Ranking Places by Intergenerational Mobility and Revealed Preference*

John Klopfer and Corbin Miller

August 20, 2026

PENDING RELEASE: DO NOT DISSEMINATE

Abstract

What is the counterfactual impact on children's outcomes of moving from one place to another? By design, standard estimates of place effects on intergenerational mobility hold parent incomes constant. In this paper, we define, estimate, and decompose counterfactual effects combining quasi-experimental exposure effects with indirect effects through parent incomes. We control for self-selection on moves across U.S. commuting zones (CZs) using a new feasible estimator equivalent to the infeasible high-dimensional fixed effects estimator, and rank the top 50 CZs with confidence intervals of 0-18 ranks. The new parent income channel doubles the effect of places on child outcomes and changes the ranking: counterfactual effects are highest in suburban and resource-rich rural CZs, and bargains are common because good outcomes associate with high incomes and low rents. Estimating revealed preference using parents' cross-CZ moves, parents prefer places with worse child outcomes but appear indifferent after accounting for rents and incomes; housing and commuter transit expansions could lower rents and persuade parents to move to places with better child outcomes.

Keywords: Intergenerational mobility, Place effects, Rankings, Revealed preference
JEL codes: D31, J62, R23

*Klopfer: Faculty of Business and Economics, The University of Hong Kong and the Hoover Institution, Stanford University; Miller: Office of Research, Applied Analytics, and Statistics, Internal Revenue Service. Feedback is appreciated. Correspondence: johnklopfer@gmail.com

We thank Philipp Barteska, Kirill Borusyak, Gregory Cox, Ash Craig, Matthias Fahn, Naijia Guo, Jasmine Hao, Anders Humlum, Julius Ilciukas, Simon Jäger, Alan Kwan, Gedeon Lim, Stelios Michalopoulos, Juan Pantano, Zhuan Pei, Kjell Salvanes, Uta Schönberg, Johannes Schmieder, Jeff Smith, Chang Sun, Betty Wang, Ashley Wong, Michael Wong, Luther Yap, Liang Zhong, and Anson Zhou for helpful discussions. Klopfer thanks the Hoover Institution for its hospitality, and the University of Hong Kong and the Research Grants Council of Hong Kong (GRF 17500922) for their support. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury or the Internal Revenue Service. All results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

The places where children grow up affect their incomes as adults (Chyn and Katz, 2021). This fact motivates programs that help families move to better places, such as the Chicago Housing Authority's Gautreaux Assisted Housing Program, HUD's Moving to Opportunity (MTO), and Seattle-King County's Creating Moves to Opportunity (CMTO) (Chyn, Collinson, and Sandler, 2025; Chetty, Hendren, and Katz, 2016; Bergman et al., 2024).

Where should parents move to raise their children's incomes? Quasi-experimental estimates of place effects, relating childhood exposures to outcomes in adulthood, may not offer a "plausible margin of choice": rural commuting zones (CZs) raise children's incomes but offer fewer jobs, and low rents suggest they are not in demand (Chetty and Hendren, 2018b). Why would poor places produce rich children? They may not: parents' moves affect their incomes, and estimates that control for parent incomes omit their effect on child outcomes by design (Wodtke, Harding, and Elwert, 2011). Given the importance of family income for children's outcomes (Loury, 1981; Becker and Tomes, 1986; M. E. Page, 2024), we ask what estimates would capture the full effect of family moves on child outcomes.

In this paper, we estimate the counterfactual effects of family moves on child outcomes, for each CZ in the United States. We allow parents' income to vary with exposures in the quasi-experimental movers design, by replacing observed parent family income ranks with latent ranks that remove place effects on income. Since existing estimates are too imprecise to compare CZs (Chetty and Hendren, 2018b), we control for self-selection on moves using a new feasible estimator equivalent to the infeasible high-dimensional fixed effects estimator, in tax data for 29 million children, and deliver estimates that permit informative inference (Mogstad et al., 2024; Andrews, Kitagawa, and McCloskey, 2024). We decompose place effects into effects on children through parent incomes, effects holding parent incomes constant, and sorting, to identify mechanisms for each place. Finally, we estimate parents' revealed preference for places based on parents' cross-CZ moves, and examine parents' valuations of child outcomes, rents, and incomes to understand margins of choice.

We find that counterfactual effects double the variance and change the geography of mobility relative to exposure effects: suburban CZs, smaller urban CZs, and resource-rich rural CZs deliver some of the best child outcomes, and the Midwest outperforms while the Southwest, South, Ozarks, and Appalachia underperform. The geography changes because indirect effects are determined by place effects on parent incomes, exposure effects are associated with rents, and these influences are approximately unrelated but coincide in suburban CZs surrounding major urban CZs. The geography of causal effects also differs from the geography of observed outcomes: relative to exposure effects, observed outcomes capture

positive selection into urban CZs and CZs in the Midwest, Great Plains, and Mountain states, and negative selection into CZs in the Southwest, South, Ozarks, and Appalachia, reinforcing the negative effects of places on incomes. Revealed preference values of places are negatively correlated with counterfactual effects, and nearly all of that relationship is explained by rents and incomes, with parents strongly valuing lower rents.

Our analysis rests on precise estimates: we reduce the noise in quasi-experimental estimates to less than one tenth of the variance. We describe the estimator before the estimates. The fundamental challenge is estimation in the high-dimensional causal model: because parents can move between any pair of CZs and each move is a choice that may be related to unobservables, the most credible estimates for CZs must control for millions of move-specific fixed effects. The size of the design matrix makes the regression computationally infeasible. Ad-hoc approaches to this problem generate considerable noise. We take a new approach. We break the estimation into a sparse residualization step for the move-specific fixed effects, and an estimation step for the exposure effects, using the Yule-Frisch-Waugh-Lovell theorem and recent results on inference to show equivalence to the efficient, infeasible estimator. In this setup, we compute estimates for an analysis population over 20 times larger than past studies, taking advantage of new data on childhood moves.

Our analysis also rests on unbiased estimates: we test for bias due to family-specific and child-specific factors, showing that few estimates are significantly different after including family fixed effects, and the variance of the bias term is two orders of magnitude smaller than the variance in place effects after including commonly suggested child-specific controls.

Third, we ask whether observed outcomes, which are currently used for most rankings and policy recommendations, are an adequate alternative to quasi-experimental estimates: we measure bias in observed outcomes using our estimates of exposure effects as a benchmark, and find that observed outcomes overstate exposure effects for major urban areas, the Upper Midwest, the Great Plains, and the Mountain states, while understating exposure effects for the South and Southwest. Thus, recommendations based on observed outcomes are distorted relative to causal estimates. Nonetheless, observed outcomes are more precise, and we ask whether efforts to predict and rank "good" CZs could be improved by relying in part on observed outcomes. Building MSE-minimizing shrinkage estimates of exposure effects as in Chetty and Hendren (2018b), we find that the estimates load almost entirely on our exposure effect estimates, rejecting prediction based on the biased observed outcomes. We find similar results for all CZ population sizes, suggesting that rankings based on quasi-experimental estimates are sufficient and preferable to rankings based on observed outcomes.

Fourth, we rank CZs based on causal place effects. We find that a majority of estimates are significant in comparisons with the national mean and with other CZs, at conventional

levels relevant for recommendations about when to move conditional on moving between any given pair of CZs. Corrections for multiple testing find that most comparisons of any meaningful magnitude are significant; this is especially true for comparisons representing the most likely relocations for families in our population. Even by Mogstad et al.'s (2024) stringent criterion holding the chance of *any* mis-ranking between any pair of CZs to less than 5 percent, it is possible with our estimates to limit the confidence intervals for most of the top 50 CZs to between 5 and 20 ranks. Given that most families will only ever rank one CZ (their origin) against any other CZs (their potential destinations), and given that the quasi-experiment speaks primarily to the benefits of earlier or later moves within pairs of CZs, we conclude that the estimates' precision is more than adequate to guide the intended recommendations.

Having answered that usable place effects can be estimated, we then ask which place effects are the right ones to estimate. Recommendations urging families to move assume that place effects identify valid counterfactuals. The only causal place effects in the literature to date are exposure effects. Counterfactual outcomes include the effects of places on parent incomes, so that exposure effects define valid counterfactuals only when places do not affect parent incomes or when parent incomes do not affect child outcomes. Neither assumption is true. How can we obtain valid counterfactuals? We define counterfactual effects and show that they can be estimated in Chetty and Hendren's (2018b) quasi-experimental framework by replacing actual parent income ranks with latent parent income ranks that remove place effects on income. We estimate place effects on parent income ranks, recover latent income ranks for every family, and estimate counterfactual effects on child outcomes. Counterfactual effects capture first-order indirect effects of places on child outcomes through parent incomes, doubling the variance of place effects. They imply a starkly different ranking, that is much-better aligned with common intuitions about geographic patterns of opportunity: notably, major urban CZs and their suburban periphery have better outcomes than suggested by exposure effects, while rural CZs have worse outcomes. Exposure effects or observed outcomes capture half or less of the outcome variation identified by counterfactual effects.

High counterfactual effects may be caused by low rents and high incomes. They capture both exposure effects and indirect effects through parent incomes, and exposure effects are strongly correlated with rents, while indirect effects are almost-entirely explained by place effects on parent incomes. These channels are approximately independent, and both point toward family resources as a determining factor for child outcomes.

Finally, we ask how to make recommendations persuasive. Parents do not choose locations based on child outcomes alone. They also weigh their own budgets and preferences.

To examine parents' budget constraints, we extend the notion of within-CZ opportunity

bargains in Chetty and Hendren (2018b) to identify bargains for cross-CZ movers, based on counterfactual effects, rents, and our estimates of place effects on parents' family incomes. Because counterfactual effects are determined by low rents and high incomes, opportunity bargains are common; they are most-often found in suburban CZs surrounding major urban CZs, or in resource-rich rural CZs. Other characteristics parents may favor, like commute times and social capital, reinforce opportunity bargains.

To examine parents' preferences, we estimate place values based on the flows of families moving from one CZ to another, and a modification of Google's PageRank (L. Page et al., 1999) algorithm proposed by Sorkin (2018), to iteratively identify "good" CZs by inflows from other "good" CZs in a form of weighted voting. These revealed preference values uncover a fundamental problem for policies encouraging moves toward CZs with positive counterfactual effects for children: counterfactual effects are negatively correlated with parents' preferences. For example, the most desirable CZ, New York, has among the most negative counterfactual effects. The negative relationship between place values and counterfactual effects may arise not because parents dislike good outcomes for their children, but because good outcomes are attached to other undesirable characteristics. Regressing place values on counterfactual effects, rents, wages, and incomes, we find that prices explain more than nine-tenths of the variance in place values due to counterfactual effects. Thus, parents are indifferent towards child outcomes, even as they conditionally value high wages, low incomes, and low rents. Opportunity bargains based on low rents may appeal to parents, but otherwise these results suggest that persuading parents to move to better places will be a challenge.

Our revealed preference findings, alongside our finding that family resources determine child outcomes, suggest giving greater attention to place-based or family-based policies raising income or lowering prices.

2 Data

We use data from full-population federal income tax records, covering the complete childhood location histories of 77 million children. This section presents population and variable definitions, followed by summary statistics.

2.1 Population Definitions

Our base population consists of all children born from 1978-1998.

To rank children by income in childhood, we restrict the population to children with (1) a

valid Social Security Number (SSN) issued by age 18 according to the sequence of issuance,[1] and (2) no record of non-citizen status.[2] These conditions ensure that children were citizens at birth, with parents subject to taxation as citizens or residents.

To rank children by income in adulthood, we exclude children deceased as of 2025, and children not observed by the IRS from 2022-25.[3]

Parent and Child Linkages.—We define parents as the first tax filers to claim each child as a dependent between 1994 and 2016, such that female filers and single male filers are aged 15-40 at the child's birth. Relative to Chetty and Hendren (2018b) who study the 7 cohorts born from 1980-86, we study the 21 cohorts born from 1978-98, using recovered tax returns from 1994-95 (Alexander et al., 2025) to extend coverage to the birth cohorts of 1978-79.[4] We link parents to 94 percent of children claimed through age 18.

Siblings.—We define siblings as children linked to the same tax filers. Studying a longer panel of birth cohorts increases the number of linked siblings more than proportionately, since nearly all siblings are found in different birth cohorts.

Stayers and Movers.—We define stayers as children whose parents lived one CZ, and movers as children whose parents lived in more than one CZ, through age 23.

We implement our definitions of stayers and movers using recovered tax returns from 1974-95 (Alexander et al., 2025) that allow us to observe parents' residence and identify moves from birth for nearly all children. We exclude the 0.02 percent of children who have no valid residence. Measuring early moves identifies a larger population of movers, increases the variance of place exposures, and reduces attenuation bias in exposure estimates.

[1]Individuals with Individual Taxpayer Identification Numbers (ITINs) are non-citizens. SSNs may be issued to foreign-born citizens and non-citizens, though non-citizen status is recorded and rarely updated. Foreign-born non-citizens are not eligible to naturalize and obtain citizenship until they reach age 18, making it unlikely that a citizen-status SSN would be issued under age 18. An exception under the Child Citizenship Act of 2000 has been available since February 27, 2001, affecting children born after the same date in 1983, for foreign-born children adopted by citizen parents or living with naturalized parents. Our age cutoff does not exclude natives: Klopfer and Miller (2024, Figure 4) show that native-born children from the 1978 birth cohort only rarely obtain an SSN after age 10, from 1983 after age 6, and from 1991 after age 1.

[2]Current citizenship status is recorded in the IRS DM-1 file and updates since 2014 are recorded in the DM-1-SSN file. Citizenship was not recorded in the DM-1 file until mid-1981, and a small number of children born and issued SSNs before mid-1981 do not have a recorded citizenship status. We include these SSNs, because most are held by citizens, and native birth is more likely the earlier the SSN is issued.

[3]Individuals are observed by the IRS if they are reported on any tax return (Form 1040) as a filer, spouse, or dependent; if they are reported on any information return (Form W-2, Form 1099, and others) or health coverage return (Form 1095); or if they are the subject of reporting from state and federal prison systems. Individuals not observed by the IRS are most often deceased or non-resident non-citizens.

[4]We do not use tax returns prior to 1994 for parent and child linkages because the 1969, 1974, 1979, 1984, and 1989 tax return files do not include dependent SSNs.

2.2 Variable Definitions

We define parent and child variables in the tax data, starting with parent geography, nativity, and income ranks before discussing child income ranks. All monetary values are reported in 2012 dollars.

Parent Geography.—In each year, we define parent ZIP codes using the first valid ZIP code found on the mother's tax return, the father's tax return, the plurality of the mother's information returns, or the plurality of the father's information returns. We translate ZIP codes to 1990 CZs using a crosswalk covering ZIP codes valid in 1974-2022, adding coverage for Puerto Rico using 2020 CZs and for overseas territories by postal abbreviation.

When defining moves, ages at moves, and years of exposure to each CZ, we treat the CZ observed closest to birth (or the later if two are equally close) as the first location. Each time a child is observed in a different CZ, we treat that CZ as the next location. To account for missing tax return data between 1974, 1979, 1984, 1989, and 1994, and for non-filing, we define the year of each move as the midpoint between the first year observed at the new location and the last year observed at the last location, rounded up. We define the age at move as the year minus the child's birth year, and consider moves from age 1 to 23. Because tax returns and information returns are filed after the end of the corresponding calendar year, we calculate exposures for the first location from age 0 to the first move, for each following location from the age a child moved in and out, and for the final location from the last move to age 24. For example, consider a child born in 1995 whose parents file 1995 taxes from one CZ, and file 1996-2019 taxes from another CZ: we would code that child as having lived in the first CZ for 1 year, and the second CZ for 23 years.[5]

Parent Family Income and Percentile Rank.—We follow Chetty et al. (2014) and subsequent papers in defining each parent's family income as the total pre-tax income by tax unit. In years when a parent files a tax return (Form 1040) as a primary or secondary filer, we define that parent's family income as the sum of Adjusted Gross Income, tax-exempt interest, and non-taxable Social Security and Disability (SSDI) benefits reported on the tax return. When no parent files, we define each parent's family income as the sum of wage earnings (from Form W-2), unemployment benefits (from Form 1099-G), and taxable and nontaxable Social Security and Disability benefits (from Form SSA-1099) for both parents.

[5]To see why this simple calculation is roughly correct, suppose the child was born on July 1, 1995, and his parents filed taxes on April 15, 1996 and moved on July 1, 1996: in that case, the child would have lived six months of 1995 in the first CZ by the time taxes are filed, and another six months of 1996 in the first CZ before moving to the second; given that the median child is born in the summer, most filers file on tax day, and most children move during the summer break, this is a fair approximation. Further, if the parents always file and move later, the child will spend at most a few months more in the first CZ, and the same few months less in the last CZ, but for intermediate moves, the errors balance.

When no parent is observed, we define family income as zero.[6]

To measure parent family income with children as the unit of observation, we average family income across both parents regardless of their marital status. For parents who file together, average family income equals their individual family incomes obtained from their joint tax return. If the parents do not file together, average family income equals the average of their new tax units' family incomes, and may be reduced by one or both of the parents being in a single-income tax unit.

To standardize the measurement of parents' income across cohorts, and to reduce the influence of annual variation in parents' family income (Mazumder, 2005), we average parent family income across all years in which the child is aged 0-23; we over-weight income from tax returns in 1979, 1984, 1989, and 1994 to account for tax years 1978-81, 1982-86, 1987-91, and 1992-94. This definition covers almost 5 times the age span available to Chetty and Hendren (2018b), and better measures family resources during early and middle childhood.

We exclude the 0.27 percent of children whose parents have zero or negative family income after averaging, since parents filing with zero or negative net income generally have positive gross income and would be mis-ranked.

We define parent family income ranks within each full birth cohort. This ensures that ranks do not vary systematically with different years of tax records for different cohorts.

Child Family Income and Percentile Rank.—We define child family income identically to parent family income. Unlike Chetty et al. (2014) and related papers, we average child family income across all available ages from 24 onward, with the final year of income measurement in 2022. We do this to reduce the influence of annual variation in children's family income: for example, for children born in 1978, we aggregate 21 years of child family income from 2002-2022 in place of the single year ranked by Chetty and Hendren (2018b). We allow for zero or negative incomes. However, zero incomes are uncommon due to the requirements that children be alive, present in the tax data (which include information returns in all outcome years), and citizens (and therefore taxed on income globally).

We define child family income ranks within each full birth cohort. Similarly to the ranking for parent family income, this ensures that ranks do not vary systematically with different age ranges and different years of tax records for different cohorts.

[6]Chetty et al. (2014) note that parents who are not observed have no taxable income. However, non-resident non-citizen parents may have no taxable income but significant non-taxable income (Chetty et al., 2014) which would distort parent ranks for immigrant children. We focus our analysis on children who were citizens at birth to ensure that parents were subject to taxation as citizens or residents.

2.3 Summary Statistics

Table 1 reports summary statistics for the full population. Panel A reports statistics for the full population, Panel B reports statistics for stayers, and Panel C reports statistics for movers who moved up to 6 times (97 percent of all movers).

The primary population for analysis is the movers population: the total count is 29.1 million, 21 times larger than the population in Chetty and Hendren (2018b, Table 1). Larger populations improve the precision of movers estimates.

3 Methods

We begin by defining our estimands, exposure effects, for each CZ. We next define our estimators, and the assumptions required to identify exposure effects. For each estimator, we discuss estimation in large populations with many nuisance parameters. Finally, we discuss variance-covariance estimators, assumptions, and estimation.

3.1 Estimands

We seek to understand how childhood exposures to locations $c \in C$ at ages $a \in A$, captured by location histories $\boldsymbol{c}_i = (c_{ia} : a \in A)'$, affect children's outcomes in adulthood. Specifically, we seek to estimate parameter vector $\boldsymbol{\mu}$ in the statistical model

$$y_i = (\boldsymbol{\mu}^0 + \boldsymbol{\mu}^1 p_i) \cdot \boldsymbol{e}_i + (\boldsymbol{\kappa}^0 + \boldsymbol{\kappa}^1 p_i) \cdot \boldsymbol{s}_i + \theta_i \tag{1}$$

where y_i is a child's family income rank or other outcome; p_i is the parents' family income rank; $\boldsymbol{\mu} = (\boldsymbol{\mu}^{0\prime}, \boldsymbol{\mu}^{1\prime})'$ summarizes causal effects of location exposures $\boldsymbol{e}_i = (\sum_A 1_c(c_{ia}) : c \in C)'$; $\boldsymbol{\kappa} = (\boldsymbol{\kappa}^{0\prime}, \boldsymbol{\kappa}^{1\prime})'$ summarizes causal effects of location sequences $\boldsymbol{s}_i = (1_{\boldsymbol{d}}(\boldsymbol{d}_i) : \boldsymbol{d} \in \boldsymbol{D})'$ where we define $\boldsymbol{d}_i = (c_{ia} : c_{ia} \neq c_{ia-1}, a \in A)'$ as the vector composed of the origin and each move's destination; and θ_i summarizes other family- and child-specific factors.

This simplified model restricts location-effect heterogeneity. First, we assume exposure and sequence effects are additively separable, ruling out order effects, complementarity, and substitutability. Second, we assume linear interactions between parental income ranks and exposure and sequence effects. Third, we assume constant (and therefore age-independent) exposure effects during childhood and zero exposure effects otherwise.[7] Fourth, we assume constant sequence effects during childhood and zero sequence effects otherwise.

[7] We rely on past validations of the second and third assumptions. Chetty et al. (2014) show linearity in parental income ranks. Chetty and Hendren (2018a) and Chetty et al. (2026a) show constant exposure effects through early adulthood and zero effects thereafter.

3.2 Estimators and Assumptions

Ideally, we would estimate location effects by randomly assigning families to location histories, as in the Moving to Opportunity experiment (see Chetty, Hendren, and Katz, 2016). However, our goal is to estimate effects for all potential moves and family income ranks, which is more readily done by observing the population of voluntary movers in tax data.

To estimate exposure effects $\boldsymbol{\mu}$, we rely on the timing of parents' moves across locations to obtain variation in children's location exposures, following Chetty and Hendren (2018b). Consider two children who have different location exposures but the same family income rank, location sequence, and other family inputs. In this case, we can isolate the exposure effects at the children's shared rank p; if we observe comparable children at multiple ranks, then we can also distinguish the intercept and slope of the exposure effects. For example, suppose the children move from New York to Los Angeles, at different ages. The child who moves at a younger age has a longer exposure in Los Angeles, and a shorter exposure in New York, so comparing outcomes identifies the rank-specific exposure effect

$$(y_i - y_j)/(e_{i,NY} - e_{j,NY}) = (\mu^0_{NY} + p\mu^1_{NY})$$

for New York relative to Los Angeles, normalizing the exposure effect of Los Angeles to zero. That is, we can identify exposure effects if we can make other family inputs comparable.

Estimator A1. A natural estimator of $\boldsymbol{\mu}$ is the sequence fixed effects specification

$$y_i = (\boldsymbol{\mu}^0 + \boldsymbol{\mu}^1 p_i) \cdot \boldsymbol{e}_i + (\boldsymbol{\alpha}^0 + \boldsymbol{\alpha}^1 p_i) \cdot \boldsymbol{s}_i + \epsilon_i \tag{2}$$

where $(\boldsymbol{\alpha}^0 + \boldsymbol{\alpha}^1 p_i) \cdot \boldsymbol{s}_i = (\boldsymbol{\kappa}^0 + \boldsymbol{\kappa}^1 p_i) \cdot \boldsymbol{s}_i + \bar{\theta}_{p_i \boldsymbol{s}_i}$ incorporates both causal effects $\boldsymbol{\kappa}$ and selection effects $\bar{\theta}_{p_i \boldsymbol{s}_i} = \mathrm{E}[\theta_i | \boldsymbol{s}_i, p_i \boldsymbol{s}_i]$, and ϵ_i captures residual variation in other family inputs. This estimator absorbs sequence effects, and identifies exposure effects from variation due to the timing of moves; similar to the estimator in Aloni and Avivi (2025), it extends the estimator in Chetty and Hendren (2018b) by allowing for multiple-move sequences.

To identify $\boldsymbol{\mu}$, the error must have a mean of zero with respect to regressors $\boldsymbol{e}_i$ and $p_i \boldsymbol{e}_i$, after absorbing the fixed effects:

Assumption A1. Conditional on location sequences and their interactions with parent family income rank $\boldsymbol{s}_i$ and $p_i \boldsymbol{s}_i$, exposures and their interactions with parent family income rank $\boldsymbol{e}_i$ and $p_i \boldsymbol{e}_i$ are jointly orthogonal to residual variation in childrens' outcomes

$$\mathrm{E}[\epsilon_i | \boldsymbol{e}_i, p_i \boldsymbol{e}_i, \boldsymbol{s}_i, p_i \boldsymbol{s}_i] = \mathrm{E}[\epsilon_i | \boldsymbol{s}_i, p_i \boldsymbol{s}_i]$$

Intuitively, Assumption A1 states that exposures cannot be related to outcomes through selection, after accounting for children's location sequences. Assumption A1 also states that the interaction of parent ranks and exposures cannot be related to outcomes through selection, after accounting for the interaction of parent ranks and children's location sequences. For example, conditional on moving to Los Angeles from New York, families should not move later when their unobservables are better. Further, it shouldn't be the case that higher-rank parents move later and lower-rank parents move earlier when their unobservables are better. That is, there should be no selection on average or by parent ranks.

We test Assumption A1 in Section 4.3, extending and formalizing the specification tests in Chetty and Hendren (2018a and 2018b) to find stronger evidence against either systematic or location-specific bias in estimates from Equation A1.

3.3 Estimation

We exploit the Yule-Frisch-Waugh-Lovell theorem to jointly estimate the 1,522 CZ intercept and slope parameters $\boldsymbol{\mu}$[8] as

$$\hat{\boldsymbol{\mu}} = (\tilde{\boldsymbol{X}}'\tilde{\boldsymbol{X}})^{-1}\tilde{\boldsymbol{X}}'\tilde{\boldsymbol{Y}} \tag{3}$$

after we obtain the residuals of childrens' outcomes and location exposures, y_i and $\boldsymbol{x}_i = (\boldsymbol{e}_i', p_i\boldsymbol{e}_i')'$ on location sequences $\boldsymbol{z}_i = (\boldsymbol{s}_i', p_i\boldsymbol{s}_i')'$.

The key technical challenge for estimation is how to obtain the residuals. Direct residualization on $\boldsymbol{Z}$ is infeasible: storing and inverting $\boldsymbol{Z}'\boldsymbol{Z}$ would require tens of terabytes of memory. Iterative residualization is inefficient, requiring days of server time and weeks for specifications that include additional slope terms for each location sequence.[9]

We develop an efficient residualization method, relying on the sparsity of $\boldsymbol{Z}$ and $\boldsymbol{X}$: first, $\boldsymbol{z}_{i\boldsymbol{d}} = (1_{\boldsymbol{d}}(\boldsymbol{d}_i), p_i 1_{\boldsymbol{d}}(\boldsymbol{d}_i))' = \boldsymbol{0}$ for $\boldsymbol{d} \neq \boldsymbol{d}_i$; second, $\boldsymbol{x}_{ic} = (e_{ic}, p_i e_{ic})' = \boldsymbol{0}$ for $c \notin \boldsymbol{d}_i$. The

[8] We observe exposures for 761 CZs: the 741 1990 CZs covering the states and the District of Columbia, 10 2020 CZs covering Puerto Rico, and 10 pseudo-CZs defined by postal abbreviation for overseas territories. We normalize $\mu^0_{LA} = \mu^1_{LA} = 0$ due to collinearity, treating Los Angeles as the base category.

Joint estimation of both $\boldsymbol{\mu}$ and $\boldsymbol{\kappa}$ by OLS is infeasible, with millions of intercept and slope parameters across all location sequences. The approach in Chetty and Hendren (2018b), two-step estimation of $\boldsymbol{\mu}$ by OLS within location sequences, is feasible but inefficient: although their second step constrains location exposure parameters to be equal across all location sequences, their first step allows for heterogeneity in exposure parameters across location sequences; consequently, they are forced to estimate orders of magnitude more parameters than necessary, and to discard many observations for which the parameters are not identified.

[9] In alternative specifications, we have also included interactions of location sequences with birth cohorts and squared birth cohorts, as in Chetty and Hendren (2018b). However, the primary reason to include cohort interactions is to address differential measurement error and selection across cohorts (Chetty and Hendren, 2018a, Online Appendix A). Doing so is not necessary when we measure location histories from birth and, because it discards up to a tenth of the population due to small cell sizes for rare location sequences, doing so is inefficient. Further, the correlation between estimates is $r = 0.99$.

sparsity of $\boldsymbol{Z}$ allows us to residualize $\boldsymbol{Y}$ and $\boldsymbol{X}$ separately for each sub-population

$$\tilde{\boldsymbol{Y}} = (\tilde{\boldsymbol{Y}}_{\boldsymbol{d}}' : \boldsymbol{d} \in \boldsymbol{D})'$$
$$\tilde{\boldsymbol{X}} = (\tilde{\boldsymbol{X}}_{\boldsymbol{d}}' : \boldsymbol{d} \in \boldsymbol{D})'$$

where we compute residuals for each sub-population, by multiplication with the annihilator matrix $\boldsymbol{M_{dd}} = \boldsymbol{I} - \boldsymbol{Z_{dd}}(\boldsymbol{Z_{dd}}'\boldsymbol{Z_{dd}})^{-1}\boldsymbol{Z_{dd}}'$, as $\tilde{\boldsymbol{Y}}_{\boldsymbol{d}} = \boldsymbol{M_{dd}}\boldsymbol{Y_d}$ and $\tilde{\boldsymbol{X}}_{\boldsymbol{d}} = \boldsymbol{M_{dd}}\boldsymbol{X_d}$. The sparsity of $\boldsymbol{X}$ allows us to compute $\tilde{\boldsymbol{X}}_{\boldsymbol{d}}$ as $\tilde{\boldsymbol{X}}_{\boldsymbol{dc}} = \boldsymbol{M_{dd}}\boldsymbol{X_{dc}}$ for $c \in \boldsymbol{d}_i$ and $\tilde{\boldsymbol{X}}_{\boldsymbol{dc}} = \boldsymbol{0}$ for $c \notin \boldsymbol{d}_i$. Our method, implemented in Stata as `residby`, runs 3-4 orders of magnitude faster than the current benchmark absorbing categorical-by-continuous interactions, `hdfe` (Correia, 2016);[10] the performance gap is larger for alternative specifications with additional slope interactions. Since many problems share this structure, our method may be of wider interest.

3.4 Variance and Covariance

Given OLS estimates $\hat{\boldsymbol{\mu}}$, we estimate the variance-covariance matrix (VCE) to rank locations by their exposure effects.

We implement the VCE estimator of Abadie et al. (2020), following their design-based approach to inference as suggested by Mogstad et al. (2024). The design-based approach is motivated by the fact that there is no sampling uncertainty in population data. The source of uncertainty is, instead, that exposures could have taken different values if children had moved at different ages. In the nomenclature of Abadie et al. (2020), we treat location exposures as causal and location sequences as descriptive.

The Abadie et al. (2020) (AAIW) VCE estimator is

$$\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}}^{AAIW} = (\tilde{\boldsymbol{X}}'\tilde{\boldsymbol{X}})^{-1}\tilde{\boldsymbol{U}}_i'\tilde{\boldsymbol{U}}_i(\tilde{\boldsymbol{X}}'\tilde{\boldsymbol{X}})^{-1} \tag{4}$$

where $\boldsymbol{u}_i = \tilde{\boldsymbol{x}}_i\hat{\epsilon}_i$, and $\hat{\epsilon}_i = \tilde{y}_i - \tilde{\boldsymbol{x}}_i'\hat{\boldsymbol{\mu}}$. We residualize $\boldsymbol{U}$ on $\boldsymbol{Z}$ separately for each sub-population

$$\tilde{\boldsymbol{U}} = (\tilde{\boldsymbol{U}}_{\boldsymbol{d}}' : \boldsymbol{d} \in \boldsymbol{D})'$$

and as before compute $\tilde{\boldsymbol{U}}_{\boldsymbol{d}}$ as $\tilde{\boldsymbol{U}}_{\boldsymbol{dc}} = \boldsymbol{M_{dd}}\boldsymbol{U_{dc}}$ for $c \in \boldsymbol{d}_i$ and $\tilde{\boldsymbol{U}}_{\boldsymbol{dc}} = \boldsymbol{0}$ for $c \notin \boldsymbol{d}_i$. Intuitively, the AAIW VCE estimator removes "uncertainty" due to estimated exposure parameter heterogeneity across location sequences.[11]

[10] We obtain these results in part by performing only the necessary computations, and in part by coding the computations efficiently: we build on the compiled-loop implementation of `regressby` (Stepner, Droste, and Townsend, 2018; Cáceres Bravo, 2018), which we modify to parallelize residualization.

[11] For comparison, the Eicker-Huber-White VCE estimator is $\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}}^{EHW} = (\tilde{\boldsymbol{X}}'\tilde{\boldsymbol{X}})^{-1}\boldsymbol{U}'\boldsymbol{U}(\tilde{\boldsymbol{X}}'\tilde{\boldsymbol{X}})^{-1}$ where $\boldsymbol{U}$

To enable secondary analyses, we will release AAIW and EHW VCE estimates: as Chetty and Hendren (2018b), Mogstad et al. (2024), and Kline (2025) observe, covariance terms are generically nonzero, so that the VCE cannot be summarized by standard errors.

4 Exposure Effects

This paper's primary objective is to estimate location exposure effects that can be relied on as an input to decisions by families and policymakers. This section reports estimated exposure effects, assesses the degree to which they reflect signal or noise, and tests the identification assumption required to interpret them as causal.

4.1 Baseline Estimates

What is the impact of location exposures on children's family income ranks in adulthood, given their parents' family income ranks?

We measure CZ exposure effects relative to the national mean, and their uncertainty, by combining estimates from equations (3) and (4),

$$\hat{\boldsymbol{\mu}}_p = \boldsymbol{Q}\boldsymbol{A}_p\hat{\boldsymbol{\mu}} \tag{5}$$

$$\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}_p} = \boldsymbol{Q}\boldsymbol{A}_p\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}}\boldsymbol{A}_p'\boldsymbol{Q}' \tag{6}$$

where $\boldsymbol{A}_p = [\boldsymbol{I}\ p\boldsymbol{I}]$ adds intercepts to slopes multiplied by parent ranks and $\boldsymbol{Q} = \boldsymbol{I} - \mathbf{1}\boldsymbol{\pi}'$ subtracts the exposure-weighted mean based on exposure shares $\boldsymbol{\pi} = \bar{\boldsymbol{e}}/\sum_C \bar{\boldsymbol{e}}_c$.

