do with movers. Take the share of the gap between the top and the bottom of the map. The slope equals

that share only if what the map misses has the same mean in the two groups. This is so even in the experiment. And by Lemma 5, the slope is not a share of the variance unless place effects and composition are uncorrelated.

Finkelstein et al. [2016] and Cantoni and Pons [2022] compute shares from estimated place effects for several partitions of areas: above and below the median, top and bottom quartiles, and so on. They find that the shares are stable across partitions. This is evidence on the second gap, the one about the question. It says that place effects are close to linear in the map on average. It is not evidence on the first gap, the one about routes. That gap depends on how the residual is related to the routes.

The mover regression itself can be pointed at another question. Regress ΔY_i on $\Delta \bar{Y}_i$ with Δc_i as the instrument. This takes more than Assumption 1: among movers, extra gains must be uncorrelated with the change in c, as they are under (24). For a group of locations, the coefficient is a ratio. Its numerator is the change in outcomes of those who move into the group minus the change of those who move out. Its denominator is the same difference in the map. The coefficient answers Q_c for all place effects under an alignment condition of its own. For a group, the condition is about the moves between the group and the rest. Within the group, and within the rest, each location takes part in these moves in proportion to its weight π_j, as the proof below shows. Such regressions are rare. What the literature reports is θ_D.

Proof of Corollary 2. Let $\hat{c} = (c - \mathbb{E}_\pi[c]) / \operatorname{Cov}_\pi(\bar{Y}_j, c_j)$ and $\hat{g} = g / \operatorname{Cov}_\pi(\bar{Y}_j, g_j)$. Both have mean zero and covariance one with the map. By (16), $Q_c = \operatorname{Cov}_\pi(\psi_j, \hat{c}_j) = \theta_\pi + \operatorname{Cov}_\pi(e_j, \hat{c}_j)$, which is (i). With $w = \hat{g} - \hat{c}$, part (i) and Theorem 1(i) give $\theta_D - Q_c = \operatorname{Cov}_\pi(e_j, w_j)$. The vector w has mean zero and is uncorrelated with the map. If $w = 0$, that is, if g is proportional to $c - \mathbb{E}_\pi[c]$, then $\theta_D = Q_c$ for all place effects. If $w \neq 0$, let $b = (t - t') / \operatorname{Var}_\pi(w_j)$ and $\psi = \{t' - b \operatorname{Cov}_\pi(w_j, \hat{c}_j)\} \bar{Y} + b\, w$. Then $Q_c = t'$ and $\theta_D = Q_c + b \operatorname{Var}_\pi(w_j) = t$, and Lemma 2 supplies the economy. This is (ii). Part (iii) is the Cauchy–Schwarz inequality, $|\operatorname{Cov}_\pi(e_j, w_j)| \leq \operatorname{sd}_\pi(e_j)\, \operatorname{sd}_\pi(w_j)$. Under Assumption 4, Lemma 5 gives $\eta^2 = \theta_{\text{var}} - \theta_\pi^2 \leq \theta_\pi(1 - \theta_\pi) \leq \frac{1}{4}$.

The regression pointed at c. Let $\operatorname{Cov}_{\text{mov}}(s_i, \Delta c_i) = 0$ and $\mathcal{L}_D(\bar{Y}, c) \neq 0$. Under Assumption 3, $\mathbb{E}_{\text{mov}}[\Delta c_i] = 0$. With Δc_i as the instrument for $\Delta \bar{Y}_i$, the coefficient is then $\operatorname{Cov}_{\text{mov}}(\Delta c_i, \Delta Y_i) / \operatorname{Cov}_{\text{mov}}(\Delta c_i, \Delta \bar{Y}_i) = \mathcal{L}_D(\psi, c) / \mathcal{L}_D(\bar{Y}, c)$. As in the proof of Proposition 3, $\mathcal{L}_D(f, c) = \operatorname{Cov}_\pi(f_j, g_j^c)$ with $g_j^c = \pi_j^{-1} \sum_k a_{jk}(c_j - c_k)$. The coefficient is therefore Q_{g^c}, and by the argument above it equals Q_c for all place effects if and only if g^c is proportional to $c - \mathbb{E}_\pi[c]$. For the indicator of a group G, let x_j be the total weight of the comparisons of j with locations on the other side. Then $g_j^c = x_j / \pi_j$ for j in G and $g_j^c = -x_j / \pi_j$ otherwise. This is proportional to $c_j - \pi(G)$ if and only if x_j / π_j is constant within G and constant outside it. □

D.3 The Policy Reading

The slope is also the return to a policy. Suppose a policy maker wants people to live where outcomes are better. She does not know which location is best for whom. She can only change the probability

with which a person lives in each location. Start from the assignment with probabilities π_j, and tilt it toward the locations whose residents do better than average:

$$\pi_j(\varepsilon) = \pi_j \left\{1 + \varepsilon \left(\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}]\right)\right\}. \tag{33}$$

For small ε these are probabilities. The expected outcome of a person is $\sum_j \pi_j(\varepsilon)\, \psi_j$. Its derivative at $\varepsilon = 0$ is $\mathrm{Cov}_\pi(\psi_j, \bar{Y}_j)$. The mean outcome of the residents of her location is, in expectation, $\sum_j \pi_j(\varepsilon)\, \bar{Y}_j$. Its derivative is $\mathrm{Var}_\pi(\bar{Y}_j)$. The slope θ_π is the ratio of the two.

For each unit by which the policy raises the mean outcome of the residents of a person's location, it raises her own outcome by θ_π. With a slope of 0.68, a policy that brings people to locations whose residents do one unit better raises their own outcomes by 0.68 on average. This reading needs no additive model. It is a statement about a small change, which leaves the locations as they are.

The mover coefficient is also the return to a policy. It is Q_c with the comparison gap as the contrast c. Tilt the assignment toward locations with a high comparison gap. This moves people along the routes that movers take, from the location with the lower mean outcome toward the location with the higher one. The mover coefficient is the return to this policy. The cross-sectional coefficient is the return to the tilt (33), toward all locations whose residents do better than average.

E Two Mover Regressions for the Cross-Sectional Coefficient

Under (24), or with $s_{jk} = 0$ on every route (Section C), the cross-sectional coefficient can be estimated without estimating the effect of each location.

Proposition 8 (Two mover regressions for the cross-sectional coefficient)**.** *Let Assumption 3 hold, and let the routes with movers connect all locations.*

(i) The cross-sectional coefficient as a mover regression. *Under (24), θ_π is the coefficient of the regression of ΔY_i on $\Delta \bar{Y}_i$ among movers, with $\Delta\varphi_i$ as the instrument. Here φ is a transformed map that solves $\sum_k a_{jk}(\varphi_j - \varphi_k) = \pi_j(\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}])$ for all j. It is computed from flows and location means. It is the map itself, up to scale and a constant, if and only if routes are aligned.*

(ii) Reweighted routes. *Let $s_{jk} = 0$ on every route. Let $w_{jk} = w_{kj} \geq 0$ be weights on the pairs of locations that exchange movers, such that the design with comparison weights $w_{jk}\, a_{jk}$ is aligned with the map. Then the regression of ΔY_i on $\Delta \bar{Y}_i$ among movers, without a constant and with weight $w_{o(i)d(i)}$ on mover i, has coefficient θ_π, whatever the place effects are. Call a set of locations closed upward if it contains, for each of its members, all the partners with higher mean outcomes. Such weights exist if and only if every set that is closed upward has a mean outcome under π at least as high as that of all locations.*

Part (i) shows how close the two-way estimator and the mover regression are. The cross-sectional coefficient is also the coefficient of a mover regression. The only change is the instrument: the gap in the map along a move is replaced by the gap in a transformed map. If place effects are

an exact linear function of the map, the two regressions estimate the same number, and the usual one is the more precise. If not, only the instrumented regression estimates θ_π, and the difference between the two is the term in Theorem 1(i).

Part (ii) keeps the regression and changes the routes. Alignment is linear in the weights on routes, so the weights can be chosen to deliver it. This takes $J-2$ restrictions on the weights. Reweighting the routes to independent pairing would take movers on every pair of locations. The weights use only flows and the map, and they are not negative. The reweighted coefficient is therefore still an average of the pair shares of Proposition 2 over routes that movers take. The condition for existence fails when a group of locations is cut off. Take a location below the mean whose partners all have lower mean outcomes. Alignment needs it to be compared mostly with locations above it, and no weights can do that. When one set of weights satisfies the condition, many do, and the choice among them affects precision. The next paragraph gives an example, and Section L the details. Proposition 7 is the coarse version of this remedy. Weights that are constant within a few groups of routes need the assumption of the design in each group, and not on each route.

Reweighted routes. If each route shows the difference of two place effects, the routes can be reweighted until they are aligned (Proposition 8(ii)). Between the 211 large zones such weights exist. The file with the routes also gives an estimate of the contrast on each pair of zones. With the weights, these estimates give a coefficient of 0.50 (s.e. 0.10). With the observed routes they give 0.53, and the cross-sectional coefficient from the published place effects is 0.47. Across all 595 zones the condition for existence fails for five small zones, and the weights that come closest reach an alignment of 0.9999. The coefficient moves from 0.60 to 0.70 (s.e. 0.10), toward the cross-sectional coefficient of 0.75. The reweighted coefficients use no estimated place effect. The weights are uneven: a few routes carry much of the weight.

Proof of Proposition 8. (i) The equation for φ is $L_D\varphi = \Pi(\bar{Y} - \mathbb{E}_\pi[\bar{Y}])$. Its right-hand side sums to zero, so it has a solution when the routes connect the locations, unique up to a constant. Under (24), $\mathbb{E}_{\text{mov}}[s_i\,\Delta z_i] = 0$ for every attribute z of locations, so the extra gains drop out of the numerator. The coefficient of the instrumented regression is

$$\frac{\text{Cov}_{\text{mov}}(\Delta\varphi_i, \Delta Y_i)}{\text{Cov}_{\text{mov}}(\Delta\varphi_i, \Delta\bar{Y}_i)} = \frac{\varphi' L_D \psi}{\varphi' L_D \bar{Y}} = \frac{(\bar{Y} - \mathbb{E}_\pi[\bar{Y}])'\Pi\,\psi}{(\bar{Y} - \mathbb{E}_\pi[\bar{Y}])'\Pi\,\bar{Y}} = \theta_\pi.$$

If $\varphi = \kappa^{-1}\bar{Y} + \text{constant}$, then $L_D\bar{Y} = \kappa\Pi(\bar{Y} - \mathbb{E}_\pi[\bar{Y}])$, which is alignment, and conversely.

(ii) With $s_{jk} = 0$, the mean of ΔY_i on the route from j to k is $\psi_k - \psi_j$. The weighted regression without a constant therefore has coefficient $\sum_{j\neq k} w_{jk}F_{jk}(\psi_k - \psi_j)(\bar{Y}_k - \bar{Y}_j) \big/ \sum_{j\neq k} w_{jk}F_{jk}(\bar{Y}_k - \bar{Y}_j)^2$. This is (11) for the design with comparison weights $w_{jk}a_{jk}$. That design is aligned, so Theorem 1(ii) applies.

For existence, let E be the set of ordered pairs (a, b) with $a_{ab} > 0$ and $\bar{Y}_a < \bar{Y}_b$, and put $\phi_{ab} = w_{ab}\,a_{ab}\,(\bar{Y}_b - \bar{Y}_a) \geq 0$. The comparison gap of the reweighted design satisfies $\pi_j g_j =$

$\sum_{a:(a,j)\in E}\phi_{aj} - \sum_{b:(j,b)\in E}\phi_{jb}$. Write the right-hand side as $(N\phi)_j$. Alignment is $N\phi = \kappa\, t$ with $t_j = \pi_j(\bar{Y}_j - \mathbb{E}_\pi[\bar{Y}])$ and $\kappa > 0$, and by scaling ϕ one can take $\kappa = 1$. Then $\sum_{(a,b)\in E}\phi_{ab}(\bar{Y}_b - \bar{Y}_a) = \bar{Y}'N\phi = \mathrm{Var}_\pi(\bar{Y}_j) > 0$, so the reweighted design has a positive denominator. By Farkas' lemma, $N\phi = t$ has a solution $\phi \geq 0$ if and only if $p't \geq 0$ for every $p \in \mathbb{R}^J$ with $N'p \geq 0$, that is, with $p_b \geq p_a$ for all $(a,b) \in E$. Call a set U of locations closed upward if $a \in U$ and $(a,b) \in E$ imply $b \in U$. The indicator of such a set is such a p, so the condition requires $\sum_{j\in U} t_j \geq 0$. Conversely, let p satisfy $p_b \geq p_a$ on E, and let its distinct values be $c_1 < \cdots < c_m$. The sets $U_l = \{j : p_j \geq c_l\}$ are closed upward, and $p = c_1\mathbf{1} + \sum_{l\geq 2}(c_l - c_{l-1})\mathbf{1}_{U_l}$. Since $\mathbf{1}'t = 0$, $p't = \sum_{l\geq 2}(c_l - c_{l-1})\sum_{j\in U_l} t_j \geq 0$. Finally, $\sum_{j\in U} t_j \geq 0$ says that the mean of the map over U, under π, is at least its mean over all locations. □

F Sorting Along Routes: Directional Isotropy

F.1 The Condition

With place effects in hand, a second condition for the equality of the two coefficients can be examined. It concerns how sorting along the routes compares with sorting across locations. Example 1 has two dimensions, and so does every map. Composition has a part that moves with the place effect. Regions whose providers treat more intensively may also have sicker residents, and neighborhoods that do more for children may attract more affluent families. It has a second part that does not. Two regions with the same practice style can have different residents because they differ in climate, housing costs, or age structure. Two numbers describe them. Across locations, let

$$\beta_\pi := \frac{\mathrm{Cov}_\pi(\psi_j, \bar{\alpha}_j)}{\mathrm{Var}_\pi(\psi_j)}, \qquad \lambda_\pi := \frac{\mathrm{Var}_\pi(\bar{\alpha}_j)}{\mathrm{Var}_\pi(\psi_j)} - \beta_\pi^2. \tag{34}$$

The *sorting slope* β_π says how much composition rises with the place effect. The *residual composition* λ_π is the variance of the part of composition that is unrelated to the place effect, relative to the variance of place effects. The same two numbers can be computed along the routes, with $\mathcal{L}_D$ in the place of the variances and the covariance:

$$\beta_D := \frac{\mathcal{L}_D(\psi, \bar{\alpha})}{\mathcal{L}_D(\psi, \psi)}, \qquad \lambda_D := \frac{\mathcal{L}_D(\bar{\alpha}, \bar{\alpha})}{\mathcal{L}_D(\psi, \psi)} - \beta_D^2. \tag{35}$$

Here β_D says how much the composition of the location changes, on average, when a move changes the place effect by one unit. λ_D measures the changes in composition that come with no change in the place effect.

Each coefficient is the same function of its two numbers,

$$\theta_D = f(\beta_D, \lambda_D), \qquad \theta_\pi = f(\beta_\pi, \lambda_\pi), \qquad f(\beta, \lambda) := \frac{1+\beta}{(1+\beta)^2 + \lambda}. \tag{36}$$

In these terms, the size of what the map misses is $\eta = \sqrt{\lambda_\pi}\,/\{(1+\beta_\pi)^2 + \lambda_\pi\}$. It is zero exactly

when there is no residual composition.

Theorem 4 (Directional isotropy). *Let* $\mathcal{L}_D(\psi,\psi)$, $\mathrm{Var}_\pi(\psi_j)$, $\mathcal{L}_D(\bar{Y},\bar{Y})$ *and* $\mathrm{Var}_\pi(\bar{Y}_j)$ *be positive.*

(i) $(\beta_D,\lambda_D)=(\beta_\pi,\lambda_\pi)$ *if and only if there is a number* $\kappa>0$ *with*

$$\mathcal{L}_D(\psi,\psi)=\kappa\,\mathrm{Var}_\pi(\psi_j),\qquad \mathcal{L}_D(\psi,\bar{\alpha})=\kappa\,\mathrm{Cov}_\pi(\psi_j,\bar{\alpha}_j),\qquad \mathcal{L}_D(\bar{\alpha},\bar{\alpha})=\kappa\,\mathrm{Var}_\pi(\bar{\alpha}_j).$$

(ii) *In that case* $\theta_D=\theta_\pi$. *The converse does not hold.*

(iii) *If composition is an exact linear function of the place effect,* $\lambda_\pi=0$, *then* (i) *holds for every design, and* $\theta_D=\theta_\pi=1/(1+\beta_\pi)$.

I call the condition in (i) *directional isotropy*. The number κ says whether moves span larger or smaller differences than the map does, and it is free. What the condition requires is that moves change place effects and composition in the same proportions as locations differ in them. In Example 1 this happens at $\rho=\frac{1}{2}$ and nowhere else. If most moves are across practice styles, composition hardly changes along them, λ_D is close to zero, and the coefficient is close to one, however much residents differ across regions. If most moves are across patient populations, the coefficient is close to zero, however much practice differs across regions.

Two familiar readings of the coefficient are special cases of the theorem. The first is the share of the variance. If composition is unrelated to the place effect, $\beta_\pi=0$, the three shares of Lemma 5 coincide, and under directional isotropy the mover coefficient equals all of them. This is the condition under which Chetty et al. [2026] read their slope as a share of the variance. It asks for more than a choice of location that is unrelated to the outcome. People who move for the climate or for family may still differ in what they would attain anywhere, and the locations they choose may differ in their effects. The reading needs zero sorting, and it needs the two coefficients to agree, as they do under isotropy or with aligned routes.