For example, children with parents at the 25th percentile of parent family income, exposed to Los Angeles for an additional year of childhood relative to the national mean, have child family income $\hat{\mu}_{LA25} = -0.0688$ (0.0043) percentiles lower. At the 75th percentile of parent family income, exposure to Los Angeles for an additional year results in child family income $\hat{\mu}_{LA75} = -0.1689$ (0.0050) percentiles lower. To obtain the full effects of location exposures, we scale effects by 24 years of exposure, as in the estimation.[12] For example, at the 25th percentile, a full childhood of exposure to Los Angeles lowers children's family incomes by $\hat{\mu}_{LA25} \times 24 = -1.650$ (0.103) percentiles; at the 75th percentile, full exposure lowers children's family incomes by $\hat{\mu}_{LA75} \times 24 = -4.053$ (0.120) percentiles.

has not been residualized. In practice, the VCE estimates are similar: AAIW standard errors for exposure effects are roughly five percent smaller than EHW standard errors. These findings support the assumption in Section 3.1 that exposure and sequence effects are additively separable.

[12] We measure childhood locations from age 0 through 23, for a total of 24 years of exposure, based on the evidence in Chetty and Hendren (2018a) and Chetty et al. (2026a) that exposure effects are approximately linear from age 3 through 23.

Figure 1 and 2 report exposure effects $\hat{\mu}_{c25}$ and $\hat{\mu}_{c75}$, and their confidence intervals, for the 50 largest CZs and for all CZs. For comparison, we also display the exposure effects and confidence intervals reported by Chetty and Hendren (2018b), which cover 166 fewer CZs out of 761, and are substantially more variable and less certain.

How large are exposure effects relative to the national mean? One relevant comparison is with parent income effects on child outcomes. We equate location exposure effects and parent income effects using the slope of child income ranks relative to parent income ranks among stayers, 0.3785. Across all CZs the standard deviation of full exposure effects weighted by population, $0.1051 \times 24 = 2.522$ for the 25th percentile of parent income or $0.1136 \times 24 = 2.727$ for the 75th percentile of parent income,[13] is equivalent to having parents $2.522/0.3785 = 6.663$ or $2.727/0.3785 = 7.204$ percentiles higher or lower in the income distribution. This evidence suggests that exposure effects are important for children's outcomes.

How significant are exposure effects relative to the national mean? Among the 50 largest CZs, at the 25th and 75th percentiles of parent income, 38/50 and 34/50 estimated exposure effects have a t-statistic greater than 1.960; accounting for multiple testing with a (conservative) Bonferroni correction, 33/50 and 23/50 have a t-statistic greater than 3.292. Among all CZs, at the 25th and 75th percentiles of parent income, 445/761 and 409/761 have a t-statistic greater than 1.960; accounting for multiple testing with a Bonferroni correction, 193/761 and 132/761 have a t-statistic greater than 3.993.[14] Since many estimates are close to zero, not all estimates will be significant in tests against the national mean.

Figure 3 and 4 map the geography of exposure effects $\hat{\mu}_{c25}$ and $\hat{\mu}_{c75}$. Intergenerational mobility is higher for children who have more exposure to CZs in the Great Plains, the Ozarks, Appalachia, and the Mountain West; and less exposure to CZs in the South Atlantic, Gulf, and Pacific Coasts, and the Central Great Lakes. Major urban areas, notably Los Angeles, New York, and Chicago, under-perform the national mean and most rural areas. Geographic patterns are similar for the 25th and 75th percentile of parent incomes. Estimates supporting these observations are generally significant after Bonferroni correction.

4.2 Validating the Estimates: Signal and Noise

Imprecision is a source of risk in decisions based on estimates. Chetty and Hendren (2018b, Section 6) raise concerns about the precision of exposure effect estimates, echoed in Mogstad

[13] As we discuss in Section 4.2, the standard deviation of the estimated exposure effects is within 5 percent of the standard deviation of the true exposure effects.

[14] For 7 CZs at the 25th percentile of parent income, and 8 CZs at the 75th percentile of parent income, we treat estimates as insignificant regardless of their t-statistics because their confidence intervals for the effect of full exposure include values more than 50 percentile ranks away from the national mean.

et al. (2024) and Andrews, Kitagawa, and McCloskey (2024). We present a formal analysis showing that our methods directly address these precision concerns.

To understand the variation of the exposure effect estimates around the national mean, we would like to partition the total variance $\sigma^2_{\hat{\boldsymbol{\mu}}_p}$ into signal $\sigma^2_{\boldsymbol{\mu}_p}$ (true variance) and noise $\sigma^2_{\hat{\boldsymbol{\mu}}_p-\boldsymbol{\mu}_p}$ (variance due to estimation), with

$$\sigma^2_{\hat{\boldsymbol{\mu}}_p} = \sigma^2_{\boldsymbol{\mu}_p} + \sigma^2_{\hat{\boldsymbol{\mu}}_p-\boldsymbol{\mu}_p}$$

We estimate these unobservable variances from demeaned estimates $\hat{\boldsymbol{\mu}}_p$ and their variances $\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}_p}$ from equations (5) and (6), as

$$\hat{\sigma}^2_{\hat{\boldsymbol{\mu}}_p} = \hat{\boldsymbol{\mu}}_p'\boldsymbol{\Pi}\hat{\boldsymbol{\mu}}_p \tag{7}$$

$$\hat{\sigma}^2_{\hat{\boldsymbol{\mu}}_p-\boldsymbol{\mu}_p} = \mathrm{tr}(\boldsymbol{\Pi}\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}_p}) \tag{8}$$

$$\hat{\sigma}^2_{\boldsymbol{\mu}_p} = \hat{\sigma}^2_{\hat{\boldsymbol{\mu}}_p} - \hat{\sigma}^2_{\hat{\boldsymbol{\mu}}_p-\boldsymbol{\mu}_p} \tag{9}$$

where we weight the estimates and their variances by exposure shares $\boldsymbol{\Pi} = \mathrm{diag}(\boldsymbol{\pi})$.[15] One way to summarize precision based on these estimates is the signal to noise ratio, $\hat{\sigma}^2_{\boldsymbol{\mu}_p}/\hat{\sigma}^2_{\hat{\boldsymbol{\mu}}_p-\boldsymbol{\mu}_p}$.

Table 2 reports the variance components. We report variances across all CZs, for the 25th and 75th percentiles of parent family income, in Panel A, Columns 1 and 2. The signal variances for exposure effects at the 25th and 75th percentiles are 0.0103 and 0.0117, or 94 and 91 percent of the total variances; the signal standard deviations are 0.101 and 0.108 percentile ranks per year of exposure; or 97 and 95 percent of the total standard deviations. In other words, the noise variances of 0.0007 and 0.0012 account for only 6 and 9 percent of the total variance, and the noise standard deviations of 0.0265 and 0.0346 are 25 and 30 percent of the total standard deviations. Summarizing the variance components, we find signal to noise ratios of 13.93 and 9.41 at the 25th and 75th percentiles, both excellent.

We build on the contribution of Chetty and Hendren (2018b) by improving the precision of their estimates. In Table 2, Panel B, Columns 1 and 2, we report estimates using the replication data from Chetty and Hendren (2018b), weighting by population for comparability because this is the weight they used to demean their estimates. This analysis reveals two new findings. First, compared to their estimated noise variances of 0.0517 and 0.0793, our estimates reduce noise by a factor of 74 and 66, respectively. Second, compared to their estimated signal variances of 0.0263 and 0.0196, our estimates suggest that the true signal variances are approximately 61 and 40 percent smaller. Re-estimating the noise in the orig-

[15] By a standard result for quadratic forms, if $\mathrm{E}[\hat{\boldsymbol{\mu}}_p] = \boldsymbol{\mu}_p$ and $\mathrm{Var}[\hat{\boldsymbol{\mu}}_p] = \boldsymbol{V}_{\hat{\boldsymbol{\mu}}_p}$, then $\mathrm{E}[\hat{\boldsymbol{\mu}}_p'\boldsymbol{\Pi}\hat{\boldsymbol{\mu}}_p] = \boldsymbol{\mu}_p'\boldsymbol{\Pi}\boldsymbol{\mu}_p + \mathrm{tr}(\boldsymbol{\Pi}\boldsymbol{V}_{\hat{\boldsymbol{\mu}}_p})$. With weights $\boldsymbol{\pi}$ and $\hat{\boldsymbol{\mu}}_p'\boldsymbol{\pi} = 0$, the quadratic forms are variances.

inal estimates by subtracting our signal variance estimates from the original total variance estimates, we find that the original signal to noise ratios may have been 0.15 and 0.13 or lower, much lower than the ratios of 0.51 and 0.25 implied by the original variance estimates. Taken together, these findings suggest signal to noise ratios 93 and 72 times greater for our estimates than for the original estimates, or two orders of magnitude greater precision.

Are exposure effect estimates more informative for larger CZs? Table 2, Columns 3 to 10 report variance components by CZ population, starting with the first 50, next 150, next 395, and final 166 CZs by population. The first three categories cover all CZs considered by Chetty and Hendren (2018b), and the fourth category covers all CZs added in our analysis. We find that estimates are most informative for the 50-largest CZs, and among high-income families in particular, with signal to noise ratios of 67 and 100 at the 25th and 75th percentiles, compared with 1.27 and 2.55 for the original estimates. For the remaining categories, signal to noise ratios are flat at 11 for the 25th percentile, while signal to noise ratios decline gradually at the 75th percentile, from 9 for CZs 51-200, to 7 for CZs 201-595, to 4 for CZs 596-761. This pattern, of greater signal at the 75th percentile for CZs 1-50, and lower signal at the 75th percentile for all remaining CZs, is consistent with the concentration of 75th percentile families in major urban CZs, and fewer high-income families in smaller CZs.

The signal to noise ratio for every population category of our estimates is substantially greater than the signal to noise ratio even for the 50-largest CZs in the original estimates. Comparing noise variance estimates directly, it is only in our smallest, newly-added CZs that the noise variance exceeds the noise variance of the original estimates for the 50-largest CZs. Thus, estimates for most CZs appear to be highly informative.

Overall, our new estimates overcome the primary limitation of the estimates in Chetty and Hendren (2018b), precision, and provide an empirical answer to the critiques in Mogstad et al. (2024) and Andrews, Kitagawa, and McCloskey (2024).

4.3 Validating the Estimates: Identification

Bias is another source of risk in decisions based on estimates. The estimates $\hat{\boldsymbol{\mu}}_p$ from equation (5) are unbiased for the true causal parameters $\boldsymbol{\mu}_p$ if Assumption A1 is valid. The assumption is valid if family- and child-specific factors θ_i are orthogonal to *all* location exposures at *all* income levels $\boldsymbol{e}_i$ and $p_i\boldsymbol{e}_i$ after accounting for location sequences and their interactions with parent family income. That is, parents should not choose the duration of a given exposure based on unobservables, conditional on having a particular sequence of exposures.

We validate the estimated exposure effects as unbiased *for each location*, first by testing for bias due to family-level factors, and second by testing for bias due to child-level factors.

4.3.1 Family-Level Factors

Because families move together, selective exposures are less of a concern for within-family comparisons than for between-family comparisons. Based on the observation that older and younger siblings have different exposures for the same family moves, Chetty and Hendren (2018a), Deutscher (2020), Alesina et al. (2021), Laliberté (2021), Britto et al. (2022), Beekers (2024), Aloni and Avivi (2025), and Chetty et al. (2026a) test for bias *on average in high- or low-mobility locations based on non-causal estimates of mobility.*[16] These analyses do not test Assumption A1, but instead a weaker assumption that stayers' outcomes do not predict selection on movers' exposures. We provide such a test for Assumption A1, taking advantage of greater statistical power relative to Chetty et al. (2026a), who report being unable to test the assumption by including family fixed effects.

With coverage of siblings across many cohorts, we use family fixed effects to test Assumption A1, modifying Estimator A1 for comparison with the sibling fixed effects Estimator A2:

Estimator A2. An estimator of $\boldsymbol{\mu}$ eliminating the influence of family-level unobservables is the nested family fixed effects specification

$$y_i = (\boldsymbol{\mu}^0 + \boldsymbol{\mu}^1 \bar{p}_{f_i \boldsymbol{d}_i}) \cdot \boldsymbol{e}_i + \boldsymbol{\beta}^0 \cdot \boldsymbol{h}_i + \boldsymbol{\beta}^1 \bar{p}_{f_i \boldsymbol{d}_i} \cdot \boldsymbol{s}_i + \tilde{\epsilon}_i \tag{10}$$

where we replace sequence intercepts $\boldsymbol{s}_i$ with nested family-by-sequence intercepts $\boldsymbol{h}_i = (1_{f\boldsymbol{d}}(f_i, \boldsymbol{d}_i) : (f, \boldsymbol{d}) \in \boldsymbol{H})'$.[17] In both Estimator A2 and A1, we replace p_i with $\bar{p}_{f_i \boldsymbol{d}_i}$, the average parent rank over all siblings with the same location sequence, to address within-family estimation noise;[18] since p_i or $\bar{p}_{f_i \boldsymbol{d}_i}$ is included in the model as a source of heterogeneity, it is not a problem that identification of slopes $\hat{\boldsymbol{\beta}}^1$, and $\hat{\boldsymbol{\mu}}^1$ comes only from variation in exposures at different parent ranks. Residualization on $\boldsymbol{\beta}^0 \cdot \boldsymbol{h}_i + \boldsymbol{\beta}^1 \bar{p}_{f_i \boldsymbol{d}_i} \cdot \boldsymbol{s}_i = (\boldsymbol{\kappa}^0 + \boldsymbol{\kappa}^1 \bar{p}_{f_i \boldsymbol{d}_i}) \cdot \boldsymbol{s}_i + \bar{\theta}_{\boldsymbol{h}_i \bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i}$ incorporates both causal effects $\boldsymbol{\kappa}$ and selection effects $\bar{\theta}_{\boldsymbol{h}_i \bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i} = \mathrm{E}[\theta_i | \boldsymbol{h}_i, \bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i]$, and $\tilde{\epsilon}_i$ captures within-family residual variation. This estimator absorbs sequence and family effects, identifies exposure effects from within-family variation in the ages at which children move, and captures heterogeneity related to family-average parent ranks.

For comparison with Estimator A1, we maintain the weaker identification assumption

[16] Specifically, these authors ask whether including family fixed effects changes the rate at which movers' outcomes converge to stayers' outcomes.

[17] Although $\boldsymbol{h}_i$ captures most variation in parent ranks, we continue to estimate sequence-by-parent-rank slopes $\bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i$ for comparability with Estimator A1.

[18] Because parent family income ranks are similar within families, but child family income ranks are not, interacting exposures with p_i in small populations (such as children with siblings in rare location sequences) makes residualization on sequence-specific slopes $\hat{\boldsymbol{\beta}}^1$, and the resulting estimates $\hat{\boldsymbol{\mu}}^0$ and $\hat{\boldsymbol{\mu}}^1$, very unstable.

Assumption A2. Conditional on family-by-sequence intercepts, $\boldsymbol{h}_i$, and sequence-by-parent-rank slopes, $\bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i$, exposures and their interactions with parent family income rank $\boldsymbol{e}_i$ and $\bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{e}_i$ are jointly orthogonal to residual variation in childrens' outcomes

$$\mathrm{E}[\tilde{\epsilon}_i | \boldsymbol{e}_i, \bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{e}_i, \boldsymbol{h}_i, \bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i] = \mathrm{E}[\tilde{\epsilon}_i | \boldsymbol{h}_i, \bar{p}_{f_i \boldsymbol{d}_i} \boldsymbol{s}_i]$$

Assumption A2 only requires that there should be no selection within families and location sequences. For example, conditional on moving to Los Angeles from New York, older siblings should not have better unobservables than younger siblings.

Comparison of Estimators A1 and A2. We test for bias in Estimator A1 under the maintained assumption that Estimator A2 is unbiased. Formally, we define bias as the difference $\hat{\boldsymbol{\mu}}_{\bar{p}}^{A1} - \hat{\boldsymbol{\mu}}_{\bar{p}}^{A2}$, and test for significant bias for all CZs jointly and for each CZ separately.

We implement coefficient comparison tests as recommended by Pei, Pischke, and Schwandt (2017), estimating both models on a common population and accounting for shared observations, by seemingly-unrelated estimation.[19]

First, we test the joint null hypothesis that all estimates from Estimator A1 are unbiased, with $\hat{\boldsymbol{\mu}}_{\bar{p}}^{A1} - \hat{\boldsymbol{\mu}}_{\bar{p}}^{A2} = \mathbf{0}$, by the Wald statistic

$$H_{\bar{p}} = (\hat{\boldsymbol{\mu}}_{\bar{p}}^{A1} - \hat{\boldsymbol{\mu}}_{\bar{p}}^{A2})' (\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}_{\bar{p}}^{A1} - \hat{\boldsymbol{\mu}}_{\bar{p}}^{A2}})^{-1} (\hat{\boldsymbol{\mu}}_{\bar{p}}^{A1} - \hat{\boldsymbol{\mu}}_{\bar{p}}^{A2})$$

and find $H_{25} = 1019$ and $H_{75} = 1041$ which, against critical value $\chi^2_{0.05,761} = 826$, suggest bias in at least one location.

Second, to determine where bias might be relevant, we test the null hypothesis for each CZ that its estimate from estimator A1 is unbiased, with $\hat{\mu}_{c\bar{p}}^{A1} - \hat{\mu}_{c\bar{p}}^{A2} = 0$, by the t-statistic. Figure 5 shows that the bias in Estimator A1 is significant only in a small number of locations: nearly all differences have a t-statistic below 1.960, and accounting for multiple testing with a Bonferroni correction, all but five have a t-statistic below 3.993. Two of the five are Los Angeles and New York, suggesting caution in comparisons with these CZs. These exceptions reinforce the value of testing Assumption A1, with a focus on location-specific selection.

Overall, the evidence suggests that Assumption A1 is violated in a handful of locations, without implicating other locations.

[19] We stack two copies of the dataset, run one regression on each, and obtain the combined VCE matrix from the Stata command `suest`, clustering on paired observations. Following the design-based interpretation of Abadie et al. (2020), we do not correct for degrees of freedom lost to sequence fixed effects, but we do correct for degrees of freedom lost to family fixed effects within sequences.

4.3.2 Child-Level Factors

Although selective exposures are less of a concern for within-family comparisons than for between-family comparisons, selective exposures for individual children within families remain a concern. Chetty and Hendren (2018b, Section 4) carry out coefficient comparison tests for specifications including controls for changes in parents' income and marital status around moves, finding that estimates are largely unchanged. Given that other researchers continue to raise concerns about identification in this setting, we believe it is useful to extend and formalize these tests with greater statistical power and additional candidate confounders.

Specifically, we consider a number of potential child-level confounders raised by Gallagher, Kaestner, and Persky (2019), Heckman and Landersø (2022) and Eshaghnia (2023): parents' marital status, parents' age when children are born, children's birth order and spacing, and family size; and one they don't, whether children grew up where parents grew up.

We emphasize, as Chetty and Hendren (2018a and 2018b) do, that location sequence fixed effects address any selection bias due to the fact of moving between locations; the remaining concern is that there may be selection bias due to the timing of moves and the duration of the resulting exposures. However, it is essential to recognize that (1) as with valid instruments, imbalance on observables *does not* necessarily indicate selection bias unless the imbalance is related to the outcome through the structural error term (e.g. Angrist and Pischke, 2009, Ch. 3.2.2); and (2) changes in exposure effect estimates due to including observables as controls may reflect mediator bias, when controls include variables on the causal path from exposures to children's outcomes, or other outcomes themselves (e.g. Rosenbaum, 1984; Angrist and Pischke, 2009, Ch. 3.2.3; Cinelli, Forney, and Pearl, 2024). The problem of over-control is not new to studies of intergenerational mobility: see for example Sampson (2008). Thus, any finding of imbalance on observables or change in estimated exposure effects would be suggestive of selection bias, but not conclusive. Further, Todd and Wolpin (2003) point out that the *ceteris paribus* treatment effect controlling for mediators is not necessarily more interesting or relevant than the total policy effect identified by Estimator A1.

We use control variables to test Assumption A1, modifying Estimator A1 for comparison with the controlled Estimator A3:

Estimator A3. An estimator of $\boldsymbol{\mu}$ eliminating the influence of child-level observables is the extended specification

$$y_i = (\boldsymbol{\mu}^0 + \boldsymbol{\mu}^1 p_i) \cdot \boldsymbol{e}_i + (\boldsymbol{\alpha}^0 + \boldsymbol{\alpha}^1 p_i) \cdot \boldsymbol{s}_i + \boldsymbol{\gamma} \cdot \boldsymbol{w}_i + \tilde{\epsilon}_i \tag{11}$$

where we add controls $\boldsymbol{w}_i$. This estimator absorbs sequence effects, and controls for differences in children's circumstances both within and between families based on the ages at which children move.

For comparison with Estimator A1, we maintain the weaker identification assumption

Assumption A3. Conditional on location sequences, their interactions with parent family income rank $\boldsymbol{s}_i$ and $p_i\boldsymbol{s}_i$, and controls $\boldsymbol{w}_i$; exposures and their interactions with parent family income rank $\boldsymbol{e}_i$ and $p_i\boldsymbol{e}_i$ are jointly orthogonal to residual variation in childrens' outcomes

$$\mathrm{E}[\tilde{\epsilon}_i|\boldsymbol{e}_i, p_i\boldsymbol{e}_i, \boldsymbol{s}_i, p_i\boldsymbol{s}_i, \boldsymbol{w}_i] = \mathrm{E}[\tilde{\epsilon}_i|\boldsymbol{s}_i, p_i\boldsymbol{s}_i, \boldsymbol{w}_i]$$

Assumption A3 relaxes the requirement that there should be no selection within location sequences, requiring only that there should be no selection after controlling for observables. For example, conditional on moving to Los Angeles from New York, and conditional on family size or birth order, one child should not have better unobservables than other children making the same move, with the same family size or birth order. Controlling for these covariates could capture confounding, for example, due to resource constraints in larger families, changes in parenting ability over time, and changes in parental labor force participation related to childcare needs for overlapping children.

Comparison of Estimators A1 and A3. We test for bias in Estimator A1 under the maintained assumption that Estimator A3 is unbiased. Formally, we define bias as the difference $\hat{\boldsymbol{\mu}}_p^{A1} - \hat{\boldsymbol{\mu}}_p^{A3}$, and test for significant bias for all CZs jointly, again by coefficient comparison tests.

First, we test the joint null hypothesis that all estimates from Estimator A1 are unbiased, with $\hat{\boldsymbol{\mu}}_p^{A1} - \hat{\boldsymbol{\mu}}_p^{A3} = \mathbf{0}$, by the Wald statistic

$$H_p = (\hat{\boldsymbol{\mu}}_p^{A1} - \hat{\boldsymbol{\mu}}_p^{A3})'(\hat{\boldsymbol{V}}_{\hat{\boldsymbol{\mu}}_p^{A1} - \hat{\boldsymbol{\mu}}_p^{A3}})^{-1}(\hat{\boldsymbol{\mu}}_p^{A1} - \hat{\boldsymbol{\mu}}_p^{A3})$$

and find $H_{25} = 15012$ and $H_{75} = 10369$ which, against critical value $\chi^2_{0.05,761} = 826$, suggest bias in at least one location.

However, we observe that the variance of exposure effects from Estimator A3 is two orders of magnitude greater than the variance of between-estimator differences due to added controls, suggesting little scope for bias to affect CZ comparisons.

Overall, Assumption A1 appears to be a reasonable approximation. With minor caveats, we conclude that our exposure effect estimates can be interpreted as causal.

5 Observed Outcomes

In this section, we compare the exposure effect estimates with children's observed outcomes. We obtain observed outcomes for both stayers and movers by modifying our benchmark Estimator A1 to remove sequence fixed effects:

Estimator A0. An estimator of $\boldsymbol{\mu}$ omitting sequence fixed effects is

$$y_i = (\boldsymbol{\mu}^0 + \boldsymbol{\mu}^1 p_i) \cdot \boldsymbol{e}_i + \eta_i \tag{12}$$

where η_i contains sequence effects $(\boldsymbol{\kappa}^0 + \boldsymbol{\kappa}^1 p_i) \cdot \boldsymbol{s}_i$ and individual unobservables θ_i. To the extent sequence effects and unobservables are associated with exposure effects, they will be captured by the resulting estimates of observed outcomes.

For stayers, exposure effects become estimable because they are no longer collinear with sequence fixed effects; the estimated exposure effects $(\hat{\boldsymbol{\mu}}^0 + \hat{\boldsymbol{\mu}}^1 p) \times 24$ are equivalent to the place-specific mean income ranks $\bar{\boldsymbol{y}}_p$. For movers, exposure effects are weighted place-specific mean income ranks, where the weights are given by exposure durations.

We find that observed outcomes for movers and stayers are similar, excepting the level differences observed in Table 1: the exposure-weighted correlation between these two measures is 0.96 for the 25th percentile of parent family income, and 0.94 for the 75th percentile.

Next, we compare observed outcomes with exposure effect estimates, to understand selection bias in observed outcomes.

5.1 Selection Bias in Observed Outcomes

We measure place-specific selection bias in observed outcomes as the difference between exposure effects and observed outcomes, under the assumption that Estimator A1 is unbiased. Formally, we define bias by the differences $\hat{\boldsymbol{\mu}}_p^{A0,M} - \hat{\boldsymbol{\mu}}_p^{A1}$ and $\hat{\boldsymbol{\mu}}_p^{A0,S} - \hat{\boldsymbol{\mu}}_p^{A1}$ for movers and stayers respectively; estimates for movers are calculated in a common population.

Selection bias accounts for one-half to two-thirds of the variance in observed outcomes. Table 3 reports the variance of the bias terms and the exposure effects, showing that the variance of the bias terms ranges from 1 to 2 times the variance of the exposure effects.

Selection bias does not reinforce or counteract exposure effects. The bias terms and the exposure effects are either uncorrelated or weakly negatively correlated, so that the variance of observed outcomes is similar to the sum of the variances.

The fact that selection and exposure effects are unrelated is an important finding. The rate of convergence of movers' outcomes to stayers' outcomes can only be used to estimate

the causal share of variance in observed outcomes when selection effects are orthogonal to exposure effects (Chetty et al., 2026a; Appendix H). We are the first to measure the selection term and validate this widely-used assumption (e.g. Chetty and Hendren, 2018a; Deutscher, 2020; Britto et al., 2022; Soria-Espín and Medina, 2025; and Chetty et al., 2026a).[20]

Where does selection bias matter? Figure 6 and 7 map the selection term in observed outcomes by CZ, at the 25th and 75th percentiles of parent family income. The bias in observed outcomes makes many large cities, such as New York, Los Angeles, and San Francisco, appear much better for children than they are in a causal sense; and makes much of the South and Southwest appear much worse for children than they are in a causal sense.

We conclude that estimates based on exposure effects may be a more reliable guide to causal effects than estimates based on observed outcomes, although there are settings where researchers may have no choice but to rely on observed outcomes.

5.2 Can We Correct for Selection Bias in Observed Outcomes?

Can we correct observed outcomes for selection bias by including control variables, to better approximate exposure effects? Altonji and Mansfield (2018) show that the underlying intuition is misleading: if the goal is to identify exposure effects, then the control variables must explain only the selection bias, not the exposure effects.

In this section, we use our estimates to expand on the analysis of correlations between exposure effects, selection effects, and covariates proposed by Chetty and Hendren (2018b), going further to compare the reliability of uncorrected observed outcomes and covariate-corrected observed outcomes as predictors of causal exposure effects.

We choose CZ covariates to match those examined in the literature testing observational estimates for selection bias (e.g. Rothbaum, 2016; Gallagher, Kaestner, and Persky, 2019; Cholli et al., 2024; and Chetty et al., 2026a): race, nativity, education, labor force participation, unemployment, marriage, divorce, single parenthood, and crime.

Table 4 reports the fraction of variance in observed outcomes, exposure effects, and selection bias that is explained by covariates. We find that these covariates explain somewhere between half and two thirds of observed outcomes, a similar proportion of the selection bias, and also a similar proportion of the exposure effects. We then residualize observed outcomes on the covariates to remove selection bias as suggested in the literature, and then correlate observed outcomes and corrected observed outcomes with our exposure effects to show how they perform as proxies for the causal variation measured by exposure effects. We find that

[20]Our finding that one-half to two-thirds of the variance in observed outcomes is bias is potentially consistent with findings in other papers that less than half of the variance is bias: other papers focus on other geographies, other subsets of the movers population, or both.

the correlation between exposure effects and corrected observed outcomes, ranging between 0.40 and 0.59, is always smaller than the comparable correlation between exposure effects and uncorrected observed outcomes, ranging between 0.58 and 0.65: for the 25th percentile of parent family income in particular, correcting observed outcomes for covariates does not increase the explanatory power of observed outcomes for exposure effects, but instead reduces explanatory power by approximately 20 percent.

Overall, controlling for covariates does not make observed outcomes better predictors of treatment effects for movers.

5.3 Can We Use Observed Outcomes to Improve Prediction?

Observed outcomes are more precise than exposure effects. In principle, it may be possible to improve prediction of true exposure effects by combining observed outcomes with estimated exposure effects, trading off bias against noise. Empirical Bayes shrinkage (Robbins, 1956) can be used to make the optimal tradeoff: exposure effects can be shrunk toward observed outcomes, by taking a weighted average giving greater weight to exposure effects when they are more precise, to observed outcomes when they are less biased, and to zero when exposure effects are noisy and observed outcomes are biased simultaneously. However, shrinkage may not be necessary when estimated exposure effects are already precise.

To test whether shrinkage improves prediction given our precise exposure effects and the bias in observed outcomes, we shrink our exposure effects toward observed outcomes as in Chetty and Hendren (2018b), obtaining the linear forecast of the true exposure effects that minimizes the mean-squared error (MSE). Figure 8 and 9 display the forecasts and confidence intervals based on their root-mean squared errors (RMSE). We find that although confidence intervals are shorter for the most noisy estimates, only the most noisy estimates are meaningfully shrunk toward observed outcomes.

To understand why shrinkage does not improve prediction, we compare the shrinkage estimates with their weighted components, the raw exposure effect estimates and the observed outcomes, in Figure 10. Overall, the forecasts are very similar to estimated exposure effects, and very different from observed outcomes. Figure 11 and 12 map the discrepancies between forecasts and the underlying data, showing where observed outcomes over- and under-estimate the effects of spending more time in a particular CZ in childhood: these results emphasize that the optimal forecasts closely match the exposure effects, avoiding reliance on the observed outcomes due to selection bias.

Our exposure effect estimates will be hard to improve using observed outcomes, so we carry out the rest of our analysis without shrinkage.

6 Comparing and Ranking Exposure Effects

In this section, we compare commuting zones' (CZs') estimated exposure effects to each other, instead of to the national mean as in Section 4. We start with pairwise comparisons, building up to compound comparisons based on pairwise comparisons.

6.1 Pairwise Comparisons

Because decisions about when (and whether) to move depend on exposure effects relative to other CZs, we measure differences $\hat{\mu}_{cp} - \hat{\mu}_{c'p}$, and their uncertainty, by combining estimates from equation (5) and (6),

$$\hat{V}_{\hat{\mu}_{cp}-\hat{\mu}_{c'p}} = \hat{V}_{\hat{\mu}_{cp}} + \hat{V}_{\hat{\mu}_{c'p}} - 2\hat{V}_{\hat{\mu}_{cp},\hat{\mu}_{c'p}} \tag{13}$$

For example, for children moving from New York to Los Angeles, with parents at the 25th percentile of parent family income, spending one more year in New York instead of Los Angeles would result in child family income $\hat{\mu}_{NY25} - \hat{\mu}_{LA25} = -0.0786$ (0.0071) percentiles lower. At the 75th percentile of parent family income, the same year in New York would result in child family income $\hat{\mu}_{NY75} - \hat{\mu}_{LA75} = -0.0497$ (0.0081) percentiles lower. These effects are typical in size: at the 25th and 75th percentiles of parent family income, 71 and 70 percent of pairwise effects, respectively, are larger than 0.0417 percentiles per year of exposure; scaled to full exposure, these effects amount to a 1 percentile change in children's family income due to choosing their origin CZ over their destination CZ, or vice versa.

How informative are pairwise comparisons of our estimates, for the typical family deciding when and where to move? What testing criteria should we use? Family decisions about when to move involve single comparisons. Confidence intervals around the benefit of the move, in terms of child percentile ranks gained per year of exposure, can be constructed based on a standard critical value of 1.96. These single comparisons between CZs are the most likely use of the estimates, and the use most consistent with the estimation strategy, which is based on variation in the timing of moves between CZs rather than variation in where families move.

Figure 13 presents a histogram of absolute t-statistics across all pairwise comparisons, weighted by the number of children making each pairwise move with their families. We find that at the 25th and 75th percentiles of parent family income, 64 and 63 percent of moves, respectively, would produce a significant gain or loss in children's outcomes, at conventional significance levels relevant for deciding when to move. Informally, using a 1 percentile cutoff for economic significance, 84 percent of economically significant differences are statistically significant at conventional levels. Conversely, among economically insignificant differences,

only 13 percent are statistically significant at conventional levels.

Family decisions about where to move involve multiple comparisons, of the family's origin CZ with up to 760 potential destination CZs, and are likely to involve identifying the best potential destination among those considered (a ranking problem as in Mogstad et al., 2024; discussed in Section 6.2 below) or identifying the gain from moving to the best potential destination (a selection problem as in Andrews, Kitagawa, and McCloskey, 2024). To give a rough idea of how many pairwise differences are significant with multiple comparisons, we consider a conservative Bonferroni correction: the critical value controlling the per-family error rate (PFER) is 3.99. By this criterion, we find that 35 percent of moves would produce a significant gain or loss. Of course, because many families consider only a few potential destinations, the appropriate critical value may be much smaller.

Are comparisons informative for all CZs, or only some CZs with extreme estimates? Figure 14 and 15 report the fraction of tests that exceed the conventional critical value of 1.96, and the corrected critical value of 3.99, for each CZ. At conventional levels we find that two-fifths or more of comparisons for a given CZ are significant, and more than four-fifths for most of the largest urban CZs. Only rarely does the fraction of significant comparisons fall below one-fifth. That is, from most CZs, most pairwise comparisons are informative for families considering when to move. Turning to tests for comparisons across all potential destinations, most of the most-populous CZs reach at least one-fifth to two-fifths of comparisons significant, with others reaching above four-fifths significant; however, it becomes more common to find origin CZs where little can be said about what pairwise differences are significant. We view this latter result as a lower bound on informativeness.

We conclude that most families making decisions based on the estimates, regardless of where they live, can consider the estimates informative both in single comparisons and multiple comparisons. However, it remains to be seen how the estimates perform when used as inputs to compound problems including ranking and selection.

6.2 Rankings

Where should a family move, if at all? Where should governments intervene to suggest that families move, or to implement place-based policies to improve neighborhood conditions?

The premise of the recent literature on place effects is that they can be ranked once they are identified (as in Chetty et al., 2014, Table 3; Chetty and Hendren, 2018b, Table 3-4; and Chetty et al., 2026a, Section 3) and used to guide recommendations such as those in Bergman et al. (2024), encouraging families to move to opportunity.

One purpose of our paper is to resolve two serious critiques of the existing rankings. First,

observed outcomes (as in Chetty et al., 2014 and Chetty et al., 2026a) are not necessarily causal, and selection bias may determine the ranking. Second, quasi-experimental estimates (as in Chetty and Hendren, 2018b) are noisy, and noise may determine the ranking. For this reason, Chetty and Hendren (2018b) do not rank their quasi-experimental estimates, instead ranking "forecasts" shrunk toward observed outcomes. Mogstad et al. (2024) argue that both the forecasts and the underlying quasi-experimental estimates are inadequate for ranking: the forecasts because they are very similar to the biased observational estimates, and the quasi-experimental estimates because the are noisy.

Our first step to resolve the critiques is to provide quasi-experimental estimates that are plausibly unbiased and reduce noise by two orders of magnitude. However, reducing bias and noise does not alone guarantee that estimates can be ranked. Mogstad et al. (2024) observe that questions about where to move require multiple comparisons among all considered destinations for a given origin, and questions about where the government should intervene multiply comparisons among all origins as well. Multiple comparisons create multiple opportunities for mis-ranking.

Our second step to resolve the critiques is to account for multiple comparisons. Mogstad et al. (2024) provide procedures to control the marginal and simultaneous family-wide error rates (FWER) for the comparisons underlying rankings. These procedures set stringent critical values for the underlying pairwise comparisons. To draw the implications of our estimates for the resulting ranking, we apply the method in Mogstad et al. (2024) to our estimates; because they did not have access to full VCE estimates from Chetty and Hendren (2018b), we modify their code to compute standard errors for pairwise comparisons using the covariance estimates, finding a small sharpening of the ranking.

Ranking results are presented in Figure 16 and 17, including "marginal" confidence sets for each CZ relative to all the other CZs, and "simultaneous" confidence sets for all possible pairwise comparisons. As in Mogstad et al. (2024), we focus on the top 50 CZs by population. The new ranking is highly informative. For children of parents at the 25th percentile of family income, we find that the top-ranked CZ, Salt Lake City, has a 95-percent marginal confidence set from rank 1-7 (families comparing Salt Lake City with possible destinations can be very sure that all but 6 other CZs have lower exposure effects); and we find that the bottom-ranked CZ, New Orleans, has a 95-percent marginal confidence set from rank 49-50 (families comparing New Orleans with possible destinations can be very sure that all but 1 other CZs have higher exposure effects). Similarly, Salt Lake City has a 95-percent simultaneous confidence set from rank 1-16 (policy makers comparing Salt Lake City with all other large CZs can be very sure that all but 15 other CZs have lower exposure effects); and we find that New Orleans has a 95-percent simultaneous confidence set from rank 44-50

(policy makers comparing New Orleans with all other large CZs can be very sure that all but 5 other CZs have higher exposure effects). Many marginal confidence sets are fewer than 10 ranks wide, and nearly all are fewer than 20 ranks wide. Simultaneous confidence sets are not much wider, typically covering fewer than half of the other CZs.

These rankings are conservative. They control for family-wide error rates (FWER), that is, for the probability that any one of the pairwise comparisons underlying a ranking is reversed. For simultaneous comparisons, two examples where such a test is relevant are (1) deciding which origin CZ's residents to target for a policy recommending moves to one or a few of the possible destination CZs, and (2) deciding which CZs to target for place-based investment, when all CZs are considered. For marginal comparisons holding the origin or destination CZ fixed, such a test is relevant primarily when all CZs are potential origins or destinations. But for a family considering when to move between a known origin and destination (for example for parents deciding when to move back home to be near grandparents), the pairwise comparisons in Section 6.1 are more appropriate.

7 Counterfactuals

Do comparisons of exposure effects offer valid counterfactuals for families deciding where to move and policymakers deciding where to invest? This section presents the assumptions required to interpret exposure effects as counterfactuals; shows how to define and estimate valid counterfactuals when those assumptions fail; presents the resulting novel, corrected counterfactual ranking; and shows that standard comparisons based on exposure effects substantially deviate from the corrected counterfactual ranking. Overall, the effect of places on child outcomes through parent income is first-order, and is missed by exposure effects.

7.1 Interpreting Exposure Effects as Counterfactuals

How do children's outcomes change as their counterfactual exposures change? We compare a child's actual rank y_i with the counterfactual rank y_i' based on the model in equation (1). For a family counterfactually spending one more year in CZ c' instead of CZ c,

$$y_i' - y_i = \underbrace{(\mu_{c'}^0 + \mu_{c'}^1 p_i) - (\mu_c^0 + \mu_c^1 p_i)}_{A} + \underbrace{\boldsymbol{\mu}^1 \cdot \boldsymbol{e}_i'(p_i' - p_i)}_{B} + \underbrace{\kappa_{\boldsymbol{d}'}^0 - \kappa_{\boldsymbol{d}}^0 + \kappa_{\boldsymbol{d}'}^1 p_i' - \kappa_{\boldsymbol{d}}^1 p_i}_{C} \tag{14}$$

where p_i', $\boldsymbol{e}_i'$, and $\boldsymbol{d}_i'$ denote the parents' counterfactual rank, exposure, and location sequence. Fixed unobservables θ_i cancel. Term A is the standard pairwise exposure effect comparison in Section 6, evaluated at the parents' actual rank. In general, Estimator A1 recovers term

A, but not terms B and C, which depend on unknowns p'_i, $\kappa^0_{\boldsymbol{d}'}$, $\kappa^0_{\boldsymbol{d}}$, $\kappa^1_{\boldsymbol{d}'}$, and $\kappa^1_{\boldsymbol{d}}$.

Under two strong assumptions that cancel terms B and C, estimates of term A provide valid counterfactuals. The first assumption is that parent ranks are identical: $p'_i = p_i$, and the second is that location sequence effects are identical: $\kappa^0_{\boldsymbol{d}'} = \kappa^0_{\boldsymbol{d}}$ and $\kappa^1_{\boldsymbol{d}'} = \kappa^1_{\boldsymbol{d}}$.

We are not aware of any paper presenting counterfactuals that states these assumptions. Are these assumptions credible?

First, parents' rank depends on their exposure.[21] When $p'_i \neq p_i$, exposure effects (term A) are invalid counterfactuals because term B and C are nonzero. When the first assumption fails, counterfactuals can only be obtained if p'_i, $\kappa^1_{\boldsymbol{d}'}$, and $\kappa^1_{\boldsymbol{d}}$ are known; or if p'_i is known and $\kappa^1_{\boldsymbol{d}'} = \kappa^1_{\boldsymbol{d}} = 0$. We are not aware of any paper correcting counterfactuals for parent rank.

Second, location sequence effects vary for many counterfactual moves. When families make the same moves earlier or later, location sequence effects cancel. When families make different moves, location sequence effects may not cancel. Existing evidence permits us to at most assume that level effects of location sequences are constant across families that make the same number of moves, $\kappa^0_{\boldsymbol{d}'} = \kappa^0_{\boldsymbol{d}}$, and lower for families that make more moves, $\kappa^0_{\boldsymbol{d}'} < \kappa^0_{\boldsymbol{d}}$;[22] and that slope effects of parent rank on child rank are positive and constant across location sequences: $\kappa^1_{\boldsymbol{d}} = \kappa^1 > 0$.[23] This weaker assumption leaves term C unknown, and requires additional measurements to obtain valid counterfactuals.

7.2 Obtaining Valid Counterfactuals

Because we cannot assume them away, obtaining valid counterfactuals requires us to estimate unknown terms B and C.

Our approach is to rewrite Estimator A1 in terms of a parent rank that is exogenous to exposures within location sequences. This approach pushes exposure-dependent parent rank variation into the error term and induces bias in estimates $\hat{\boldsymbol{\mu}}^0$ and $\hat{\boldsymbol{\mu}}^1$. We show that under mild assumptions, the bias exactly recovers terms B and C, providing valid counterfactuals within our existing estimation framework.

First, we model parents' rank as a function of exposures and latent rank q_i,

$$p_i = \boldsymbol{\rho} \cdot \boldsymbol{e}_i + q_i \tag{15}$$

[21] Places including CZs affect workers' income (Combes, Duranton, and Gobillon, 2008; De la Roca and Puga, 2017; Dauth et al., 2022; Card, Rothstein, and Yi, 2025; Sprung-Keyser and Porter, 2026; Tan, 2026), in estimates based on modified Abowd, Kramarz, and Margolis (1999) panel estimators.

[22] Moving hurts children (Chetty, Hendren, and Katz, 2016). We generalize the assumption in Chetty and Hendren (2018a) that $\kappa^0_{\boldsymbol{d}} = \kappa^0 < 0$ for once-movers, and $\kappa^0_{\boldsymbol{d}} = 0$ for stayers.

[23] All else equal, raising parents' income raises children's income (Akee et al., 2010; Aizer et al., 2016; Hilger, 2016; Barr, Eggleston, and Smith, 2022; Sikhova, 2023; Bailey et al., 2024; Dustmann, Landersø, and Andersen, 2024; Britto et al., 2025).

where rank premium $\rho_{c'} - \rho_c$ gives the effect on parents' rank of spending one more year in CZ c' instead of CZ c.

Second, we recover the parents' latent rank by estimating the rank premia due to each family's exposures. We assume that the latent rank does not depend on exposures conditional on parents' location sequences. We control for sequence fixed effects, estimating

$$p_i = \boldsymbol{\rho} \cdot \boldsymbol{e}_i + \boldsymbol{\beta} \cdot \boldsymbol{s}_i + \tilde{q}_i \tag{16}$$

where $\boldsymbol{\beta} \cdot \boldsymbol{s}_i$ captures selection on location sequences, and $\tilde{q}_i$ captures residual variation in parent rank. We calculate the latent rank from the estimated parameters, $\hat{q}_i = p_i - \hat{\boldsymbol{\rho}} \cdot \boldsymbol{e}_i$.

Third, the estimand in equation 1 can be rewritten as a function of the true latent rank q_i and the exposures,

$$y_i = (\boldsymbol{\mu}^0 + \kappa^1 \boldsymbol{\rho} + \boldsymbol{\mu}^1(\boldsymbol{\rho} \cdot \boldsymbol{e}_i)) \cdot \boldsymbol{e}_i + \boldsymbol{\mu}^1 \cdot q_i \boldsymbol{e}_i + \boldsymbol{\kappa}^0 \cdot \boldsymbol{s}_i + \boldsymbol{\kappa}^{\mathbf{1}} \cdot q_i \boldsymbol{s}_i + \theta_i \tag{17}$$

where we maintain the assumption that $\kappa^1_{\boldsymbol{d}} = \kappa^1$.

Next, we obtain parameters for candidate counterfactuals, by estimation:

Estimator A4 Building an estimator with the same regressors, plugging in the population estimates of latent ranks $\hat{q}_i$ for q_i,[24] we obtain

$$y_i = \boldsymbol{\mu}^0 \cdot \boldsymbol{e}_i + \boldsymbol{\mu}^1 \cdot \hat{q}_i \boldsymbol{e}_i + \boldsymbol{\alpha}^0 \cdot \boldsymbol{s}_i + \boldsymbol{\alpha}^1 \cdot \hat{q}_i \boldsymbol{s}_i + \epsilon_i \tag{18}$$

where it is already clear that by pushing parts of term B and C into the error term, where those parts are multiplied by exposures, $\hat{\boldsymbol{\mu}}^0$ and $\hat{\boldsymbol{\mu}}^1$ may recover the indirect effects of exposure on child rank through parent rank.

Finally, we verify that the estimates can be used to construct valid counterfactuals. The biased estimates from Estimator A4 capture key parts of term B and C that are needed to construct the counterfactual differences,

$$\begin{aligned}
\hat{\mu}^{0,A4}_{c'} - \hat{\mu}^{0,A4}_{c} &\to \mu^0_{c'} - \mu^0_c + \kappa^1(\rho_{c'} - \rho_c) + 24(\mu^1_{c'}\rho_{c'} - \mu^1_c \rho_c) \\
\hat{\mu}^{1,A4}_{c'} - \hat{\mu}^{1,A4}_{c} &\to \mu^1_{c'} - \mu^1_c
\end{aligned}$$

where we rely on differencing to eliminate nuisance terms.

In the estimated counterfactual, since Estimator A4 does not control for the part of

[24] Our approach is similar to Farber and Gibbons (1996) but opposite in effect: they regress on residuals to block pre-treatment confounding and we regress on residuals to unblock post-treatment mediation.

parent family income rank that moves with exposures, bias terms B' and C' capture the impact of exposures on child ranks through their impact on parent ranks:

$$\begin{aligned}\widehat{y_i' - y_i} &= \hat{\mu}_{c'}^{0,A4} - \hat{\mu}_{c}^{0,A4} + (\hat{\mu}_{c'}^{1,A4} - \hat{\mu}_{c}^{1,A4})p_i \\ &\to \underbrace{(\mu_{c'}^0 + \mu_{c'}^1 p_i) - (\mu_c^0 + \mu_c^1 p_i)}_{A} + \underbrace{24(\mu_{c'}^1 \rho_{c'} - \mu_c^1 \rho_c)}_{B'} + \underbrace{\kappa^1(\rho_{c'} - \rho_c)}_{C'} \end{aligned} \tag{19}$$

where terms B' and C' from the estimated counterfactual in equation 19 are analogues to terms B and C from the true counterfactual in equation 14.

Under what conditions does the Estimator A4 counterfactual coincide with the true counterfactual? Both term B and term C must cancel. Term B, the influence of parent rank premia through exposure effect heterogeneity at counterfactual exposures,

$$\begin{aligned} B' - B &= 24(\mu_{c'}^1 \rho_{c'} - \mu_c^1 \rho_c) - (\boldsymbol{\mu}^1 \cdot \boldsymbol{e}_i'(\rho_{c'} - \rho_c)) \\ &= \mu_{c'}^1 \rho_{c'}(24 - e_{ic'}' - e_{ic}') - \mu_c^1 \rho_c (24 - e_{ic'}' - e_{ic}') \\ &\quad + \sum_{k \neq c', c} \mu_k^1 e_{ik}' (\rho_{c'} - \rho_c) \end{aligned}$$

cancels either if all exposure is on the counterfactual origin and destination, $e_{ic'}' + e_{ic}' = 24$; or if parent rank has the same effect on child rank in both CZs, $\mu_{c'}^1 \rho_{c'} = \mu_c^1 \rho_c$, and the weighted average effect of parent rank through other CZs is zero. Term C, the influence of parent rank premia through location sequence effects,

$$C' - C = -\kappa_{\boldsymbol{d}'}^0 + \kappa_{\boldsymbol{d}}^0$$

cancels under our maintained assumption $\kappa_{\boldsymbol{d}'}^0 = \kappa_{\boldsymbol{d}}^0$ if the counterfactual does not change the family's number of moves.

In summary, under mild assumptions, Estimator A4 makes it possible to construct valid counterfactuals that combine both the direct effect of places on child ranks, and indirect effects through the effect of places on parent ranks.

7.3 Counterfactual Estimates

What are the counterfactual effects of places on children's outcomes? We now present our novel counterfactual estimates, which differ from the conventional estimates of exposure effects presented in Section 4. These are the relevant effects for families and policymakers.

We measure CZs' counterfactual effects relative to the national mean, and their uncertainty, by combining estimates from equations (3) and (4), for parents at the 25th and 75th

percentiles of parent family income.

Counterfactual estimates lead to different conclusions about which places are the best places for children's outcomes. For example, for parents at the 25th percentile of the national family income distribution, each additional year a family spends in Atlanta (relative to the national mean) raises its child's family income rank by 0.17 percentiles more than predicted by the exposure effects alone. Ranked by exposure effects, for parents at the 25th percentile, Atlanta is 29th among the 50 largest CZs, and 671th nationally; ranked by counterfactual total effects, Atlanta is 1st among the 50 largest CZs, and 211th nationally.

Large differences are not unusual. Figure 18 reports counterfactuals $\hat{\mu}_{c25}^{A4}$ and $\hat{\mu}_{c75}^{A4}$, and their confidence intervals, for the 50 largest CZs; the original estimates from Figure 1 are underlaid for comparison. The width of the counterfactuals' confidence intervals is similar to those of the exposure effects. However, for most of the 50 largest CZs, the confidence intervals of the counterfactuals do not overlap the confidence intervals of the exposure effects. Large urban CZs generally produce higher child family income ranks than the exposure effects suggest: 35/50 produce higher ranks for children of parents at the 25th percentile, and 34/50 produce higher ranks for children of parents at the 75th percentile.

Counterfactuals sharpen the ranking of commuting zones. Figures 19 and 20 present ranking results for counterfactual total effects. The new ranking is even more informative than the ranking of exposure effects in Section 6.2. For children of parents at the 25th percentile of income, we find that the top-ranked CZ, Atlanta, has a 95-percent marginal confidence set from rank 1-5; and we find that the bottom-ranked CZ, New Orleans, has a 95-percent marginal confidence set of rank 50 only. Results for simultaneous confidence sets are similar. It is now the case that several CZs' ranks can be distinguished from all other CZs, and most confidence sets are 5-10 ranks long. The improvement is especially notable for children of parents at the 75th percentile.

Counterfactuals change the geography of mobility. Figures 22 and 23 map the geography of counterfactual total effects $\hat{\mu}_{c25}^{A4}$ and $\hat{\mu}_{c75}^{A4}$; and Figure 24 maps counterfactual indirect effects, that is, the differences $\hat{\mu}_{c25}^{A4} - \hat{\mu}_{c25}^{A1}$ and $\hat{\mu}_{c75}^{A4} - \hat{\mu}_{c75}^{A1}$ between counterfactual total effects and exposure effects, on children's family income ranks. Allowing places to affect child outcomes through parent incomes resolves counterintuitive patterns in the exposure effects. Notably, we find that the Ozarks, Appalachia, and the rural South, Southwest, and Pacific Northwest look worse for children in the counterfactuals, consistent with common intuition. Many urban areas, particularly in the Northeast, Midwest, Texas, the South, and the Pacific coast, look better in counterfactuals than in exposure effects. Interestingly, New York is no better, but adjacent commuting zones Bridgeport, Poughkeepsie, Newark, and Toms River (the Jersey Shore), all connected to the city by commuter rail, outperform in the

counterfactual estimates compared to exposure effects.

What other measures of place effects most closely match the estimated counterfactuals? Table 5 correlates counterfactual total effects $\hat{\mu}_{c25}^{A4}$ and $\hat{\mu}_{c75}^{A4}$ against exposure effects $\hat{\mu}_{c25}^{A1}$ and $\hat{\mu}_{c75}^{A1}$, and observed outcomes $\hat{\mu}_{c25}^{A0,S}$ and $\hat{\mu}_{c75}^{A0,S}$, as well as the exposure effects, observed outcomes, and forecasts estimated by Chetty and Hendren (2018b). We find that exposure effects are the best proxy for counterfactuals, with correlations of 0.751 and 0.737 respectively for the 25th and 75th percentile of parent family income; by this measure, exposure effects (the state of the art) miss nearly half the true variation in counterfactual effects even if precisely estimated. Counterfactual correlations with observed outcomes are lower at 0.615 and 0.693 respectively; by this measure, observed outcomes (the standard proxy) miss half or more of the true variation.[25] Finally, we confirm Chetty and Hendren's (2018b) conclusion that their shrinkage forecasts and observed outcomes outperform their exposure effect estimates because the latter are noisy. In summary, we show that our improved exposure effect estimates are the best proxy for counterfactuals, but estimated counterfactuals are preferable because they capture effects of places on child outcomes through parent incomes.

8 Mechanisms

What characteristics explain the patterns in counterfactual effects and exposure effects? Because the counterfactuals are so different from the exposure effects already characterized in Chetty and Hendren's 2018b correlational analysis, this section focuses on the counterfactuals, why they are different from the exposure effects, and what this means for the availability of "opportunity bargains" offering good outcomes for children at low cost.

8.1 Income, Prices, and Place Effects

What explains place effects, and what in particular explains the substantial differences between counterfactual effects and exposure effects?

The causal effects of places on parent family income explain more than half of the full causal effect of places on child ranks: Table 5, Column 2 shows that counterfactual effects are strongly correlated with parent family income rank premia, at 0.751 and 0.719 respectively for the 25th and 75th percentiles of parent family income. The rank-premium effect is first-order, and is not documented in the literature. It offers a simple answer to the question of how to improve places to improve outcomes. First, it suggests that resources may matter more

[25] Interestingly, observed outcomes are equally or more-strongly correlated with counterfactuals than they are with exposure effects, because both counterfactuals and observed outcomes capture the causal effects of parent rank premia, while exposure effects are almost unrelated to parent rank premia.

than fixed demographics and entrenched social structures. Second, it suggests that directing resources to families rather than to places may improve both targeting and efficiency.

Indirect effects, the difference between counterfactual effects and exposure effects, are almost entirely explained by the causal effects of places on parent ranks. Figure 21 scatters indirect effects on children's family income ranks against parent rank premia $\hat{\rho}_c$. Parent rank premia explain nine-tenths of the indirect effects of CZs on child ranks. Because the indirect effects approximate the sum of terms B' and C' ($24\mu_c^1\rho_c$ and $\kappa^1\rho_c$) relative to the base category, and our estimate of term B' is weakly correlated with the indirect effects, we infer that indirect effects are driven by term C' and the slope in the regression of the indirect effects on $\hat{\rho}_c$ is approximately κ^1, the causal effect of parent rank on child rank: children gain approximately 0.35 percentiles per percentile of parent rank.

Exposure effects are weakly related to parent rank effects, and more-strongly related to rents. Table 5, Column 3 reports weak correlations between exposure effects and parent rank premia of 0.145 and 0.087 respectively. Exposure effects are negatively related to rents, with correlations of -0.392 and -0.732 respectively. The relationship between exposure effects and rents may be causal: low rents increase resources available for investment in children, and lower the price of services including daycare, schools, and enrichment activities.

Thus, indirect effects and exposure effects identify two distinct channels through which geographic mobility affects family resources and child outcomes: the new, indirect channel operates through incomes; and the direct channel operates through prices including rents. We return to the relationship between exposure effects and rents in the next section.

Overall, incomes and rents explain between three and four fifths of the variance in counterfactual effects: if budget constraints are not the whole story behind child outcomes, they are nearly the whole story. Place effects are not mysterious.

One implication of our findings is that parents' geographic mobility matters far more for children's outcomes than previously understood. When indirect effects are orthogonal to exposure effects, their variances are additive. Taken together, the variance of counterfactual effects is more than double the variance of conventional exposure effects.

8.2 Opportunity Bargains Across Commuting Zones

Where should parents move? Chetty and Hendren (2018b) propose that parents should move not based on place effects, but instead based on "opportunity bargains," which they define as counties within commuting zones that have good child outcomes and low rents.

There is an unmet need for estimates that identify opportunity bargains across rather than only within commuting zones. Chetty and Hendren (2018b) do not define opportunity

bargains across CZs, because parents' labor market opportunities may vary across CZs. We define opportunity bargains across CZs as those that produce good child outcomes with low rents and (to account for the labor market) high parent family incomes.

Counterfactual effects identify candidate opportunity bargains across commuting zones; because they include both direct and indirect effects, they fully capture place effects on child outcomes, and they predict lower rents and higher parent family incomes.

We confirm opportunity bargains by comparing counterfactual effects with rents and parent rank premia. Comparing counterfactual effects with rents is necessary because they are only weakly correlated: Table 5, Column 2 reports that counterfactual effects and rents are uncorrelated at the 25th percentile of parent family income, at 0.011, and negatively correlated at the 75th percentile, at -0.311.[26] Thus, not all high-opportunity places have low rents, but both low-income and high-income families can find places with both high outcomes and low rents either through chance or through negative correlation. Comparing counterfactual effects with parent family income rank premia is less necessary because they are strongly positively correlated. Few places have both high counterfactual effects and low rank premia. Regardless, our estimated rank premia make comparisons possible.

There is no "best" opportunity bargain for every family, because families must choose how to weigh outcomes, rents, and incomes when making location choices. However, some CZs dominate others across all dimensions. For example, among the 50-most-populous CZs, Atlanta has the highest counterfactual effect, the highest parent rank premium, and median rents of 723 or 852 dollars (2000 USD) for below-median and above-median households respectively; by contrast, Miami has counterfactual effects ranked 47th or 48th out of 50, a parent rank premium at 48th out of 50, and median rents of 723 or 913 dollars.

What characterizes opportunity bargains across commuting zones? Using counterfactual effects instead of exposure effects makes it clear that, for outcomes as well as for costs, high family incomes matter as much as low rents. Consequently, many opportunity bargains are not in rural areas, but instead in CZs that contain suburbs and exurbs of major urban CZs; such CZs allow parents to save on rent while maintaining a high family income.[27] For Boston,

[26]The correlations between rents and indirect effects (not reported in Table 5) are 0.410 and 0.362 respectively for the 25th and 75th percentiles of parent family income rank. In Table 5, Column 3 and 7, we find that rents and exposure effects are similarly or more-negatively correlated across CZs, at -0.392 and -0.732 respectively. The contributions of indirect effects and exposure effects to the variance of counterfactual total effects are similar, so the rent correlations approximately offset each other at the 25th percentile, while the negative rent correlation with exposure effects dominates at the 75th percentile.

[27]Low-rent, high-income commuting zones are common. In Table 5, we find that the correlations between rents and parent family income rank premia are only 0.379 and 0.321 (Column 5), respectively, for low-income and high-income families. This occurs despite the fact that rents and the wage premia estimated by Card, Rothstein, and Yi (2025) are strongly correlated at 0.857 and 0.879 (Column 6), because the correlation between wage premia and rank premia is only 0.434 (Column 5). Wage premia and rank premia are weakly

rents are lower and family incomes and child outcomes are higher in Manchester; Providence has lower rents and similar incomes and child outcomes, qualifying as an unambiguous opportunity bargain. For New York, rents, incomes, and outcomes are better in Bridgeport and Poughkeepsie. For Philadelphia, Reading and Allentown are bargains. For Washington, D.C., Fredericksburg is a bargain; Baltimore has lower rents and similar incomes, but only has better outcomes at the 75th percentile of parent family income, therefore qualifying as a high-income opportunity bargain but not a low-income opportunity bargain. For Chicago, Gary and Kenosha are bargains. For Minneapolis, St. Cloud is a bargain. For Houston, Pearland is a bargain. For San Francisco, Sacramento is a bargain. For San Jose, Modesto is a bargain, with lower rents, similar incomes, and better outcomes.

Not all opportunity bargains are near major urban CZs. There are economic opportunities elsewhere. For example, for low-income families, natural resource extraction generates opportunity bargains comparable to Atlanta in Valdez, Alaska (the terminus and oil loading port for the Trans-Alaska Pipeline), Sitka, Alaska (a major fishing port), Williston, North Dakota (the center of the Williston Basin fracking boom), and Midland, Texas (the center of the Permian Basin fracking boom).

Opportunity bargains may involve tradeoffs along dimensions beyond child outcomes, rents, and parent family incomes. One important dimension is commute times. Theoretically, longer commute times may offset the lower rents of opportunity bargain CZs. Empirically, we find the opposite is true for cross-CZ opportunity bargains: the fraction of individuals with a commute under 15 minutes, obtained from Chetty and Hendren (2018b), is strongly negatively correlated with rents. Every suburban opportunity bargain described above has a higher fraction of short commutes than the neighboring major urban CZ, and every small-city or rural opportunity bargain described above has an even higher fraction of short commutes. Thus, commute times reinforce, rather than compensate, cross-CZ opportunity bargains.[28] Another important dimension is social capital. We find that the social capital index of Rupasingha and Goetz (2008), obtained from Chetty and Hendren (2018b), is negatively correlated with rents. Thus, amenities provided by civic, religious, charitable, recreational, and other organizations also reinforce cross-CZ opportunity bargains.

Overall, opportunity bargains are common, well-predicted by counterfactual effects, and readily confirmed with additional estimates of rents and parent family income rank premia.

correlated because they measure different things. Wage premia are measured for workers by establishment location, while rank premia are measured for families by residential location. Thus, rank premia capture additional effects on employment, self-employment, and cross-CZ commuting.

[28] This evidence is correlational, and the purpose of this paper is not to evaluate spatial equilibrium models. However, the low rents and high incomes we identify may reflect the fact that married couples can specialize to mitigate the worker-level tradeoff between rents, wages, and commute times.

9 Revealed Preference

Do parents value the commuting zone characteristics summarized by counterfactual effects and the underlying direct and indirect effects? In this section, we estimate parents' preferences over commuting zones, and factors related to their preferences.

9.1 Methods

We use flows of families between locations to identify parents' revealed preferences for places. Formally, we value locations by network centrality, computed as in Sorkin (2018): each move from one place to another is a vote for the destination, recursively weighted by the value of the origin. Our mapping of CZ-to-CZ flows to parents' preferences is similar to Sorkin's (2018) mapping of employer-to-employer flows to workers' preferences.[29]

To understand how the valuation works, it is worth stating the underlying equations of Sorkin's estimator. Suppose that families assign each location c an unobserved common value v_c, which drives their observable CZ-to-CZ moves. Denote the number of families moving to destination c' from origin c as $M_{c'c}$. The values are implicitly defined by the equation

$$v_{c'} = \frac{\sum_c M_{c'c} v_c}{\sum_c M_{cc'}} \tag{20}$$

where the value of destination CZ c' is the value of origin CZs c weighted by how many families move to the destination from each origin, normalized by how many families move out of the destination to all origins. Normalizing each weighted row total (entry, CZ c' as destination) by the corresponding column total (exit, CZ c' as origin) reduces the impact of gross flows on estimated values. This normalization differentiates Sorkin's approach from Google's PageRank (L. Page et al., 1999) algorithm which intends for gross flows, in that case the number of backlinks from one website to another, to influence the valuation.

We rely on Result 1 and 2 of Sorkin (2018) to compute our CZ values.[30] First, we aggregate moves in the mobility matrix $\boldsymbol{M}$, and we sum moves across all destinations (rows) for a given origin (column), in the diagonal matrix $\boldsymbol{S}$, where $S_{cc} = \sum_{c'} M_{c'c}$. Second, we solve by recursion for the unique fixed point $\boldsymbol{v}$ of the system of equations

$$\boldsymbol{S}^{-1}\boldsymbol{M}\boldsymbol{v} = \boldsymbol{v} \tag{21}$$

[29]The interpretation in our setting is simplified by the fact that families need to live somewhere, even if workers don't need to be employed anywhere.

[30]These propositions establish uniqueness of the solution and invariance to the number of families that do not move. Relevant to the first proposition, we verify that all CZs are in the strongly-connected set of CZs where at least one family moves in and at least one family moves out.

where each row corresponds to one of the rows in equation 20. Computationally, $\boldsymbol{v}$ is the first eigenvector of the matrix $\boldsymbol{S}^{-1}\boldsymbol{M}$.

Some simple algebra and examples may help to fix intuition. In the two-CZ example, it is clear that a CZ is only revealed preferred if it attracts more residents than it loses. The closed-form solution to the two-CZ problem, $\frac{v_1}{v_2} = \frac{M_{12}}{M_{21}}$, makes it clear that it is the ratio of the flows that determines the relative values, not the scale of the flows.

In the three-CZ example, the same intuition largely holds but is complicated by recursion. Consider the closed-form solution to the three-CZ problem

$$\boldsymbol{M} = \begin{bmatrix} M_{11} & M_{12} & M_{13} \\ M_{21} & M_{22} & M_{23} \\ M_{31} & M_{32} & M_{33} \end{bmatrix}, \quad \boldsymbol{v} = \begin{bmatrix} \frac{M_{12}M_{13}+M_{12}M_{23}+M_{13}M_{32}}{M_{12}M_{31}+M_{21}M_{32}+M_{31}M_{32}} \\ \frac{M_{12}M_{21}+M_{21}M_{23}+M_{23}M_{31}}{M_{12}M_{31}+M_{21}M_{32}+M_{31}M_{32}} \\ 1 \end{bmatrix}$$

where $v_3 = 1$ is a scale normalization.

Example 1 Our first three-CZ example shows that in a migration network with zero net flows, gross flows don't affect the values. Population and migration rates don't matter, in themselves. In this case, CZs 1 and 2 have greater gross flows but the larger numerator from each row is offset by a larger denominator from each column:

$$\boldsymbol{M} = \begin{bmatrix} 0 & 2 & 1 \\ 2 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}, \quad \boldsymbol{v} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

Example 2 Our second three-CZ example shows that CZs with equal net flows may be valued differently due to their position in the migration network. In this case, CZs 1 and 2 have equal net flows (+1) but CZ 2 has a higher value because its flows come from higher-valued CZ 1 while CZ 1's flows come from lower-valued CZ 3:

$$\boldsymbol{M} = \begin{bmatrix} 0 & 1 & 3 \\ 2 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}, \quad \boldsymbol{v} = \begin{bmatrix} \frac{7}{4} \\ \frac{9}{4} \\ 1 \end{bmatrix}$$

Example 3 Our last three-CZ example shows that the CZ with the most-negative net flow is not necessarily the lowest-valued CZ, again due to its position in the migration network. In this case, CZ 1 has the most-negative net flow (-3) but out-ranks CZ 3 because more of

its flows come from top-ranked CZ 2:

$$\boldsymbol{M} = \begin{bmatrix} 0 & 2 & 1 \\ 4 & 0 & 4 \\ 2 & 1 & 0 \end{bmatrix}, \quad \boldsymbol{v} = \begin{bmatrix} \frac{11}{10} \\ \frac{14}{5} \\ 1 \end{bmatrix}$$

In summary, the revealed preference ranking is insensitive to the scale of the flows, but depends strongly on the ratio and origin of the flows.

9.2 Place Values

We compute parents' revealed preference separately for above-median and below-median income families, and rank the values for the 50 most-populous CZs in Figure 25. We find that among large urban CZs, New York is most-preferred: under the ratio interpretation in Sorkin (2018), below-median-income parents are about 15 percent more likely to choose New York than the next-most-preferred CZ, San Jose, in a pairwise contest; above-median-income parents are about 72 percent more likely to choose New York than San Jose. Los Angeles, Newark, and Chicago are all highly preferred, but lower-ranked than New York and San Jose. These five CZs are in the top 10 among all CZs for below-median-income parents, and in the top 60 for above-median-income parents. However, not all large urban CZs are preferred: for example, Orlando, Nashville, and Charlotte rank at the bottom of the 50 most-populous CZs, and in the bottom 80 and 70 among all CZs, respectively.