The second is the ratio of standard deviations. Suppose composition is an exact linear function of the place effect. A stylized case is that families sort across neighborhoods on what a neighborhood does for children, and on nothing else. Then by part (iii) the coefficient is $1/(1+\beta_\pi)$ whatever the routes are. Routes cannot matter, because moves have no second dimension to select on. For $\beta_\pi>-1$ the coefficient is then the ratio of the standard deviation of place effects to that of the map, and the variance share is its square. A coefficient of 0.20, often described as locations explaining twenty percent of the variation, is then a variance share of four percent.

Directional isotropy is a statement about place effects and composition. It can be tested when place effects can be estimated (Section F.2).

Alignment and isotropy. Independent pairing implies both alignment and directional isotropy, and each of the two implies $\theta_D=\theta_\pi$. Neither implies the other. Alignment restricts routes and the map, and it delivers equality for all place effects. Isotropy restricts routes, place effects, and composition, and it delivers more than equality: movers then also reproduce the sorting slope and the residual composition of the cross-section.

Proof of Theorem 4. Since $\bar{Y} = \psi + \bar{\alpha}$, $\mathcal{L}_D(\psi, \bar{Y}) = \mathcal{L}_D(\psi, \psi) + \mathcal{L}_D(\psi, \bar{\alpha})$ and $\mathcal{L}_D(\bar{Y}, \bar{Y}) = \mathcal{L}_D(\psi, \psi) + 2\mathcal{L}_D(\psi, \bar{\alpha}) + \mathcal{L}_D(\bar{\alpha}, \bar{\alpha})$. Dividing both by $\mathcal{L}_D(\psi, \psi)$ gives $\theta_D = f(\beta_D, \lambda_D)$. The same algebra with Cov_π gives $\theta_\pi = f(\beta_\pi, \lambda_\pi)$.

(i) If the three moments are proportional, the ratios in (34) and (35) coincide. Conversely, let $\kappa = \mathcal{L}_D(\psi, \psi) / \mathrm{Var}_\pi(\psi_j) > 0$. Then $\beta_D = \beta_\pi$ gives $\mathcal{L}_D(\psi, \bar{\alpha}) = \kappa \, \mathrm{Cov}_\pi(\psi_j, \bar{\alpha}_j)$, and $\lambda_D = \lambda_\pi$ gives $\mathcal{L}_D(\bar{\alpha}, \bar{\alpha}) = \kappa \, \mathrm{Var}_\pi(\bar{\alpha}_j)$.

(ii) follows from (36). For the converse, take Example 1 with $c = 1$, so that $(\beta_\pi, \lambda_\pi) = (0, 1)$ and $\theta_\pi = \frac{1}{2}$. The design that only compares the locations $(-1, -1)$ and $(1, 1)$ has $(\beta_D, \lambda_D) = (1, 0)$ and $\theta_D = f(1, 0) = \frac{1}{2}$.

(iii) If $\bar{\alpha}_j = a + \beta\psi_j$ for all j, then for every design $\mathcal{L}_D(\psi, \bar{\alpha}) = \beta \, \mathcal{L}_D(\psi, \psi)$ and $\mathcal{L}_D(\bar{\alpha}, \bar{\alpha}) = \beta^2 \mathcal{L}_D(\psi, \psi)$, so $(\beta_D, \lambda_D) = (\beta, 0)$. The same holds in the cross-section, and $f(\beta, 0) = 1/(1 + \beta)$. □

F.2 The Test

Where place effects can be estimated, both coefficients can be estimated, and directional isotropy can be tested. By Theorem 4, isotropy is two restrictions,

$$g_\beta := \beta_D - \beta_\pi = 0, \qquad g_Q := (\lambda_D - \lambda_\pi) \, \mathcal{L}_D(\psi, \psi) \, \mathrm{Var}_\pi(\psi_j) = 0. \tag{37}$$

The first asks whether composition rises with the place effect along moves as much as it does across locations. The second asks whether moves change the residual part of composition as much as locations differ in it.

Both sides of each restriction are built from six second moments: the variances and the covariance of place effects and composition across locations, and the same three moments along routes. Place effects and location means are estimated with noise, and the square of a noisy estimate is too large. This is the same inflation that produces limited-mobility bias [Andrews et al., 2008, Kline et al., 2020]. Here it is removed by cross-fitting. Individuals are split at random into two halves. In each half, place effects are estimated from that half's movers, and location means from that half's residents. Every product in the six moments then takes one factor from each half, so the errors of the two factors are independent.

Every estimator in this construction is linear in outcomes. The influence of each individual on the six moments is therefore available in closed form, and a multiplier bootstrap over individuals needs no re-estimation. The same six moments give θ_π, the two shares of the variance, and the difference $\theta_D - \theta_\pi$, each with a confidence interval. Let $\hat{g} = (\hat{g}_\beta, \hat{g}_Q)'$, let $\hat{V}_g$ be its bootstrap variance, and let $W_{\text{iso}} = \hat{g}' \hat{V}_g^{-1} \hat{g}$.

Proposition 9 (Inference with estimated place effects)**.** *Let J be fixed and $N \to \infty$. Suppose individuals are independent and identically distributed, and the split is an independent fair coin for each individual. Outcomes have finite moments of order $4 + \epsilon$ for some $\epsilon > 0$. Every location has a positive share of residents in each half, and the movers of each half connect all locations. Suppose* $\mathrm{Var}_\pi(\psi_j) > 0$, $\mathcal{L}_D(\psi, \psi) > 0$, *and*

the limiting variance Ω *of* $\sqrt{N}\,\hat{g}$ *is nonsingular. Then*

$$\sqrt{N}\,\{\hat{g} - g\} \xrightarrow{d} \mathcal{N}(0, \Omega), \qquad N\hat{V}_g \xrightarrow{p} \Omega,$$

and under directional isotropy W_{iso} *converges in distribution to a chi-squared variable with two degrees of freedom.*

The proof is a smooth-map argument. With a fixed number of locations, the estimates of place effects and location means are smooth functions of finitely many sample averages. Section F.3 gives the estimators, the influence functions, and the bootstrap. The result excludes the case $\lambda_\pi = 0$, where the variance of the second restriction is singular. In that case no test is needed: by Theorem 4(iii) the coefficient equals θ_π for every design.

When the test rejects, the two estimates say which way the routes differ. The first, $\hat{g}_\beta$, says whether sorting is stronger or weaker along moves than across locations. The second, $\hat{g}_Q$, says whether moves change residual composition more or less than locations differ in it.

F.3 Estimators, Influence Functions, and the Bootstrap

Throughout this subsection the weights π are mover incidence, the shares of moves that involve each location, and are written μ. The six moments are named $a, \ldots, f$: the first three are $\text{Var}_\pi(\psi_j)$, $\text{Cov}_\pi(\psi_j, \bar{\alpha}_j)$ and $\text{Var}_\pi(\bar{\alpha}_j)$, and the last three are their counterparts along routes, $\mathcal{L}_D(\psi, \psi)$, $\mathcal{L}_D(\psi, \bar{\alpha})$ and $\mathcal{L}_D(\bar{\alpha}, \bar{\alpha})$. In this subsection J is the number of locations that are retained in both halves of the sample.

This subsection states first-order inference for a fixed calendar window $\mathcal{T}$. The paper's two-period estimator is the case $T = 2$; for $T > 2$, the result is its fixed-window implementation extension under (59). The length T is fixed as $N \to \infty$, the additive model (4) is understood period by period over the window, and the fixed-window residual-exogeneity condition of Assumption 5 is maintained. A one-time mover enters the first stage through $\Delta Y_i = \Delta\psi_i + \bar{\varepsilon}_{i,d} - \bar{\varepsilon}_{i,o}$: the structural signal is the paper's two-period contrast, while its residual is averaged over the two spells. Exogenous mobility and Assumption 3, for the limiting retained analysis population, are maintained throughout. The sampling unit is the individual: individual histories are independent and identically distributed draws from a common population—place effects are fixed parameters, and a draw comprises the individual's type, assignment path, eligibility indicators, and outcomes—and observations may be arbitrarily dependent within a individual. The split is data-independent: the analysis takes it to be an independent fair coin per individual, and Remark 1 covers the exactly balanced partition used in the implementation as well as user-supplied fixed splits. All moments along routes are centered over retained movers, matching the intercept of the mover regression.

Split samples and retained locations. Let R_{is} indicate that the complete history of individual i is assigned to split $s \in \{1, 2\}$; the split is assigned independently of the individual data. The set of

retained locations is $\mathcal{J} = \{1, \dots, J\}$. Let M_i indicate a retained mover, write $o(i)$ and $d(i)$ for its origin and destination, and define

$$D_i := \mathbf{e}_{d(i)} - \mathbf{e}_{o(i)}, \qquad \iota_i := \tfrac{1}{2}\{\mathbf{e}_{o(i)} + \mathbf{e}_{d(i)}\}, \qquad m := \sum_{i=1}^{N} M_i.$$

The estimated mover-incidence weights are

$$\widehat{\mu} := \frac{1}{m}\sum_{i=1}^{N} M_i \iota_i, \qquad \mathbf{1}'\widehat{\mu} = 1. \tag{38}$$

The split-s place-effect regression may use a larger connected estimation universe $\mathcal{J}_s^E \supseteq \mathcal{J}$, of size J_s. Let $D_{is}^E \in \mathbb{R}^{J_s}$ be mover i's incidence vector there, write $o_s^E(i), d_s^E(i)$ for its endpoint coordinates, let E_{is} indicate inclusion in that split regression, and let $S_s \in \{0,1\}^{J \times J_s}$ select the coordinates of the retained locations. With ΔY_i equal to the destination-spell outcome average minus the origin-spell average, set

$$L_s := \sum_{i=1}^{N} E_{is} D_{is}^E D_{is}^{E\prime}, \qquad b_s := \sum_{i=1}^{N} E_{is} D_{is}^E \Delta Y_i, \qquad \widehat{\psi}_E^s := L_s^+ b_s, \qquad \widehat{\psi}^s := S_s \widehat{\psi}_E^s. \tag{39}$$

The Moore–Penrose normalization is used for notation only; every reported object is invariant to an additive constant in $\widehat{\psi}_E^s$.

Common full-window means and the canonical leave-one-out gap. Fix a calendar window $\mathcal{T}$ with $T := |\mathcal{T}|$ fixed. Let A_{it} indicate eligibility of observation (i,t) for a location mean. In the paper's two-period case, $A_{it} \equiv 1$: its two observations are the structural origin and destination periods. In the multi-period implementation, the convention excludes a mover's transition-period observation from the location means while stayers are always retained, independently of whether that observation appears in the event-study outcome sample. For each core location, period, and split, define

$$n_{jt}^s := \sum_i R_{is} A_{it} \mathbf{1}\{j(i,t) = j\}, \qquad \overline{Y}_{jt}^s := \frac{\sum_i R_{is} A_{it} \mathbf{1}\{j(i,t) = j\} Y_{it}}{n_{jt}^s},$$

and the equally weighted window average

$$\widehat{z}_j^s := \frac{1}{T}\sum_{t \in \mathcal{T}} \overline{Y}_{jt}^s, \qquad \widehat{\bar{\alpha}}_j^s := \widehat{z}_j^s - \widehat{\psi}_j^s. \tag{40}$$

By equation (59), the population counterpart of $\widehat{z}_j^s$ is the same $\bar{Y}_j = \psi_j + \bar{\alpha}_j$ used by the two-period target. The order of aggregation follows Finkelstein et al. [2016]: individuals are averaged within location-period first, and the periods are then averaged with equal weights. A location is retained only if $n_{jt}^s > 0$ for every $t \in \mathcal{T}$ and both splits; otherwise the location and its movers are

removed. The window is never shortened or renormalized. All six moments therefore use common, non-leave-one-out, split-specific full-window means.

Only the canonical mover regression and the event study use leave-one-out. With pooled full-sample counts

$$n_{jt} := \sum_{\ell} A_{\ell t} \mathbf{1}\{j(\ell, t) = j\}, \qquad \delta_{ijt} := A_{it} \mathbf{1}\{j(i, t) = j\},$$

define

$$\overline{Y}_{jt,-i} := \frac{\sum_{\ell} A_{\ell t} \mathbf{1}\{j(\ell, t) = j\} Y_{\ell t} - \delta_{ijt} Y_{it}}{n_{jt} - \delta_{ijt}}, \qquad \widehat{z}_{j,-i} := \frac{1}{T} \sum_{t \in \mathcal{T}} \overline{Y}_{jt,-i},$$

and $\widehat{\delta}_i^{(-i)} := \widehat{z}_{d(i),-i} - \widehat{z}_{o(i),-i}$. If any of the $2T$ endpoint-period denominators is nonpositive, the entire canonical gap is missing; no averaging over the remaining periods is performed. Writing $\overline{Y}_{jt}$ for the pooled mean,

$$\overline{Y}_{jt,-i} - \overline{Y}_{jt} = \frac{\delta_{ijt}(\overline{Y}_{jt} - Y_{it})}{n_{jt} - \delta_{ijt}}.$$

Let $n_{\min} := \min_{j \in \mathcal{J}, t \in \mathcal{T}} n_{jt}$. Under a uniform $4 + \eta$ moment bound, $\max_{i,t} |Y_{it} - \overline{Y}_{j(i,t),t}| = O_p(N^{1/(4+\eta)})$, so $n_{\min} / N^{1/2+1/(4+\eta)} \to \infty$ suffices for the leave-one-out perturbation to be $o_p(N^{-1/2})$ uniformly in i. Under the fixed-J conditions of Proposition 10, retained cell counts grow proportionally with N, so this requirement holds automatically.

Six cross-fitted moments. Let

$$\widehat{P}_{\mu} := I - \mathbf{1}\widehat{\mu}', \qquad p_s := \widehat{P}_{\mu} \widehat{\psi}^s, \qquad r_s := \widehat{P}_{\mu} \widehat{\alpha}^s.$$

For retained mover i, define

$$x_{is} := D_i' \widehat{\psi}^s, \qquad v_{is} := D_i' \widehat{\alpha}^s,$$

and let x_{is}^c and v_{is}^c denote centering over the m retained movers. The primitive estimators are

$$\begin{aligned}
\widehat{a} &:= \sum_j \widehat{\mu}_j p_{1j} p_{2j}, & \widehat{b} &:= \frac{1}{2} \sum_j \widehat{\mu}_j (p_{1j} r_{2j} + p_{2j} r_{1j}), \\
\widehat{c} &:= \sum_j \widehat{\mu}_j r_{1j} r_{2j}, & \widehat{d} &:= \frac{1}{m} \sum_i M_i x_{i1}^c x_{i2}^c, \\
\widehat{e} &:= \frac{1}{2m} \sum_i M_i (x_{i1}^c v_{i2}^c + x_{i2}^c v_{i1}^c), & \widehat{f} &:= \frac{1}{m} \sum_i M_i v_{i1}^c v_{i2}^c.
\end{aligned} \tag{41}$$

Every quadratic moment multiplies one split-1 estimate by one split-2 estimate, so first-stage estimation errors, independent across splits, contribute no leading-order plug-in bias [Kline, 2025].

Write $\widehat{\vartheta} = (\widehat{a}, \widehat{b}, \widehat{c}, \widehat{d}, \widehat{e}, \widehat{f})'$. For $a > 0$ and $d > 0$, let

$$Q_\pi := c - \frac{b^2}{a}, \qquad Q_D := f - \frac{e^2}{d},$$

and define

$$g(\vartheta) := \begin{pmatrix} g_\beta \\ g_Q \end{pmatrix} = \begin{pmatrix} e/d - b/a \\ (f - e^2/d)a - (c - b^2/a)d \end{pmatrix}. \tag{42}$$

Direct contribution and gradients. For a retained mover, set

$$h_i := \begin{pmatrix} \frac{1}{2}\sum_{\ell \in \{o(i), d(i)\}} p_{1\ell} p_{2\ell} \\ \frac{1}{4}\sum_{\ell \in \{o(i), d(i)\}} (p_{1\ell} r_{2\ell} + p_{2\ell} r_{1\ell}) \\ \frac{1}{2}\sum_{\ell \in \{o(i), d(i)\}} r_{1\ell} r_{2\ell} \\ x^c_{i1} x^c_{i2} \\ \frac{1}{2}(x^c_{i1} v^c_{i2} + x^c_{i2} v^c_{i1}) \\ v^c_{i1} v^c_{i2} \end{pmatrix}. \tag{43}$$

Then

$$\widehat{\vartheta} = \frac{1}{m}\sum_i M_i h_i, \qquad \widehat{\phi}_i^{\text{edge}} := \frac{M_i}{m}(h_i - \widehat{\vartheta}). \tag{44}$$

This contribution includes estimation of $\widehat{\mu}$ and the mover centering: derivatives of the weighted centers vanish because the opposite split's centered vector has weighted mean zero.