We also map the revealed preference values in Figure 26. The places that parents prefer are generally in California, the Northeast Corridor between Washington, D.C. and Boston, the Great Plains, Alaska, and Hawaii; and along the Great Lakes, the lower Mississippi, and the Rio Grande. The places that parents are less likely to choose include central Florida, southern Appalachia, the Ozarks, inland Wisconsin and Minnesota, the northern Rocky Mountains, and the Great Basin Desert around Las Vegas. Finally, values are more uniform for below-median-income parents than for above-median-income parents.

9.3 What Do Parents Value?

Parents do not appear to value places that generate good outcomes for their children. Table 5, Column 1, correlates place values with a range of mobility measures and covariates. Strikingly, we find that revealed preference values are negatively correlated with counterfactual total effects, and exposure effects, at the 25th percentile (Panel A) and the 75th percentile (Panel B) of parent family income. We also find that place values are weakly correlated with

observed outcomes, suggesting that parents who move to highly valued places are positively selected, or value positively selected peers.[31]

Parents appear to value places based on their own earnings potential, rather than income in itself. Table 5, Column 1, presents evidence that values are positively correlated with the CZ-level wage premia reported in Card, Rothstein, and Yi (2025). Positive correlations with wage premia combined with weak negative correlations with rank premia suggest that rank premia, which capture the employment, self-employment, and commuting margins, capture something that parents value negatively; that something may be the cost of employment, the cost of commuting, or perhaps the cost of a dispersed suburban community.

Places that parents value also have high rents. This is not surprising, given that rents are in themselves a revealed preference measure. But the correlations are not very strong, revealing the obvious: high rents are also a disamenity.

Do parents value the impact of places on child outcomes? It is unclear whether the negative correlation of counterfactual effects with place values captures parents' feelings about child outcomes, or just the other factors that are correlated with child outcomes, such as rents, wages, family incomes, commute times, or population density.

To understand how much of the negative correlation between counterfactual effects (and other place effects) and place values is due to preferences over other characteristics, in Table 6 we regress place values on place effects, controlling for other characteristics to isolate preferences for child outcomes. Column 1, 4, and 7 present the univariate regression coefficients; Column 2, 5, and 8 present the coefficients after controlling for the parent family income rank premium, wage premium, and median rent; and column 3, 6, and 9 add controls for the fraction of individuals with a commute under 15 minutes, and for the natural log of population density. All controls enter with the expected signs: parents prefer places with high wages, avoid places with high incomes conditional on wages (which may indicate disutility from working or commuting), and avoid places with high rents; they also prefer places with short commutes and high population density.

Parents do not appear to dislike places because they generate positive outcomes for children, but they do not appear to like such places, either. Adding the first set of controls for incomes, wages, and rents reduces the negative coefficients on counterfactual effects and on exposure effects by a factor of 2.5 up to 5.3; adding the second set of controls for commutes and density reduces the negative coefficients by a factor of 3.9 up to 5.6. The coefficients generally remain significant, but given that explained variance scales in the square of the

[31] We thank Jeffrey Smith for the observation that our findings, at the CZ scale, complement past findings that parents sort to neighborhoods and schools where students have high test scores, but not those that are effective according to value added (e.g. Rothstein, 2006; Bayer, Ferreira, and McMillan, 2007; Imberman and Lovenheim, 2016; MacLeod and Urquiola, 2019; Abdulkadiroğlu et al., 2020).

coefficient, they show that the influence of child outcomes on parent preferences is at least 15 times smaller than the uncontrolled coefficients suggest. Therefore, child outcomes account for 1 percent or less of the variance in place values.

These revealed preference estimates offer both good news and bad news for policies intended to move families to opportunity. The good news is that places with great child outcomes are not so popular that congestion would get in the way of policy implementation. The bad news is that not all opportunity bargains are ones that parents are willing to accept: in particular, opportunity bargains that offer greater counterfactual effects are not attractive on their own, and the fact that many such opportunity bargains derive from greater rank premia is not attractive, either. The more promising opportunity bargains arise from lower rents, where lower rents do not imply higher wage premia, but these are few.

We conclude that other policies, apart from recommendations about where to move, might encounter less resistance from parents and might be more effective. Such policies might include place-based or family-based tax credits and transfers, initiatives to increase housing supply and lower rents in high-wage CZs, and transit improvements to speed up commutes and make high-wage CZs more accessible to commuters.

10 Conclusion

In this paper, we show that usable place effects on child outcomes can be estimated, as an input to recommendations about where families should move. We provide the first estimates that capture valid counterfactuals, meaning children's likely outcomes after following recommendations. And we find that despite the fact that most places that generate good outcomes are bargains for the families that move there, families with children tend to move elsewhere. Nonetheless, there are three optimistic ways to read our results:

First, many families do move to opportunity bargains, and more information about child outcomes may persuade others who were not persuaded by low rents or high incomes. Further, we identify a large number of opportunity bargains so that families may be given a menu of options, some of which they may prefer for idiosyncratic reasons even when they do not value opportunity bargains in general.

Second, not all families need to move to achieve good results, and moving too many families may be counterproductive: moving many families to the same location may affect prices and family budget constraints, which appear to determine child outcomes. Fortunately, with many opportunity bargains, it is not necessary to recommend unsustainable migration to a small number of places. Further, moving willing families out of bad places may make bad places better for families that do not move, by lowering rents.

Finally, families may not need to move. If high incomes and low prices create opportunity, many alternative policies should succeed. Raising incomes or lowering rents need not be costly: proposals as varied as *Abundance* (Klein and Thompson, 2025) and the *Strategic Plan 2026-2030* (U.S. Department of the Treasury, 2026) seek to do so by removing constraints such as housing supply restrictions (e.g. E. L. Glaeser and Gyourko, 2003 and 2018; Quigley and Raphael, 2004 and 2005; Hsieh and Moretti, 2019).

11 References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2020). Sampling-based vs. design-based uncertainty in economics and econometrics. *Econometrica*, *88*(2), 265–296 (cit. on pp. 11, 17).

Abdulkadiroğlu, A., Pathak, P. A., Schellenberg, J., & Walters, C. R. (2020). Do parents value school effectiveness? *American Economic Review*, *110*(5), 1502–1539 (cit. on p. 38).

Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, *67*(2), 251–333 (cit. on p. 27).

Aizer, A., Eli, S., Ferrie, J., & Lleras-Muney, A. (2016). The long-run impact of cash transfers to poor families. *American Economic Review*, *106*(4), 935–971 (cit. on p. 27).

Akee, R. K. Q., Copeland, W. E., Keeler, G., Angold, A., & Costello, E. J. (2010). Parents' incomes and children's outcomes: A quasi-experiment using transfer payments from casino profits. *American Economic Journal: Applied Economics*, *2*(1), 86–115 (cit. on p. 27).

Alesina, A., Hohmann, S., Michalopoulos, S., & Papaioannou, E. (2021). Intergenerational mobility in Africa. *Econometrica*, *89*(1), 1–35 (cit. on p. 16).

Alexander, J. T., Bleckley, D., Fisher, J., Genadek, K., Leonard, S. H., & Magganas, A. (2025). Newly-available individual-level U.S. tax data from 1969-1994. *National Tax Journal*, *78*(4) (cit. on p. 5).

Aloni, T., & Avivi, H. (2025). One land, many promises: Assessing the consequences of unequal childhood location effects (Working Paper). (Cit. on pp. 9, 16).

Altonji, J. G., & Mansfield, D. (2018). Estimating group effects using averages of observables. *American Economic Review*, *108*(4-5), 974–1005 (cit. on p. 21).

Andrews, I., Kitagawa, T., & McCloskey, A. (2024). Inference on winners. *Quarterly Journal of Economics*, *139*(1), 305–358 (cit. on pp. 1, 14, 15, 24).

Angrist, J. D., & Pischke, J.-S. (2009). Mostly harmless econometrics: An empiricist's companion. Princeton University Press. (Cit. on p. 18).

Bailey, M. J., Hoynes, H. W., Rossin-Slater, M., & Walker, R. (2024). Is the social safety net a long-term investment? large-scale evidence from the Food Stamps program. *Review of Economic Studies*, *91*(3), 1291–1330 (cit. on p. 27).

Barr, A., Eggleston, J., & Smith, A. A. (2022). Investing in infants: The lasting effects of cash transfers to new families. *The Quarterly Journal of Economics*, *137*(4), 2539–2583 (cit. on p. 27).

Bayer, P., Ferreira, F., & McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, *115*(4), 588–638 (cit. on p. 38).

Becker, G. S., & Tomes, N. (1986). Human capital and the rise and fall of families. *Journal of Labor Economics*, *4*(3, Part 2), S1–S39 (cit. on p. 1).

Beekers, L. (2024, July). Sources of regional variation in intergenerational mobility: Evidence from the Netherlands (CentER Discussion Paper No. 2024-015). Tilburg University, CentER. (Cit. on p. 16).

Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L., & Palmer, C. (2024). Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. *American Economic Review*, *114*(5), 1281–1337 (cit. on pp. 1, 24).

Britto, D. G. C., Fonseca, A., Pinotti, P., Sampaio, B., & Warwar, L. (2025). Do CCTs create conditions to thrive? Bolsa Família and social mobility in Brazil (tech. rep. No. 2025/655-1). UNU-WIDER. (Cit. on p. 27).

Britto, D. G. C., Fonseca, A., Pinotti, P., Sampaio, B., & Warwar, L. (2022). Intergenerational mobility in the land of inequality (IZA Discussion Paper No. 15611). IZA Institute of Labor Economics. (Cit. on pp. 16, 21).

Cáceres Bravo, M. M. (2018). `regressby`: Super fast grouped regressions in Stata (fork). (Cit. on p. 11).

Card, D., Rothstein, J., & Yi, M. (2025). Location, location, location. *American Economic Journal: Applied Economics*, *17*(1), 297–336 (cit. on pp. 27, 33, 38, 49–51).

Chetty, R., Friedman, J., Hendren, N., Jones, M., & Porter, S. (2026a). The opportunity atlas: Mapping the childhood roots of social mobility. *American Economic Review*, *116*(1), 1–51 (cit. on pp. 8, 12, 16, 21, 24, 25).

Chetty, R., & Hendren, N. (2018a). The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects. *Quarterly Journal of Economics*, *133*(3), 1107–1162 (cit. on pp. 8, 10, 12, 16, 18, 21, 27).

Chetty, R., & Hendren, N. (2018b). The impacts of neighborhoods on intergenerational mobility II: County-level estimates. *Quarterly Journal of Economics*, *133*(3), 1163–1228 (cit. on pp. 1–5, 7–10, 12–15, 18, 21, 22, 24, 25, 31, 32, 34, 47, 49–51, 53, 54).

Chetty, R., Hendren, N., & Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment. *American Economic Review*, *106*(4), 855–902 (cit. on pp. 1, 9, 27).

Chetty, R., Hendren, N., Kline, P., & Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the United States. *Quarterly Journal of Economics*, *129*(4), 1553–1623 (cit. on pp. 6–8, 24, 25).

Cholli, N. A., Durlauf, S. N., Landersø, R., & Navarro, S. (2024, October). Understanding the heterogeneity of intergenerational mobility across neighborhoods (NBER Working Paper No. 33035). National Bureau of Economic Research. (Cit. on p. 21).

Chyn, E., Collinson, R., & Sandler, D. (2025). The long-run effects of America's largest residential racial desegregation program: Gautreaux (tech. rep.). SSRN. (Cit. on p. 1).

Chyn, E., & Katz, L. F. (2021). Neighborhoods matter: Assessing the evidence for place effects. *Journal of Economic Perspectives*, *35*(4), 197–222 (cit. on p. 1).

Cinelli, C., Forney, A., & Pearl, J. (2024). A crash course in good and bad controls. *Sociological Methods & Research*, *53*(3), 1071–1104 (cit. on p. 18).

Combes, P.-P., Duranton, G., & Gobillon, L. (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics*, *63*(2), 723–742 (cit. on p. 27).

Correia, S. (2016, March). A feasible estimator for linear models with multi-way fixed effects. (Cit. on p. 11).

Dauth, W., Findeisen, S., Moretti, E., & Suedekum, J. (2022). Matching in cities. *Journal of the European Economic Association*, *20*(4), 1478–1521 (cit. on p. 27).

De la Roca, J., & Puga, D. (2017). Learning by working in big cities. *The Review of Economic Studies*, *84*(1), 106–142 (cit. on p. 27).

Deutscher, N. (2020). Place, peers, and the teenage years: Long-run neighborhood effects in Australia. *American Economic Journal: Applied Economics*, *12*(2), 220–249 (cit. on pp. 16, 21).

Dustmann, C., Landersø, R., & Andersen, L. (2024). Unintended consequences of welfare cuts on children and adolescents (tech. rep.). SSRN. (Cit. on p. 27).

Eshaghnia, S. M. S. (2023, August). Is Zip Code destiny? (Working paper). University of Chicago. (Cit. on p. 18).

Farber, H. S., & Gibbons, R. (1996). Learning and wage dynamics. *The Quarterly Journal of Economics*, *111*(4), 1007–1047 (cit. on p. 28).

Gallagher, R. M., Kaestner, R., & Persky, J. J. (2019). The geography of family differences and intergenerational mobility. *Journal of Economic Geography*, *19*(3), 589–618 (cit. on pp. 18, 21).

Glaeser, E., & Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, *32*(1), 3–30 (cit. on p. 40).

Glaeser, E. L., & Gyourko, J. (2003). The impact of building restrictions on housing affordability. *Economic Policy Review*, *9*(2), 21–39 (cit. on p. 40).

Heckman, J. J., & Landersø, R. (2022). Lessons for Americans from Denmark about inequality and social mobility. *Labour Economics*, *77*, 101999 (cit. on p. 18).

Hilger, N. G. (2016). Parental job loss and children's long-term outcomes: Evidence from 7 million fathers' layoffs. *American Economic Journal: Applied Economics*, *8*(3), 247–283 (cit. on p. 27).

Hsieh, C.-T., & Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, *11*(2), 1–39 (cit. on p. 40).

Imberman, S. A., & Lovenheim, M. F. (2016). Does the market value value-added? Evidence from housing prices after a public release of school and teacher value-added. *Journal of Urban Economics*, *91*, 104–121 (cit. on p. 38).

Klein, E., & Thompson, D. (2025, March). Abundance. Avid Reader Press. (Cit. on p. 40).

Kline, P. (2025, August). Branching fixed effects: A proposal for communicating uncertainty (Working paper). University of California, Berkeley. (Cit. on p. 12).

Klopfer, J., & Miller, C. (2024). Using SSNs to identify immigration, selection, and censoring in administrative data (tech. rep. No. 2024/009). HKU Jockey Club Enterprise Sustainability Global Research Institute. (Cit. on pp. 5, 46).

Laliberté, J.-W. (2021). Long-term contextual effects in education: Schools and neighborhoods. *American Economic Journal: Economic Policy*, *13*(2), 336–377 (cit. on p. 16).

Loury, G. C. (1981). Intergenerational transfers and the distribution of earnings. *Econometrica*, *49*(4), 843–867 (cit. on p. 1).

MacLeod, W. B., & Urquiola, M. (2019). Is education consumption or investment? Implications for school competition. *Annual Review of Economics*, *11*, 563–589 (cit. on p. 38).

Mazumder, B. (2005). Fortunate sons: New estimates of intergenerational mobility in the United States using social security earnings data. *Review of Economics and Statistics*, *87*(2), 235–255 (cit. on p. 7).

Mogstad, M., Romano, J., Shaikh, A., & Wilhelm, D. (2024). Inference for ranks with applications to mobility across neighbourhoods and academic achievement across countries. *Review of Economic Studies*, *91*(1), 476–518 (cit. on pp. 1, 3, 11–13, 15, 24, 25).

Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The PageRank citation ranking: Bringing order to the web (tech. rep. No. 1999-66). Stanford InfoLab. (Cit. on pp. 4, 35).

Page, M. E. (2024). New advances on an old question: Does money matter for children's outcomes? *Journal of Economic Literature*, *62*(3), 891–947 (cit. on p. 1).

Pei, Z., Pischke, J.-S., & Schwandt, H. (2017). Poorly measured confounders are more useful on the left than on the right (tech. rep. No. 23232). National Bureau of Economic Research. (Cit. on p. 17).

Quigley, J. M., & Raphael, S. (2004). Is housing unaffordable? why isn't it more affordable? *Journal of Economic Perspectives*, *18*(1), 191–214 (cit. on p. 40).

Quigley, J. M., & Raphael, S. (2005). Regulation and the high cost of housing in california. *American Economic Review*, *95*(2), 323–328 (cit. on p. 40).

Robbins, H. (1956). An empirical Bayes approach to statistics. In J. Neyman (Ed.), Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability (pp. 157–163, Vol. 1). University of California Press. (Cit. on p. 22).

Rosenbaum, P. R. (1984). The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society: Series A (General)*, *147*(5), 656–666 (cit. on p. 18).

Rothbaum, J. (2016). Sorting and geographic variation in intergenerational mobility, U.S. Census Bureau, Social, Economic, and Housing Statistics Division. (Cit. on p. 21).

Rothstein, J. (2006). Good principals or good peers? Parental valuation of school characteristics, Tiebout equilibrium, and the incentive effects of competition among jurisdictions. *American Economic Review*, *96*(4), 1333–1350 (cit. on p. 38).

Rupasingha, A., & Goetz, S. J. (2008). US county-level social capital data, 1990–2005. (Cit. on p. 34).

Sampson, R. J. (2008). Moving to inequality: Neighborhood effects and experiments meet social structure. *American Journal of Sociology*, *114*(1), 189–231 (cit. on p. 18).

Sikhova, A. (2023). Understanding the effect of parental education and financial resources on the intergenerational transmission of income. *Journal of Labor Economics*, *41*(3), 771–811 (cit. on p. 27).

Soria-Espín, J., & Medina, O. (2025, May). Intergenerational mobility in Spain: Geographic analysis and causal neighborhood effects (Working paper). World Inequality Lab. (Cit. on p. 21).

Sorkin, I. (2018). Ranking firms using revealed preference. *Quarterly Journal of Economics*, *133*(3), 1331–1393 (cit. on pp. 4, 35, 37).

Sprung-Keyser, B., & Porter, S. (2026, May). The geography of lifecycle human capital accumulation (NBER Working Paper No. 35203). National Bureau of Economic Research. (Cit. on p. 27).

Stepner, M., Droste, M., & Townsend, R. M. (2018). `regressby`: Stata module for fast grouped regressions. (Cit. on p. 11).

Tan, H. R. (2026). Shifting the land of opportunity: The changing role of local labor markets. (Cit. on p. 27).

Todd, P. E., & Wolpin, K. I. (2003). On the specification and estimation of the production function for cognitive achievement. *The Economic Journal*, *113*(485), F3–F33 (cit. on p. 18).

U.S. Department of the Treasury. (2026). Strategic plan 2026-2030 (Strategic Plan). U.S. Department of the Treasury. Washington, DC. (Cit. on p. 40).

Wodtke, G. T., Harding, D. J., & Elwert, F. (2011). Neighborhood effects in temporal perspective. *American Sociological Review*, *76*(5), 713–736 (cit. on p. 1).

12 Tables

Table 1: Summary Statistics

	(1) Total	(2) Stayers	(3) Movers
N	76,790,109 (100.0%)	46,917,409 (61.1%)	29,060,530 (37.8%)
Parent Income	68,322 (362,048)	70,759 (392,008)	65,188 (312,609)
Parent Rank	50.047 (28.856)	51.736 (28.596)	47.756 (29.090)
Child Income	46,018 (2,001,774)	47,752 (2,543,439)	43,705 (379,589)
Child Rank	50.000 (28.870)	51.153 (28.889)	48.374 (28.753)
Filer Age SSN Issued	14.576 (16.211)	14.994 (17.924)	14.000 (13.165)
Filer Age Child Born	28.713 (7.918)	29.211 (8.043)	27.986 (7.673)
Child Birth Year	1988.473 (5.917)	1988.320 (5.915)	1988.605 (5.914)
Child Age SSN Issued	1.856 (2.851)	1.886 (2.884)	1.839 (2.816)
Child Age Linked	6.931 (5.304)	7.035 (5.311)	6.859 (5.293)
Between-CZ Moves	0.869 (1.483)	0.000 (0.000)	2.069 (1.281)
Age at First Move	8.381 (6.423)	. (.)	8.500 (6.451)

Summary statistics for analysis populations used to estimate exposure effects and observed outcomes by CZ. Column 1 reports the population of children who (1) were born from 1978-98, (2) had a valid Social Security Number (SSN) by age 18, (3) were never recorded as a non-citizen from 2014-25, (4) were alive as of 2025, (5) were observed by the IRS at least once from 2022-25, and (6) were claimed as a child dependent at least once from 1994-2016; Column 2 reports the sub-population of "stayers" whose parents were observed in only one CZ from child age 0-23, and Column 3 reports the sub-population of "movers" whose parents were observed in more than one CZ from child age 0-23, excluding children who moved more than 6 times. Parent income is the average pre-tax income across both parents' tax units and all years from child age 0-23. Child income is the average from the year each child turns 24 through 2024. Parent rank and child rank are defined within birth cohorts. Filer age SSN issued is defined based on the primary filer parent's birth month and SSN issue month from the crosswalk in Klopfer and Miller (2024). Filer age child born is defined based on the primary filer's birth year and the child's birth year. Child age SSN issued is defined based on the child's birth month and SSN issue month. Child age linked is based on the child's birth year and the year the child was first claimed as a dependent. Between-CZ moves reports the number of times a child is observed in a new CZ, from age 1-23. Age at first move reports the child's age in the year the child was first observed in a new CZ. See Section 2 for detailed population and variable definitions.

Table 2: Variance Components of Exposure Effect Estimates

	Total		By Population							
			1-50		51-200		201-595		596-761	
	$p = 25$ (1)	$p = 75$ (2)	$p = 25$ (3)	$p = 75$ (4)	$p = 25$ (5)	$p = 75$ (6)	$p = 25$ (7)	$p = 75$ (8)	$p = 25$ (9)	$p = 75$ (10)
Panel A: This Paper										
Total Variance	0.0110	0.0129	0.0062	0.0089	0.0047	0.0051	0.0224	0.0249	0.1093	0.1086
Noise Variance	0.0007	0.0012	0.0001	0.0001	0.0004	0.0005	0.0019	0.0030	0.0095	0.0214
Signal Variance	0.0103	0.0117	0.0061	0.0088	0.0043	0.0046	0.0204	0.0219	0.0998	0.0872
Signal to Noise Ratio	13.93	9.41	66.80	99.85	10.65	9.14	10.52	7.25	10.55	4.07
Number of CZs	755	754	50	50	150	150	395	395	160	159
Panel B: Chetty and Hendren (2018b)										
Total Variance	0.0780	0.0989	0.0169	0.0223	0.0477	0.0538	0.3629	0.4700		
Noise Variance	0.0517	0.0793	0.0075	0.0063	0.0381	0.0465	0.2417	0.4137		
Signal Variance	0.0263	0.0196	0.0095	0.0160	0.0096	0.0073	0.1212	0.0563		
Signal to Noise Ratio	0.51	0.25	1.27	2.55	0.25	0.16	0.50	0.14		
Number of CZs	595	595	50	50	150	150	395	395		

Variance components of exposure effect estimates by CZ. We estimate exposure effects based on the sequence fixed effects specification in equation (2), for the population of children who moved up to six times, reported in Table 1. We drop move sequences with too few observations to estimate the fixed effects (N=21,848,889). Estimation is carried out as in equation (3) and (4). We calculate rank-specific exposure effects and their variances by linear combination of the main exposure effect and the interaction with parent family income rank, demeaned using CZ exposure weights, as in equations (5) and (6). Panel A reports variance components based on our exposure effect estimates and standard errors. Column 1 and 9 exclude 6 CZs, and column 2 and 10 exclude 7 CZs, for which exposure effect confidence intervals for 24 years of exposure exceed 50 percentile ranks from the mean. The total variance and noise variance are estimated as in equations (7) and (8) as the mean of squared exposure effect estimates and the mean of squared standard errors respectively, weighted by exposure shares. The signal variance is the total variance minus the noise variance estimated in equation (9). Variance component estimates for sub-populations in columns 3-10 share the same national mean, and sum to the variance component estimates for the total population in columns 1-2, if weighted proportionately. Panel B reports variance components based on the exposure effect estimates and standard errors from Chetty and Hendren (2018b), `online_table3-2.dta`. Exposure effect estimates are `causal_p25_czkr26` and `causal_p75_czkr26`, and standard errors are `causal_p25_czkr26_bs` and `causal_p75_czkr26_bs` from coefficient and variance estimates demeaned using population shares from the 2000 census. Because the authors do not report exposure shares, the total variance and noise variance are weighted by population shares for comparability with our estimates; the weighted correlation between exposure shares and population shares is $r = 0.98$. Our calculation is different from the calculation in Chetty and Hendren (2018b) which is weighted by the inverse variance of the estimates. As the authors observe, weighting by population shares gives larger variance estimates than weighting by the inverse variance, because weighting by the inverse variance down-weights observations with the highest estimated variance.

Table 3: Variance Components of Observed Outcomes

	Observed Outcomes $\hat{\boldsymbol{\mu}}_p^{A0}$		Exposure Effects $\hat{\boldsymbol{\mu}}_p^{A1}$	Selection Bias $\hat{\boldsymbol{\mu}}_p^{A0} - \hat{\boldsymbol{\mu}}_p^{A1}$		Covariance $\hat{\boldsymbol{\mu}}_p^{A0} - \hat{\boldsymbol{\mu}}_p^{A1}, \hat{\boldsymbol{\mu}}_p^{A1}$		Selection Share $(\hat{\boldsymbol{\mu}}_p^{A0} - \hat{\boldsymbol{\mu}}_p^{A1})/\hat{\boldsymbol{\mu}}_p^{A0}$	
	Stayers (1)	Movers (2)	Movers (3)	Stayers (4)	Movers (5)	Stayers (6)	Movers (7)	Stayers (8)	Movers (9)
$p = 25$	19.520	13.363	6.361	12.059	8.053	0.550	-0.525	0.618	0.413
$p = 75$	17.176	12.177	7.434	11.592	7.210	-0.925	-1.233	0.675	0.420

Variance components of observed outcomes by CZ. Exposure effects estimated as in Table 2 (N=21,848,889). We estimate observed outcomes (exposure effects without sequence fixed effects) as in equation (12), for the population of children who moved up to six times (Movers), using the same estimation sample as the fixed effects specification; and for the population of children who did not move (Stayers), using their full population (N=46,917,409). For all specifications and populations, we calculate rank-specific exposure effects and their variances by linear combination of the main exposure effect and the interaction with parent family income rank, demeaned using CZ exposure weights, as in equations (5) and (6). The first row shows results for below-median parent family income rank $p = 25$, and the second row shows results for above-median parent family income rank $p = 75$. We estimate selection bias in column 4-5 by the difference between observed outcomes in column 1-2 with exposure effects in column 3; we estimate their covariances in column 6-7; and we estimate the share of selection bias in the variance in column 8-9 by the values in column 4-5 divided by the values in column 1-2. We find that both for below-median and above-median parents, selection bias explains the majority of the variance in observed outcomes for stayers, and a near-majority for movers.

Table 4: Explained Variance in Observed Outcomes

	Explained Variance (R^2)					Exposure-Observed Correlation (ρ)			
	Observed Outcomes $\hat{\boldsymbol{\mu}}_p^{A0}$		Exposure Effects $\hat{\boldsymbol{\mu}}_p^{A1}$	Selection Bias $\hat{\boldsymbol{\mu}}_p^{A0} - \hat{\boldsymbol{\mu}}_p^{A1}$		Observed Uncorrected $\hat{\boldsymbol{\mu}}_p^{A0}$		Observed Corrected $\tilde{\boldsymbol{\mu}}_p^{A0}$	
	Stayers (1)	Movers (2)	Movers (3)	Stayers (4)	Movers (5)	Stayers (6)	Movers (7)	Stayers (8)	Movers (9)
$p = 25$	0.684	0.711	0.528	0.716	0.738	0.620	0.633	0.395	0.401
$p = 75$	0.544	0.596	0.658	0.422	0.466	0.576	0.652	0.557	0.587

Explained variance in observed outcomes by CZ. We report R^2 from exposure-weighted regressions of observed outcomes, exposure effects, and the selection bias term on CZ-level covariates drawn from Card, Rothstein, and Yi (2025; CRY) and Chetty and Hendren (2018b; CH): from CRY, share with some college education or more; from CH, share black, share foreign born, share in the labor force, unemployment rate, share married, share divorced, share single mother, violent crime rate, and total crime rate. Observed outcomes in column 1-2, exposure effects in column 3, and the selection bias term in column 4-5 are as defined in Table 3. For the regressions in column 1-5 we omit CZs with missing covariates (N=663). We find that the covariates explain somewhere between half and two thirds of the variance in observed outcomes, exposure effects, and the selection bias term. We then residualize observed outcomes on the covariates to remove selection bias as suggested in the literature, and then correlate observed outcomes in column 6-7 (N=755 for $p = 25$ and N=754 for $p = 75$) and corrected observed outcomes in column 8-9 (N=663) with our exposure effects to show how they perform as proxies for the causal variation measured by exposure effects. We find that the correlation between exposure effects and observed outcomes is greater for the uncorrected observed outcomes than for the regression-corrected observed outcomes.

Table 5: Pairwise Correlations of Key Measures

	This Paper					External		Chetty and Hendren (2018b)		
	Revealed Pref. $\hat{\boldsymbol{v}}_p$	Counterf. Effects $\hat{\boldsymbol{\mu}}_p^{A4}$	Exposure Effects $\hat{\boldsymbol{\mu}}_p^{A1}$	Observed Outcomes $\hat{\boldsymbol{\mu}}_p^{A0,S}$	Rank Premium $\hat{\boldsymbol{\rho}}$	Wage Premium $\hat{\boldsymbol{\psi}}$	Median Rent $\bar{\boldsymbol{r}}_p$	Exposure Effects $\hat{\boldsymbol{\mu}}_p^{A1}$	Observed Outcomes $\hat{\boldsymbol{\mu}}_p^{A0,S}$	Forecast Effects $\boldsymbol{\mu}_p^{f}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: $p = 25$										
KM: $\hat{\boldsymbol{v}}_p$	1.000									
KM: $\hat{\boldsymbol{\mu}}_p^{A4}$	-0.350	1.000								
KM: $\hat{\boldsymbol{\mu}}_p^{A1}$	-0.395	0.751	1.000							
KM: $\hat{\boldsymbol{\mu}}_p^{A0,S}$	0.156	0.615	0.620	1.000						
KM: $\hat{\boldsymbol{\rho}}$	-0.159	0.751	0.145	0.296	1.000					
Ext: $\hat{\boldsymbol{\psi}}$	0.662	0.049	-0.385	0.226	0.434	1.000				
Ext: $\bar{\boldsymbol{r}}_p$	0.517	0.011	-0.392	0.185	0.379	0.857	1.000			
CH: $\hat{\boldsymbol{\mu}}_p^{A1}$	-0.115	0.368	0.470	0.377	0.087	-0.151	-0.199	1.000		
CH: $\hat{\boldsymbol{\mu}}_p^{A0,S}$	0.015	0.608	0.708	0.923	0.163	0.002	-0.047	0.460	1.000	
CH: $\boldsymbol{\mu}_p^{f}$	-0.154	0.681	0.738	0.820	0.254	-0.070	-0.112	0.604	0.924	1.000
Panel B: $p = 75$										
KM: $\hat{\boldsymbol{v}}_p$	1.000									
KM: $\hat{\boldsymbol{\mu}}_p^{A4}$	-0.415	1.000								
KM: $\hat{\boldsymbol{\mu}}_p^{A1}$	-0.379	0.737	1.000							
KM: $\hat{\boldsymbol{\mu}}_p^{A0,S}$	-0.007	0.693	0.576	1.000						
KM: $\hat{\boldsymbol{\rho}}$	-0.255	0.719	0.087	0.403	1.000					
Ext: $\hat{\boldsymbol{\psi}}$	0.433	-0.148	-0.611	-0.172	0.434	1.000				
Ext: $\bar{\boldsymbol{r}}_p$	0.311	-0.311	-0.732	-0.405	0.321	0.879	1.000			
CH: $\hat{\boldsymbol{\mu}}_p^{A1}$	0.024	0.249	0.332	0.301	0.028	-0.160	-0.250	1.000		
CH: $\hat{\boldsymbol{\mu}}_p^{A0,S}$	-0.024	0.570	0.651	0.887	0.158	-0.314	-0.554	0.334	1.000	
CH: $\boldsymbol{\mu}_p^{f}$	-0.037	0.611	0.663	0.881	0.203	-0.293	-0.539	0.393	0.982	1.000

Pairwise correlations of key measures of place preferences, place effects, and covariates by CZ. Column 1: parents' revealed preference valuation for each CZ, estimated as the first eigenvector of the row-normalized migration matrix for the population of movers (N=21,848,889). Column 2: counterfactual total effects, estimated based on exposure-corrected latent parent family income rank as in equation 18. Columns 3-4: exposure effects for movers and observed outcomes for stayers as defined in Table 3. Column 5: parent family income rank premia estimated as in equation 16. Column 6: workers' wage premium for each CZ from Card, Rothstein, and Yi (2025). Column 7: median rent for below- and above-median income households reported from the 2000 decennial census, via Chetty and Hendren (2018b; CH). Columns 8-10: exposure effects for movers, observed outcomes for stayers, and forecast exposure effects from CH. Correlations weighted by exposures of the movers population. Panel A: 25th percentile of parent family income (or below-median family income) except for columns 5-6 which do not depend on parent family income. Panel B: 75th percentile of parent family income (or above-median family income).

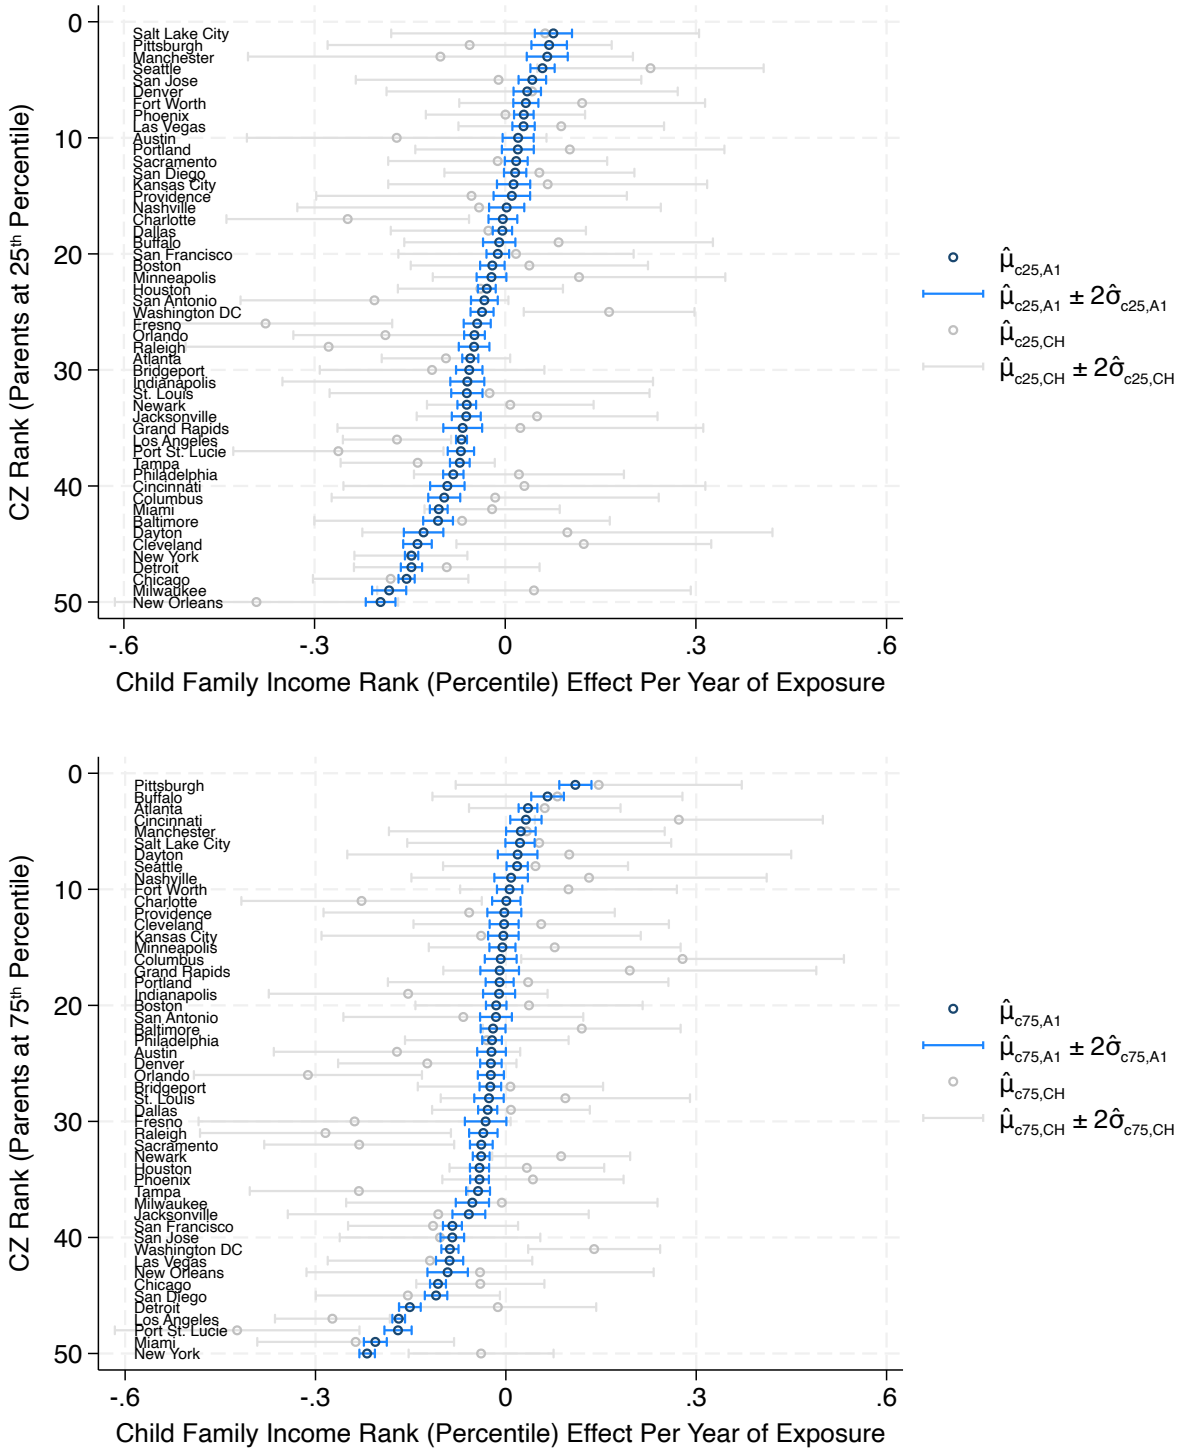
Table 6: Predictors of Revealed Preference

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: $p = 25$									
Counterf. Effect $\hat{\mu}_{cp}^{A4}$	-0.0271 (0.0028)	-0.0071 (0.0029)	-0.0070 (0.0033)						
Exposure Effect $\hat{\mu}_{cp}^{A1}$				-0.0463 (0.0040)	-0.0087 (0.0029)	-0.0083 (0.0035)			
Observed Outcome $\hat{\mu}_{cp}^{A0,S}/24$							0.0152 (0.0024)	0.0088 (0.0015)	0.0060 (0.0017)
Rank Premium $\hat{\rho}_c$		-0.0170 (0.0017)	-0.0177 (0.0016)		-0.0194 (0.0011)	-0.0201 (0.0009)		-0.0208 (0.0010)	-0.0206 (0.0009)
Wage Premium $\hat{\psi}_c$		0.0964 (0.0044)	0.0859 (0.0039)		0.0962 (0.0044)	0.0864 (0.0040)		0.0957 (0.0043)	0.0807 (0.0041)
Median Rent $\bar{r}_{cp}/1000$		-0.0155 (0.0037)	-0.0044 (0.0032)		-0.0161 (0.0037)	-0.0047 (0.0032)		-0.0143 (0.0036)	-0.0076 (0.0033)
p(travel<15 mins.)			0.0657 (0.0042)			0.0659 (0.0042)			0.0540 (0.0048)
ln(pop. density)			0.0045 (0.0003)			0.0044 (0.0003)			0.0048 (0.0003)
N	690	690	690	690	690	690	690	690	690
adj. R^2	0.115	0.644	0.750	0.165	0.646	0.750	0.053	0.658	0.753
Panel B: $p = 75$									
Counterf. Effect $\hat{\mu}_{cp}^{A4}$	-0.0395 (0.0032)	-0.0139 (0.0055)	-0.0089 (0.0050)						
Exposure Effect $\hat{\mu}_{cp}^{A1}$				-0.0464 (0.0043)	-0.0185 (0.0054)	-0.0104 (0.0049)			
Observed Outcome $\hat{\mu}_{cp}^{A0,S}/24$							0.0107 (0.0040)	0.0384 (0.0035)	0.0230 (0.0036)
Rank Premium $\hat{\rho}_c$		-0.0200 (0.0029)	-0.0220 (0.0026)		-0.0244 (0.0016)	-0.0250 (0.0014)		-0.0313 (0.0015)	-0.0288 (0.0013)
Wage Premium $\hat{\psi}_c$		0.1180 (0.0070)	0.0906 (0.0062)		0.1182 (0.0070)	0.0911 (0.0062)		0.0935 (0.0069)	0.0784 (0.0063)
Median Rent $\bar{r}_{cp}/1000$		-0.0348 (0.0045)	-0.0182 (0.0040)		-0.0374 (0.0045)	-0.0191 (0.0041)		-0.0045 (0.0043)	-0.0061 (0.0039)
p(travel<15 mins.)			0.0900 (0.0062)			0.0898 (0.0062)			0.0666 (0.0069)
ln(pop. density)			0.0068 (0.0004)			0.0067 (0.0005)			0.0060 (0.0005)
N	690	690	690	690	690	690	690	690	690
adj. R^2	0.183	0.453	0.607	0.146	0.457	0.608	0.009	0.531	0.628

Regressions of parents' revealed preference valuation for each CZ, estimated as the first eigenvector of the row-normalized migration matrix for the population of movers (N=21,848,889), on intergenerational mobility measures and explanatory covariates. The main coefficient of interest is the conditional effect of mobility measures on revealed preference valuations. Column 1-3: counterfactual total effects. Column 4-6: exposure effects. Column 7-9: observed outcomes. The first column for each measure is the univariate regression. The second column adds controls for the parent family income rank premium, workers' wage premium for each CZ from Card, Rothstein, and Yi (2025), and median rent for below- and above-median income households reported from the 2000 decennial census, via Chetty and Hendren (2018b; CH); these are the key explanatory variables for place preferences and migration, summarizing income and prices. The third column adds controls for commute distance and for population density, two key amenities that may trade off with rent and earnings. CZ observations weighted by exposures of the movers population, excluding CZs for which covariates are missing in any regression. Panel A: 25th percentile of parent family income (or below-median income). Panel B: 75th percentile of parent family income (or above-median income).

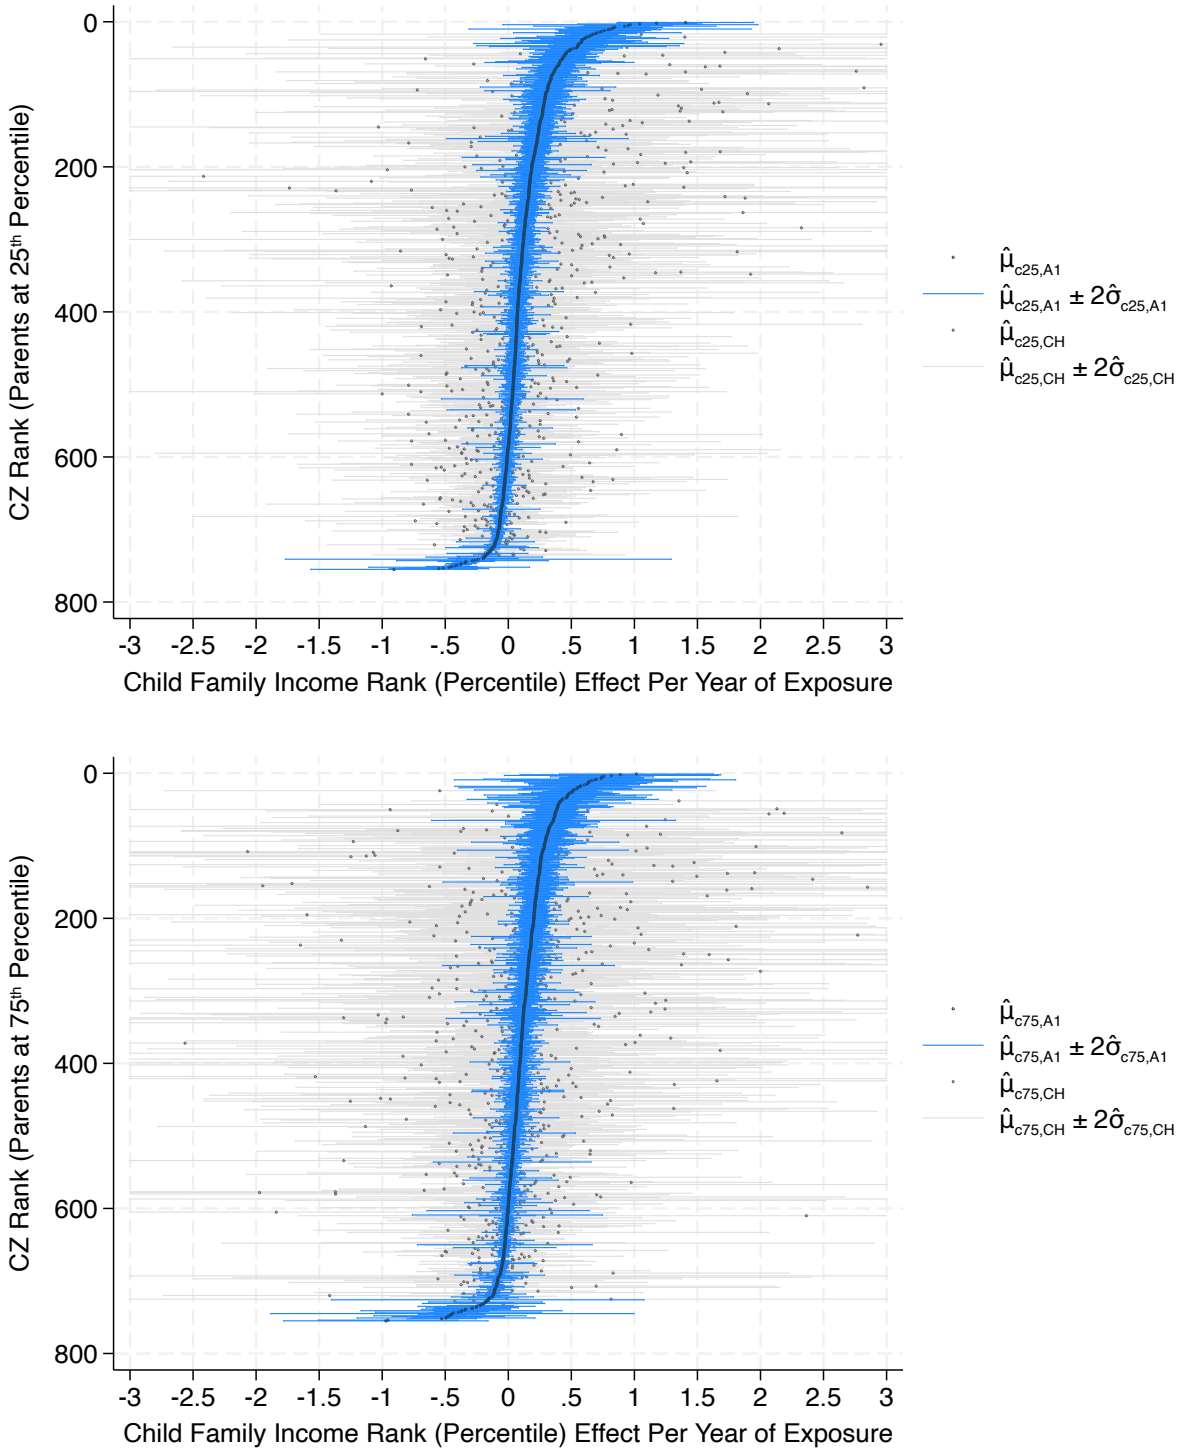
13 Figures

Figure 1: Exposure Effects and Confidence Intervals, Top 50 CZs by Population



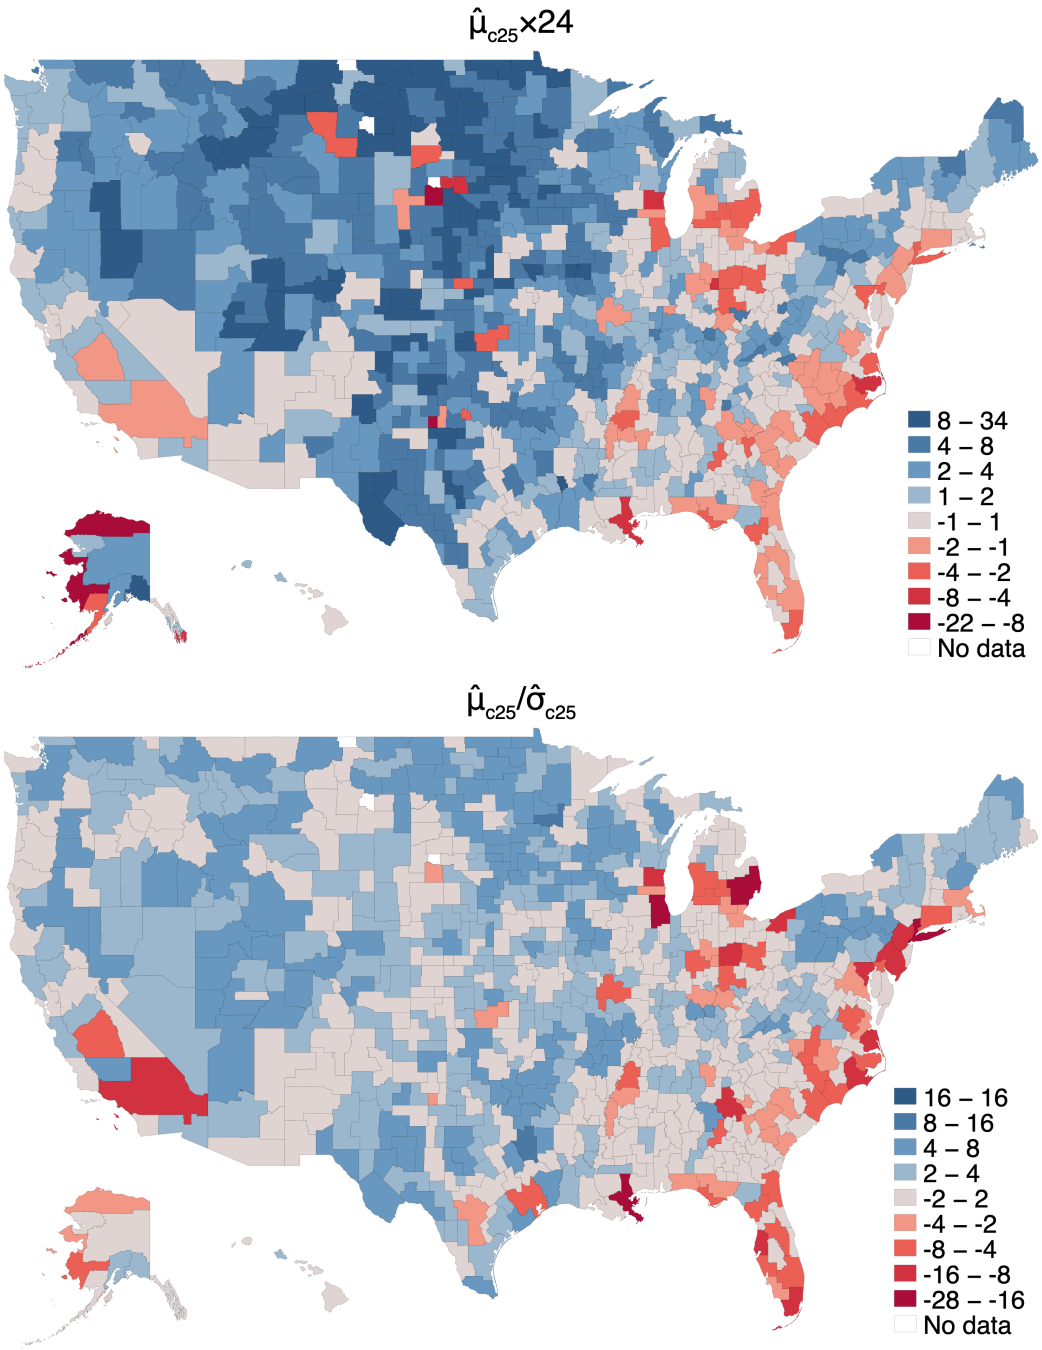
Effects of one year of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889). Background: estimates in gray are the comparable exposure effect estimates and standard errors from Chetty and Hendren (2018b). The vertical axes plot the rank of each CZ among the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. Top: effect for children at the 25th percentile of parent family income. Bottom: effect for children at the 75th percentile of parent family income.

Figure 2: Exposure Effects and Confidence Intervals, All CZs



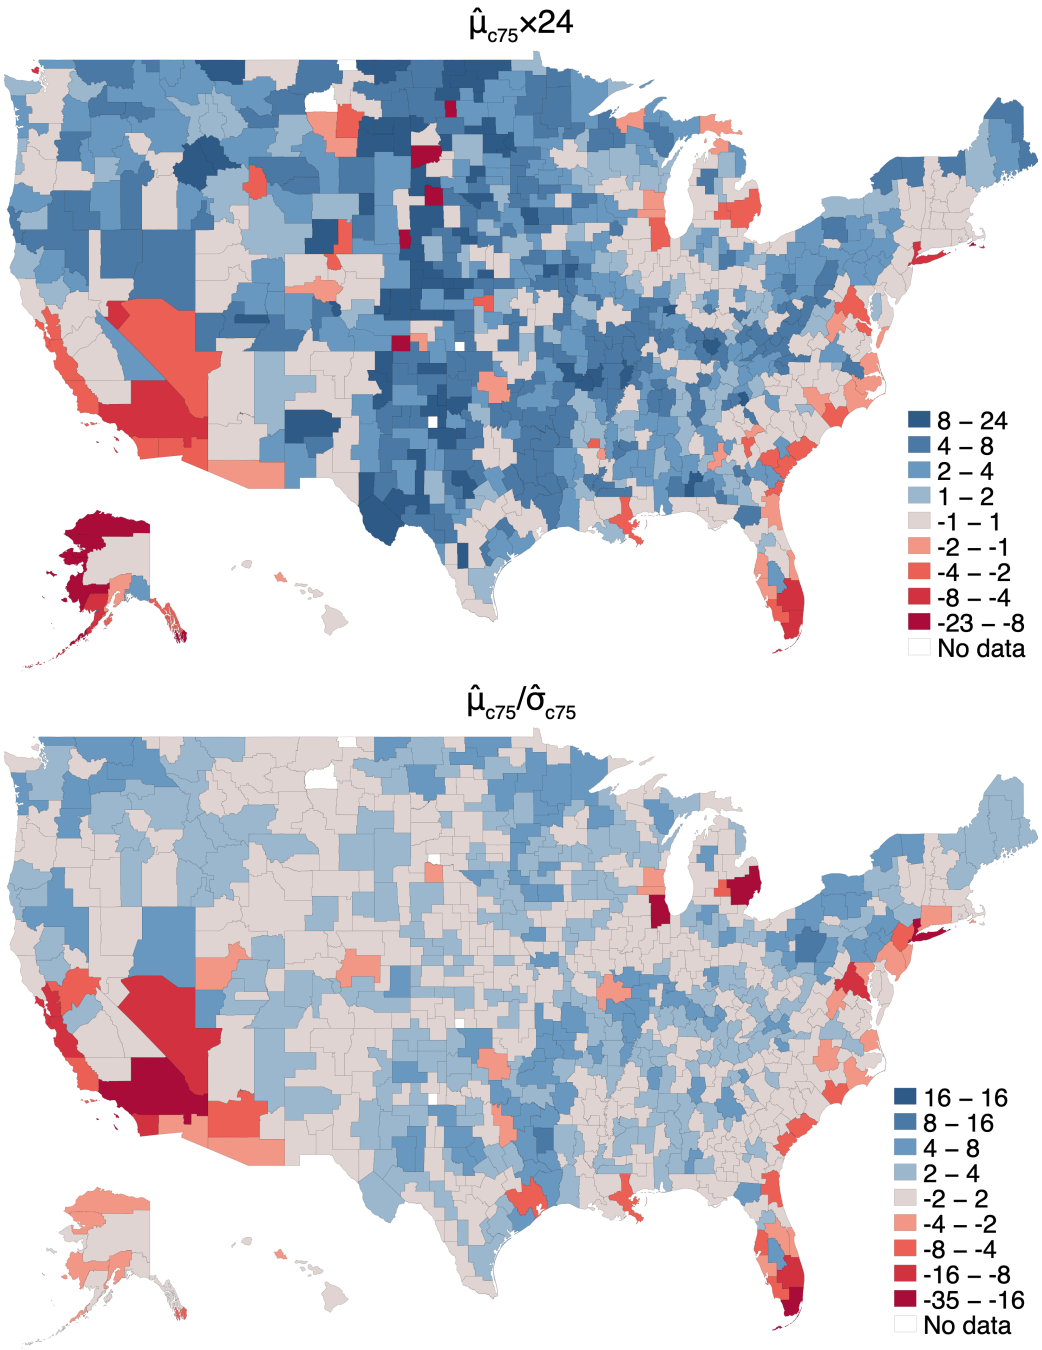
Effects of one year of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889). Background: estimates in gray are the comparable exposure effect estimates and standard errors from Chetty and Hendren (2018b), covering the 595 most populous CZs out of 761; estimates and confidence intervals falling outside the display range are truncated. The vertical axes plot the rank of all CZs such that the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean (all CZs in Table 2, Column 1-2). Top: effect for children at the 25th percentile of parent family income. Bottom: effect for children at the 75th percentile of parent family income.

Figure 3: Exposure Effects on Children's Income by Commuting Zone, $p = 25$



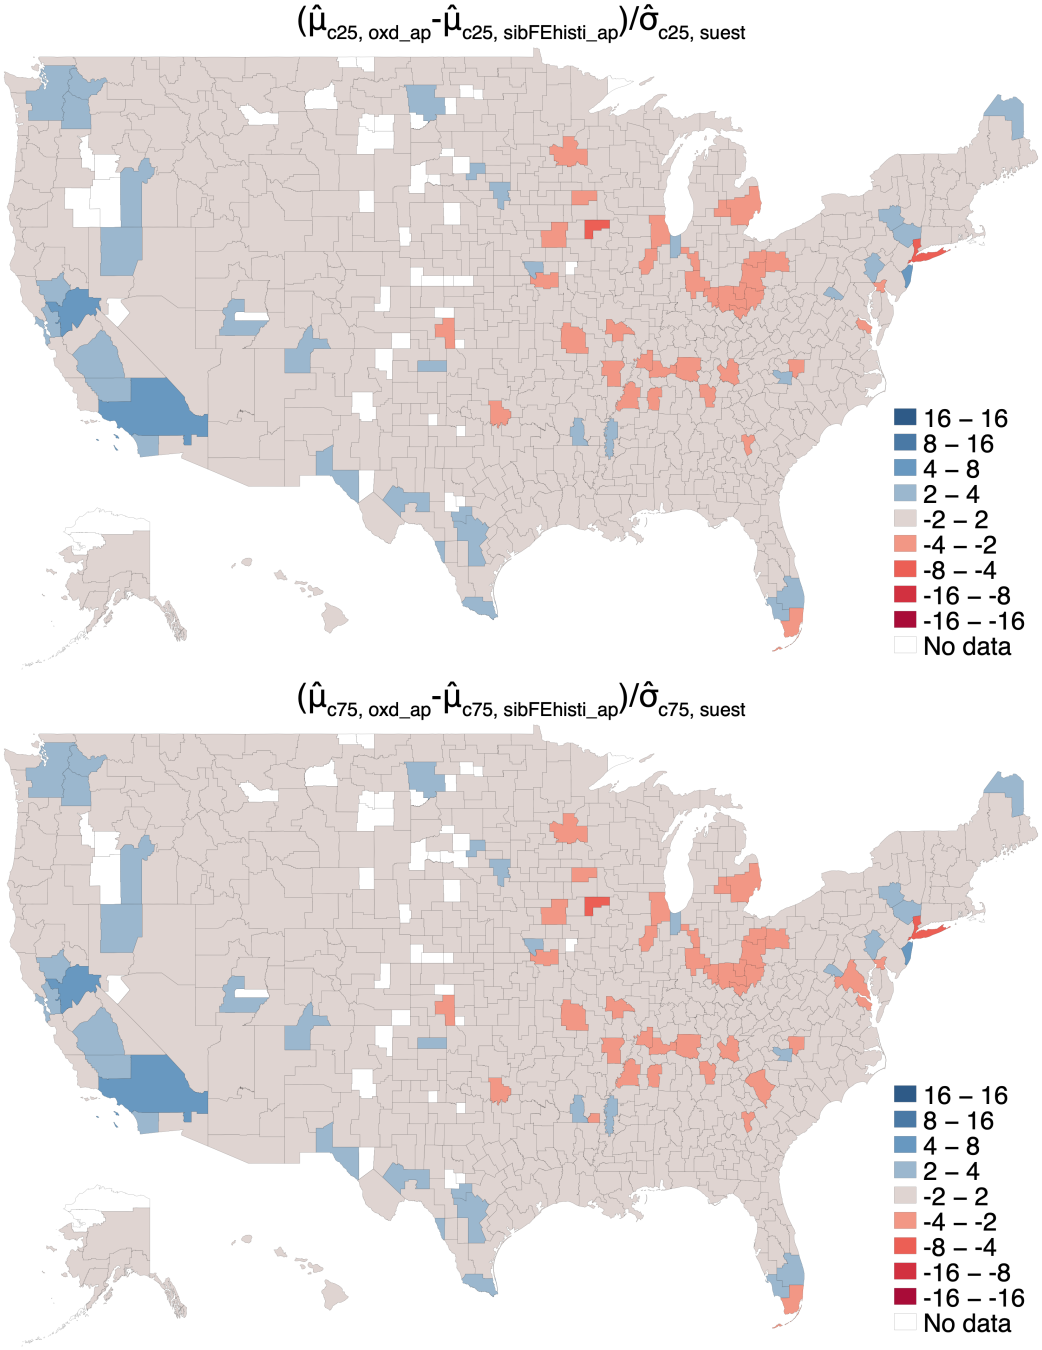
Effects of 24 years of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889). Estimates displayed only if the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean (all CZs in Table 2, Column 1-2). Overseas territories CZs omitted. Top: effect for children at the 25th percentile of parent family income. Bottom: t-statistic of effect relative to national mean.

Figure 4: Exposure Effects on Children's Income by Commuting Zone, $p = 75$



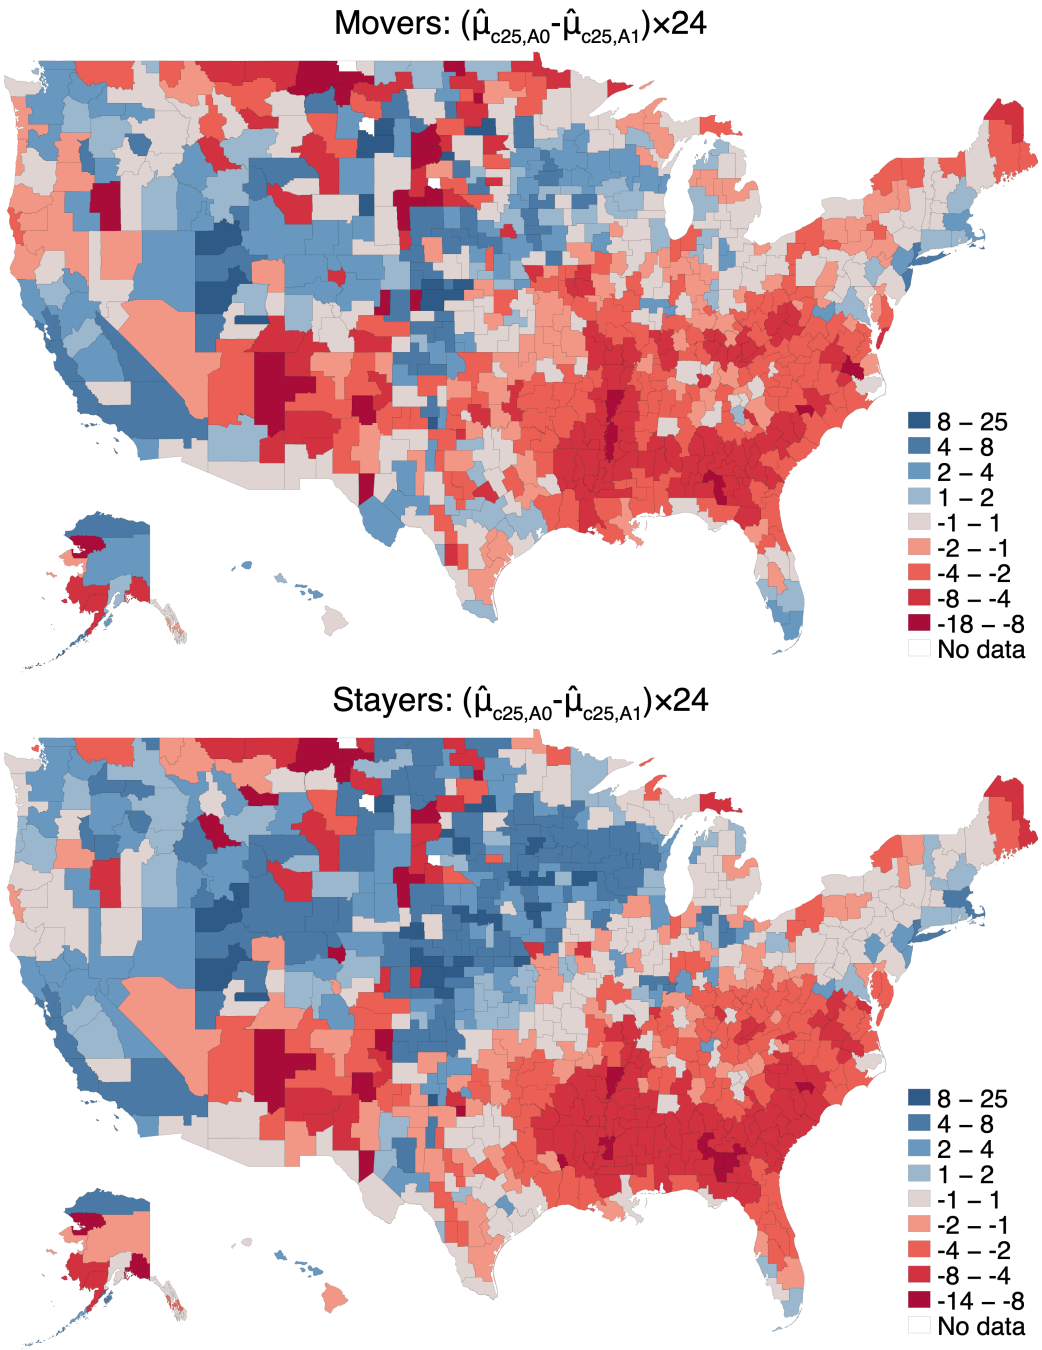
Effects of 24 years of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889). Estimates displayed only if the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean (all CZs in Table 2, Column 1-2). Overseas territories CZs omitted. Top: effect for children at the 75th percentile of parent family income. Bottom: t-statistic of effect relative to national mean.