Suppress hats temporarily, let $D_\mu := \operatorname{diag}(\mu)$, and for a mover-centered scalar u_i define

$$\tau(u) := \frac{1}{m}\sum_i M_i D_i u_i.$$

The gradients with respect to the core vectors (ψ^1, z^1) are

$$G^1_\psi = \begin{pmatrix} (D_\mu p_2)' \\ \frac{1}{2}\{D_\mu (r_2 - p_2)\}' \\ -(D_\mu r_2)' \\ \tau(x^c_2)' \\ \frac{1}{2}\{\tau(v^c_2) - \tau(x^c_2)\}' \\ -\tau(v^c_2)' \end{pmatrix}, \qquad G^1_z = \begin{pmatrix} 0' \\ \frac{1}{2}(D_\mu p_2)' \\ (D_\mu r_2)' \\ 0' \\ \frac{1}{2}\tau(x^c_2)' \\ \tau(v^c_2)' \end{pmatrix}. \tag{45}$$

The matrices G^2_ψ, G^2_z interchange subscripts 1 and 2. Their rows sum to zero, which is the normalization invariance. Rows five and six of G^s_z are the transition response to the same common full-window mean used in the cross-sectional moments. A hat on any gradient below means evaluation at the corresponding estimates.

The restriction Jacobian is

$$G_g(\vartheta) = \begin{pmatrix} b/a^2 & -1/a & 0 & -e/d^2 & 1/d & 0 \\ f - e^2/d - b^2 d/a^2 & 2bd/a & -d & ae^2/d^2 - c + b^2/a & -2ae/d & a \end{pmatrix}, \qquad \widehat{G}_g := G_g(\widehat{\vartheta}). \tag{46}$$

For reported shares, write $V_\pi := a + 2b + c$ and $V_D := d + 2e + f$. On $V_\pi, V_D > 0$, define

$$\theta_{\text{var}} := a/V_\pi, \qquad \theta_\pi := (a+b)/V_\pi, \qquad S_{\text{eq}} := (a+2b)/V_\pi, \qquad \theta_D := (d+e)/V_D,$$

with Jacobian, in this order,

$$G_\theta(\vartheta) = \begin{pmatrix} (2b+c)/V_\pi^2 & -2a/V_\pi^2 & -a/V_\pi^2 & 0 & 0 & 0 \\ (b+c)/V_\pi^2 & (c-a)/V_\pi^2 & -(a+b)/V_\pi^2 & 0 & 0 & 0 \\ c/V_\pi^2 & 2c/V_\pi^2 & -(a+2b)/V_\pi^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & (e+f)/V_D^2 & (f-d)/V_D^2 & -(d+e)/V_D^2 \end{pmatrix}. \tag{47}$$

When $\theta_{\text{var}} > 0$, the derivative of $\theta_{\text{sd}} = \sqrt{\theta_{\text{var}}}$ is the first row of (47) divided by $2\sqrt{\theta_{\text{var}}}$.

Individual influence. For a split-s estimation mover, define

$$\widehat{u}_i^s := \Delta Y_i - D_{is}^{E\prime} \widehat{\psi}_E^s, \qquad \dot{\psi}_{E,i}^s := E_{is} L_s^+ D_{is}^E \widehat{u}_i^s.$$

The full-window cell-mean contribution on the core is

$$\dot{z}_i^s := R_{is} \sum_{t \in \mathcal{T}} \frac{1}{T} A_{it} \mathbf{1}\{j(i,t) \in \mathcal{J}\} \mathbf{e}_{j(i,t)} \frac{Y_{it} - \overline{Y}^s_{j(i,t),t}}{n^s_{j(i,t),t}}.$$

The estimated first-order contribution of individual i is

$$\widehat{\phi}_i = \widehat{\phi}_i^{\text{edge}} + \sum_{s=1}^{2} \left(\widehat{G}_\psi^s S_s \dot{\psi}_{E,i}^s + \widehat{G}_z^s \dot{z}_i^s \right). \tag{48}$$

Equivalently, with

$$W_s := L_s^+ S_s' \widehat{G}_\psi^{s\prime},$$

the place-effect term is

$$\widehat{G}_\psi^s S_s \dot{\psi}_{E,i}^s = E_{is} \{ W_s[d_s^E(i), \cdot] - W_s[o_s^E(i), \cdot] \}' \widehat{u}_i^s,$$

so one six-right-hand-side Laplacian solve is required per split. The normal equations, the within-cell residual identities, and (44) imply $\sum_i \widehat{\phi}_i = 0$. The finite-sample covariance estimate and its

root-N scaling are

$$\widehat{\mathrm{Var}}(\widehat{\vartheta}) := \sum_i \widehat{\phi}_i \widehat{\phi}_i', \qquad \widehat{\Sigma}_\vartheta := N \sum_i \widehat{\phi}_i \widehat{\phi}_i'. \tag{49}$$

Finite-dimensional reduction. The proposition below exploits that, with J, J_s, and T fixed, the entire construction is one smooth map of a fixed-dimensional vector of sample means. Collect for each individual the graph block, the cell block, and the endpoint block,

$$\begin{aligned} W_i^{\mathrm{g}} &:= \left\{ E_{is} \mathbf{1}\{o_s^E(i) = j,\, d_s^E(i) = k\} \,(1,\ \Delta Y_i) \right\}_{j \neq k,\, s}, \\ W_i^{\mathrm{c}} &:= \left\{ R_{is} A_{it} \mathbf{1}\{j(i,t) = j\} \,(1,\ Y_{it}) \right\}_{j \in \mathcal{J},\, t \in \mathcal{T},\, s}, \qquad W_i^{\mathrm{e}} := \left\{ M_i \mathbf{1}\{o(i) = j,\, d(i) = k\} \right\}_{j \neq k}, \end{aligned} \tag{50}$$

and set $W_i := (W_i^{\mathrm{g}}, W_i^{\mathrm{c}}, W_i^{\mathrm{e}})$, $\overline{W} := N^{-1} \sum_i W_i$, and $\varpi := \mathbb{E}[W_i]$, common to all individuals under the sampling frame below. Then $\widehat{\vartheta} = \Phi(\overline{W})$, where Φ is infinitely differentiable on the open set $\mathcal{O}$ on which all retained cell counts are positive, the share of retained movers is positive, and both split estimation graphs are connected: on $\mathcal{O}$ the split Laplacian has constant rank $J_s - 1$, so L_s^+ is a smooth function of the edge shares, and every remaining operation is a ratio, product, or centering. Define the population target and influence

$$\vartheta := \Phi(\varpi), \qquad \phi_i := \frac{1}{N} \nabla\Phi(\varpi)\,(W_i - \varpi), \tag{51}$$

so that $\sum_i \phi_i = \nabla\Phi(\varpi)(\overline{W} - \varpi)$ and $\mathbb{E}\phi_i = 0$. Evaluating the three blocks of $\nabla\Phi$—the endpoint shares, the split first stages, and the cell sums and counts—yields exactly the three terms of (48) with population values in place of estimates: the endpoint block gives the edge kernel, the chain rule through $\widehat{\psi}_E^s = L_s^+ b_s$ gives the $E_{is} L_s^+ D_{is}^E u_i^s$ term, and the ratio expansion of each cell mean gives $\dot{z}_i^s$, with the gradient matrices (45).

Proposition 10 (Fixed-J asymptotic linearity)**.** *Fix $\mathcal{J}$, $\mathcal{J}_s^E$, S_s, and $\mathcal{T}$; if any of these is selected from the data, it equals the fixed object with probability approaching one. Suppose:*

(i) individual histories—including the split indicator, drawn as an independent Bernoulli($\frac{1}{2}$) assignment per individual, independently of the remaining individual data—are independent and identically distributed across individuals, with $\sup_t \mathbb{E}|Y_{it}|^{4+\eta} < \infty$ for some $\eta > 0$;

(ii) the share of retained movers and the share of every retained cell are positive: $\pi_M := \Pr(M_i = 1) > 0$ and $\pi_{jt}^s := \Pr\left(R_{is} A_{it} \mathbf{1}\{j(i,t) = j\} = 1\right) > 0$;

(iii) $L_s^0 := \mathbb{E}\left[E_{is} D_{is}^E D_{is}^{E\prime}\right]$ is the Laplacian of a connected graph on $\mathcal{J}_s^E$: $\operatorname{rank} L_s^0 = J_s - 1$ and $\lambda_2(L_s^0) > 0$.

Set $\vartheta := \Phi(\varpi)$ and $\Sigma_\vartheta := \mathrm{Var}\{\nabla\Phi(\varpi)\, W_i\} = \mathbb{E}[(N\phi_i)(N\phi_i)']$, which is finite under (i). Then

$$\widehat{\vartheta} - \vartheta = \sum_{i=1}^N \phi_i + O_p(N^{-1}), \qquad \sqrt{N}(\widehat{\vartheta} - \vartheta) \xrightarrow{d} \mathcal{N}(0, \Sigma_\vartheta), \tag{52}$$

and moreover

$$N \max_i \|\widehat{\phi}_i - \phi_i\| \xrightarrow{p} 0, \qquad N \sum_i \widehat{\phi}_i \widehat{\phi}_i' \xrightarrow{p} \Sigma_\vartheta.$$

Proof. By the law of large numbers, $\overline{W} \xrightarrow{p} \varpi$ and $N^{-1} L_s \xrightarrow{p} L_s^0$; by (ii)–(iii), $\varpi \in \mathcal{O}$. Fix compacts $K \subset K' \subset \mathcal{O}$ with ϖ interior to K, so that $\Pr\{\overline{W} \in K'\} \to 1$ and $\nabla\Phi$ is bounded and Lipschitz with a bounded Hessian on K'. The entries of W_i are indicators or indicators times Y_{it} or ΔY_i, so they inherit finite $4+\eta$ moments from (i), and $\|\overline{W} - \varpi\| = O_p(N^{-1/2})$. On the event $\{\overline{W} \in K'\}$, a second-order Taylor expansion gives

$$\Phi(\overline{W}) - \Phi(\varpi) = \nabla\Phi(\varpi)(\overline{W} - \varpi) + O_p\big(\|\overline{W} - \varpi\|^2\big) = \sum_i \phi_i + O_p(N^{-1}),$$

which is the linear representation; the identification of $\nabla\Phi(\varpi)(W_i - \varpi)/N$ with the population counterpart of (48) is the block computation described above. The scaled scores $N\phi_i$ are independent and identically distributed, mean zero, with finite variance, so the classical central limit theorem gives the normal limit, the $O_p(N^{-1})$ remainder being $o_p(N^{-1/2})$.

For the influence accuracy, write $\widehat{\phi}_i = N^{-1}\nabla\Phi(\overline{W})(W_i - \overline{W})$, which is the estimated-gradient, estimated-centering version of ϕ_i and coincides term by term with (48). Under the sampling frame (i) there is no centering gap—$\mathbb{E}W_i = \varpi$ for every individual—so

$$N(\widehat{\phi}_i - \phi_i) = \{\nabla\Phi(\overline{W}) - \nabla\Phi(\varpi)\}(W_i - \overline{W}) + \nabla\Phi(\varpi)(\varpi - \overline{W}).$$

The gradient difference is $O_p(N^{-1/2})$ by the Lipschitz property, while $\max_i \|W_i\| = O_p(N^{1/(4+\eta)})$ under the moment bound, so $N \max_i \|\widehat{\phi}_i - \phi_i\| = O_p(N^{1/(4+\eta)-1/2}) + O_p(N^{-1/2}) = o_p(1)$. Consequently $N\sum_i \|\widehat{\phi}_i - \phi_i\|^2 \le \big(N \max_i \|\widehat{\phi}_i - \phi_i\|\big)^2 \to_p 0$, and by Cauchy–Schwarz it suffices that $N\sum_i \phi_i\phi_i' = N^{-1}\sum_i (N\phi_i)(N\phi_i)' \to_p \Sigma_\vartheta$, which is the law of large numbers for independent and identically distributed matrices with a finite mean. □

Two consequences are worth recording. First, no remainder condition is assumed anywhere: at fixed dimension the $O_p(N^{-1})$ remainder is delivered by Taylor's theorem, and the cross-split structure is needed only for the *centering* of the quadratic moments, not for their rates. Second, the leave-one-out refinement of the canonical regressor is covered by the earlier bound, since (ii) implies $n_{\min} \asymp N$.

Remark 1 (Balanced and conditioned splits). The implementation assigns an exactly balanced split without replacement, and a user-supplied split is fixed rather than random. Neither departure alters the influence representation, for a structural reason: Φ is degree-zero homogeneous in each split block of $\overline{W}$ separately—$\widehat{\psi}_E^s = L_s^+ b_s$ is invariant to a common scaling of the split-s edge counts and sums, the cell means are within-split ratios, and the endpoint block enters only through m-normalized shares—so by Euler's identity $\nabla\Phi(\varpi)$ annihilates the mean direction of each block. The difference between an exactly balanced partition and independent fair coins, and likewise the between-stratum mean gap that arises when conditioning on the realized balanced partition, lie

exactly along those block-mean directions and are therefore removed by the gradient: the influence functions, the asymptotic covariance, and the bootstrap are unchanged. The same holds for any fixed partition whose stratum shares converge to one half. For general fixed shares the argument applies verbatim with the block masses of ϖ set to those shares: the target $\Phi(\varpi)$ is unchanged by block homogeneity, while Σ_ϑ then reflects the split proportions.

Interior test and Gaussian multiplier bootstrap. The restriction test requires the population moments to satisfy $a, d > 0$; inference for the reported share ratios additionally requires V_π and V_D to be bounded away from zero along the asymptotic sequence. Define

$$\Omega := G_g(\vartheta)\, \Sigma_\vartheta\, G_g(\vartheta)',$$

and assume $\Omega \succ 0$ for the joint test. Draw independent $\xi_i^{(b)} \sim \mathcal{N}(0,1)$ and form

$$\widehat{\vartheta}^{*(b)} := \widehat{\vartheta} + \sum_i \xi_i^{(b)} \widehat{\phi}_i, \qquad \widehat{g}^{*(b)} := g(\widehat{\vartheta}^{*(b)}). \tag{53}$$

Conditional on the data, the primitive perturbation is Gaussian with covariance $\sum_i \widehat{\phi}_i \widehat{\phi}_i'$; the implementation draws this six-dimensional Gaussian directly, which is distributionally identical to using Gaussian individual multipliers.

Let $\mathcal{B}_v$ index draws for which the positive a, d checks pass and $\widehat{g}^{*(b)}$ is finite, let $B_v := |\mathcal{B}_v|$, and let $\overline{g}^*$ be their mean. If $B_v < 3$, the joint test is reported missing; otherwise define

$$\widehat{V}_g := \frac{1}{B_v - 1} \sum_{b \in \mathcal{B}_v} (\widehat{g}^{*(b)} - \overline{g}^*)(\widehat{g}^{*(b)} - \overline{g}^*)',$$

so that, as $B_v \to \infty$, $\widehat{V}_g = \widehat{G}_g \big(\sum_i \widehat{\phi}_i \widehat{\phi}_i' \big) \widehat{G}_g' + o_p(N^{-1})$. Set $\widehat{g} := g(\widehat{\vartheta})$ and

$$\mathcal{W} := \widehat{g}' \widehat{V}_g^{-1} \widehat{g}, \qquad \mathcal{W}^{*(b)} := (\widehat{g}^{*(b)} - \widehat{g})' \widehat{V}_g^{-1} (\widehat{g}^{*(b)} - \widehat{g}).$$

Theorem 5 (Interior validity)**.** *Under Proposition 10, the interior conditions, and $\Omega \succ 0$, as $N, B_v \to \infty$,*

$$N \widehat{V}_g \xrightarrow{p} \Omega.$$

If, in addition, $H_0 : g(\vartheta) = 0$,

$$\mathcal{W} \xrightarrow{d} \chi^2_2,$$

and, conditionally on the data, the multiplier law of $\mathcal{W}^{*(b)}$ consistently estimates the null law. The finite-draw value

$$\widehat{p}_B := \frac{1 + \sum_{b \in \mathcal{B}_v} \mathbf{1}\{\mathcal{W}^{*(b)} \geq \mathcal{W}\}}{B_v + 1}$$

is asymptotically valid as $N, B_v \to \infty$. Scalar restriction tests use the analogous plus-one rule for the absolute centered perturbation.

Proof. Proposition 10 and the delta method give $\sqrt{N}\{\widehat{g} - g(\vartheta)\} \xrightarrow{d} \mathcal{N}(0, \Omega)$. For the conditional part, the two inputs of the Gaussian multiplier central limit theorem are consequences of the proposition rather than assumptions: covariance consistency is $N\sum_i \widehat{\phi}_i\widehat{\phi}_i' \to_p \Sigma_\vartheta$, and maximal negligibility follows from $\sqrt{N}\max_i \|\widehat{\phi}_i\| \le \sqrt{N}\max_i \|\phi_i\| + \sqrt{N}\max_i \|\widehat{\phi}_i - \phi_i\| = N^{-1/2}\,O_p(N^{1/(4+\eta)}) + o_p(N^{-1/2}) \to_p 0$. Hence, conditionally on the data, $\sqrt{N}\sum_i \xi_i^{(b)}\widehat{\phi}_i \xrightarrow{d} \mathcal{N}(0, \Sigma_\vartheta)$ in probability, the smooth map g and the interior conditions transfer the limit to $\widehat{g}^{*(b)} - \widehat{g}$, and the continuous mapping theorem gives the two quadratic-form limits. The plus-one correction delivers validity of the finite-draw p-value as $B_v \to \infty$, with Monte Carlo error of order $B_v^{-1/2}$ if the number of draws is held fixed. □

Numerical interior rule. Let ϵ_{mach} denote machine precision and set $s_\pi := |a| + 2|b| + |c|$, $s_D := |d| + 2|e| + |f|$. The implementation requires

$$a > \sqrt{\epsilon_{\text{mach}}}\, s_\pi, \qquad d > \sqrt{\epsilon_{\text{mach}}}\, s_D,$$

and applies the analogous conditions to V_π, V_D for share ratios. Invalid bootstrap maps are left missing rather than truncated; in particular, negative variance-share draws are not truncated before taking a square root. Under the interior assumptions their share is $o_p(1)$. The joint covariance is accepted only when its diagonal is finite and positive and the smallest eigenvalue of its correlation matrix exceeds 10^{-10}; otherwise the joint statistic and its p-values are reported missing. No generalized inverse is used for the joint test.