Figure 5: Comparison of Estimators A1 and A2 by CZ: Very Rare Significant, Location-Specific Selection in Estimator A1



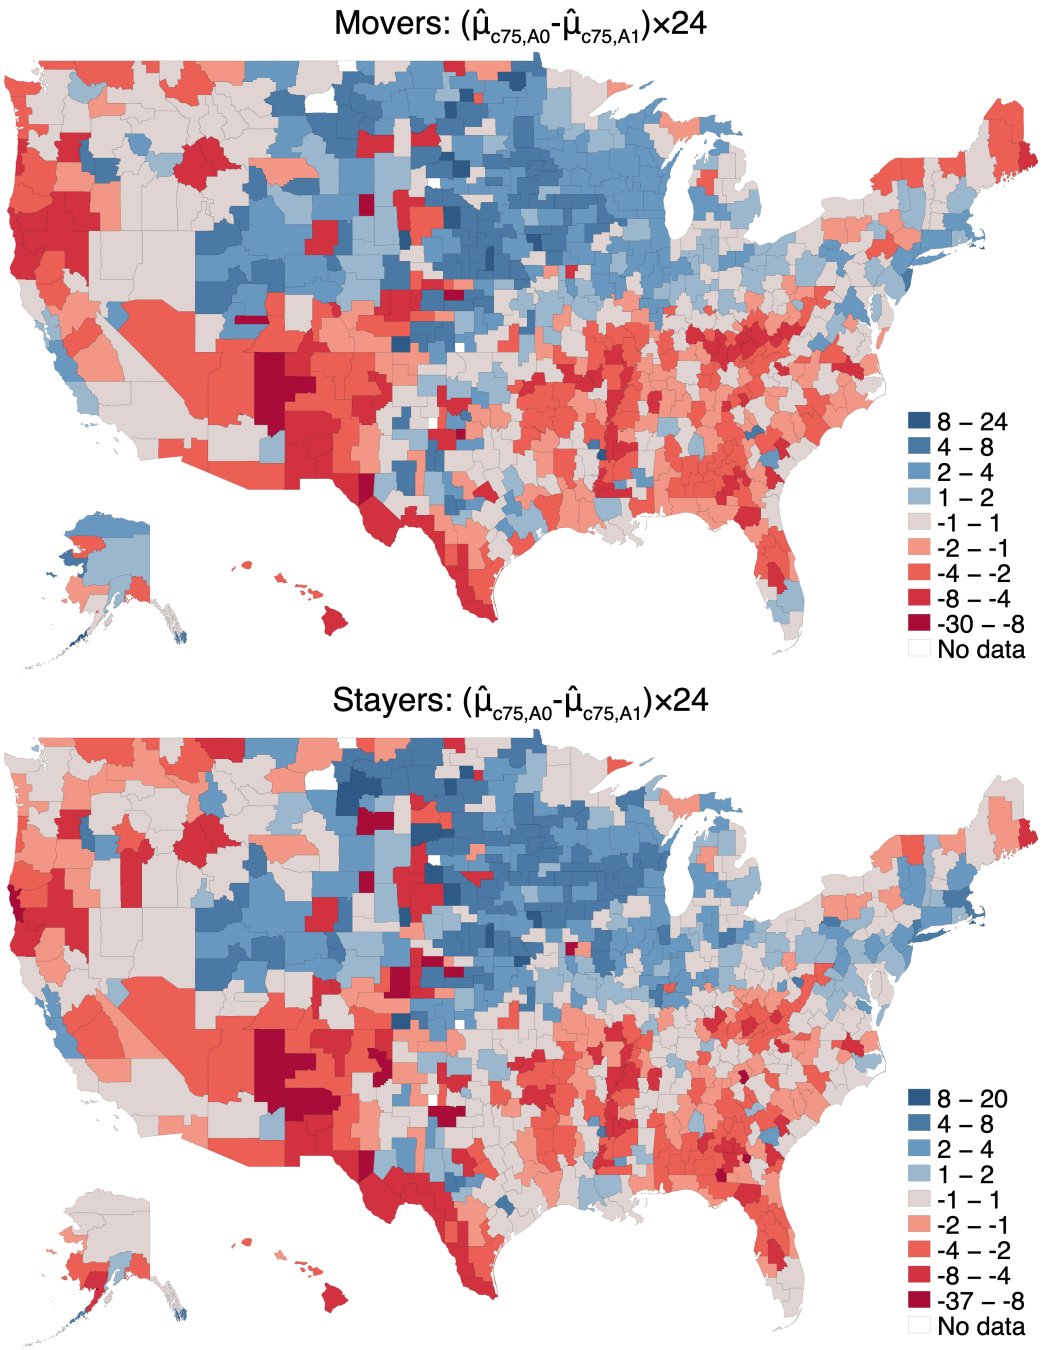
t-statistic of the difference in point estimates for equation (2) (Estimator A1) and (10) (Estimator A2), replacing parent family income rank with the within-family mean rank; standard errors derived by seemingly-unrelated regression. Initial estimation as in Table 2, on a common estimation population; we drop move sequences with too few observations to estimate the sequence fixed effects and parent income rank interactions, and family-by-sequence fixed effects, in Estimator A2 (N=10,618,414). Estimates displayed only if the confidence interval of the difference in effects of 24 years of exposure is contained within 50 percentile ranks of the national mean. Overseas territories CZs omitted. Top: t-statistic for children at the 25th percentile of parent family income. Bottom: t-statistic for children at the 75th percentile of parent family income.

Figure 6: Selection Bias in Observed Outcomes by Commuting Zone, $p = 25$



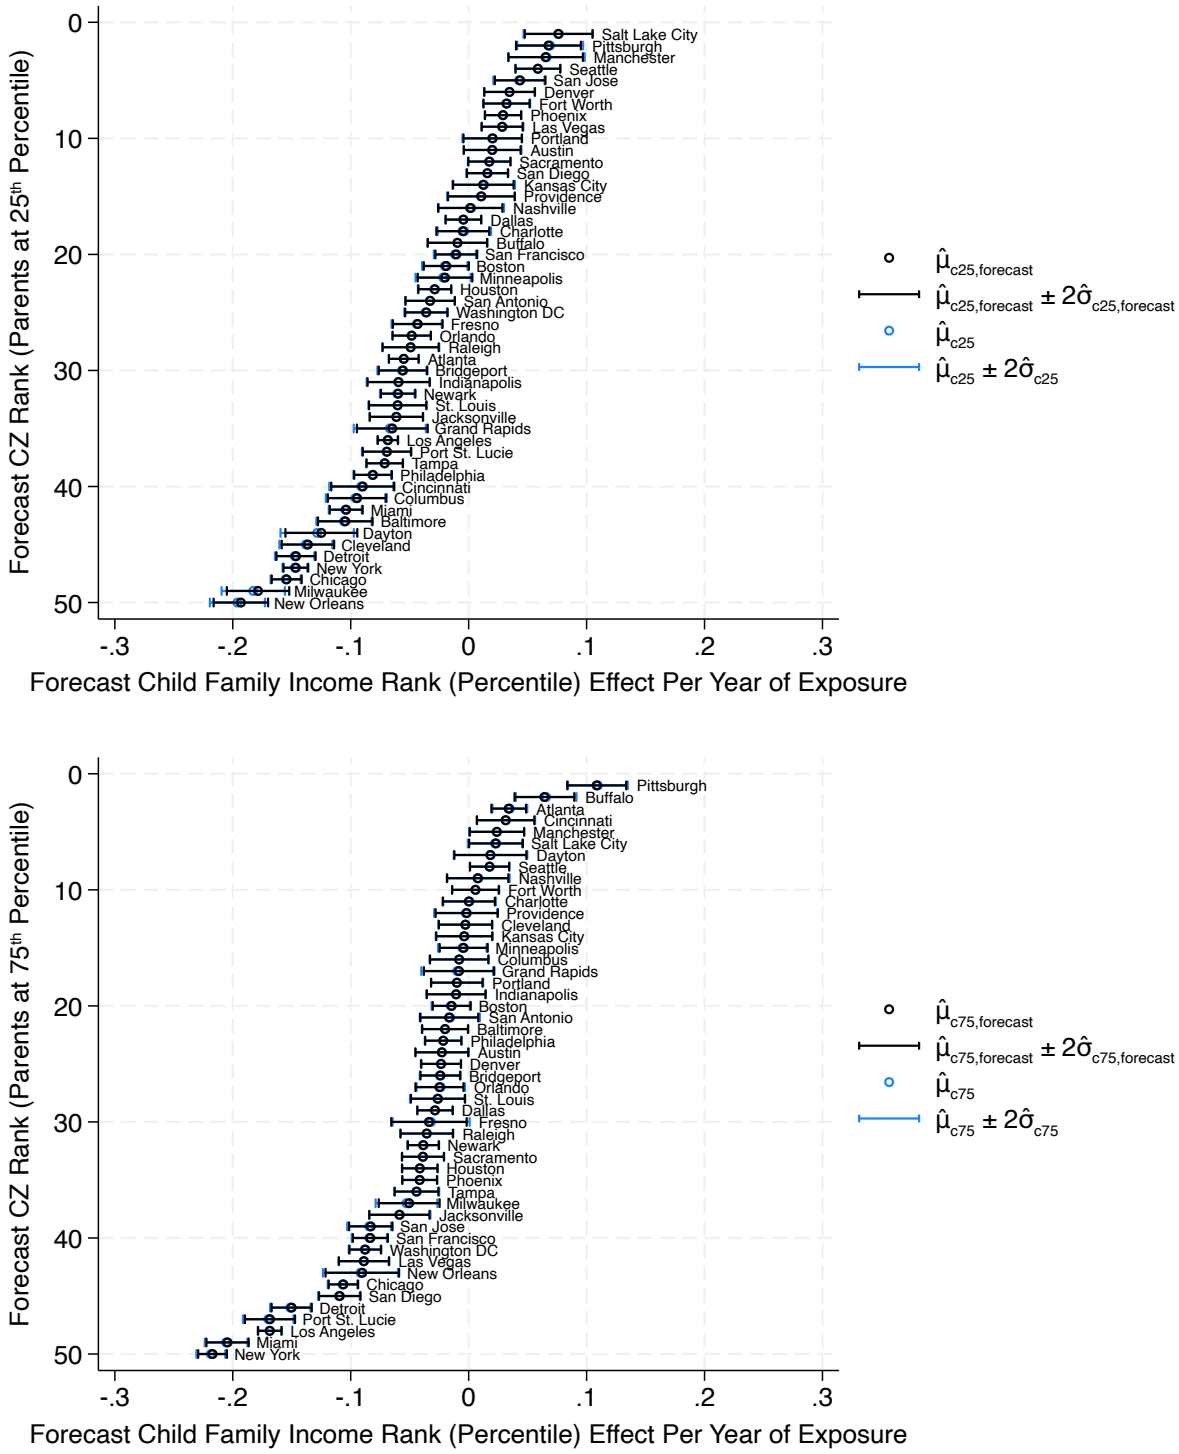
Difference in point estimates for equation (12) (Estimator A0) observed outcomes and (2) (Estimator A1) exposure effects, scaled up to effect of 24 years of exposure, for children at the 25th percentile of parent family income. Exposure effects estimated as in Table 2 (N=21,848,889). Observed outcomes estimated as in Table 3, on a common estimation population with exposure effects for movers, or on the population of stayers (N=46,917,409). Estimates displayed only if the confidence interval of the difference in effects of 24 years of exposure is contained within 50 percentile ranks of the national mean. Overseas territories CZs omitted. Top: Difference for movers. Bottom: Difference for stayers.

Figure 7: Selection Bias in Observed Outcomes by Commuting Zone, $p = 75$



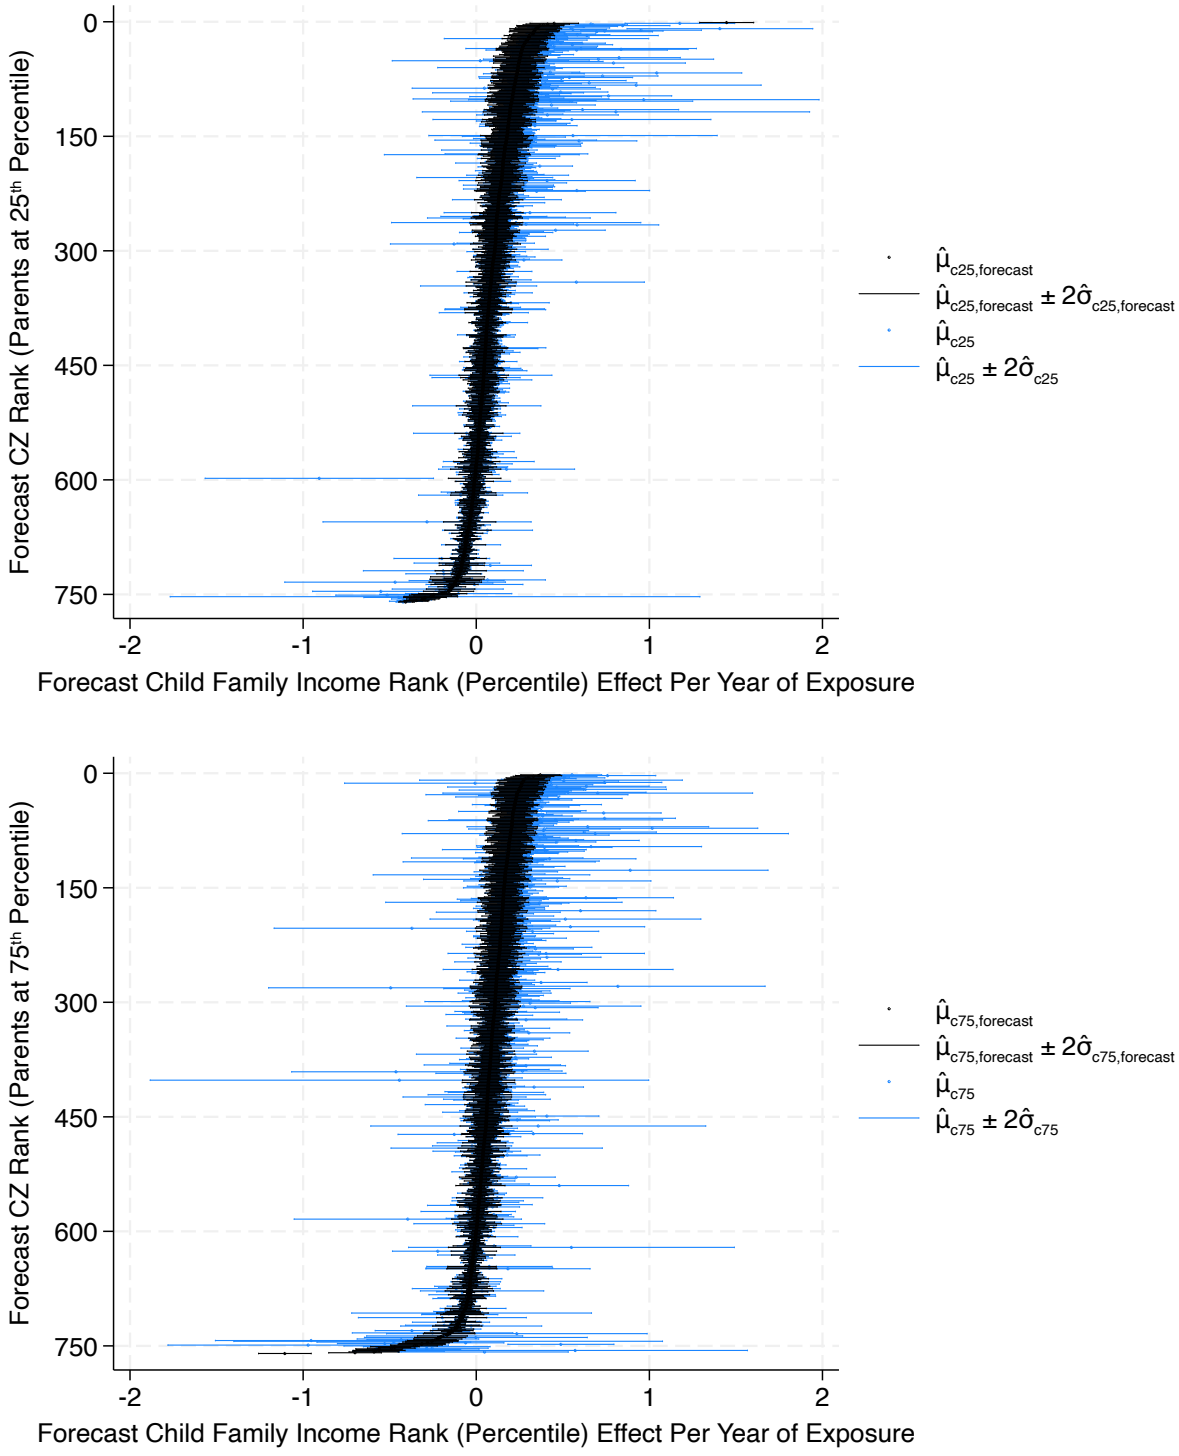
Difference in point estimates for equation (12) (Estimator A0) observed outcomes and (2) (Estimator A1) exposure effects, scaled up to effect of 24 years of exposure, for children at the 75th percentile of parent family income. Exposure effects estimated as in Table 2 (N=21,848,889). Observed outcomes estimated as in Table 3, on a common estimation population with exposure effects for movers, or on the population of stayers (N=46,917,409). Estimates displayed only if the confidence interval of the difference in effects of 24 years of exposure is contained within 50 percentile ranks of the national mean. Overseas territories CZs omitted. Top: Difference for movers. Bottom: Difference for stayers.

Figure 8: Forecast Exposure Effects and Confidence Intervals (RMSE), Top 50 CZs by Population



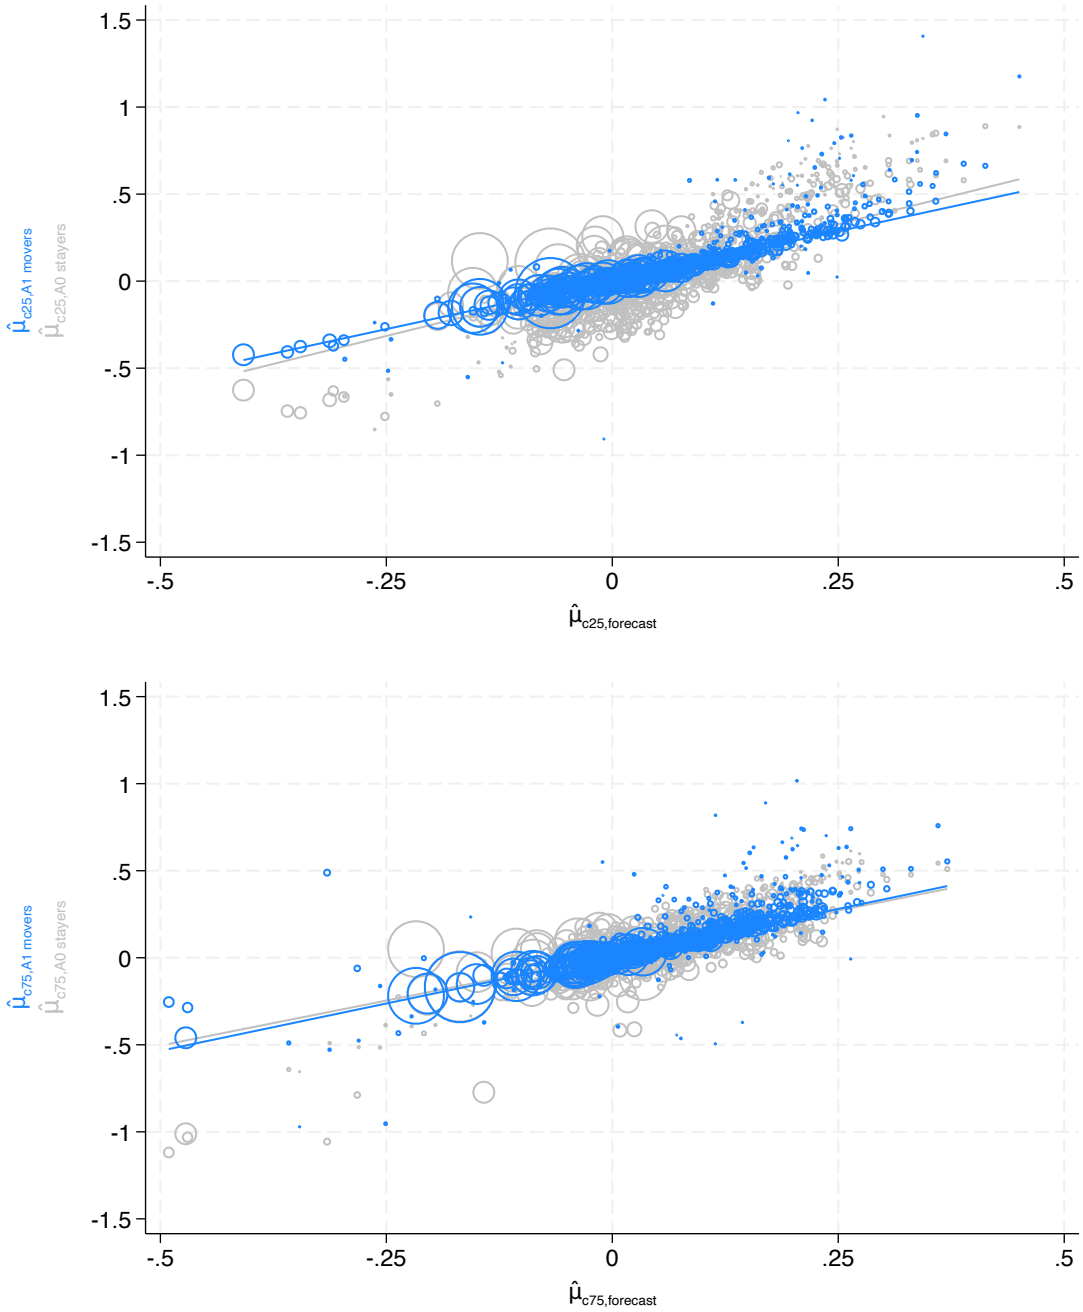
Forecast effects of one year of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, compared with estimated exposure effects as in Figure 1. Forecast effects are obtained by empirical Bayes shrinkage: we regress the exposure effects on the observed outcomes for stayers in Table 2, weighted by exposure, and obtain the MSE-minimizing forecast as a weighted average between the exposure effect and the predicted exposure effect conditional on the observed outcome, where weights are based on the standard error of the exposure effect and the standard error of the prediction. The vertical axes plot the forecast rank of each CZ among the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. Top: forecast for children at the 25th percentile of parent family income. Bottom: forecast for children at the 75th percentile of parent family income.

Figure 9: Forecast Exposure Effects and Confidence Intervals (RMSE), All CZs



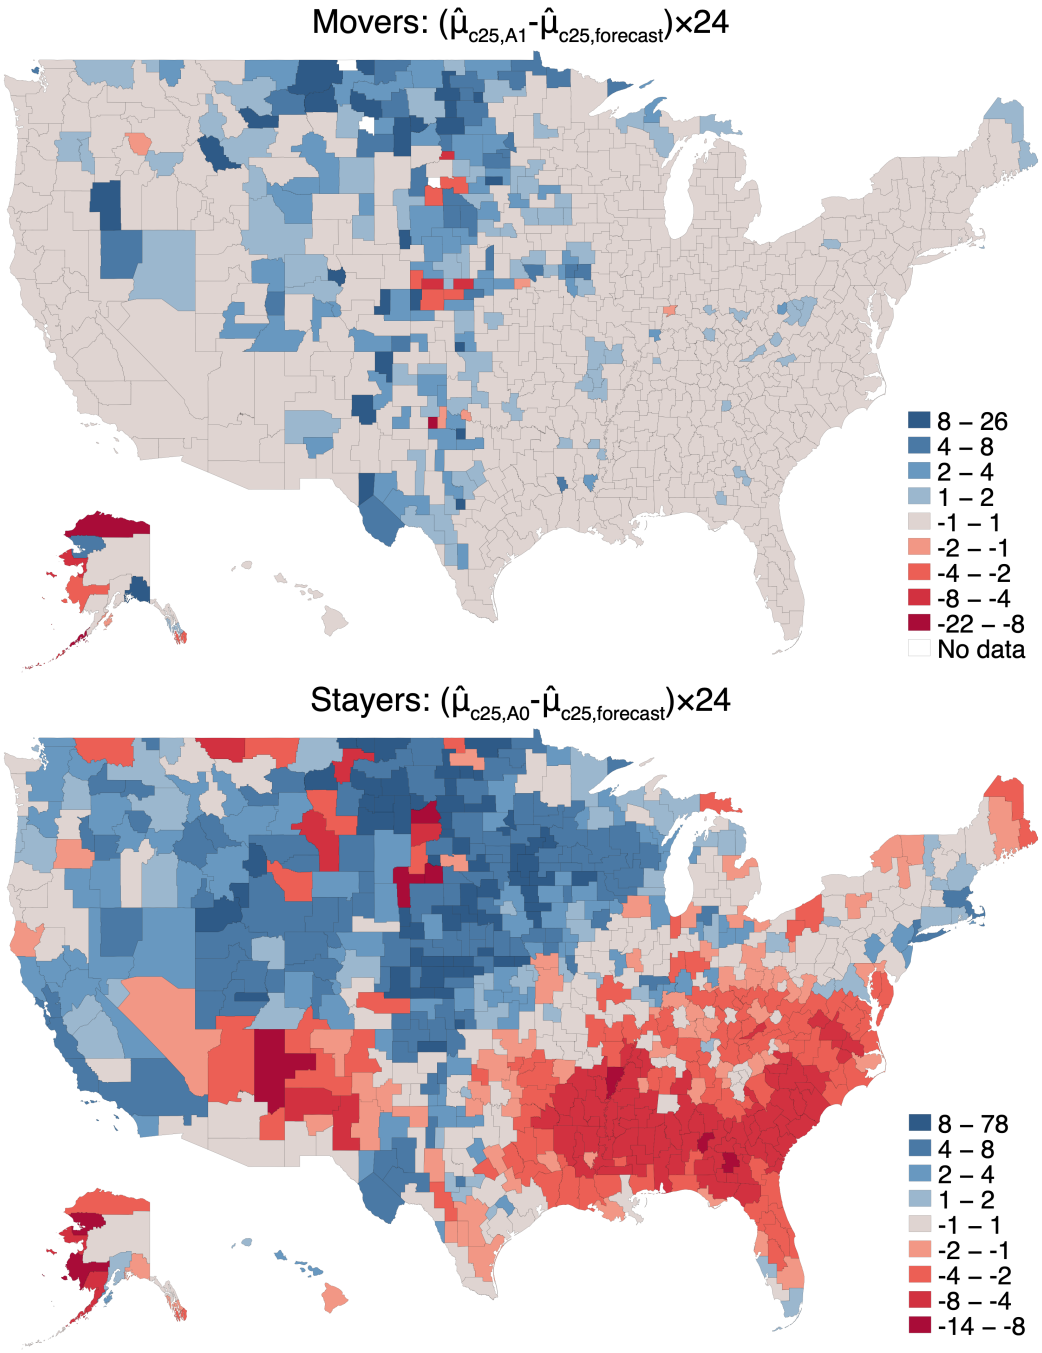
Forecast effects of one year of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, compared with estimated exposure effects as in Figure 2. Forecast effects are obtained by empirical Bayes shrinkage: we regress the exposure effects on the observed outcomes for stayers in Table 2, weighted by exposure, and obtain the MSE-minimizing forecast as a weighted average between the exposure effect and the predicted exposure effect conditional on the observed outcome, where weights are based on the standard error of the exposure effect and the standard error of the prediction. The vertical axes plot the forecast rank of all CZs such that the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean. Top: forecast for children at the 25th percentile of parent family income. Bottom: forecast for children at the 75th percentile of parent family income.

Figure 10: Forecast Exposure Effects Are More Similar to Estimated Exposure Effects than to Observed Outcomes



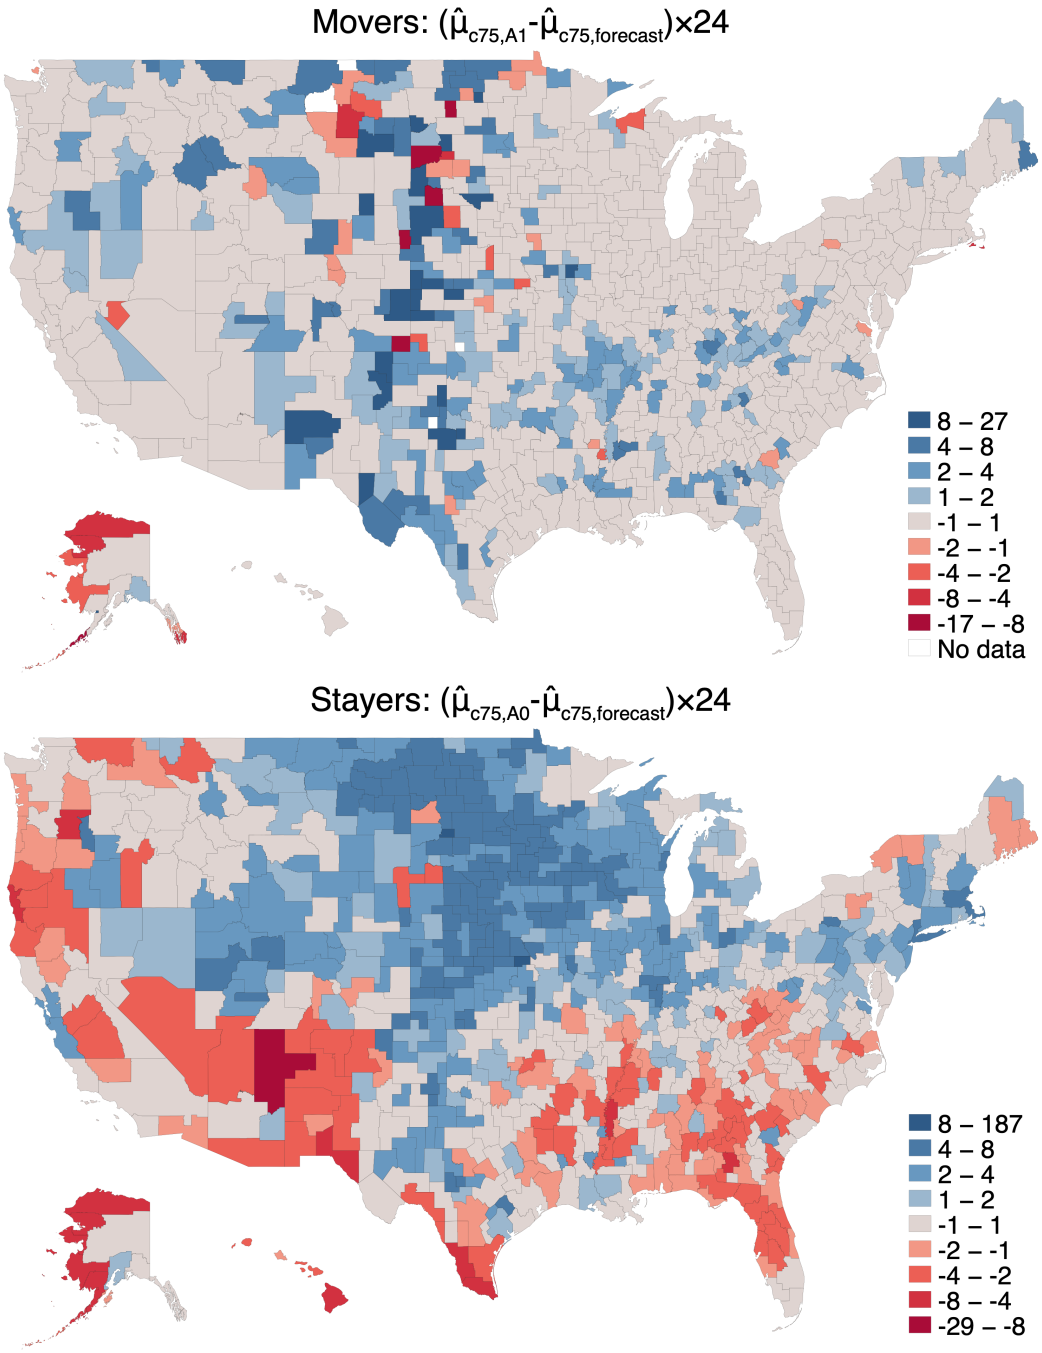
Exposure effects and observed outcomes compared with forecast exposure effects as in Figure 9. We plot lines for the regressions of exposure effects (blue) and observed outcomes (gray), respectively, on forecast exposure effects. We scatter exposure effects and observed outcomes, with circle sizes proportional to exposure weights, to show that all large CZs' forecasts are essentially identical to exposure effects, small CZs' forecasts are very similar, and observed outcomes are very different from the forecasts. Further, because exposure effects are estimated with error for smaller CZs and observed outcomes are uncertain due to bias, their weights may sum to less than one and forecasts may shrink toward zero relative to both exposure effects and observed outcomes. Top: comparison for children at the 25th percentile of parent family income. Bottom: comparison for children at the 75th percentile of parent family income.

Figure 11: Forecast Exposure Effects Compared to Estimated Exposure Effects and Observed Outcomes by Commuting Zone, $p = 25$



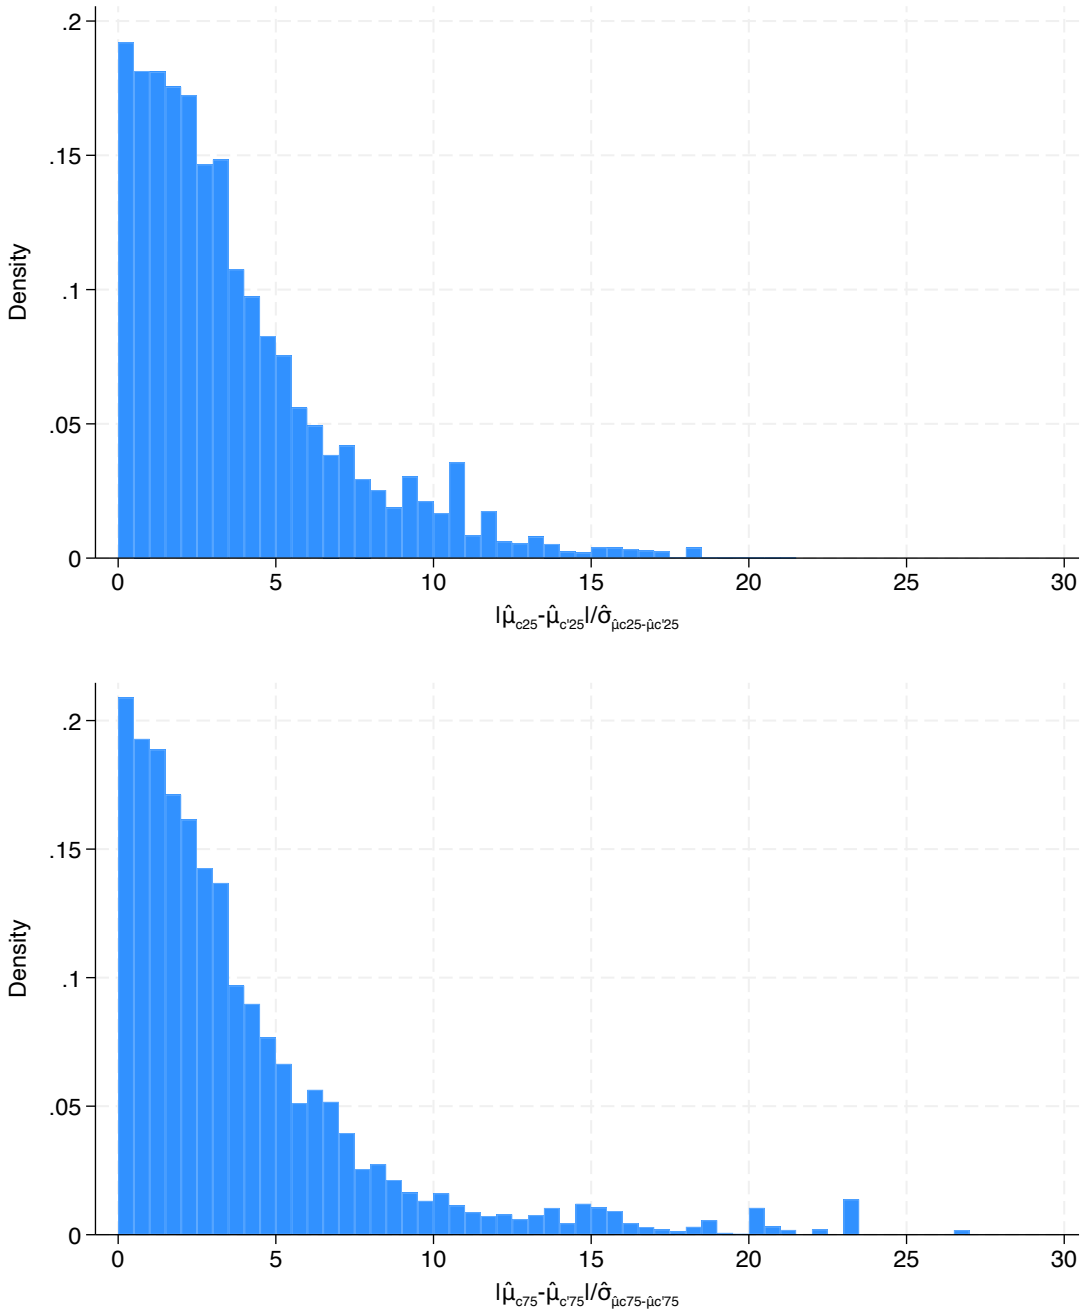
Difference in exposure effects for movers (top) and observed outcomes for stayers (bottom), respectively, with forecast exposure effects as in Figure 9, scaled up to effect of 24 years of exposure, for children at the 25th percentile of parent family income. Exposure effects estimated as in Table 2 (N=21,848,889). Observed outcomes estimated as in Table 3, on the population of stayers (N=46,917,409). Estimates displayed only if the confidence interval of the difference in effects of 24 years of exposure is contained within 50 percentile ranks of the national mean. Overseas territories CZs omitted.