Remark 2 (Growing number of locations). Proposition 10 is a fixed-dimension result: with J fixed, every estimator is a smooth function of finitely many sample means and the second-order remainder is $O_p(N^{-1})$ by Taylor's theorem. When J grows with N, that argument no longer applies. Each $\widehat{\psi}_j$ then carries estimation noise of order one over the square root of its degree, and the six quadratic forms accumulate J such terms; the cross-split construction still removes the leading plug-in bias, but controlling the *variance* of the second-order remainder at the $\sqrt{N}$ scale requires conditions on the mover network—uniformly vanishing effective resistance (leverage) of mover edges, an expansion condition on the degree-normalized split Laplacians, and no dominant mover-incidence weight—of the kind developed for variance components in sparse two-way models by Kline et al. [2020]; see also Andrews et al. [2008] and Bonhomme et al. [2023].

Remark 3 (Affine boundary). The theorem does not cover exact affine composition, where the first-order covariance of the residual restriction is singular. No boundary law or coverage result is derived here; in the affine simulation designs reported in Table 3, the ordinary joint procedure under-rejects.

For applied work, one can first distinguish the boundary from the interior. For any fixed $\lambda_\pi > 0$, residual intensity is an interior parameter; provided the other interior conditions hold, a consistent interval excludes zero with probability approaching one and the joint test applies. At exact affinity, if $\beta_\pi \ge 0$ and $d = \text{Var}_{\text{mov}}(\Delta\psi_i) > 0$, Theorem 4(iii) gives $\theta_D = \theta_{\text{sd}}$ for every mobility pattern.

Conditional on exact affinity, no isotropy test is needed for that SD-share interpretation.

Near-affinity has deterministic content. Under Assumption 3, the symmetrized-flow representation and $(q_j - q_k)^2 \le 2q_j^2 + 2q_k^2$ give $\mathrm{Var}_{\mathrm{mov}}(\Delta q_i) \le 4Q_\pi$, hence, by projection optimality and Cauchy–Schwarz,

$$Q_D \le 4Q_\pi, \qquad |\beta_D - \beta_\pi| \le 2\sqrt{Q_\pi / d}.$$

Let $U \ge Q_\pi$, $s := 1 + \beta_\pi$, and $r := 2\sqrt{U/d}$. Then $Q_D/d \le r^2$ and $|\beta_D - \beta_\pi| \le r$. If $\beta_\pi \ge 0$ and $r < s$, the denominators of f are bounded away from zero and

$$|\theta_D - \theta_{\mathrm{sd}}| \le \frac{r}{s(s-r)} + \frac{r^2}{(s-r)^3} + \frac{U}{2as^3}, \qquad a = \mathrm{Var}_\pi(\psi_j).$$

The first two terms bound the projection along routes relative to $1/s$; the last uses $\theta_{\mathrm{sd}} = \{s^2 + Q_\pi/a\}^{-1/2}$. This is a deterministic population sensitivity bound.

In practice there are three outcomes. If a reported U satisfies $r < s$ and the displayed bound is substantively small, report the share of the standard deviation with that bound and do not interpret the joint p-value. If the lower endpoint of an interval for λ_π is positive and the remaining interior and numerical covariance conditions hold, report the joint isotropy test. If neither condition holds, report the evidence as inconclusive between the near-affine and interior regimes.

G The Exposure Design

A second design measures the outcome once, in adulthood, and compares children who moved at different ages [Chetty and Hendren, 2018a]. Child i moves once, at age $m(i)$, from $o(i)$ to $d(i)$. Her outcome in adulthood is

$$Y_i = \alpha_i + \psi_{o(i)} + \omega_{m(i)}\,\Delta\psi_i + \varepsilon_i. \tag{54}$$

Here ψ_j is the effect of a whole childhood in location j. The share of childhood that remains at age m is ω_m, a number between zero and one. The model is as favorable as the additive model of the first design: a location does the same for every child, and exposure adds up. Assumption 1 has a counterpart here: among children who leave the same origin at the same age, extra gains are unrelated to the map of their destinations. I state the design in the additive form that the literature uses. The map is the mean outcome of children who never moved. Among children who move at age m, the researcher regresses

$$Y_i = c_{o(i)} + b_m\,\Delta\bar{Y}_i + u_i, \tag{55}$$

with an intercept for each origin, and reads the decline of b_m with the age at move as the effect of exposure.[13] The outcome is observed once, so the person component is not differenced out. The

[13] This is the semi-parametric specification of Chetty and Hendren [2018a], their equations (4) and (5), within one cell of parental income and birth cohort. Their fixed effects are for the origin by parental income decile by birth cohort, and by age at move when ages are pooled. Their map is the predicted outcome of permanent residents at the parents' income rank and for the child's cohort, and they add interactions of the gap with cohort. Their parametric specifications, equations (6) and (7), replace the origin fixed effects by a control for the outcomes of permanent residents in the origin.

origin is absorbed by its intercept, so the regression compares children from the same origin who went to different destinations. The pairs it compares are pairs of destinations.

Proposition 11 (The coefficient at each age)**.** *Let the shocks in* (54) *have mean zero given person components, locations, and ages at move, and let the map differ between two destinations that age-m movers reach from the same origin. The coefficient of regression* (55) *is* $b_m = \omega_m \, \theta_{D_m} + \xi_m$*, which is* (12)*. Here* θ_{D_m} *is* (11) *for the design* D_m *with comparison weights* $a_{jk} \propto \sum_o P_m(o) \, p_m(j \mid o) \, p_m(k \mid o)$*, where* $P_m(o)$ *is the share of age-m movers who leave o and* $p_m(\cdot \mid o)$ *is the distribution of their destinations. The term* ξ_m *is the coefficient of the same regression with* α_i *as the outcome.*

The coefficient at age m has three parts: the share of childhood that is left, a mover coefficient on the comparisons made at that age, and the selection of families. The results of the paper apply to θ_{D_m} as they stand. What is specific to this design is the comparison across ages.

Chetty and Hendren [2018a] define selection relative to an experiment. Let β_m be the coefficient when children are assigned to locations at random at age m. The coefficient in the data is $b_m = \beta_m + \delta_m$, and δ_m is their selection effect. Their identifying assumption is that δ_m is the same at all ages. The decline of b_m with the age at move is then the decline of β_m.

Proposition 12 (The exposure design)**.** *Let the conditions of Proposition 11 hold.*

(i) Two parts of selection. *If children are assigned to locations with probabilities* π_j*, then* $\beta_m = \omega_m \theta_\pi$*, and*

$$\delta_m = \underbrace{\xi_m}_{\text{selection of families}} + \underbrace{\omega_m \left(\theta_{D_m} - \theta_\pi\right)}_{\text{comparisons of movers, not of the experiment}} . \tag{56}$$

Let ξ_m *be the same at all ages, and let* $\omega_m = 0$ *at some age. Then* δ_m *is the same at all ages if and only if* $\theta_{D_m} = \theta_\pi$ *at every age with* $\omega_m > 0$*.*

(ii) Age profile. *For adjacent ages,*

$$b_m - b_{m+1} = \underbrace{\left(\omega_m - \omega_{m+1}\right) \theta_{D_m}}_{\text{extra exposure}} + \underbrace{\omega_{m+1} \left(\theta_{D_m} - \theta_{D_{m+1}}\right)}_{\text{changing comparisons}} + \underbrace{\xi_m - \xi_{m+1}}_{\text{selection of families}} . \tag{57}$$

(iii) Siblings. *Let each family move once, and regress* Y_i *on family intercepts and* $\omega_{m(i)} \Delta \bar{Y}_i$*. Suppose that, among the families that move between the same two locations, differences in exposure between siblings are uncorrelated with differences in their person components and shocks. Then the coefficient is* θ_D *for a design with the following weights. The pair* $\{j, k\}$ *has weight proportional to* $\mathbb{E}[H \, \mathbf{1}\{\text{the family moves between } j \text{ and } k\}]$*, where* H *is the sum of squared deviations of* $\omega_{m(i)}$ *from its mean within the family.*

In the additive model the selection effect has two parts. The first is the selection of families: children who move to better locations differ in ways that would show anywhere. This is what

The coefficient is then still linear in place effects, with weights that the destinations and the map determine, but it is not of the form (11). I state the results for the specification with origin fixed effects.

the assumption is meant for, and what the tests in the literature address. Comparisons of siblings, moves caused by displacement shocks, and placebo outcomes all speak to it.

The second part is not about families. It is the difference between the comparisons that movers make and the comparisons of the experiment, scaled by the share of childhood that remains. It shrinks with the age at move because exposure does, unless it is zero. Suppose the selection of families is the same at all ages, and exposure has ended at some age. Then the assumption of constant selection holds exactly when the mover coefficient at each age is the cross-sectional coefficient. Theorem 1 gives the condition on the comparisons: at each age they are aligned with the map. The assumption therefore covers more than its name suggests, and the second part can be examined with flows and location means.

The Opportunity Atlas illustrates the two parts [Chetty et al., 2026]. It estimates the slope twice. The first estimate compares the three arms of a housing experiment within each of five cities. It is 0.68, with a standard error of 0.33. Vouchers were assigned at random, so the selection of families is absent by design, and the comparisons are those between the neighborhoods where the three groups lived. The second estimate uses children who move between census tracts at different ages, in order to use all tracts in the United States. It is 0.57, and it rests on constant selection. By part (i) and Theorem 1, 0.57 is the cross-sectional coefficient if the comparisons are aligned with the map. This can be examined with flows between tracts and the map.

G.1 The Age Profile and Its Unit

Part (ii) shows what the profile contains when only the selection of families is constant. The first term is the value of an extra year. The second bends the profile where the comparisons change. Suppose two adjacent ages have the same remaining exposure, so that moving a year earlier adds nothing. In Example 1, let the destinations compared among the younger movers differ in p, and those compared among the older movers differ in z. The younger coefficient is ω_m and the older is zero. An exposure effect appears where there is none, and no selection of families is involved. Whether comparisons change with age can be checked: θ_{D_m} can be estimated for groups of ages.

If the comparisons are the same at all ages, the profile shows that exposure matters, and at which ages. Its shape is that of ω_m. Its level is another matter. An extra year of exposure is valued at θ_D units of place effect per unit of the map, and this number is a property of the comparisons. If exposure is linear over a childhood of T years, b_m falls by θ_D/T per year of age, for any θ_D. One reading is that a child who moves at birth picks up this share of the map, and that this share is the part of the map that is place. That reading requires $\theta_D = \theta_\pi$. Theorems 1 and 3 say when that holds and what remains otherwise.

Siblings experience the same move at different ages, and family intercepts remove both the family component and the route. They settle whether the response is a response to exposure. They do not change the unit. By part (iii) their coefficient is again a θ_D.

G.2 *Place Effects in the Exposure Design*

The exposure design also has its counterpart of the two-way estimator. Chetty and Hendren [2018b] estimate the effect of each commuting zone and county from families who move with children of different ages. They then regress these effects on the outcomes of permanent residents. In the terms of this paper, that regression is the cross-sectional coefficient, estimated in levels as in Section VI. Across commuting zones it is 0.032 per year of exposure for children of parents at the 25th percentile, against about 0.04 per year for the mover coefficient of Chetty and Hendren [2018a]. The authors attribute the difference to the samples and to the income group. They also report that the outcomes of permanent residents and the effects of commuting zones have a correlation of 0.80. The map therefore misses more than a third of the variance of place effects, which is the residual of Section IV. Their two papers contain the target, the residual, and the mover coefficient. The results here say how the three are linked: through the comparison weights.

Proof of Propositions 11 and 12. Coefficient at age m. Fix m and let a tilde denote the deviation from the mean among age-m movers with the same origin. By (54), $\tilde{Y}_i = \tilde{\alpha}_i + \omega_m \tilde{\psi}_{d(i)} + \tilde{\varepsilon}_i$, and the regressor is $\tilde{\bar{Y}}_{d(i)}$, so

$$b_m = \omega_m \frac{\mathbb{E}[\tilde{\psi}_{d(i)} \tilde{\bar{Y}}_{d(i)}]}{\mathbb{E}[\tilde{\bar{Y}}^2_{d(i)}]} + \frac{\mathbb{E}[\tilde{\alpha}_i \tilde{\bar{Y}}_{d(i)}]}{\mathbb{E}[\tilde{\bar{Y}}^2_{d(i)}]}.$$

The second ratio is ξ_m. For attributes f and h, $\mathbb{E}[\tilde{f}_{d(i)} \tilde{h}_{d(i)}]$ is the average over origins of the covariance of f and h across destinations. By Lemma 1 applied to each distribution $p_m(\cdot \mid o)$,

$$\mathbb{E}[\tilde{f}_{d(i)} \tilde{h}_{d(i)}] = \sum_{j<k} \Big\{ \sum_o P_m(o)\, p_m(j \mid o)\, p_m(k \mid o) \Big\} (f_j - f_k)(h_j - h_k),$$

which is proportional to $\mathcal{L}_{D_m}(f, h)$. The first ratio is therefore θ_{D_m}.

Proposition 12. (i) Under random assignment at age m with probabilities π_j, the destination is independent of the person component, the origin, and the shock. By (54), the coefficient of Y_i on $\bar{Y}_{d(i)}$ is $\omega_m \operatorname{Cov}_\pi(\psi_j, \bar{Y}_j) / \operatorname{Var}_\pi(\bar{Y}_j) = \omega_m \theta_\pi$. Subtracting from (12) gives (56). If $\xi_m = \xi$ for all m and $\omega_{m_0} = 0$, then $\delta_{m_0} = \xi$, so δ_m is constant if and only if $\omega_m(\theta_{D_m} - \theta_\pi) = 0$ for all m.

(ii) By (12), $b_m - b_{m+1} = \omega_m \theta_{D_m} - \omega_{m+1} \theta_{D_{m+1}} + \xi_m - \xi_{m+1}$. Add and subtract $\omega_{m+1} \theta_{D_m}$.

(iii) Let family q move once, with $\Delta \bar{Y}_q$ and $\Delta \psi_q$ common to its children, and let $\bar{\omega}_q$ be the mean of $\omega_{m(i)}$ among them. Deviations from family means satisfy $Y_i - \bar{Y}_{q\cdot} = (\omega_{m(i)} - \bar{\omega}_q) \Delta \psi_q + v_i$, where v_i collects the deviations of person components and shocks. The regressor is $(\omega_{m(i)} - \bar{\omega}_q) \Delta \bar{Y}_q$. Under the stated condition the coefficient is $\mathbb{E}[H_q \Delta \bar{Y}_q \Delta \psi_q] / \mathbb{E}[H_q \Delta \bar{Y}_q^2]$ with $H_q = \sum_i (\omega_{m(i)} - \bar{\omega}_q)^2$. Collecting families by the pair of locations they move between gives the stated design. □

H The Event Study and the Two-Period Coefficient

In the appendices that follow, $\mathcal{J}_0$ denotes the set of all locations and $J_0 = J$ its size. Moments with the subscript mov are taken over movers; under Assumption 3, $\mathrm{Cov}_{\mathrm{mov}}(\Delta f_i, \Delta h_i) = \mathcal{L}_D(f, h)$.

The event study used in applications is the multi-period specification of Finkelstein et al. [2016] with base period $k = -1$. This section makes its relation to the paper's two-period objects precise in two steps. First, on a balanced event-time window the event-study coefficients coincide, by projection algebra alone, with a sequence of per-event-time long-difference regressions (Proposition 13). Second, under the additive model with residual exogeneity, those per-horizon coefficients converge—pre-event to zero and post-event to a common first-difference limit (Theorem 6). Under mean stationarity for the compressed origin–destination endpoints and the eligible-window bridge below, that limit equals the paper's two-period estimand.

Consider movers $i \in M$ observed in periods $t \in \{1, \dots, T\}$, $T \geq 4$, each changing location exactly once at t_i^*, with event time $\tau_{it} := t - t_i^*$ and $o(i) := j(i, t_i^* - 1)$, $d(i) := j(i, t_i^*)$.

Assumption 5 (Residual exogeneity)**.** $Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}$ with $\psi_{j(i,t)} = \psi_{o(i)}$ for $t < t_i^*$ and $\psi_{d(i)}$ thereafter. Let $\mathcal{F}^T$ be the sigma-field generated by all individuals' individual effects and realized location paths and by the place effects. Then $\mathbb{E}[\varepsilon_{it} \mid \mathcal{F}^T] = 0$; the residual vectors $(\varepsilon_{i1}, \dots, \varepsilon_{iT})'$ are conditionally uncorrelated across individuals, with unrestricted dependence across t within a individual [Card et al., 2013, 2018].

Assumption 6 (Balanced event-time window)**.** Fix $\mathcal{K} := \{-K_-, \dots, -2\} \cup \{0, \dots, K_+\}$ and restrict move dates to $K_- + 1 \leq t_i^* \leq T - K_+$, so that every mover's event window $\mathcal{K} \cup \{-1\}$ lies inside the calendar panel. Estimation uses only the event-window observations, $\tau_{it} \in \mathcal{K} \cup \{-1\}$.