Figure 12: Forecast Exposure Effects Compared to Estimated Exposure Effects and Observed Outcomes by Commuting Zone, $p = 75$



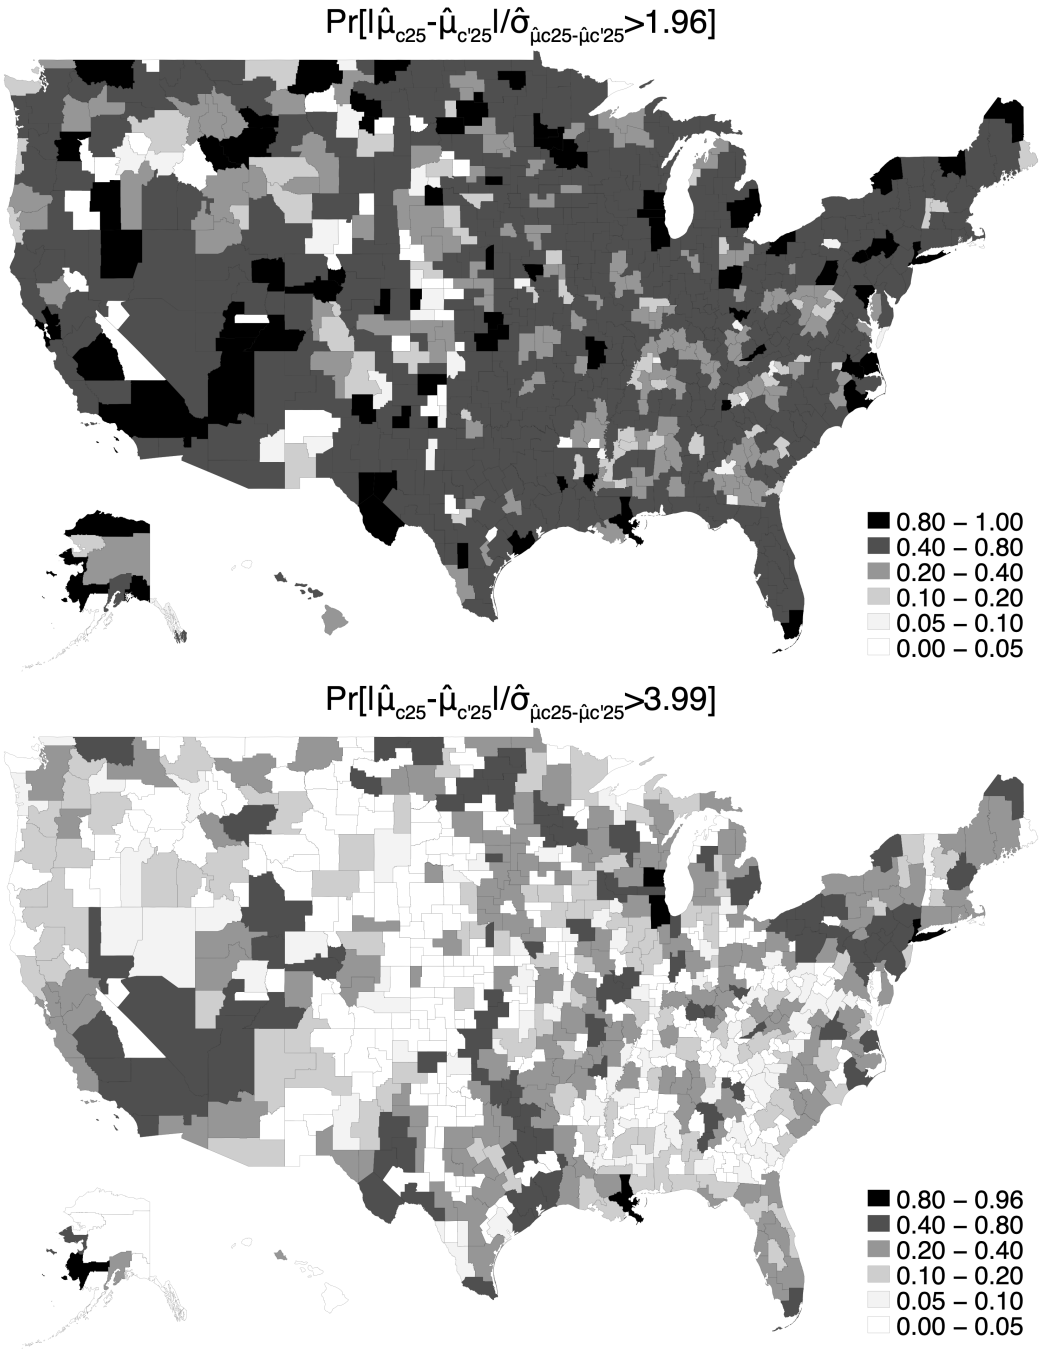
Difference in exposure effects for movers (top) and observed outcomes for stayers (bottom), respectively, with forecast exposure effects as in Figure 9, scaled up to effect of 24 years of exposure, for children at the 75th percentile of parent family income. Exposure effects estimated as in Table 2 (N=21,848,889). Observed outcomes estimated as in Table 3, on the population of stayers (N=46,917,409). Estimates displayed only if the confidence interval of the difference in effects of 24 years of exposure is contained within 50 percentile ranks of the national mean. Overseas territories CZs omitted.

Figure 13: Pairwise Differences in Exposure Effects Between Commuting Zones, Test Statistics Weighted by Destination Frequencies



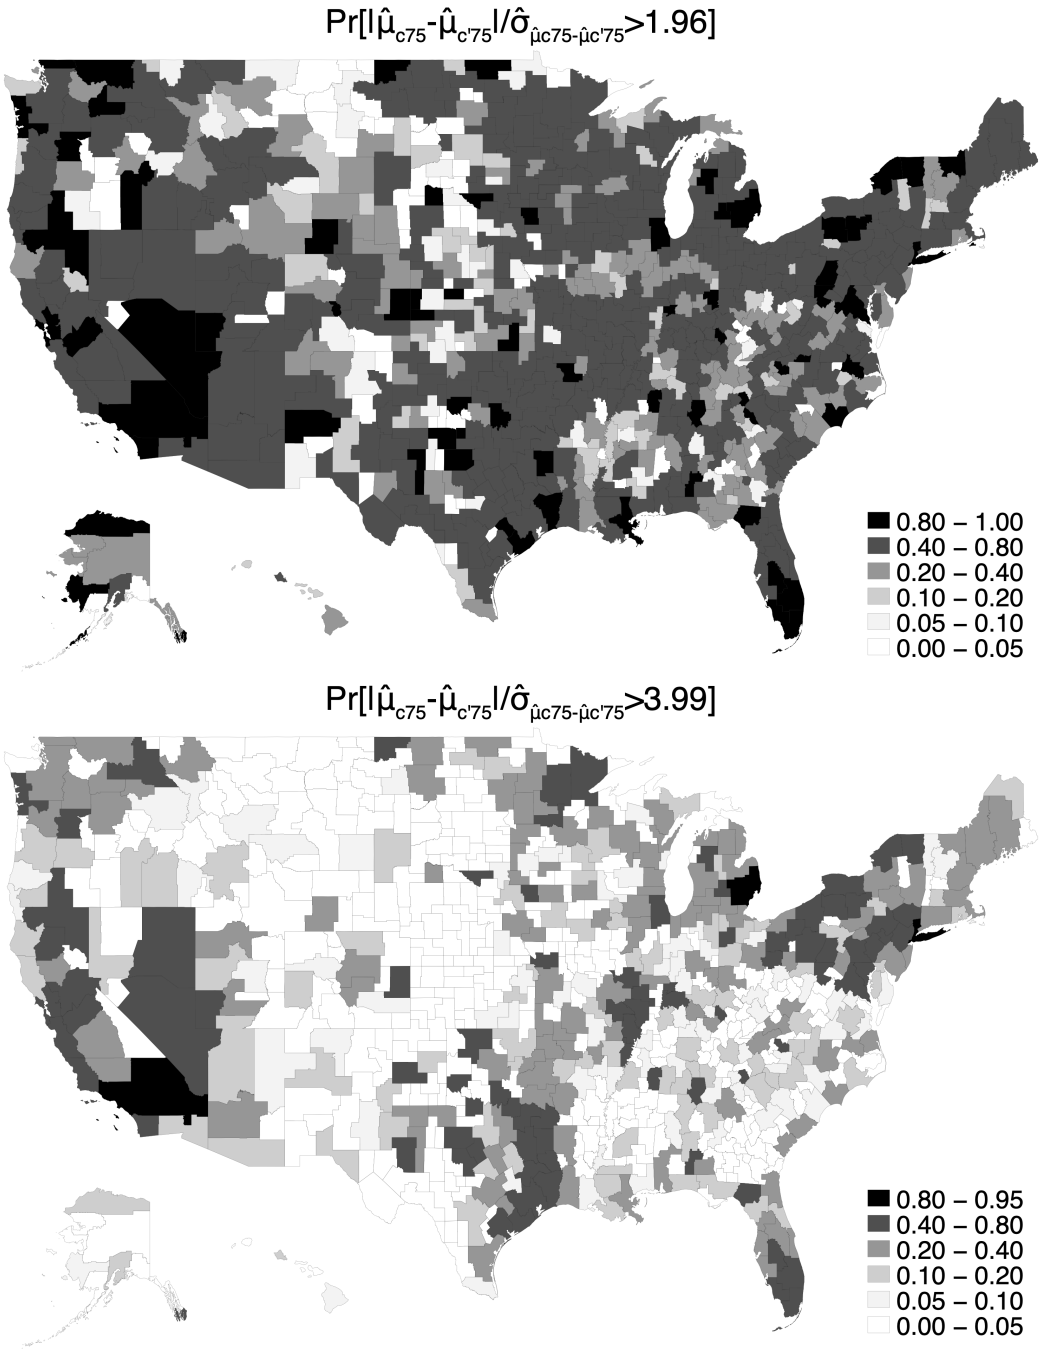
Distribution of t-statistics of the absolute difference in exposure effects for each pair of CZs, weighted by destination frequencies. Estimation as in Table 2 (N=21,848,889). The distribution is informative about the real-world impact (or lack of impact) of families' moves from origin to destination on children's outcomes. Top: distribution for children at the 25th percentile of parent family income. Bottom: distribution for children at the 75th percentile of parent family income.

Figure 14: Pairwise Differences in Exposure Effects Between Commuting Zones, Null Rejections Weighted by Destination Frequencies, by Commuting Zone, $p = 25$



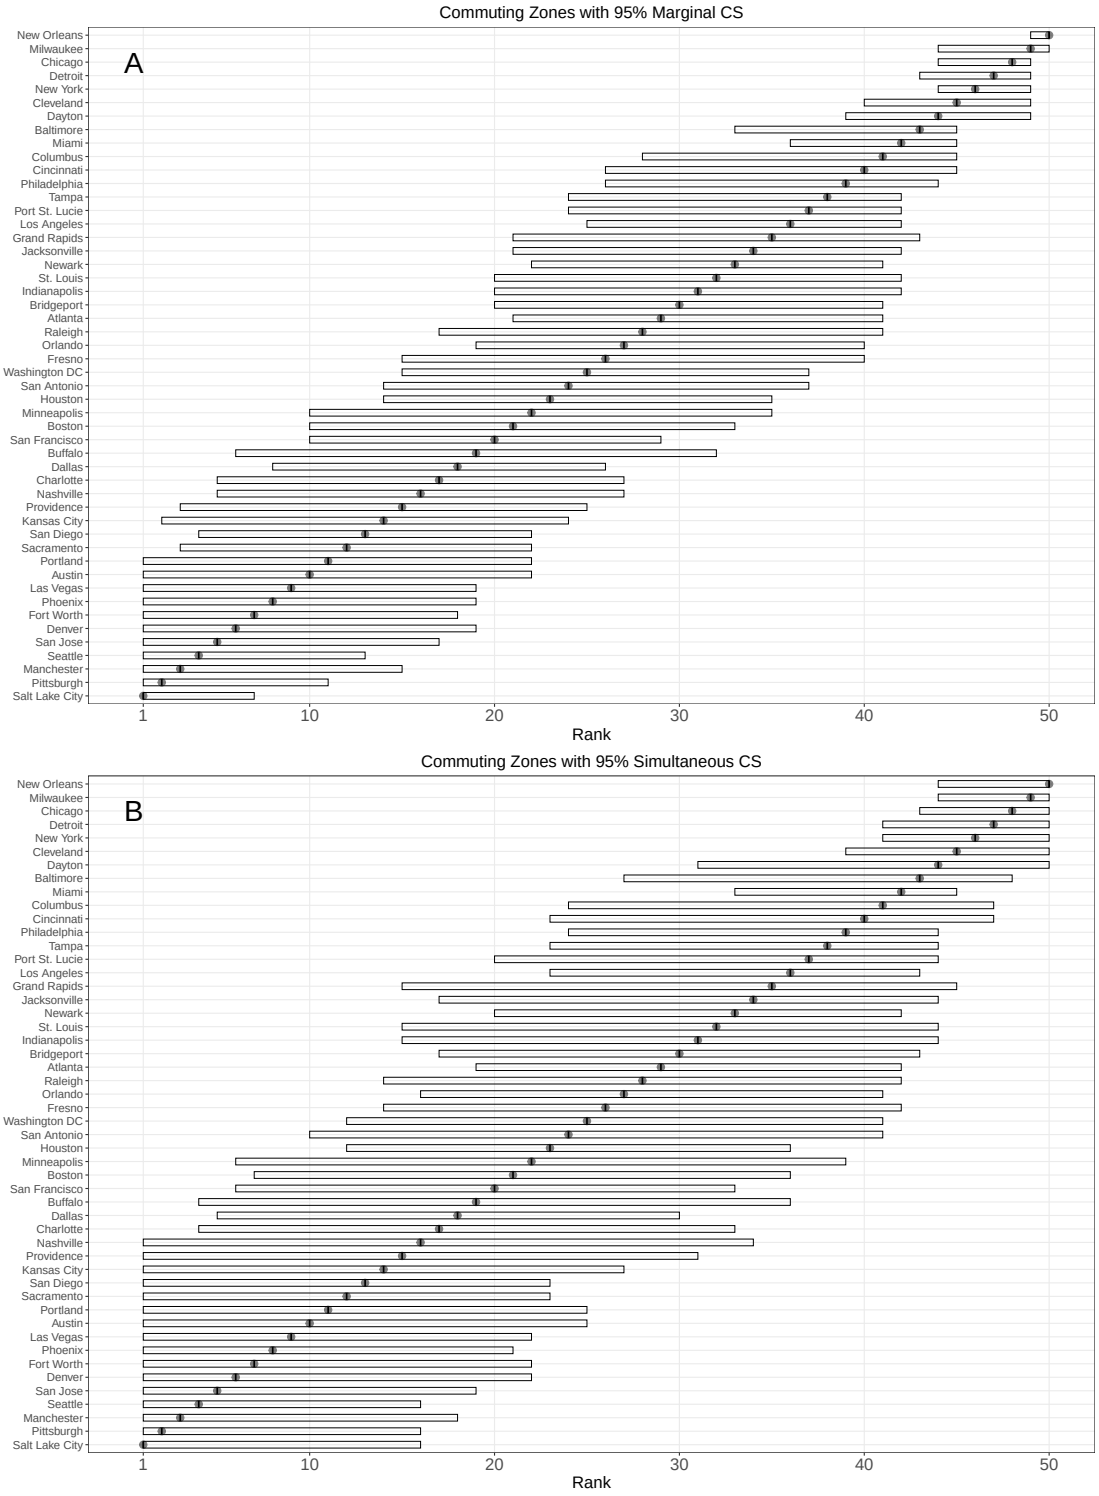
Probability that a typical family's move from each origin CZ produces a significant exposure effect relative to the origin, for children at the 75th percentile of parent family income. Calculation based on t-statistics of the absolute difference in exposure effects for each pair of CZs, weighted by destination frequencies. Estimation as in Table 2 (N=21,848,889). The distribution is informative about the real-world impact (or lack of impact) of families' moves from origin to destination on children's outcomes, by origin. Top: tests for significance at the 95-percent confidence level, appropriate when families decide whether to move or when to move, but not where. Bottom: tests for significance at the Bonferroni-corrected 95-percent confidence level, appropriate when families decide where to move, comparing all potential destination CZs.

Figure 15: Pairwise Differences in Exposure Effects Between Commuting Zones, Null Rejections Weighted by Destination Frequencies, by Commuting Zone, $p = 75$



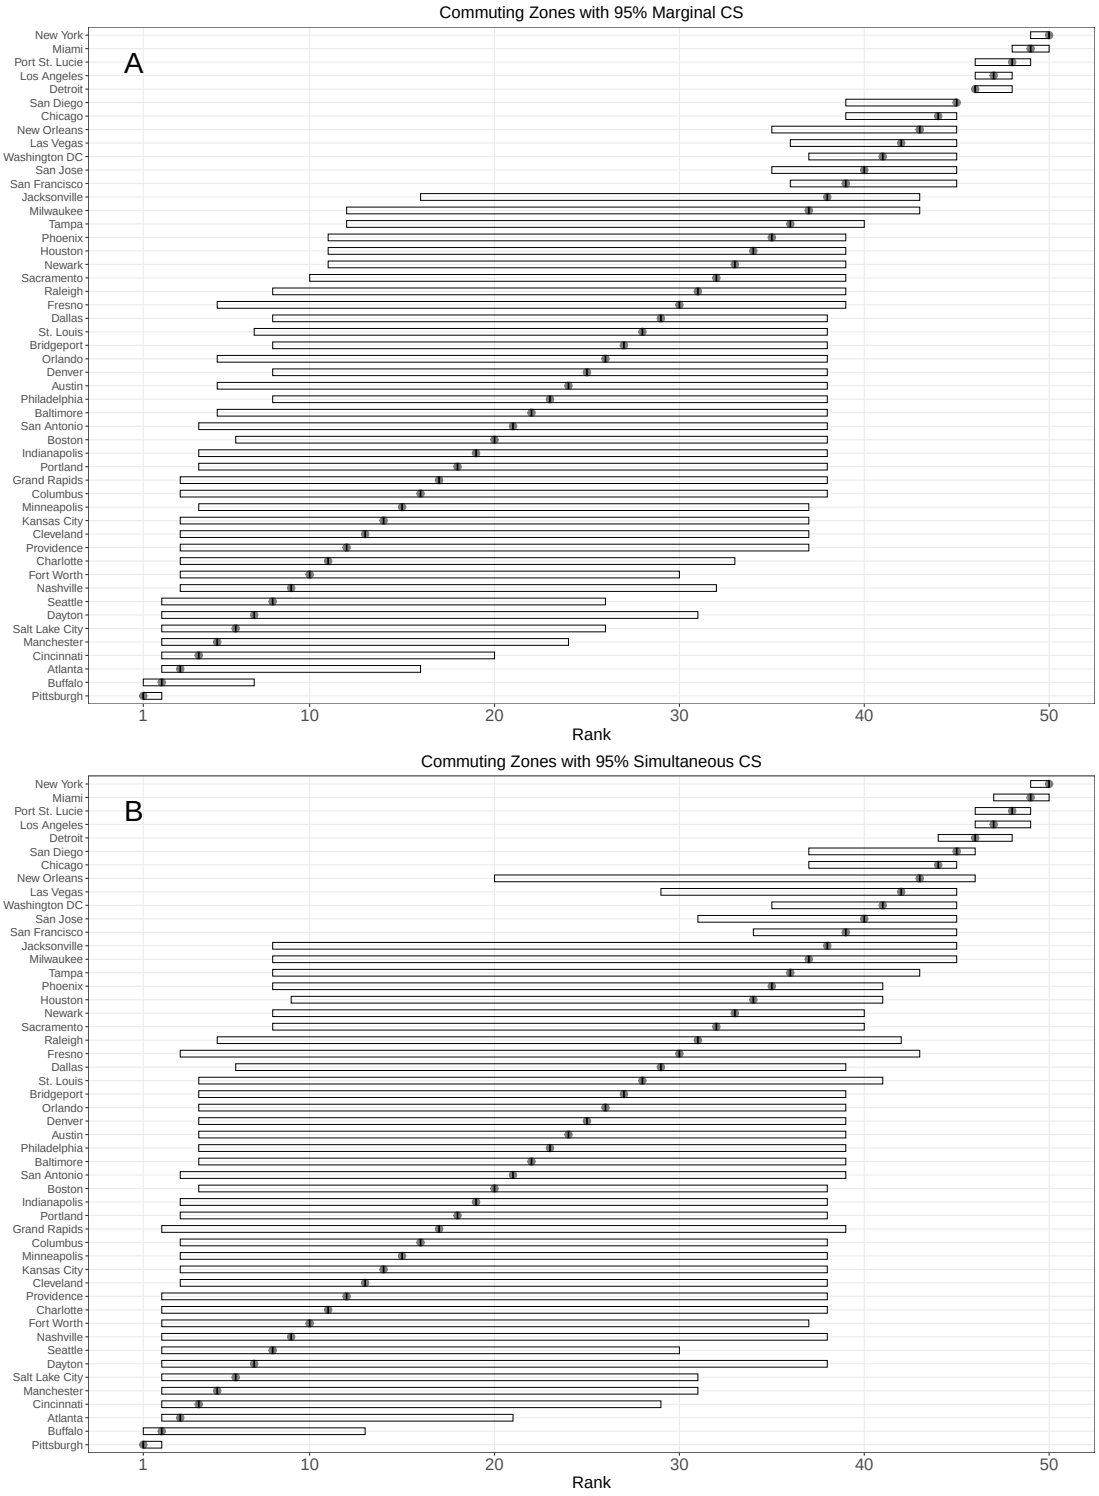
Probability that a typical family's move from each origin CZ produces a significant exposure effect relative to the origin, for children at the 75th percentile of parent family income. Calculation based on t-statistics of the absolute difference in exposure effects for each pair of CZs, weighted by destination frequencies. Estimation as in Table 2 (N=21,848,889). The distribution is informative about the real-world impact (or lack of impact) of families' moves from origin to destination on children's outcomes, by origin. Top: tests for significance at the 95-percent confidence level, appropriate when families decide whether to move or when to move, but not where. Bottom: tests for significance at the Bonferroni-corrected 95-percent confidence level, appropriate when families decide where to move, comparing all potential destination CZs.

Figure 16: Rankings by Exposure Effects, Top 50 CZs by Population, $p = 25$



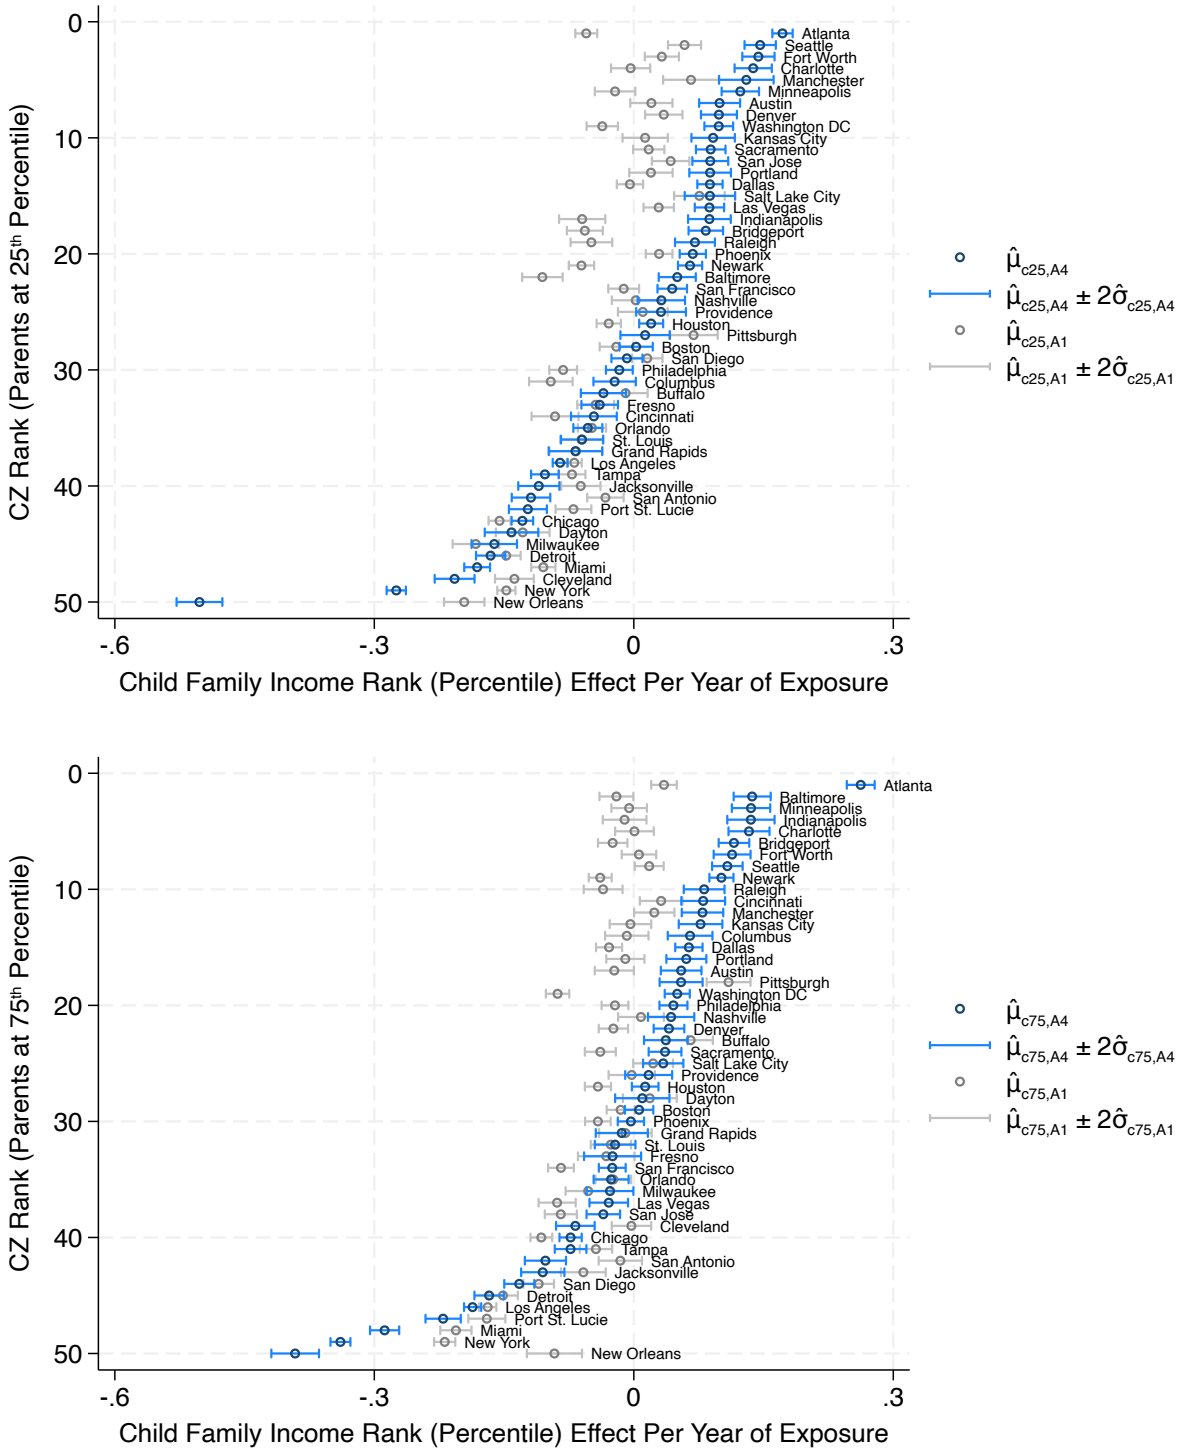
Confidence sets for the ranking of CZs by exposure effects presented in Figure 1, for children at the 25th percentile of parent family income. The top-ranked CZ has the most-positive estimated exposure effect on child family income rank in adulthood. The confidence sets compare CZs within the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. We implement the inference procedures in Mogstad et al. (2024), modified to account for the covariance of individual exposure effect estimates as in the pairwise comparisons in Figure 13. Top: marginal confidence sets for pairwise comparisons between each origin CZ and all destination CZs. Bottom: simultaneous confidence sets for all pairwise comparisons.

Figure 17: Rankings by Exposure Effects, Top 50 CZs by Population, $p = 75$



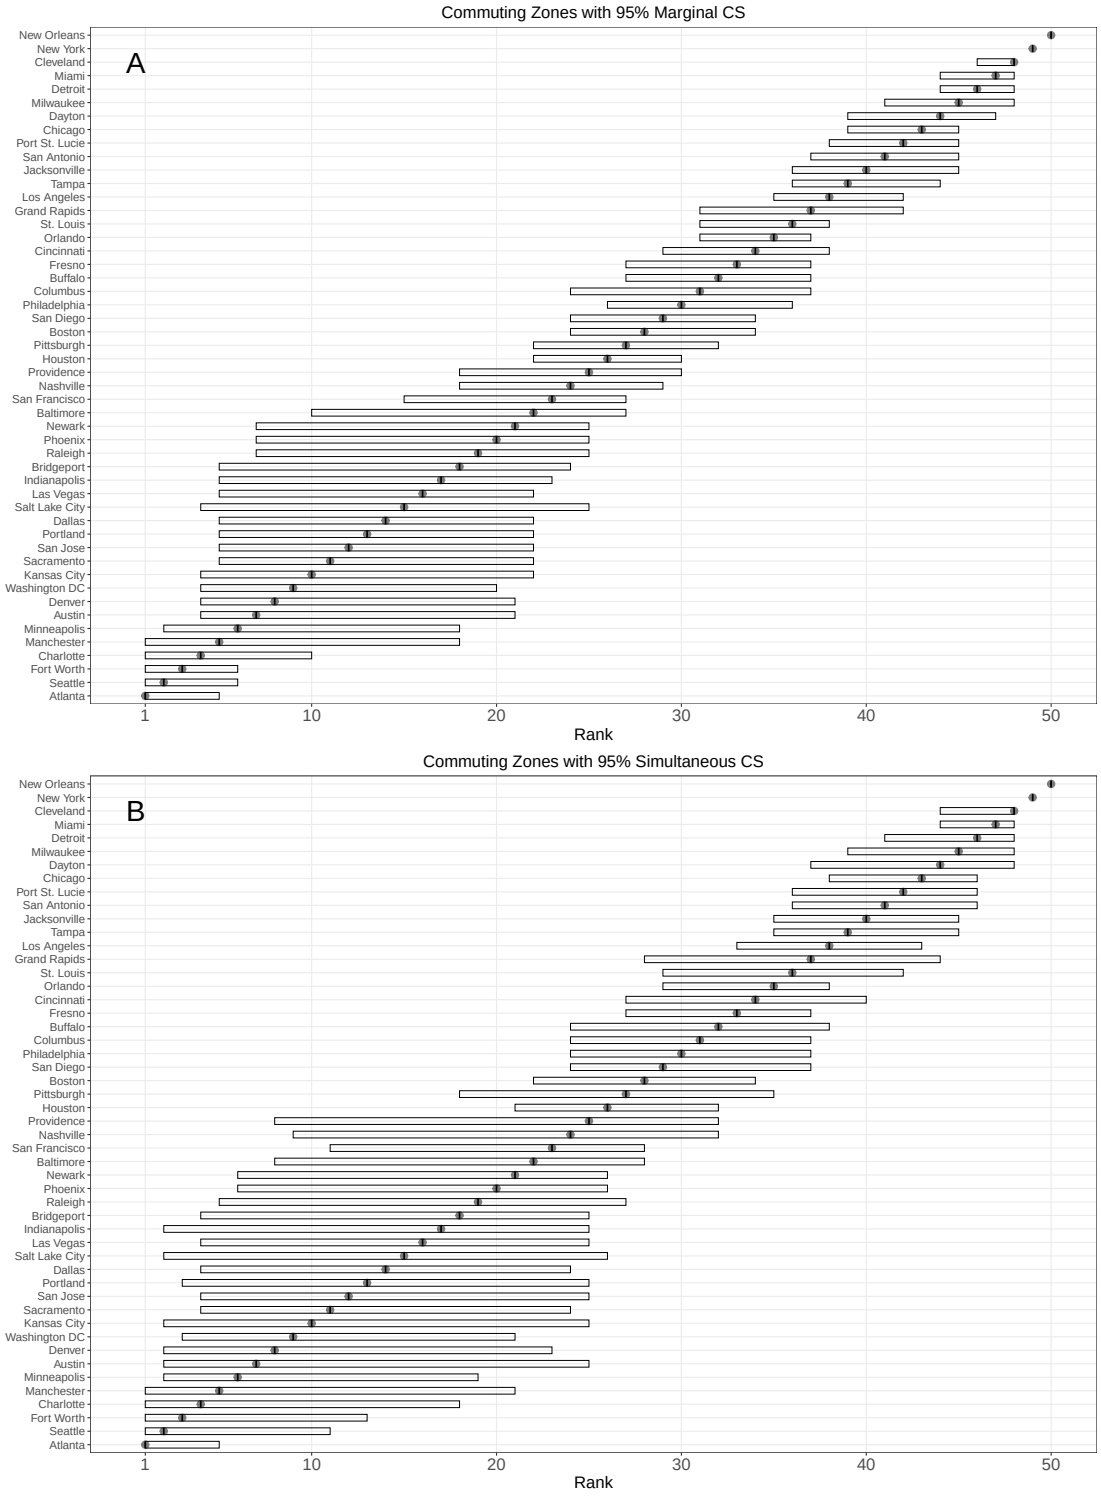
Confidence sets for the ranking of CZs by exposure effects presented in Figure 1, for children at the 75th percentile of parent family income. The top-ranked CZ has the most-positive estimated exposure effect on child family income rank in adulthood. The confidence sets compare CZs within the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. We implement the inference procedures in Mogstad et al. (2024), modified to account for the covariance of individual exposure effect estimates as in the pairwise comparisons in Figure 13. Top: marginal confidence sets for pairwise comparisons between each origin CZ and all destination CZs. Bottom: simultaneous confidence sets for all pairwise comparisons.