Let $\mathcal{T} := \{1, \dots, T\}$ denote the fixed calendar window and let $A_{\ell t}$ indicate that individual ℓ's observation is eligible for the location mean at t. In particular, $A_{\ell t} = 0$ for a mover's transition-year observation, while stayers remain eligible. Define

$$\overline{Y}_j^{(-i)} := \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \frac{\sum_{\ell \neq i} A_{\ell t} \mathbf{1}\{j(\ell, t) = j\} Y_{\ell t}}{\sum_{\ell \neq i} A_{\ell t} \mathbf{1}\{j(\ell, t) = j\}}, \qquad \delta_i^{(-i)} := \overline{Y}_{d(i)}^{(-i)} - \overline{Y}_{o(i)}^{(-i)}. \tag{58}$$

Thus individuals are averaged within location–year first and the resulting location–year means are then averaged equally over the fixed years, as in Finkelstein et al. [2016]. The superscript $(-i)$ is the leave-one-out refinement used here; Finkelstein et al. [2016] use the same aggregation order. If any required post-LOO endpoint-year denominator is zero, the complete gap is left undefined rather than averaging a shorter window. Write $\bar{Y}_j$ for the large-cell probability limit of this full-window location attribute and $\Delta \bar{Y}_i := \bar{Y}_{d(i)} - \bar{Y}_{o(i)}$.

The common-window construction is a measurement convention within the maintained structural regime. Maintain the eligible-window condition

$$\bar{Y}_j := \frac{1}{T} \sum_{t \in \mathcal{T}} \mathbb{E}[Y_{it} \mid A_{it} = 1, j(i, t) = j] = \psi_j + \bar{\alpha}_j, \qquad j \in \mathcal{J}_0. \tag{59}$$

Although the event study itself requires a longer panel, the aggregation formula specializes to the paper's two-period location mean when $T = 2$ and $A_{it} \equiv 1$, the eligibility convention for its structural origin and destination periods. For $T > 2$, the empirical eligibility convention targets the same location object under (59).

Following equation (6) of Finkelstein et al. [2016][14], the event-study specification is

$$Y_{it} = \tilde{a}_i + \theta_{\tau_{it}} \delta_i^{(-i)} + \rho_{\tau_{it}} + e_{it}, \tag{60}$$

estimated by OLS on the mover sample over the window of Assumption 6: individual effects $\tilde{a}_i$, event-time main effects ρ_k, and the common full-window leave-one-out destination–origin gap $\delta_i^{(-i)}$ interacted with event time, normalized at the base period ($\theta_{-1} = \rho_{-1} = 0$).[15] The coefficients of interest are the θ_k.

The reference point is the per-k long difference: for $k \in \mathcal{K}$ define $Y_{ik}^* := Y_{i,t_i^*+k} - Y_{i,t_i^*-1}$ and estimate $Y_{ik}^* = c_k + \widetilde{\theta}_k \delta_i^{(-i)} + e_{ik}$ by OLS. The intercept is retained, so $\widetilde{\theta}_k$ is the centered sample slope $\widehat{\mathrm{Cov}}_{\mathrm{mov}}(\delta_i^{(-i)}, Y_{ik}^*) / \widehat{\mathrm{Var}}_{\mathrm{mov}}(\delta_i^{(-i)})$. The equivalence below does not require the mean gap to vanish. These per-k long differences are the algebraic reference for the event-study coefficients; the next result shows that the two coincide.

Proposition 13 (The event-study coefficients are long differences)**.** *If* $\widehat{\mathrm{Var}}_{\mathrm{mov}}(\delta_i^{(-i)}) > 0$*, then the sample OLS coefficients in equation* (60) *on the balanced window of Assumption 6 satisfy* $\theta_k = \widetilde{\theta}_k$ *and* $\rho_k = c_k$ *exactly for every* $k \in \mathcal{K}$*.*

Proof. Let $n = |\mathcal{K}| + 1$, and let $Y_i \in \mathbb{R}^n$ collect the outcomes of mover i at the event times in $\mathcal{K} \cup \{-1\}$. Write $\rho, \theta \in \mathbb{R}^n$ for the event-time effects and the coefficients on the gap, with entries zero at $k = -1$. Let $\mathbf{1}$ be the vector of ones and $M = I - \mathbf{1}\mathbf{1}'/n$. On the balanced window, least squares minimizes $\sum_i \|Y_i - \tilde{a}_i \mathbf{1} - \rho - \delta_i^{(-i)}\theta\|^2$. Minimizing over $\tilde{a}_i$ leaves $\sum_i \|MY_i - M\rho - \delta_i^{(-i)} M\theta\|^2$. The map $\rho \mapsto M\rho$ is one-to-one from the vectors with $\rho_{-1} = 0$ onto the range of M, and likewise for θ. The problem is therefore the least-squares regression of MY_i on $(1, \delta_i^{(-i)})$, one coordinate at a time. Let c_k^0 and b_k be the intercept and the slope of the regression of Y_{i,t_i^*+k} on $(1, \delta_i^{(-i)})$, which exist because $\widehat{\mathrm{Var}}_{\mathrm{mov}}(\delta_i^{(-i)}) > 0$. Then $M\theta = Mb$ and $M\rho = Mc^0$, so $\theta = b - b_{-1}\mathbf{1}$ and $\rho = c^0 - c_{-1}^0 \mathbf{1}$. Hence $\theta_k = b_k - b_{-1}$ and $\rho_k = c_k^0 - c_{-1}^0$. These are the slope and the intercept of the regression of $Y_{ik}^* = Y_{i,t_i^*+k} - Y_{i,t_i^*-1}$ on $(1, \delta_i^{(-i)})$. □

[14] Their specification additionally includes calendar-time effects and age controls; both are covered by the time-varying-controls remark below. Estimating calendar effects *jointly* within equation (60) requires *cohort exchangeability*—that the joint distribution of $(\Delta\psi_i, \delta_i^{(-i)})$ not vary with the move date t_i^*: after partialling the individual effects, calendar dummies are functions of (k, t_i^*), so without exchangeability they load on cohort-varying moments of $(\Delta\psi_i, \delta_i^{(-i)})$ and the θ_k mix cohorts in the manner of the staggered-adoption critiques [Sun and Abraham, 2021, De Chaisemartin and d'Haultfoeuille, 2020, Borusyak et al., 2024]. The per-k population projections remain defined cohort by cohort either way, and Theorem 6 applies within each cohort under its conditions.

[15] Following Finkelstein et al. [2016], a mover's transition-year observation is excluded when constructing the location–year means because it may combine exposure to the origin and destination. This exclusion concerns only the construction of location means: the transition-year outcome remains in the event-study regression, so the coefficient at $k = 0$ is reported.

Theorem 6. *Under Assumptions 5 and 6, suppose the T-period individual panels are independently and identically sampled across individuals,* $|M|/N \xrightarrow{p} \pi_M > 0$, $\sup_t \mathbb{E}[Y_{it}^2] < \infty$, *every eligible location–year cell used in (58) has a positive limiting share, and* $\mathrm{Var}_{\mathrm{mov}}(\Delta\bar{Y}_i) > 0$, *in the fixed-*$J_0$*, large-location regime of the paper, as* $N \to \infty$*: (i) for every* $k \in \mathcal{K}$ *with* $k \le -2$, $\widetilde{\theta}_k \xrightarrow{p} 0$*; and (ii) for every* $k \in \mathcal{K}$ *with* $k \ge 0$,

$$\widetilde{\theta}_k \xrightarrow{p} \theta_{\mathrm{FD}} := \frac{\mathrm{Cov}_{\mathrm{mov}}(\Delta\psi_i, \Delta\bar{Y}_i)}{\mathrm{Var}_{\mathrm{mov}}(\Delta\bar{Y}_i)},$$

independently of k. If Assumption 3, applied to the compressed origin–destination endpoints, and the fixed-window bridge (59) are also maintained, then $\bar{Y}_j = \psi_j + \bar{\alpha}_j$ *and* $\theta_{\mathrm{FD}} = \theta_D$ *in equation (9).*

Proof. Under Assumption 5, $Y_{it} = a_i + \Delta\psi_i \mathbf{1}\{\tau_{it} \ge 0\} + \varepsilon_{it}$, with $a_i := \alpha_i + \psi_{o(i)}$, so $Y_{ik}^* = \Delta\psi_i \mathbf{1}\{k \ge 0\} + \Delta\varepsilon_{ik}$. Because J_0 and T are fixed and every eligible location–year cell used in (58) grows proportionally with N, the cell-mean law of large numbers and the exact leave-one-out identity give

$$\frac{1}{|M|} \sum_{i \in M} \left(\delta_i^{(-i)} - \Delta\bar{Y}_i\right)^2 \xrightarrow{p} 0.$$

Conditional mean zero and cross-individual residual orthogonality in Assumption 5, together with the omission of the focal individual from every cell mean, make the sample covariance between $\Delta\varepsilon_{ik}$ and $\delta_i^{(-i)}$ $o_p(1)$. Substitution in the centered OLS slope yields zero before the move and θ_{FD} after it. The fixed-window bridge gives $\bar{Y}_j = \psi_j + \bar{\alpha}_j$; together with Assumption 3, this yields $\theta_{\mathrm{FD}} = \theta_D$. □

Time-varying controls. Applied two-way specifications include time-varying controls, $Y_{it} = \alpha_i + \psi_{j(i,t)} + X_{it}'\gamma + \varepsilon_{it}$, with calendar-time effects the leading case. Every object in the paper is then defined on the control-adjusted outcome $\widetilde{Y}_{it} := Y_{it} - X_{it}'\gamma$: location means, the mover regression, the moments along routes and across locations, and the shares of Section D. The interpretation adjusts accordingly—the recovered shares are shares of the variance of *control-adjusted* location means (for instance, healthcare spending adjusted for age and year)—and Theorem 6 holds unchanged with $\widetilde{Y}$ in place of Y.

Unbalanced event windows. Assumption 6 delivers the exact finite-sample projection identity above. In a conventional unbalanced event window, equation (60) instead forms cohort–horizon-weighted projections. If the cohort-exchangeability condition stated above holds and, conditional on move cohort, availability at each horizon is mean-independent of $(\Delta\psi_i, \Delta\bar{Y}_i)$ and the corresponding outcome innovation, those weights do not change the limits in Theorem 6: pre-move coefficients converge to zero and post-move coefficients to the common θ_{FD}. Without these conditions, the coefficients remain well-defined weighted projections.

I Without a Steady State

Assumption 3 gives the paper one map and flows that balance. This section drops the assumption. Locations may grow or shrink, and the mean outcome of their residents may change from period to period. The results of Sections III to VI then hold with one more term in the comparison weights. The other sections of this Supplement maintain the steady state.

I.1 The Map of the Window

Individuals are observed in periods $t = 1, \dots, T$, and the paper has $T = 2$. A mover changes location once, and ΔY_i is the change in her outcome between the period before the move and the period after it. Equation (2) defines place effects and extra gains as before, for the movers of all periods together. Let $\bar{Y}_{jt}$ be the mean outcome of the residents of location j in period t. The map is the mean over the periods,

$$\bar{Y}_j := \frac{1}{T} \sum_{t=1}^{T} \bar{Y}_{jt}, \tag{61}$$

and composition is $\bar{\alpha}_j := \bar{Y}_j - \psi_j$, as in (5). Finkelstein et al. [2016] construct their regressor in this way: the authors average the mean outcome of each area over the years of their sample. The gap $\Delta \bar{Y}_i$ is the gap in this map. The weights π of the question do not depend on the period, and the cross-sectional coefficient is (6) with this map.

The map of the window answers a question about the window. Let $\theta_{\pi t}$ be the cross-sectional coefficient of period t, which is (6) with $\bar{Y}_{jt}$ in the place of the map.

Proposition 14 (The window and its periods)**.**

$$\theta_\pi \operatorname{Var}_\pi(\bar{Y}_j) = \frac{1}{T} \sum_{t=1}^{T} \theta_{\pi t} \operatorname{Var}_\pi(\bar{Y}_{jt}), \qquad \operatorname{Var}_\pi(\bar{Y}_j) \le \frac{1}{T} \sum_{t=1}^{T} \operatorname{Var}_\pi(\bar{Y}_{jt}). \tag{62}$$

The inequality holds with equality if and only if the maps of the periods differ by constants. In that case $\theta_{\pi t} = \theta_\pi$ *for every* t*.*

Proof. Place effects do not depend on the period, so $\operatorname{Cov}_\pi(\psi_j, \bar{Y}_j) = T^{-1} \sum_t \operatorname{Cov}_\pi(\psi_j, \bar{Y}_{jt})$, which is the first statement. For the second, $\operatorname{sd}_\pi(\bar{Y}_j) \le T^{-1} \sum_t \operatorname{sd}_\pi(\bar{Y}_{jt})$ by the triangle inequality, and the square of a mean is at most the mean of the squares. Both hold with equality if and only if the vectors $\bar{Y}_{jt} - \mathbb{E}_\pi[\bar{Y}_{jt}]$ are the same for all t. □

The coefficient of the window is an average of the coefficients of the periods, weighted by the variances of their maps, times a factor that is at least one. The factor is the mean variance of the maps of the periods over the variance of the map of the window. Composition that changes from period to period adds variance to the map of each period, and averaging over periods removes part of this variance. In a steady state the factor is one, and all the coefficients are equal.

I.2 The Coefficient and the Results

Let $\delta_j := \sum_k F_{kj} - \sum_k F_{jk}$ be the net inflow of location j, as a share of all movers. It is zero in a steady state. Let $\mu := \mathbb{E}_{\text{mov}}[\Delta \bar{Y}_i] = \sum_j \delta_j \bar{Y}_j$ be the mean gap among movers. Let $\tilde{D}$ be the design with the comparison weights

$$\tilde{a}_{jk} := a_{jk} + \delta_j \delta_k. \tag{63}$$

The weight of a pair of locations rises if both gain movers or both lose movers. It falls if one gains and the other loses, and it can then be negative.

Proposition 15 (The mover coefficient without a steady state)**.** *Let Assumption 1 hold for the gap in the map (61), and let this gap vary among movers. The coefficient of regression (8) is*

$$\theta_{\tilde{D}} = \frac{\mathcal{L}_{\tilde{D}}(\psi, \bar{Y})}{\mathcal{L}_{\tilde{D}}(\bar{Y}, \bar{Y})}. \tag{64}$$

For all attributes f and h of locations, $\mathcal{L}_{\tilde{D}}(f,h) = \text{Cov}_{\text{mov}}(\Delta f_i, \Delta h_i)$. The comparison gap of $\tilde{D}$ is

$$g_j = \frac{1}{\pi_j}\Big\{ \sum_k a_{jk} \, (\bar{Y}_j - \bar{Y}_k) - \delta_j \, \mu \Big\}. \tag{65}$$

Proof. For attributes f and h, $\mathbb{E}_{\text{mov}}[\Delta f_i \, \Delta h_i] = \mathcal{L}_D(f,h)$, as in the proof of Proposition 1, and $\mathbb{E}_{\text{mov}}[\Delta f_i] = \sum_j f_j \, \delta_j$. Because $\sum_j \delta_j = 0$,

$$\sum_{j<k} \delta_j \delta_k \, (f_j - f_k)(h_j - h_k) = -\Big(\sum_j \delta_j f_j \Big) \Big(\sum_j \delta_j h_j \Big).$$

Hence $\mathcal{L}_{\tilde{D}}(f,h) = \mathcal{L}_D(f,h) - \mathbb{E}_{\text{mov}}[\Delta f_i] \, \mathbb{E}_{\text{mov}}[\Delta h_i] = \text{Cov}_{\text{mov}}(\Delta f_i, \Delta h_i)$. The gap $\Delta \bar{Y}_i$ is a function of the two locations of i. By (2), $\text{Cov}_{\text{mov}}(\Delta Y_i, \Delta \bar{Y}_i) = \text{Cov}_{\text{mov}}(\Delta \psi_i, \Delta \bar{Y}_i) + \text{Cov}_{\text{mov}}(s_i, \Delta \bar{Y}_i)$, and the last term is zero by Assumption 1. Regression (8) has a constant, so its coefficient is $\text{Cov}_{\text{mov}}(\Delta \psi_i, \Delta \bar{Y}_i) / \text{Var}_{\text{mov}}(\Delta \bar{Y}_i)$, which is (64). For (65), $\sum_k \delta_j \delta_k \, (\bar{Y}_j - \bar{Y}_k) = -\delta_j \, \mu$. □

The proposition asks nothing of the sizes of locations, their shares of movers, or their mean outcomes in single periods. In a steady state $\delta = 0$, and $\tilde{D}$ is the design D of the paper.

The term $\delta_j \delta_k$ comes from the constant of the regression. With a constant, the regression measures each gap from the mean gap among movers, and this mean is zero when flows balance. The constant takes up any change at the move that is the same on all routes. The relative-year effects of Finkelstein et al. [2016] are constants of this kind, one for each year relative to the move (Proposition 13). The authors include these effects to allow changes in the utilization of movers around the move that are the same for every origin and destination.

The proofs of Proposition 3, of Theorems 1, 2 and 3, and of Propositions 4 and 5 use two properties of comparison weights. The weights are symmetric, and $\mathcal{L}_D(f,f) \geq 0$ for every attribute f. The weights $\tilde{a}_{jk}$ have both properties, because $\mathcal{L}_{\tilde{D}}(f,f) = \text{Var}_{\text{mov}}(\Delta f_i)$. The proofs also use economies with given place effects, a given map, and given flows, and Lemma 6 supplies these

economies without a steady state. The results therefore hold for $\tilde{D}$: with $\tilde{a}_{jk}$ in the place of a_{jk}, with the comparison gap (65), and with the alignment $r_{\tilde{D}} = \text{Corr}_\pi(g_j, \bar{Y}_j)$ of this comparison gap. The statements below change.

Pair shares. Part (i) of Proposition 2 holds for $\tilde{D}$: the mover coefficient is a weighted sum of pair shares, and the weights sum to one. Some $\tilde{a}_{jk}$ can be negative, so the coefficient can lie outside the range of the pair shares, and part (ii) of that proposition does not apply.

The sign of a comparison gap. By (65), a location whose mean outcome is above that of its partners can have a negative comparison gap. This requires that the location gain movers while movers go up the map on average, or lose movers while movers go down. Theorem 2 holds as stated in terms of the sign of the comparison gap, with $\pi_j g_j$ in the place of $a_j(\bar{Y}_j - \bar{Y}_j^P)$ in (19).

Independent pairing. Part (iv) of Theorem 1 reads: the comparison gap is proportional to the map for every map if and only if $\tilde{a}_{jk} \propto \pi_j \pi_k$.

The two-way regression. The regression of the route contrasts on indicators of destination and origin includes a constant, which takes up a change at the move that is common to all routes. In a steady state the constant changes no estimate. The normal equations are (25) with $\tilde{a}_{jk}$ in the place of a_{jk} and with each route contrast measured from the mean contrast among movers. Condition (24) and the net extra gain in (26) are stated for extra gains measured from their mean among movers. Proposition 5 needs routes that connect all locations and one pair of locations with movers in both directions. The constant is then not a combination of the indicators. Finkelstein et al. [2016] make this point for their relative-year effects: movers in one direction between two areas do not identify the difference between the place effects of the two areas, and movers in both directions do.

Lemma 6 (Representation without a steady state)**.** *Take J locations with any place effects ψ_j and any composition $\bar{\alpha}_j$. Take any flows $F_{jk} \geq 0$ with $F_{jj} = 0$ that sum to one. There is an additive economy with exogenous mobility and two periods that has these place effects and these flows, and whose map is $\psi_j + \bar{\alpha}_j$. The map can be the mean over the two periods of the mean outcome of residents, or the mean outcome of permanent residents. Let the routes connect $J \geq 3$ locations, and let χ be an attribute of locations with $\mathcal{L}_{\tilde{D}}(\chi, \bar{Y}) = 0$. There is then an economy with these place effects, this map, and these flows that satisfies Assumptions 1 and 2 and has the extra gain $s_{jk} = \chi_k - \chi_j$ on every route.*

Proof. Let the movers have mass one, with F_{jk} of them on the route from j to k, and let location j have a mass $u_j > 0$ of stayers. Its size is $n_{j1} = u_j + \sum_k F_{jk}$ in the first period and $n_{j2} = u_j + \sum_k F_{kj}$ in the second. Every mover draws her person component from one distribution with mean zero, whatever her route. The stayers of j draw theirs from a distribution with mean $\bar{\alpha}_j^s$. The mean person component of the residents of j is $u_j \bar{\alpha}_j^s / n_{j1}$ in the first period and $u_j \bar{\alpha}_j^s / n_{j2}$ in the second. Let $\bar{\alpha}_j^s = 2\bar{\alpha}_j / \{u_j(1/n_{j1} + 1/n_{j2})\}$, or $\bar{\alpha}_j^s = \bar{\alpha}_j$ if the map is the mean outcome of permanent residents. Outcomes are $Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}$, with shocks that have mean zero and are independent of person components and locations. Mobility is exogenous, and the map is $\psi_j + \bar{\alpha}_j$.

For the second statement, give the movers the outcomes of the proof of Lemma 3, and give the stayers the outcome $\alpha_i + \psi_l + \varepsilon_{it}$ in every location l. As in that proof, ψ is the vector of place effects,

$s_{jk} = \chi_k - \chi_j$, and Assumption 2 holds. By Proposition 15, $\mathrm{Cov}_{\mathrm{mov}}(s_i, \Delta\bar{Y}_i) = \mathcal{L}_{\tilde{D}}(\chi, \bar{Y}) = 0$, which is Assumption 1. The movers who live in j have the mean outcome $\psi_j + \chi_j$ there. Let $\bar{\alpha}_j^{\mathrm{s}}$ solve

$$\tfrac{1}{2}\, u_j\, \bar{\alpha}_j^{\mathrm{s}} \Big(\frac{1}{n_{j1}} + \frac{1}{n_{j2}} \Big) + \tfrac{1}{2}\, \chi_j \Big(\frac{\sum_k F_{jk}}{n_{j1}} + \frac{\sum_k F_{kj}}{n_{j2}} \Big) = \bar{\alpha}_j,$$

or let $\bar{\alpha}_j^{\mathrm{s}} = \bar{\alpha}_j$ if the map is the mean outcome of permanent residents. The map is then $\psi_j + \bar{\alpha}_j$. □

I.3 Assumption 1 and the Periods

Assumption 1 is one condition, on the regressor that is used. With the map of the window, extra gains are uncorrelated with the gap in that map among the movers of all periods together. The regressor is the same in every period, so no version of the assumption for each period is needed.

A regressor built from the mean outcomes of single periods is another matter. Take two periods, and take the mean outcome of the destination after the move minus the mean outcome of the origin before the move. Then

$$\bar{Y}_{d(i)2} - \bar{Y}_{o(i)1} = \Delta\bar{Y}_i + \tfrac{1}{2}\Big\{ \big(\bar{Y}_{d(i)2} - \bar{Y}_{d(i)1}\big) + \big(\bar{Y}_{o(i)2} - \bar{Y}_{o(i)1}\big) \Big\}.$$

The second term adds the changes over time of the two locations. It is a sum over the two locations and not a difference between them. The coefficient is then not a ratio of the form (64), and the results do not apply to it. The map of the window has no term of this kind.

An event study has a regression for each period after the move (Supplement H). Let Y_{ik}^* be the change in the outcome of mover i from the last period before the move to period k after the move. Define the extra gain of that period by $\mathbb{E}[Y_{ik}^* \mid o(i) = j,\, d(i) = l] = \psi_l - \psi_j + s_{jl}^k$. The regression of Y_{ik}^* on the gap, with a constant, has the coefficient

$$\theta_{\tilde{D}} + \frac{\mathrm{Cov}_{\mathrm{mov}}(s_i^k, \Delta\bar{Y}_i)}{\mathrm{Var}_{\mathrm{mov}}(\Delta\bar{Y}_i)}. \tag{66}$$

Each period after the move therefore needs Assumption 1 for its own extra gain. For a period before the move there is no change of place, and the coefficient is the second term alone, with the mean change before the move on each route in the place of the extra gain. The coefficients before the move are the evidence on trends of Section II.1.

I.4 The Flows between U.S. Commuting Zones

The flows of Section VII do not balance. Half the sum of the absolute net inflows of the 595 zones is 25 percent of all moves on the listed pairs. Pairs with fewer than 25 children are not listed, and 3,139 of the 11,017 listed routes have no listed route in the other direction. Net inflows enter the results through the mean gap among movers. This mean is -0.12 rank points, and the standard deviation of the gap is 3.4.

Table 1 computes the coefficient for origin against destination without a constant, which is the form of the steady state. The coefficient is 0.609. With a constant, as in (64), it is 0.606. The alignment with population weights is 0.766 in both forms. The bounds are 0.247 and 0.881 in the form of the steady state, and 0.245 and 0.880 with the weights $\tilde{a}_{jk}$. At least 105 of the 595 zones have a comparison gap of the wrong sign at every level of the map, against 106 in the form of the steady state. Of the 7,078 pairs of zones with movers, 60 have a negative weight $\tilde{a}_{jk}$.

J Estimation and Inference

Alignment can be checked with flows and location means only. In this section they are computed in a sample of N individuals, drawn as in Section II.2. The number of locations is fixed, and N grows. The approximation is good when each location has many movers and, for the estimates that use flows, when each route does. Where cells are thin, the estimates use leave-one-out means and split samples.

J.1 Estimates and a Test

The sample comparison weights are the shares of movers between each pair of locations. The sample map is the vector of location means. Put them into (14) to (18), with the shares of moves that involve each location as the weights π. This gives the estimates $\hat{g}_j$, $\hat{g}_j^{\perp}$ and $\hat{r}_D$. The plot behind them is simple: for each location, the mean outcome of the locations it exchanges movers with, against its own mean outcome. Under alignment the points lie on a line.

The estimate $\hat{r}_D$ is biased toward misalignment. A location with few movers has a noisy comparison gap, and noise is not proportional to the map. A second estimate avoids this. Split the individuals at random into two halves. Compute $\hat{g}^{\perp}$ in each half, from that half's flows and location means. Then replace the square of $\hat{g}^{\perp}$ by the product of the two halves. I call the result the split-sample estimate, $\tilde{r}_D$.

For a test of alignment, let $\hat{v}_j = \hat{\pi}_j \hat{g}_j^{\perp}$ and let $\hat{V}$ be the estimate of the variance of the vector $\hat{v}$ from Section J.3. The test statistic is $W_{\text{al}} = \hat{v}' \hat{V}^{+} \hat{v}$, where $\hat{V}^{+}$ is the generalized inverse.

Proposition 16 (Inference on alignment)**.** *Let J be fixed. Suppose individuals are independent and identically distributed, outcomes have finite fourth moments, every location has a positive share of residents and of movers, and $\text{Var}_{\pi}(\bar{Y}_j)$ and $\mathcal{L}_D(\bar{Y}, \bar{Y})$ are positive.*

(i) *If $r_D < 1$, then $\sqrt{N}(\hat{r}_D - r_D)$ and $\sqrt{N}(\tilde{r}_D - r_D)$ are asymptotically normal, and the bootstrap that resamples individuals is consistent for their distributions.*

(ii) *If $r_D = 1$ and the limiting variance of $\sqrt{N}\hat{v}$ has rank $J - 2$, then W_{al} converges in distribution to a chi-squared variable with $J - 2$ degrees of freedom. In this case $N(1 - \hat{r}_D^2)$ converges to a positive random variable. Its split-sample counterpart, before the truncation at zero in $\tilde{r}_D$, converges to a variable with mean zero.*

Part (ii) matters in practice. Under alignment, $1 - \hat{r}_D^2$ is centered at a positive number of the order of J divided by the number of movers. Sparse flows therefore look like misalignment. The split-sample estimate removes this bias. It cannot exceed one, so under alignment its mean is slightly below one. With many locations and few movers on each route, sample flows are far from the population flows, and an approximation with a fixed number of locations is poor. Remark 2 says what a theory for that case would require.

The bounds of Theorem 3 are smooth functions of the coefficient and the alignment when $r_D < 1$ and $|1 - 2\theta_D| < 1/r_D$. The coefficient and the alignment are jointly asymptotically normal, by the argument of part (i). The estimates of the bounds are then asymptotically normal too, and confidence intervals for θ_π follow from Imbens and Manski [2004] and Stoye [2009].

J.2 With Estimates of Place Effects

Where place effects can be estimated, the cross-sectional coefficient and the difference $\theta_D - \theta_\pi$ can be estimated too. Place effects and location means are estimated with noise, and the square of a noisy estimate is too large. This is the inflation that produces limited-mobility bias [Andrews et al., 2008, Kline et al., 2020]. Cross-fitting removes it. Individuals are split at random into two halves. In each half, place effects are estimated from that half's movers, and location means from that half's residents. Every product of two estimates then takes one factor from each half, so the errors of the two factors are independent.

Every estimator in this construction is linear in outcomes. The influence of each individual is therefore available in closed form, and a multiplier bootstrap over individuals needs no re-estimation. Section F gives the estimators, their limiting distribution with a fixed number of locations, and the test of directional isotropy.

J.3 Estimators and Proof of Proposition 16

This subsection gives the estimators of alignment and proves Proposition 16. The weights are the shares of moves that involve each location.

Estimators. Let n be the number of movers. For mover i let $D_i = \mathbf{e}_{d(i)} - \mathbf{e}_{o(i)}$ and $\iota_i = \frac{1}{2}\{\mathbf{e}_{o(i)} + \mathbf{e}_{d(i)}\}$, where $\mathbf{e}_j$ is the jth unit vector. Let $\hat{x}$ be the vector of sample location means, each the average of the two period means of the location. Define

$$\hat{L} = \frac{1}{n}\sum_i D_i D_i', \qquad \hat{\pi} = \frac{1}{n}\sum_i \iota_i, \qquad \hat{M} = \operatorname{diag}(\hat{\pi}) - \hat{\pi}\hat{\pi}', \qquad \hat{\kappa} = \frac{\hat{x}'\hat{L}\hat{x}}{\hat{x}'\hat{M}\hat{x}}, \qquad \hat{v} = \hat{L}\hat{x} - \hat{\kappa}\,\hat{M}\hat{x}.$$

The entries of $\hat{L}\hat{x}$ are $\hat{\pi}_j \hat{g}_j$, so $\hat{v}_j = \hat{\pi}_j \hat{g}_j^{\perp}$. By construction $\mathbf{1}'\hat{v} = 0$ and $\hat{x}'\hat{v} = 0$. The estimate of alignment is

$$\hat{r}_D = (1 + \hat{\Lambda}^2)^{-1/2}, \qquad \hat{\Lambda}^2 = \frac{\hat{x}'\hat{M}\hat{x}}{(\hat{x}'\hat{L}\hat{x})^2}\sum_j \frac{\hat{v}_j^2}{\hat{\pi}_j},$$

and $\hat{\Lambda}^2$ is the sample counterpart of $(1 - r_D^2)/r_D^2$. For the split-sample estimate, assign each individual to one of two halves by a fair coin. Compute $\hat{v}^A$ and $\hat{x}^{A\prime}\hat{L}^A\hat{x}^A$ from the flows and the location means of half A, and likewise for half B. Then

$$\tilde{\Lambda}^2 = \hat{x}'\hat{M}\hat{x}\sum_j \frac{1}{\hat{\pi}_j}\,\frac{\hat{v}_j^A}{\hat{x}^{A\prime}\hat{L}^A\hat{x}^A}\,\frac{\hat{v}_j^B}{\hat{x}^{B\prime}\hat{L}^B\hat{x}^B}, \qquad \tilde{r}_D = \big(1 + \max\{\tilde{\Lambda}^2, 0\}\big)^{-1/2}.$$

Influence functions and the test. Let N_{jt} be the number of residents of location j in period t. For individual ℓ, let $\rho_{\ell t}$ be her outcome in period t minus the sample mean of her location in that period. Her contribution to the location means is the vector

$$\zeta_\ell = \tfrac{1}{2}\sum_{t=1,2} \mathbf{e}_{j(\ell,t)}\,\frac{\rho_{\ell t}}{N_{j(\ell,t),t}}.$$

For a mover i, let

$$A_i = D_i\,(D_i'\hat{x}) - \hat{\kappa}\left\{\iota_i \circ (\hat{x} - \bar{x}\mathbf{1}) - \hat{\pi}\,\iota_i'(\hat{x} - \bar{x}\mathbf{1})\right\}, \qquad \bar{x} = \hat{\pi}'\hat{x},$$

where $\circ$ is the entrywise product, and let $\bar{A}$ be the mean of the A_i over movers. With $\hat{P} = I - \hat{M}\hat{x}\hat{x}'/(\hat{x}'\hat{M}\hat{x})$ and $\hat{B} = \hat{L} - \hat{\kappa}\hat{M}$, the estimated influence of individual ℓ on $\hat{v}$ is

$$\hat{\phi}_\ell = \frac{\mathbf{1}\{\ell \text{ is a mover}\}}{n}\,\hat{P}\,(A_\ell - \bar{A}) + \hat{B}\,\zeta_\ell, \qquad \hat{V} = \sum_\ell \hat{\phi}_\ell\hat{\phi}_\ell'.$$

The statistic is $W_{\text{al}} = \hat{v}'\hat{V}^+\hat{v}$. Here $\hat{V}^+$ is the generalized inverse that discards the two smallest eigenvalues of $\hat{V}$, which correspond to the two linear restrictions on $\hat{v}$.

Proof of Proposition 16. With J fixed, the entries of $\hat{L}$ and $\hat{\pi}$ are shares of movers on each route, and the entries of $\hat{x}$ are cell means. All are smooth functions of finitely many sample averages over individuals, so $(\hat{L}, \hat{\pi}, \hat{x})$ is consistent and jointly asymptotically normal at rate $\sqrt{N}$.

(i) Where $\hat{x}'\hat{L}\hat{x}$ and $\hat{x}'\hat{M}\hat{x}$ are positive, $\hat{\Lambda}^2$ is a smooth function of $(\hat{L}, \hat{\pi}, \hat{x})$, and $\hat{r}_D$ is a smooth function of $\hat{\Lambda}^2$. The delta method gives the limit, and the bootstrap is consistent for smooth functions of sample averages. For the split-sample estimate, write $\hat{d}^A$ and $\hat{d}^B$ for the two normalized vectors in the sum and d for their common limit. Then $\hat{d}^{A\prime}\hat{\Pi}^{-1}\hat{d}^B = d'\hat{\Pi}^{-1}d + d'\hat{\Pi}^{-1}\{(\hat{d}^A - d) + (\hat{d}^B - d)\} + O_p(N^{-1})$. The linear term is asymptotically normal, and it does not vanish when $d \neq 0$.

(ii) Let $r_D = 1$, so that $v = L_D x - \kappa M_\pi x = 0$ at the population values $x = \bar{Y}$. A first-order expansion of $\hat{v}$ in $(\hat{L} - L_D, \hat{M} - M_\pi, \hat{x} - x)$ gives

$$\hat{v} = P\left\{(\hat{L} - L_D) - \kappa(\hat{M} - M_\pi)\right\}x + (L_D - \kappa M_\pi)(\hat{x} - x) + O_p(N^{-1}), \qquad P = I - \frac{M_\pi x\, x'}{x' M_\pi x}.$$

The projection P appears because $\hat{\kappa} - \kappa = x'\{(\hat{L} - L_D) - \kappa(\hat{M} - M_\pi)\}x/(x'M_\pi x) + O_p(N^{-1})$ when $v = 0$; the term in $\hat{x} - x$ drops out of $\hat{\kappa} - \kappa$ for the same reason. The first term of the expansion

is an average over movers of the vectors A_i at the population values, centered. The second is an average over all individuals of the vectors ζ_ℓ. Hence $\hat{v} = \sum_\ell \phi_\ell + O_p(N^{-1})$, where ϕ_ℓ is $\hat{\phi}_\ell$ at the population values, and $\sqrt{N}\hat{v}$ converges to a normal vector with variance $V = \lim N \sum_\ell \mathbb{E}[\phi_\ell \phi_\ell']$. The estimated influence functions replace population values by consistent estimates in a smooth map, so $N\hat{V} \to V$ in probability, as in the proof of Proposition 10. V has rank $J-2$ by assumption, so the two smallest eigenvalues of $N\hat{V}$ converge to zero and the others stay away from zero. W_{al} therefore converges to a chi-squared variable with $J-2$ degrees of freedom.

Finally, $N\hat{\Lambda}^2$ is a positive definite quadratic form in $\sqrt{N}\hat{v}$, up to factors that converge in probability, so its limit is a positive random variable. $N\tilde{\Lambda}^2$ is a bilinear form in $\sqrt{N}\hat{v}^A$ and $\sqrt{N}\hat{v}^B$. The two halves share no individuals, so the two vectors are asymptotically independent with mean zero, and the limit has mean zero. □

K Simulations

A simulated economy shows the results at work. It is a larger version of Example 1, with 256 locations indexed by (p, z) and 16 values of each. The place effect varies with p. Composition rises with the place effect, with a sorting slope of 0.6, and has a residual part that varies with z. The cross-sectional coefficient is 0.49. The economy is additive with exogenous mobility, so each route shows the difference of two place effects, and nothing below is caused by extra gains. Only the routes vary. A share ρ of moves change p and keep z, and the others change z and keep p.

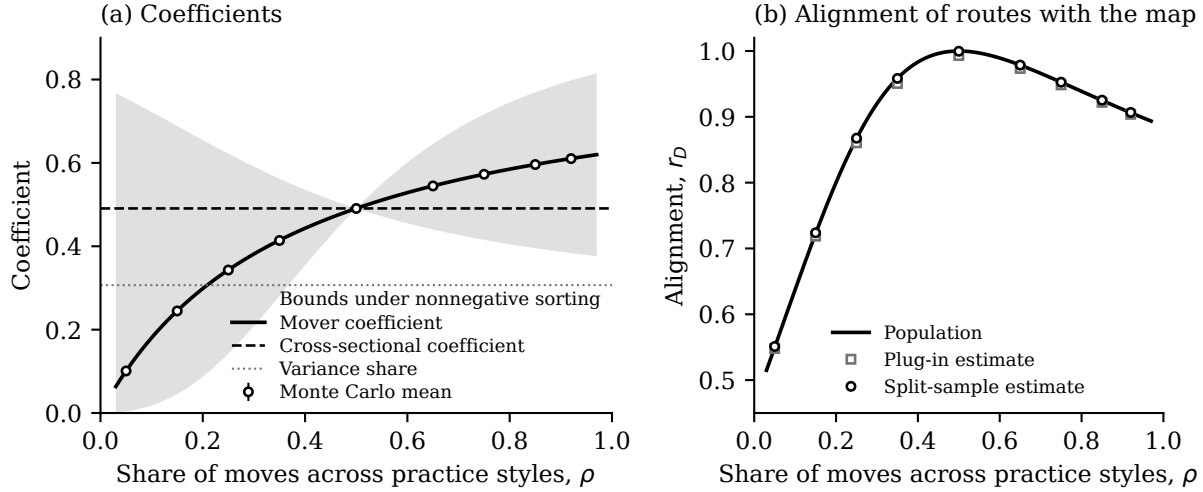

Figure 1: The Mover Coefficient and Alignment as Routes Change

Notes. The economy of this section, which is Panel B of Table 3. Lines are population values. Circles and squares are means over simulated samples (500 at the four values of ρ in the table, 200 elsewhere), and bars in panel (a) are two standard deviations across samples. The shaded band in panel (a) shows the bounds of Theorem 3.

Panel (a) of Figure 1 shows the mover coefficient. It is 0.25 at $\rho = 0.15$, where most moves change z, and 0.61 at $\rho = 0.92$, where most change p. It equals the cross-sectional coefficient only at

$\rho = \frac{1}{2}$. The locations, their place effects, and their residents are the same throughout. A two-way decomposition of the variance of outcomes into people and places gives the same answer for all these routes. Panel (b) shows the alignment, which is computed from flows and location means alone. It is one at $\rho = \frac{1}{2}$ and falls on both sides, to 0.93 at $\rho = 0.85$ and to 0.72 at $\rho = 0.15$.

The shaded band in panel (a) shows the bounds of Theorem 3. They contain the cross-sectional coefficient for every ρ, and they are wide. At $\rho = 0.85$, a coefficient of 0.596 and an alignment of 0.925 give bounds of 0.40 and 0.77. The rest of this section reports the estimates and tests of Section J in samples from these economies.

The simulated economies hold the cross-section fixed and change only the routes. Every economy is additive, with exogenous mobility. Each route shows the difference of two place effects, so Assumption 1 holds in its strongest form. What the simulations show is therefore present in the model that is most favorable to the design. It is not caused by extra gains or by effects that differ across people.

Design. The economy is a larger version of Example 1. There are 256 locations on a grid, indexed by (p, z) with 16 values each. The place effect varies with p and has unit variance. Composition is $\bar{\alpha}_j = \beta_\pi \psi_j + q_j$, where the residual q_j varies with z. The sorting slope is $\beta_\pi = 0.6$ and the residual composition is $\lambda_\pi = 0.7$. Hence $\theta_\pi = 0.49$, between a variance share of 0.31 and an equalization share of 0.67. Each location has 1,200 stayers, and there are 25,000 movers, so about 332,000 individuals and a mover share of 7.5 percent. Outcomes follow (4). Person components have a standard deviation of one within each location, and shocks a standard deviation of 0.2. Each location has the same composition in both periods under every pattern of routes, so Assumption 3 holds exactly, as the paragraph on individuals below explains.

Routes are the only thing that varies. In Panel A of Table 3, destinations are drawn independently of origins, which is independent pairing. In Panel B, a share ρ of moves are across practice styles, which change p and keep z. The others are across patient populations, which change z and keep p. At $\rho = \frac{1}{2}$ the routes are aligned with the map and directionally isotropic. Away from $\rho = \frac{1}{2}$ they are neither. Each design is simulated 500 times, and every estimate of Section J is computed on each sample.

The coefficient follows the routes. The four economies of Panel B have the same locations, the same place effects, and the same composition. A two-way variance decomposition would call them one economy: the variance of person components, the variance of place effects, and their correlation agree to about one percent. Yet the mover coefficient ranges from 0.25 to 0.61. With most moves across patient populations, $\rho = 0.15$, it is half of θ_π. With most moves across practice styles it is 0.60 to 0.61, well above θ_π. Figure 1 shows the whole range. The cross-sectional coefficient is flat, because the cross-section does not change. The mover coefficient rises from near zero to above 0.6.

Panel A confirms the cases in which the two coefficients agree by construction. Under independent pairing the mover coefficient equals θ_π in every row. It equals the variance share only in the second row, where composition is unrelated to the place effect. In the third row and in Panel C, composition is an exact linear function of the place effect. The coefficient is then $1/(1 + \beta_\pi) = 0.625$

Table 3: Simulated Mover Designs: One Cross-Section, Many Routes

	Two-way decomposition			Mover	Cross-sectional		Align-	Bounds on θ_π,	Rejection rate	
Design	Var α	Var ψ	Corr	$\hat{\theta}_D$	θ_π	$\hat{\theta}_\pi$	ment $\tilde{r}_D$	Theorem 3	align.	isotropy
Panel A. Destinations drawn independently of origins; the cross-section varies										
General ($\beta_\pi > 0, \lambda_\pi > 0$)	2.14	1.00	0.41	0.491	0.491	0.491	1.000	[0.49, 0.50]	0.048	0.054
Orthogonal ($\beta_\pi = 0$)	1.75	1.00	0.00	0.588	0.588	0.588	1.000	[0.58, 0.59]	0.068	0.040
Affine ($\lambda_\pi = 0$)	1.39	1.00	0.51	0.625	0.625	0.625	1.000	[0.62, 0.63]	0.052	0.020
Panel B. One cross-section; only the routes vary										
Balanced ($\rho = 0.50$)	2.09	1.00	0.42	0.491	0.491	0.491	1.000	[0.48, 0.50]	0.066	0.048
Practice styles ($\rho = 0.85$)	2.09	1.00	0.42	0.596	0.491	0.491	0.925	[0.40, 0.77]	1.000	1.000
Patient populations ($\rho = 0.15$)	2.11	1.00	0.41	0.245	0.491	0.491	0.724	[0.05, 0.69]	1.000	1.000
Practice styles, near pure ($\rho = 0.92$)	2.09	1.00	0.42	0.610	0.491	0.491	0.907	[0.38, 0.80]	1.000	1.000
Panel C. Composition is an exact linear function of the place effect										
Affine, patient populations ($\rho = 0.15$)	1.36	1.00	0.51	0.624	0.625	0.625	0.998	[0.61, 0.64]	0.088	0.002

Notes. Means over 500 simulated samples of an economy with 256 locations, 1,200 stayers per location and 25,000 movers. The first three columns are the variance of person components, the variance of place effects, and their correlation, pooled over individuals and periods. $\hat{\theta}_D$ is the coefficient of regression (8) with leave-one-out location means. θ_π is the population value, with the shares of moves that involve each location as weights, and $\hat{\theta}_\pi$ its cross-fitted estimate. $\tilde{r}_D$ is the split-sample estimate of alignment. The bounds are (23) at $\hat{\theta}_D$ and $\tilde{r}_D$. The last two columns are the shares of samples in which the test of alignment (Proposition 16) and the test of directional isotropy (Proposition 9) reject at the 5 percent level. In Panel B the cross-section is the same in every row, and ρ is the share of moves across practice styles; the other moves are across patient populations.

under both patterns of routes, as Theorem 4(iii) says.

Alignment sees the problem from flows and location means. Panel (b) of Figure 1 shows the alignment of the routes with the map. It is one at $\rho = \frac{1}{2}$ and falls on both sides, to 0.93 at $\rho = 0.85$ and to 0.72 at $\rho = 0.15$. No place effect is estimated to compute it. The split-sample estimate is on the population line. The plug-in estimate is a little below it. At $\rho = \frac{1}{2}$ its mean is 0.993, where the truth is one.

Both tests detect the problem in every sample. At the 5 percent level, the test of alignment rejects in every sample of every design that is not aligned. In the five aligned designs it rejects in 5 to 9 percent of samples, which is a little too often. It tests 254 restrictions at once. The test of directional isotropy uses estimated place effects. It rejects in every sample of the three designs that are not isotropic. In the three isotropic designs with residual composition it rejects in 4 to 5 percent of samples. When composition is an exact linear function of the place effect, routes cannot matter, and this test rejects in 2 percent of samples or fewer.

The bounds contain the truth and say how much is known. The bound of Theorem 3(i) holds with equality in this economy, as in Example 1: the routes have only two dimensions to choose from. The bounds of Theorem 3 are the edges of the shaded band in the figure. They always contain θ_π, and they coincide when routes are aligned. In the designs that are not aligned, the estimated bounds contain θ_π in every sample. They are also wide. At $\rho = 0.85$, a coefficient of 0.596 and an alignment

of 0.925 give bounds of 0.40 and 0.77. At $\rho = 0.15$, a coefficient of 0.25 and an alignment of 0.72 give bounds of 0.05 and 0.69. An alignment above 0.9 looks high, and it still leaves θ_π anywhere between 0.4 and 0.8. This is the content of the anything-goes result in numbers. Where place effects can be estimated, the cross-fitted estimate $\hat{\theta}_\pi$ is on target in every design. The confidence interval for the difference of the two coefficients covers the true difference in 93 to 97 percent of samples, and in nearly every sample of the affine design with two kinds of moves (Table 4).

Locations. There are 256 locations, indexed by (p, z) with p and z each taking 16 values. The 16 values are the means of 16 equally likely bins of a standard normal variable, standardized to have mean zero and unit variance across locations. The place effect is the value of p. The residual q_j is $\sqrt{\lambda_\pi}$ times the value of z, and composition is $\bar{\alpha}_j = \beta_\pi \psi_j + q_j$. The general design has $\beta_\pi = 0.6$ and $\lambda_\pi = 0.7$, the orthogonal design $\beta_\pi = 0$, and the affine design $\lambda_\pi = 0$.

Routes. Under independent pairing, every pair of locations is equally likely. In the designs with two kinds of moves, a share ρ of the comparison weight is spread evenly over the pairs of locations that share z, and the rest evenly over the pairs that share p. The direction of each move is drawn at random. All locations have the same share of moves, so the weights π are equal.

Individuals. Each location has 1,200 stayers, whose person components are normal with mean $\bar{\alpha}_j$ and unit variance. There are 25,000 movers. A mover between j and k has a person component that is normal with mean $\varphi_j + \varphi_k$ and unit variance, where the vector φ solves $\varphi_j + \sum_k W_{jk} \varphi_k = \bar{\alpha}_j$ and W_{jk} is the share of the moves that involve j whose other end is k. The mean person component of the movers who leave j is then $\bar{\alpha}_j$, and so is that of the movers who arrive in j. Each location therefore has the same composition in both periods, for every pattern of routes, and Assumption 3 holds exactly. This is the economy with sorted movers of Section C. Outcomes follow (4), with normal shocks of standard deviation 0.2, independent across individuals and periods.

Estimates. The mover coefficient is the coefficient of regression (8), where the mean outcome of a location in a period leaves out the mover herself. Alignment and its test are computed as in Section J.3, with the average of the two period means of each location as the map. The bounds of Theorem 3 are evaluated at the estimated coefficient and the split-sample alignment. The cross-fitted estimates and the test of directional isotropy follow Section F.3, with place effects estimated by least squares on the movers of each half and 1,000 bootstrap draws.

Additional results. Table 4 reports means and standard deviations across samples. The plug-in estimate of alignment is below the population value in every design. The split-sample estimate is within 0.002 of it. It cannot exceed one, so with aligned routes its mean is slightly below one. The estimates of residual composition show how the routes differ from the cross-section: $\hat{\lambda}_D$ is near four when most moves are across patient populations and near zero when most are across practice

styles, while $\hat{\lambda}_\pi$ stays at 0.7. Confidence intervals for θ_π and for the difference $\theta_D - \theta_\pi$ cover the population values in about 95 percent of samples. With affine composition and two kinds of moves, the interval for the difference covers too often.

Table 4: Simulated Mover Designs: Additional Results

			Alignment			Residual composition		Coverage of 95% CI	
Design	$\hat{\theta}_D$	$\hat{\theta}_\pi$	r_D	$\hat{r}_D$	$\tilde{r}_D$	$\hat{\lambda}_\pi$	$\hat{\lambda}_D$	θ_π	$\theta_D - \theta_\pi$
General ($\beta_\pi > 0$, $\lambda_\pi > 0$)	0.491 (0.002)	0.491 (0.001)	1.000	0.998	1.000 (0.000)	0.701	0.700 (0.009)	0.95	0.95
Orthogonal ($\beta_\pi = 0$)	0.588 (0.003)	0.588 (0.002)	1.000	0.997	1.000 (0.000)	0.700	0.700 (0.009)	0.94	0.95
Affine ($\lambda_\pi = 0$)	0.625 (0.001)	0.625 (0.001)	1.000	0.997	1.000 (0.000)	0.000	0.000 (0.000)	0.96	0.93
Balanced ($\rho = 0.50$)	0.491 (0.002)	0.491 (0.002)	1.000	0.993	1.000 (0.001)	0.700	0.699 (0.015)	0.94	0.97
Practice styles ($\rho = 0.85$)	0.596 (0.001)	0.491 (0.002)	0.925	0.922	0.925 (0.002)	0.699	0.123 (0.004)	0.95	0.95
Patient populations ($\rho = 0.15$)	0.245 (0.004)	0.491 (0.002)	0.722	0.719	0.724 (0.006)	0.701	3.969 (0.114)	0.95	0.94
Practice styles, near pure ($\rho = 0.92$)	0.610 (0.001)	0.491 (0.002)	0.906	0.904	0.907 (0.002)	0.700	0.061 (0.003)	0.95	0.95
Affine, patient populations ($\rho = 0.15$)	0.624 (0.002)	0.625 (0.002)	1.000	0.969	0.998 (0.003)	0.000	0.000 (0.001)	0.94	1.00

Notes. The designs of Table 3, in the same order. Entries are means over 500 samples, with standard deviations across samples in parentheses. r_D is the population alignment, $\hat{r}_D$ the plug-in estimate, and $\tilde{r}_D$ the split-sample estimate. $\hat{\lambda}_\pi$ and $\hat{\lambda}_D$ are the cross-fitted estimates of residual composition across locations and along routes. The last two columns are the shares of samples in which the bootstrap confidence interval contains the population value.

L Details for the Evidence

L.1 Further Results for the U.S. Commuting Zones

Groups of comparisons. Proposition 7 needs the assumption of the design group by group, and not route by route. Take the 211 large zones and the comparisons between destinations. I sort the pairs of zones by how common their comparison is relative to independent pairing, and form groups of pairs with equal comparison weight. With population weights, the alignment of the groups together is 0.97 with three groups and 0.98 with five, against 0.85 for the single design. Groups that are formed by the position of the two zones on the map add nothing. With three such groups the alignment is 0.86, as it is with groups drawn at random.

The group coefficients can be estimated without place effects. A comparison is a pair of destinations of one origin. Its estimate is the difference between the published estimates for the two pairs of origin and destination. The coefficient of a group is the weighted regression of these differences on the gaps in the map, over the comparisons of the group. With five groups the coefficients are 0.28, 0.39, 0.38, 0.53 and 0.66, from the least common comparisons to the most common. Their standard errors are between 0.15 and 0.20. The bounds of Proposition 7 are 0.23 and 0.40. The single design has a coefficient of 0.42, with a standard error of 0.12, and bounds of 0.18 and 0.70. With three groups the bounds are 0.24 and 0.47.

The groups narrow the bounds more than the confidence interval. Five groups reduce the width of the bounds from 0.52 to 0.17. The 95 percent confidence interval for the cross-sectional

coefficient is 0.05 to 0.86 for the single design and −0.01 to 0.70 with five groups. Five coefficients are estimated in place of one, and the coefficient of the least common comparisons has a weight of 0.79 out of 0.97. With the group coefficients computed from the published place effects, the bounds with five groups are 0.36 and 0.52. The single design then gives 0.25 and 0.78, at the coefficient that the same place effects imply. The cross-sectional coefficient computed from these place effects is 0.47. It is above the bounds from the estimates for pairs and inside the confidence interval.

With mover weights, the alignment of these groups together is 0.98 with three groups and 0.99 with five, against 0.96 for the single design. The bounds from the estimates for pairs are 0.34 and 0.47 with five groups, against 0.29 and 0.57 for the single design. The routes are not misaligned because the mean outcome of a zone's partners is a curve in its own mean outcome. They are misaligned zone by zone. A group is there a set of comparisons between destinations, and not a set of children. In the comparisons of origin against destination, where a group is a set of routes, five groups raise the alignment from 0.84 to 0.90. Most pairs of zones have no published route there, and no weights can use comparisons that are missing. With three groups the bounds from the estimates for pairs are 0.29 and 0.62, against 0.26 and 0.80 for the single design. With five groups the coefficients range from 0.18 to 0.94, and no value satisfies the condition of Proposition 7(iii) at the point estimates.

Tax flows. The county-to-county migration flows of all tax filers, which the Internal Revenue Service publishes, give the same picture as the routes of the sample. The alignment is between 0.76 and 0.97, and the two coefficients computed from the place effects differ by between −0.10 and 0.09. The construction is described below.

Public data for the United States. Table 1 uses two files from the replication files of Chetty and Hendren [2018b]. Online data table 3 gives, for each commuting zone, the raw estimate of the effect of one year of childhood on the income rank at age 26, its standard error, the mean rank of permanent residents, for parents at the 25th and the 75th percentile, and the population in 2000. I multiply the effects and their standard errors by twenty. Online data table 5 lists each ordered pair of commuting zones with at least 25 children in the sample of movers. It gives the number of those children, and the estimate for the pair from which the place effects are computed: the coefficient on age at move, as a linear function of the rank of the parents. I use the children whose income is observed at age 26. There are 1.39 million of them on 11,017 ordered pairs between the 595 zones that have an estimated effect, and 1.07 million between the 211 zones with at least 250,000 people. The counts pool all parental incomes. The files give two standard errors for each effect, a raw one and a bootstrapped one. The standard errors of Table 1 and the sorting slopes of Table 6 use the raw one. For parents at the 75th percentile the alignment is between 0.80 and 0.98, and the differences between the two coefficients are between −0.08 and −0.03.

What the map misses and the sign of sorting. Three numbers need the variance of place effects across zones: what the map misses, η, the correlation implied by (21), and the sorting slope. The

Table 5: What the Map Misses and the Sign of Sorting, with Each Published Standard Error

	Raw standard errors			Bootstrapped standard errors		
Comparisons, weights on locations	η	Corr.	Sorting slope	η	Corr.	Sorting slope
A. All 595 commuting zones						
Origin and destination, population	0.62	−0.23	−0.21	—	—	1.52
Origin and destination, mover weights	0.63	−0.18	−0.20	0.29	−0.38	0.22
Two destinations, population	0.67	−0.28	−0.25	—	—	0.38
Two destinations, mover weights	0.56	−0.29	−0.11	0.30	−0.54	0.25
B. The 211 commuting zones with at least 250,000 people						
Origin and destination, population	0.48	0.29	0.05	0.37	0.38	0.33
Origin and destination, mover weights	0.47	0.22	0.08	0.39	0.27	0.25
Two destinations, population	0.48	0.17	0.05	0.37	0.22	0.33
Two destinations, mover weights	0.41	−0.06	0.18	0.33	−0.08	0.38

Notes. The rows are those of Table 1, and the place effects are the published two-way estimates. η is the size of what the map misses about place effects, relative to the map. Corr. is the correlation between what the map misses and the part of the comparison gap that the map does not explain, implied by (21). The sorting slope is $\mathrm{Cov}_\pi(\bar{\alpha}_j, \psi_j)/\mathrm{Var}_\pi(\psi_j)$. All three use the variance of the published effects across zones less $\sum_j \pi_j(1-\pi_j)s_j^2$, where s_j is the raw or the bootstrapped standard error of zone j in online data table 3 of Chetty and Hendren [2018b]. A dash marks a row in which this variance is below θ_π^2 times the variance of the map. The slope of −0.21 in the first row has a standard error of 0.16 by the delta method.

variance of place effects is the variance of the published effects net of their sampling error. The sampling error is most of that variance: between 0.62 and 0.68 of it with the raw standard errors, and between 0.70 and 0.89 with the bootstrapped ones. Table 5 reports the three numbers with each standard error, for the rows of Table 1. Across all zones the two standard errors give opposite signs of sorting. With the raw standard errors the slope is negative in all four rows, and η is above the largest value that nonnegative sorting allows, $\sqrt{\theta_\pi(1-\theta_\pi)}$. With the bootstrapped standard errors the slope is positive in all four rows. In two of these rows the corrected variance is below θ_π^2 times the variance of the map, which no place effects can produce, and η cannot be computed. Across the 211 large zones the slope is positive in every row with both standard errors, and η is below that largest value. The script `ch_public_sorting_two_se.py` in the folder `us_public` computes the table.

The pairs that are not listed. Chetty and Hendren [2018b, footnote 16] report that the 11,216 listed pairs account for 75 percent of the moves across zones in their sample. About 466,000 children are therefore on pairs with fewer than 25 children. The coefficient, the alignment, and the published place effects all use the listed pairs. The place effects use all of them. The coefficient and the alignment use the 11,017 between the 595 zones, which carry 99.4 percent of the children. To see what the missing moves could change, I allocate them in two ways and compute the alignment of all moves. In both, the missing moves are a third of the listed moves between the zones of the panel, 463,000 across all zones and 357,000 across the large zones. They enter as flows, and the comparison weights and the mover weights are computed anew. The zones that are compared with no other destination stay out. In the first, the missing quarter is an exactly aligned component: independent pairing with the weights of the row, on all pairs. In the order of the rows of Table 1,

the alignment is then 0.89, 0.91, 0.92 and 0.98 across all zones, and 0.92, 0.94, 0.92 and 0.99 across the large zones. In the second, the missing moves are paired independently on the pairs that are not listed, with at most 24 children on each, which is what the truncation allows. The number on a pair is the smaller of 24 and a multiple of the product of the weights of its two zones, and the multiple is set to place all missing moves. The alignment is then 0.92, 0.94, 0.94 and 0.99 across all zones, and 0.93, 0.96, 0.92 and 0.99 across the large zones. In seven of the eight rows it is higher than in the first case, because the listed pairs are the common ones and the pairs that are not listed are those that independent pairing would add. The smallest number of zones with a comparison gap of the wrong sign, over all levels of the map, is 106, 106, 59 and 59 of the zones in the first four rows and 32, 32, 20 and 20 in the last four. In the second case it is 54, 55, 12 and 17, and 19, 17, 2 and 2. Neither case is a bound: missing moves that offset the listed ones would raise the alignment further, and missing moves that resemble the listed ones would leave it unchanged. Confidence intervals in Table 1 follow Imbens and Manski [2004]. The end points move with the coefficient by the delta method. The sampling error of the alignment changes no digit: resampling the children across the listed pairs gives it a standard error below 0.001.

The mover coefficient from the estimates for pairs. For the mover coefficient without place effects, I evaluate the estimate for each ordered pair at the 25th percentile, reverse its sign, and express it in percentile ranks over twenty years of childhood. I then regress it on the gap in the map between destination and origin, without a constant and with the number of children as weights. The coefficient is 0.61 (s.e. 0.07) across all zones and 0.53 (s.e. 0.09) across the large zones. The place effects on the same routes give 0.63 and 0.56. For two destinations, I regress the same estimate on the map of the destination, with an intercept for each origin and the same weights. The coefficient is 0.47 (s.e. 0.10) across all zones and 0.42 (s.e. 0.12) across the large zones, where the place effects on the same comparisons give 0.60 and 0.52. The standard errors treat the estimates for pairs as independent. For parents at the 75th percentile, the coefficient for origin against destination is 0.72 (s.e. 0.09) across all zones. With an alignment of 0.80, the bounds of Theorem 3 are then 0.36 and 0.92. The script `ch_public_pairs_coefficient.py` in the folder `us_public` computes these numbers.

Group coefficients from the estimates for pairs. For two destinations j and k of an origin o, the comparison has weight $P(o)\,p(j \mid o)\,p(k \mid o)$. Its estimate is the estimate for the pair of o and j less the estimate for the pair of o and k. The coefficient of a group is the regression of these differences on $\bar{Y}_j - \bar{Y}_k$ with these weights, over the pairs of destinations in the group. For origin against destination, a comparison is a route, its estimate is the estimate for the pair, and each group has its own constant. The standard errors treat the estimates for pairs as independent.

The number v_G of Proposition 7 is quadratic in the group coefficients, so their sampling error raises it. I subtract the part that the published sampling variances imply. I keep v_G at or above t_G^2 / R_D^2, the smallest value that v_G can take for given t_G and R_D. With one group the two are equal,

and the bounds are those of Theorem 3. The confidence interval follows Imbens and Manski [2004]. The end points of the bounds move with the estimates for pairs and with the flows. The estimates are drawn from normal distributions with the published variances. The flows are drawn by resampling the children across the listed pairs. The groups and the map are held fixed. The script `ch_public_groups_direct.py` in the same folder computes these numbers.

Reweighted routes. For each pair of zones that exchange children, the contrast between the zone with the higher and the zone with the lower value of the map combines the estimates of the two directions with precision weights. With the observed routes this gives 0.60 across all zones, a little below the 0.61 of the previous paragraph. The weights on pairs solve the alignment condition of Proposition 8(ii), with population weights on locations. Among the solutions I take the one with the smallest variance of the reweighted coefficient, treating the estimates for pairs as independent. The standard errors treat the weights as fixed. Relative to the number of movers, the median weight is 0.50 between the large zones and 0.37 across all zones, the 99th percentile is 16 and 18, and the largest weights are several hundred.

Tax flows. The county outflow files of the Statistics of Income division of the Internal Revenue Service, at https://www.irs.gov/statistics/soi-tax-stats-migration-data, give the number of returns and of exemptions that leave each county for each other county, for 2004–2005 to 2010–2011. I count exemptions, which approximate persons, sum over the seven pairs of years, assign counties to commuting zones with online data table 4 of Chetty and Hendren [2018b], and drop moves within a zone. The files disclose a pair of counties only when enough returns move between them. The disclosed pairs carry 76 percent of all persons who leave a county. This leaves 38.9 million moves between the 595 zones, on 17,900 pairs of zones. In the order of the rows of Table 1, the alignment with these routes is 0.76, 0.82, 0.83 and 0.93 across all zones, and 0.82, 0.88, 0.88 and 0.97 across the large zones. The differences between the two coefficients are -0.10, -0.06, -0.10 and -0.05, and 0.09, 0.05, 0.07 and 0.01.

With flows F_{jk} from j to k, the comparison weights are $a_{jk} \propto F_{jk} + F_{kj}$ for origin and destination, and $a_{jk} \propto \sum_o P(o)\, p(j \mid o)\, p(k \mid o)$ for two destinations, where $P(o)$ is the share of movers who leave o and $p(\cdot \mid o)$ is the distribution of their destinations. The alignment is $r_D = \mathrm{Corr}_\pi(g_j, \bar{Y}_j)$ with $g_j = \pi_j^{-1} \sum_k a_{jk}(\bar{Y}_j - \bar{Y}_k)$. With tens of millions of moves its sampling error is negligible, and across the seven pairs of years it varies by less than 0.02. The two coefficients are linear in the place effects, so their standard errors and that of their difference follow from the published standard errors, with the locations, the flows, and the map held fixed. This treats the estimates of different locations as independent. They are estimated jointly, and their covariances are not published, so the standard errors are approximate. The variance of place effects is the weighted variance of the estimates minus $\sum_j \pi_j(1-\pi_j)s_j^2$, where s_j is the standard error of location j, and η^2 is that variance less θ_π^2 times the variance of the map, relative to the variance of the map. Here θ_π^2 is taken net of its squared standard error.

Table 6: The Sign of Sorting Where Place Effects Have Been Estimated

Outcome	Locations	Sorting slope $\hat{\beta}_\pi$
Medicare utilization[a]	306 hospital referral regions, United States	0.32
Voter turnout[a]	States, United States	0.08
Income rank at 26[b]	211 large commuting zones, United States	0.05
Income rank at 26[b]	595 commuting zones, United States	−0.21 (0.16)

Notes. The sorting slope is $\mathrm{Cov}_\pi(\bar{\alpha}_j, \psi_j) / \mathrm{Var}_\pi(\psi_j)$. Standard error in parentheses. [a]Computed from the published variance share and equalization share as $(S_{\text{eq}} - \theta_{\text{var}})/(2\,\theta_{\text{var}})$: 0.44 and 0.72 in Finkelstein et al. [2016], 0.36 and 0.42 in Cantoni and Pons [2022]. [b]Computed from the place effects of Chetty and Hendren [2018b] for parents at the 25th percentile, with population weights and the variance of the effects net of sampling error. The sampling error is measured by the standard errors that the files label raw, and the standard error of the slope is by the delta method. With the bootstrapped standard errors of the same files the two slopes are 0.33 and 1.52. No standard error is computed for the large zones.

Three checks. Counting returns in place of persons changes the alignment by less than 0.02. For parents at the 75th percentile the alignment is between 0.80 and 0.98, and the differences between the two coefficients are between −0.07 and 0.00. Across the 2,872 counties with a map and disclosed flows, the alignment is between 0.73 and 0.91. The script and its output are in the folder `us_public` of the replication files.

L.2 The Sign of Sorting

The bounds of Theorem 3 are for a researcher who assumes that sorting is not negative. Under Assumption 1 the data on movers cannot establish the sign. Two-way estimates of place effects do, under the condition that they require. Table 6 reports the sorting slope β_π, the slope of composition on place effects across locations, in four settings: two published studies and the two sets of commuting zones of Section VII. The slope is positive in three. It is negative across all commuting zones of the United States, at −0.21 with a standard error of 0.16. With the bootstrapped standard errors of the same files it is positive there as well (Table 5). Under that condition the sign has support in these settings, and it is not a law. A study that reports a mover coefficient can say which sign it assumes, and a study that estimates place effects can report the sign.

L.3 The Variance Share

In published estimates, the mover coefficient is not the variance share of place effects. Finkelstein et al. [2016] report that equalizing place effects would remove 72 percent of the variance of Medicare utilization across regions, and that equalizing patient characteristics would remove 56 percent. The variance share of place effects is then 0.44 and the equalization share 0.72. Their midpoint is 0.58. The shares of gaps between groups of regions that the authors compute from estimated place effects are 0.54 to 0.61. Their event-study coefficient is about 0.5, and the slope of movers' changes in utilization on the gap is 0.63. For voter turnout across states, the two shares in Cantoni and Pons [2022] are 0.36 and 0.42, with a midpoint of 0.39. Their share of the gap between states above and below the median is 0.37, and their event-study coefficient is 0.40.

Take childhood moves between commuting zones in the United States. The cross-sectional coefficient, estimated from the effects of areas, is 0.032 per year of exposure [Chetty and Hendren, 2018b]. The mover coefficient is about 0.04 per year [Chetty and Hendren, 2018a]. The two come from samples that differ. In the first two studies the mover coefficient is numerically close to the midpoint, and not to the variance share. In the third it is a quarter above the cross-sectional coefficient, and the authors attribute the difference to the samples and to the income group. Whether the routes in those data are aligned with the map cannot be said from the published numbers. That takes the flows, and Section VII uses public ones.