Figure 18: Counterfactual Total Effects and Confidence Intervals, Top 50 CZs by Population



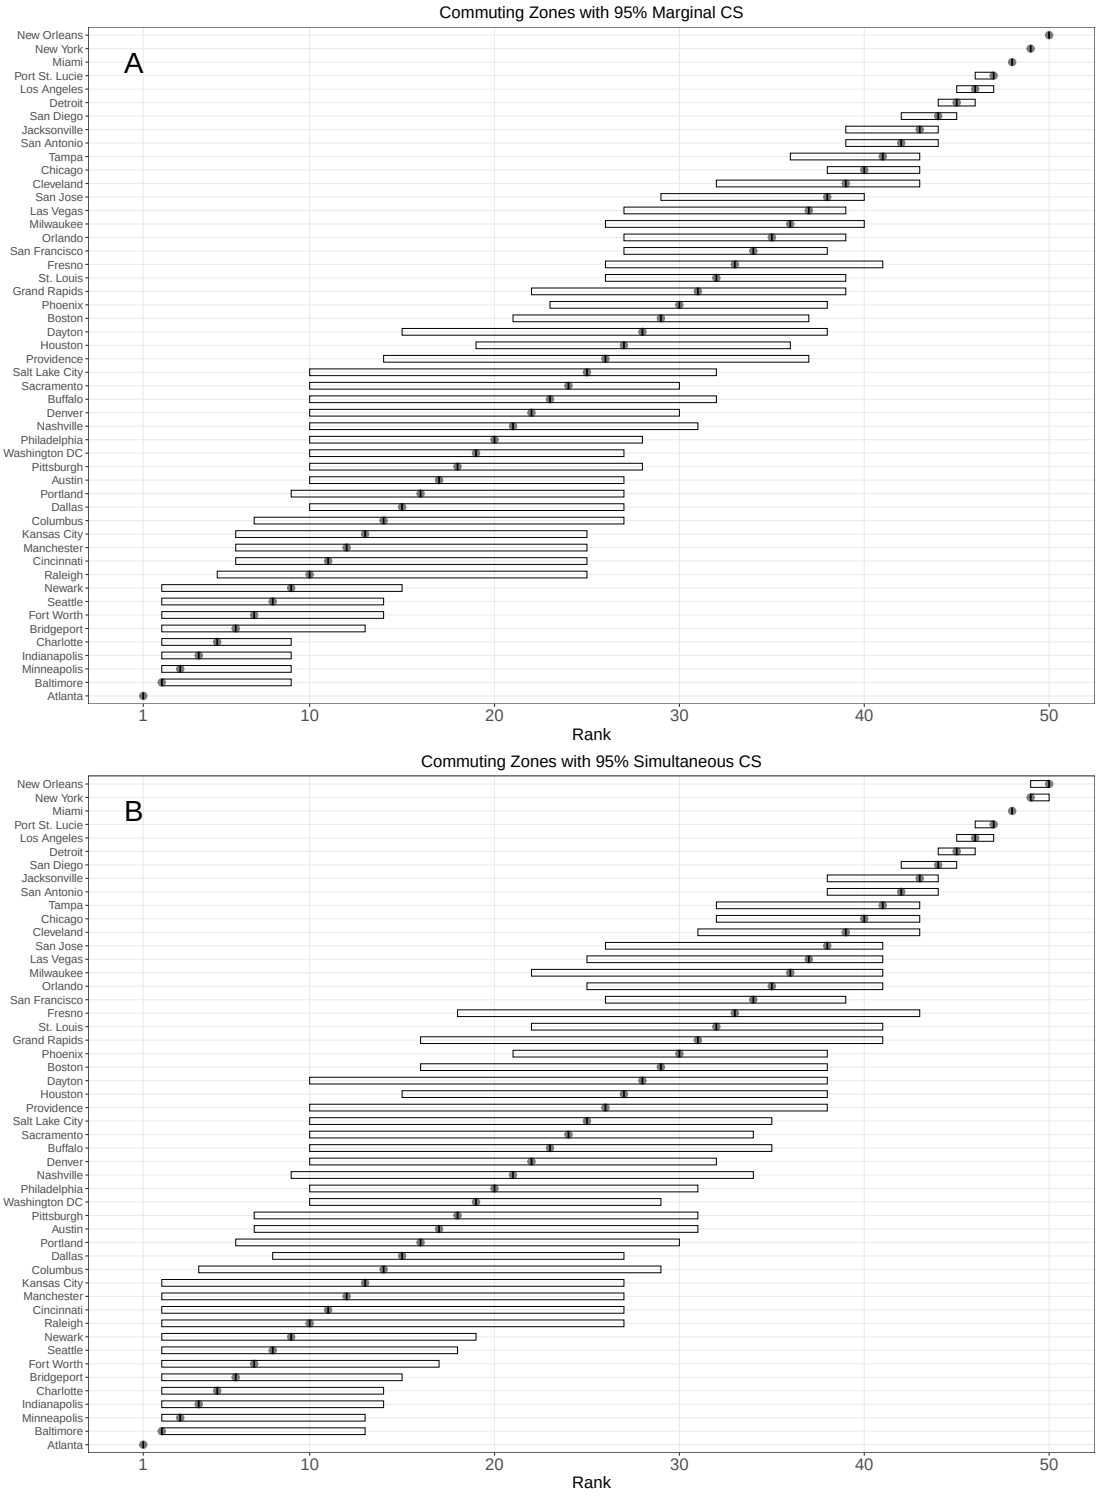
Counterfactual total effects of one year of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889), but replacing parent rank p_i with latent parent rank $\hat{q}_i$ for Estimator A4. Exposure effect estimates from Estimator A1 reproduced from Figure 1 for comparison. The vertical axes plot the rank of each CZ among the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. Top: effect for children at the 25th percentile of parent family income. Bottom: effect for children at the 75th percentile of parent family income.

Figure 19: Rankings by Counterfactual Total Effects, Top 50 CZs by Population, $p = 25$



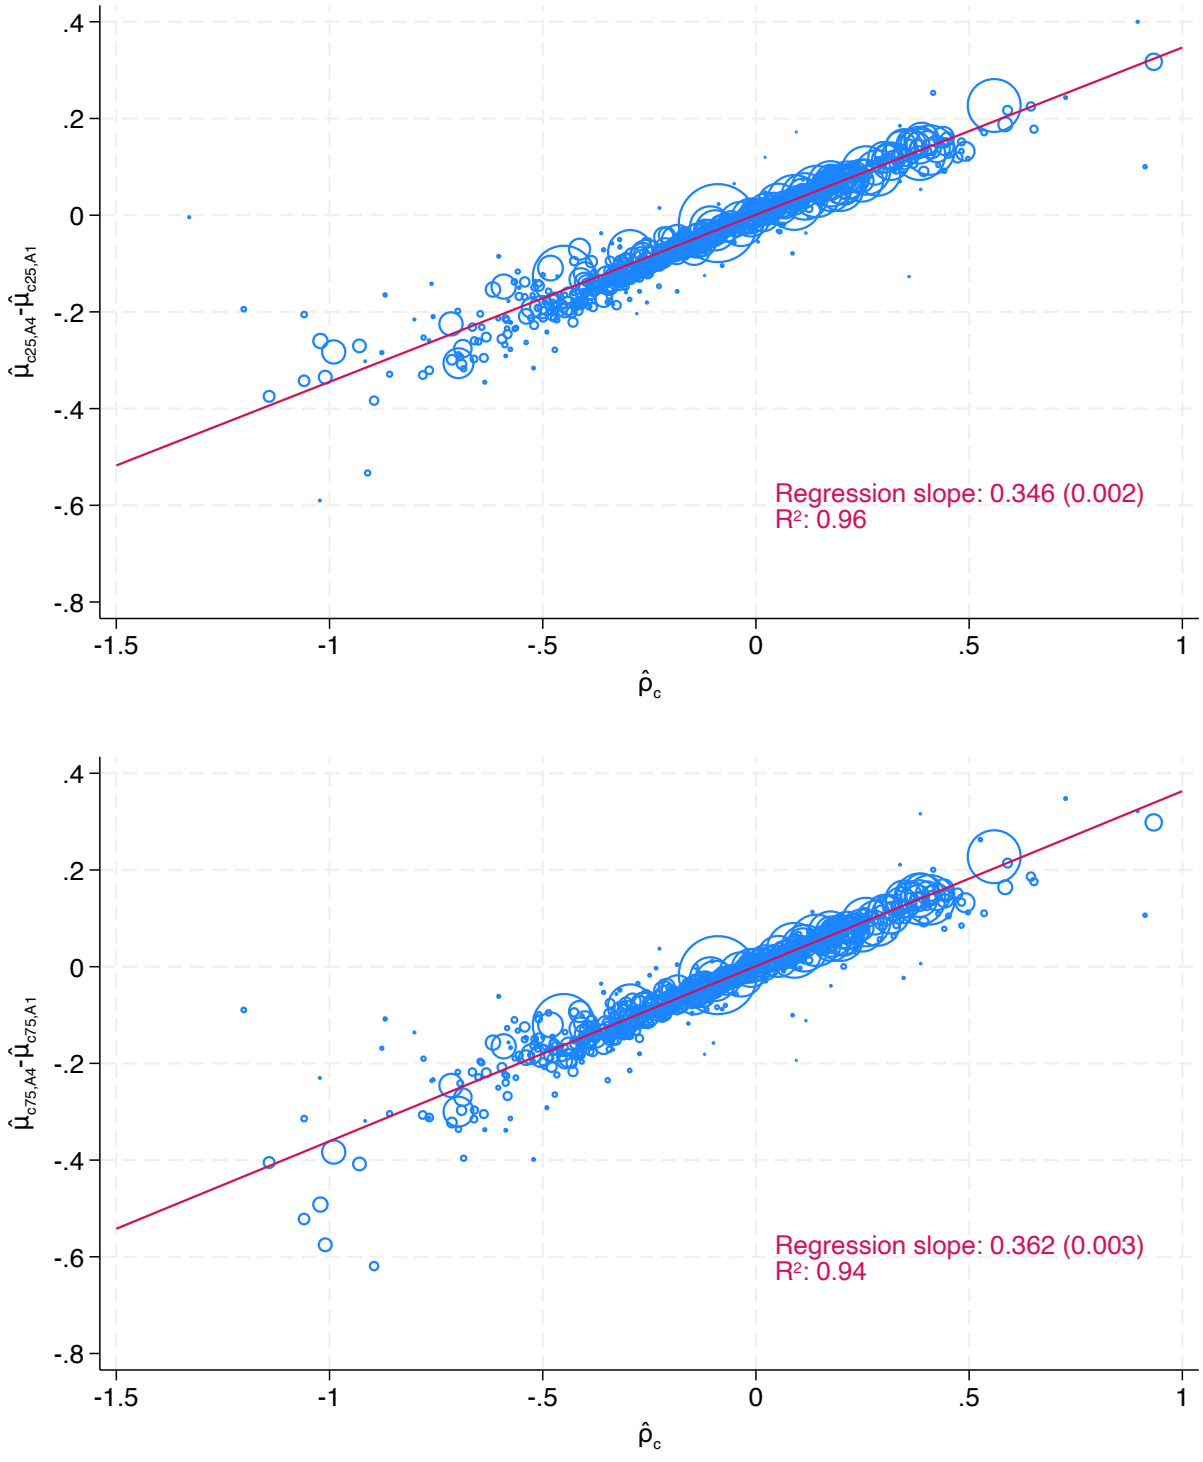
Confidence sets for the ranking of CZs by counterfactual total effects presented in Figure 18, for children at the 25th percentile of parent family income. The top-ranked CZ has the most-positive estimated exposure effect on child family income rank in adulthood. The confidence sets compare CZs within the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. We implement the inference procedures in Mogstad et al. (2024), modified to account for the covariance of individual counterfactual total effect estimates as in the pairwise comparisons in Figure 13. Top: marginal confidence sets for pairwise comparisons between each origin CZ and all destination CZs. Bottom: simultaneous confidence sets for all pairwise comparisons.

Figure 20: Rankings by Counterfactual Total Effects, Top 50 CZs by Population, $p = 75$



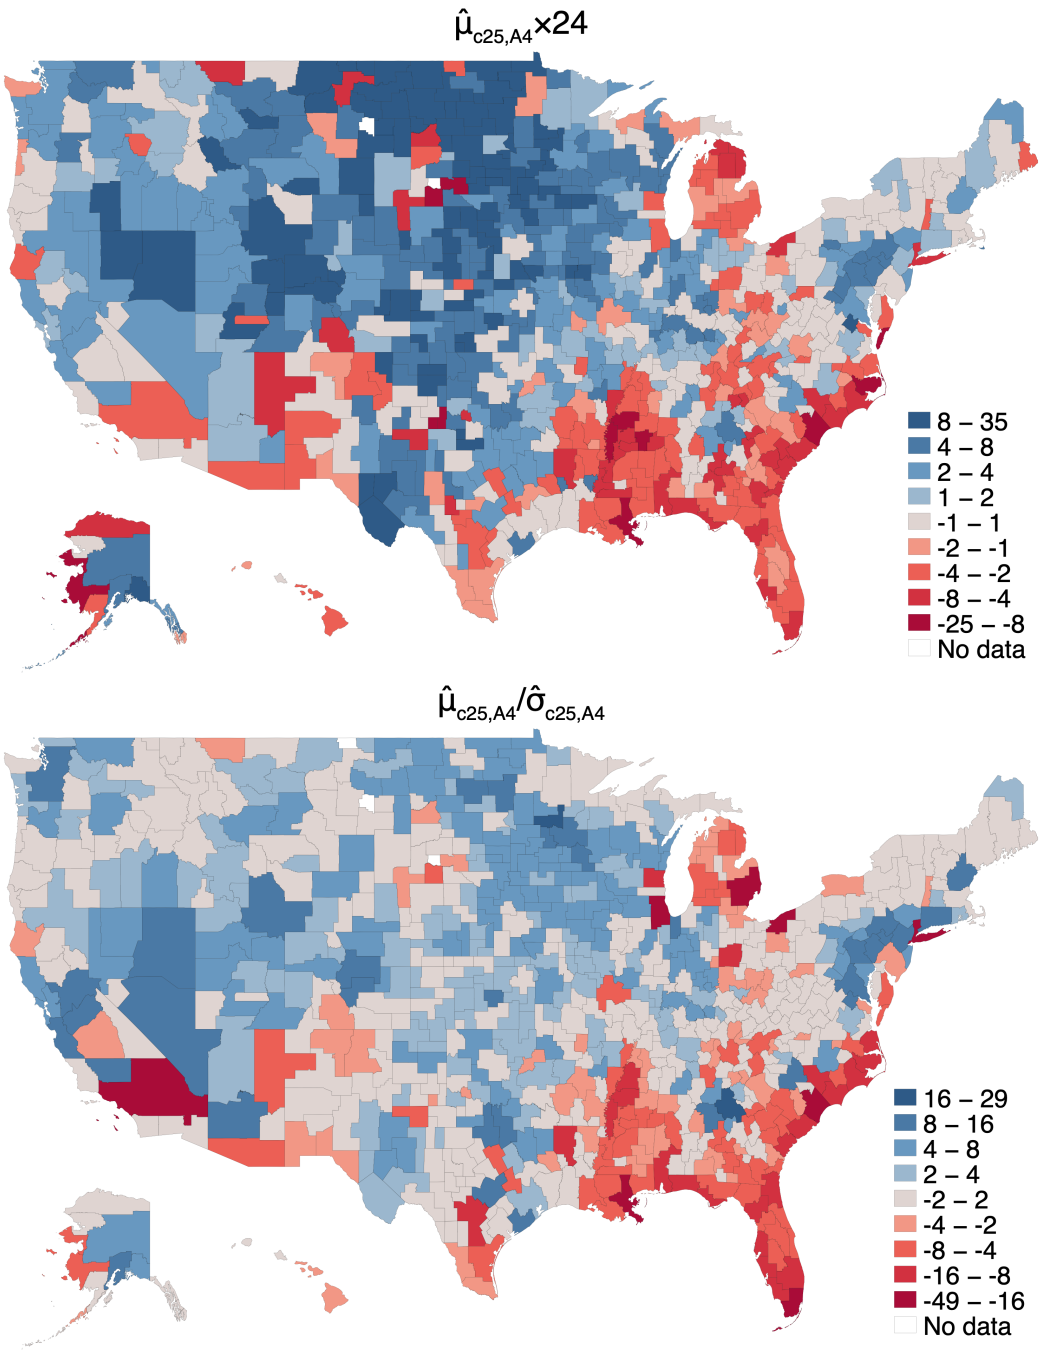
Confidence sets for the ranking of CZs by counterfactual total effects presented in Figure 18, for children at the 75th percentile of parent family income. The top-ranked CZ has the most-positive estimated exposure effect on child family income rank in adulthood. The confidence sets compare CZs within the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. We implement the inference procedures in Mogstad et al. (2024), modified to account for the covariance of individual counterfactual total effect estimates as in the pairwise comparisons in Figure 13. Top: marginal confidence sets for pairwise comparisons between each origin CZ and all destination CZs. Bottom: simultaneous confidence sets for all pairwise comparisons.

Figure 21: Counterfactual Indirect Effects on Children's Income, and Parent Rank Premia



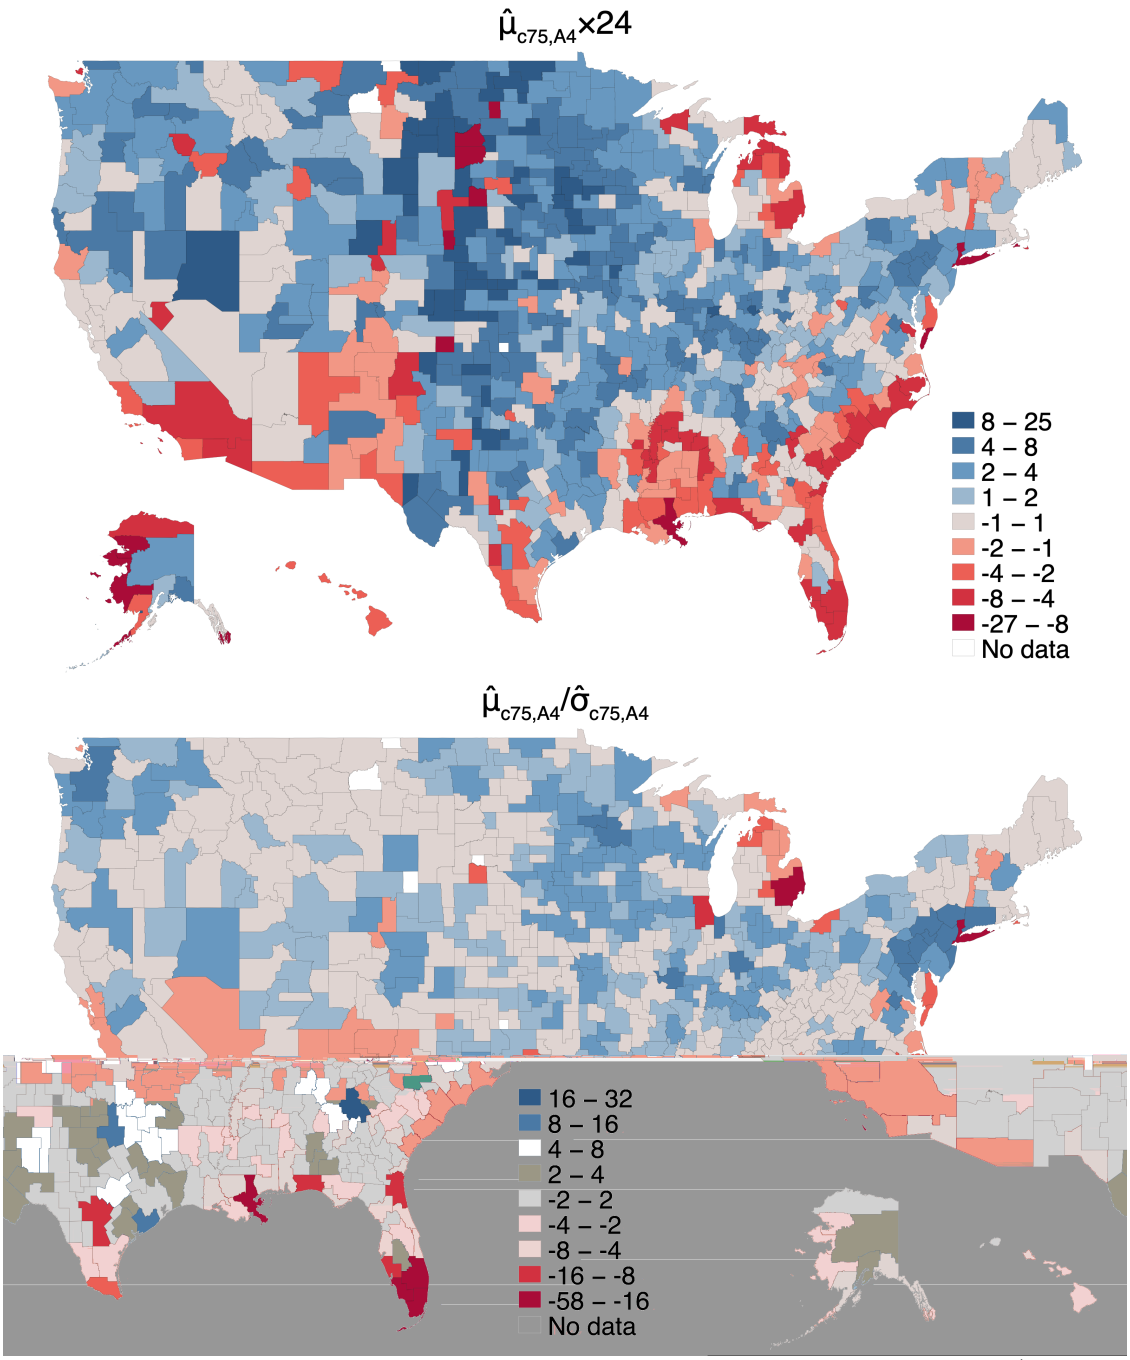
Counterfactual indirect effects (vertical axis) defined as the difference between total effects from Estimator A4 and direct exposure effects from Estimator A1 as in Figure 18, compared with parent rank premia (horizontal axis) estimated from equation 16. Estimation as in Table 2 (N=21,848,889), but replacing parent rank p_i with latent parent rank $\hat{q}_i$ for Estimator A4, and replacing child rank y_i with parent rank p_i and dropping interaction terms to estimate parent rank premia. We plot observations in blue, with circle sizes proportional to exposure weights. We plot the exposure-weighted regression line in red, reporting the regression slope, standard error, and explained variance. Top: comparison for children at the 25th percentile of parent family income. Bottom: comparison for children at the 75th percentile of parent family income.

Figure 22: Counterfactual Total Effects on Children's Income by Commuting Zone, $p = 25$



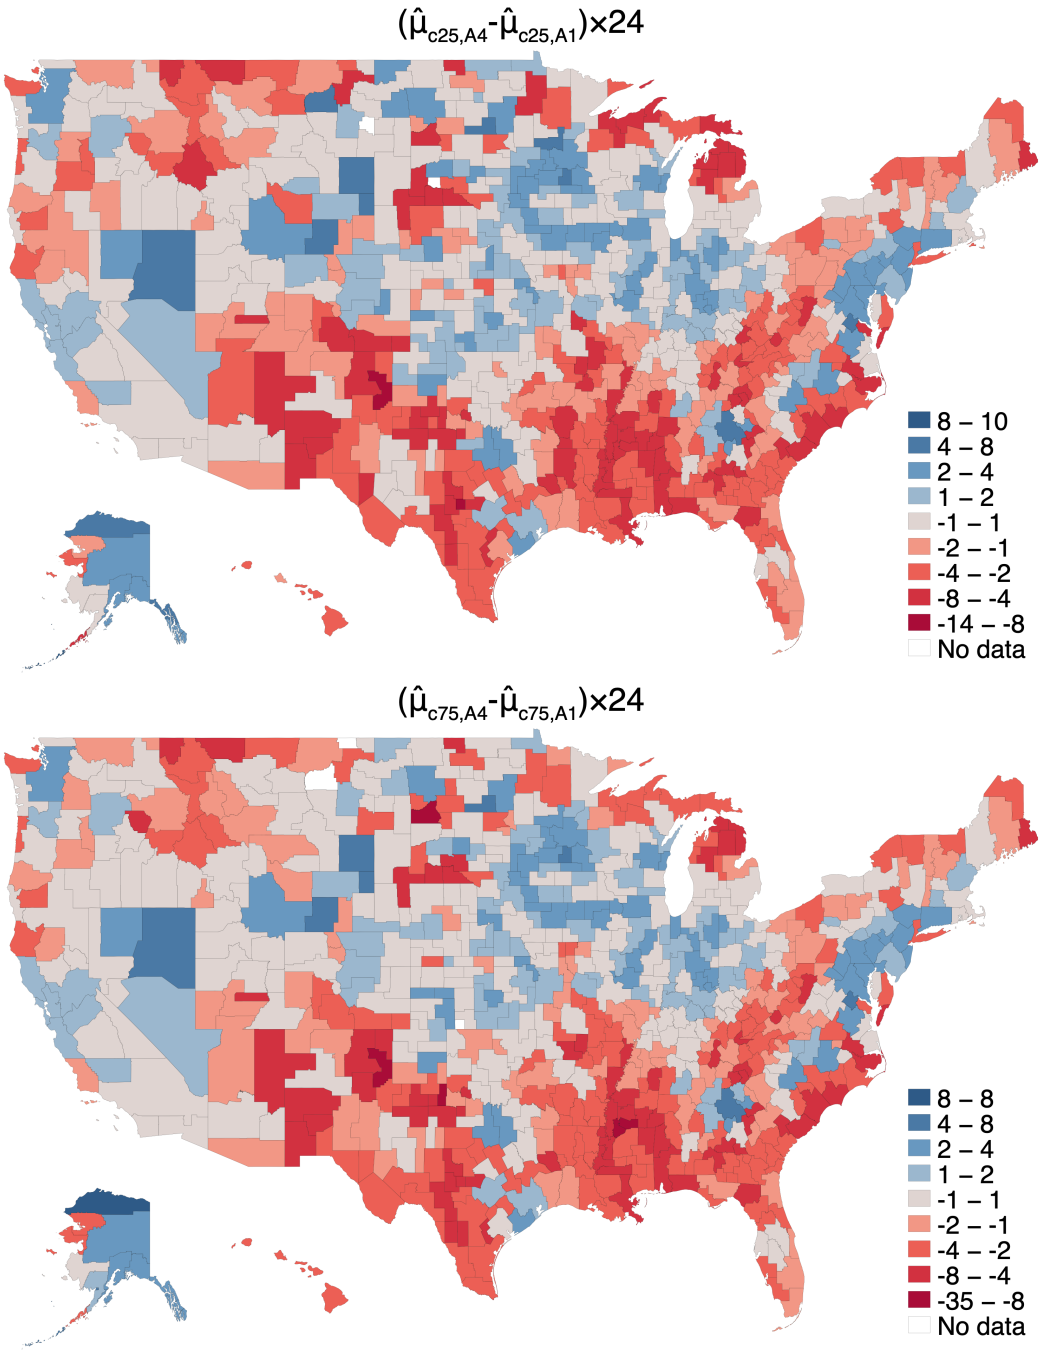
Counterfactual total effects of 24 years of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889), but replacing parent rank p_i with latent parent rank $\hat{q}_i$. Estimates displayed only if the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean (all CZs in Table 2, Column 1-2). Overseas territories CZs omitted. Top: effect for children at the 25th percentile of parent family income. Bottom: t-statistic of effect relative to national mean.

Figure 23: Counterfactual Total Effects on Children's Income by Commuting Zone, $p = 75$



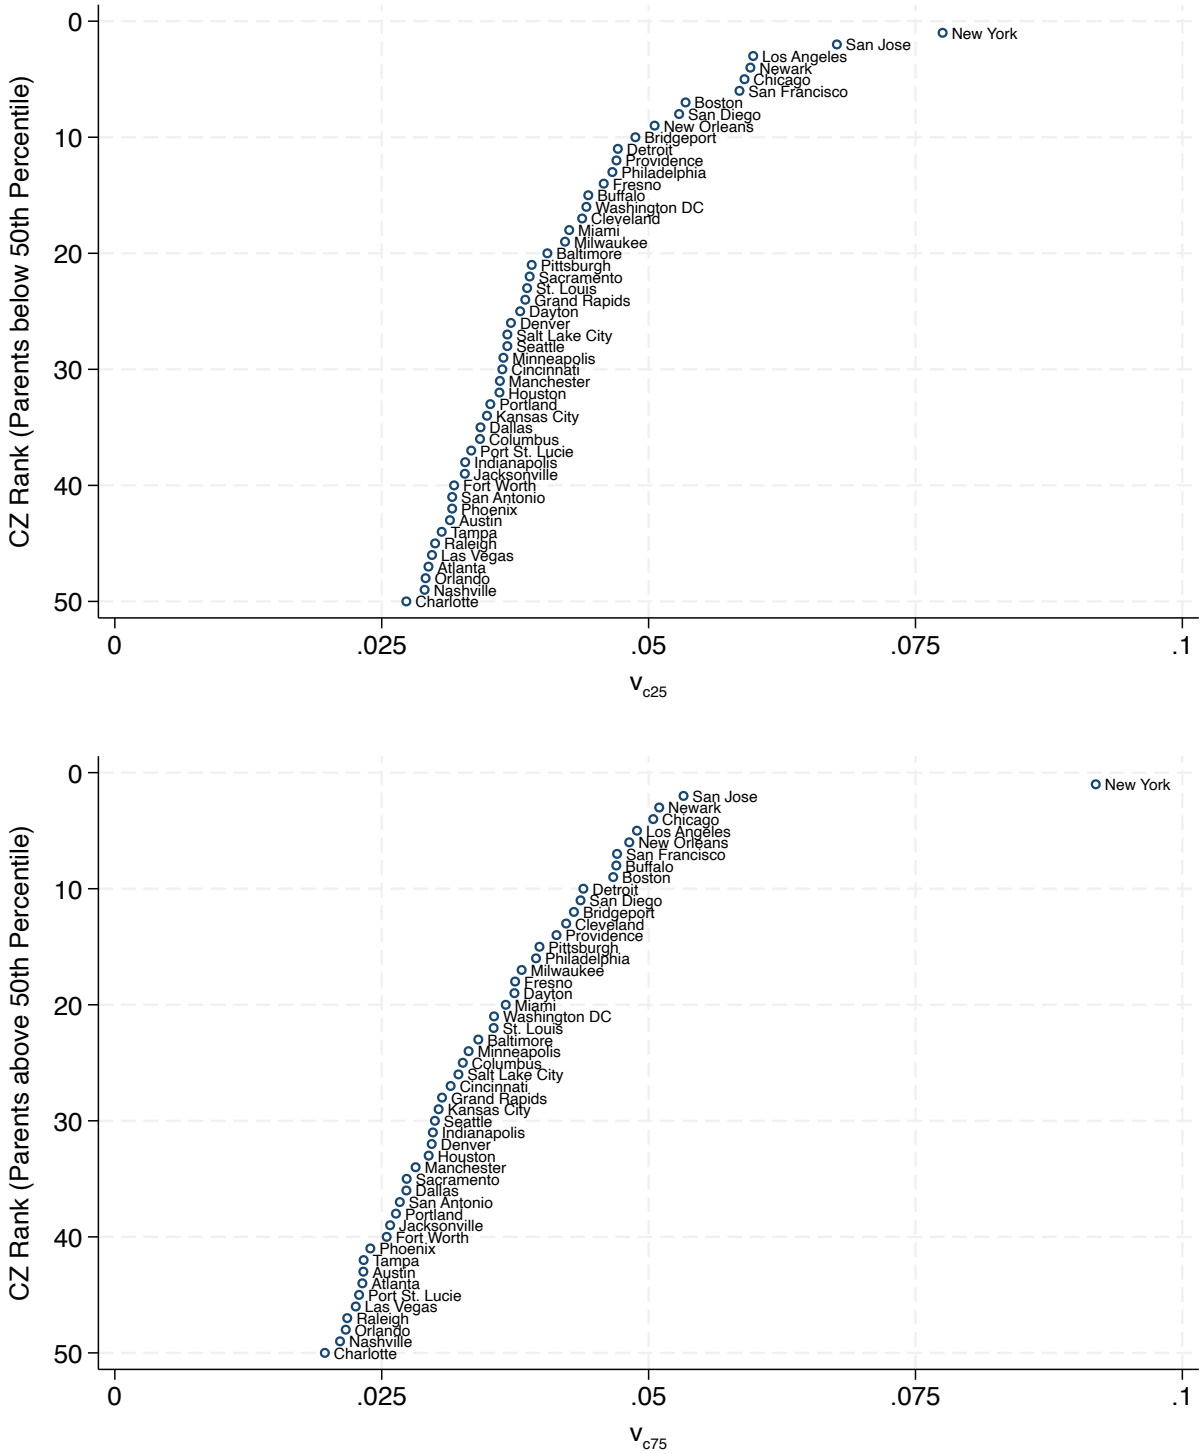
Counterfactual total effects of 24 years of exposure to each CZ on child family income rank (in percentiles) in adulthood, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889), but replacing parent rank p_i with latent parent rank $\hat{q}_i$. Estimates displayed only if the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean (all CZs in Table 2, Column 1-2). Overseas territories CZs omitted. Top: effect for children at the 75th percentile of parent family income. Bottom: t-statistic of effect relative to national mean.

Figure 24: Counterfactual Indirect Effects on Children's Income by Commuting Zone



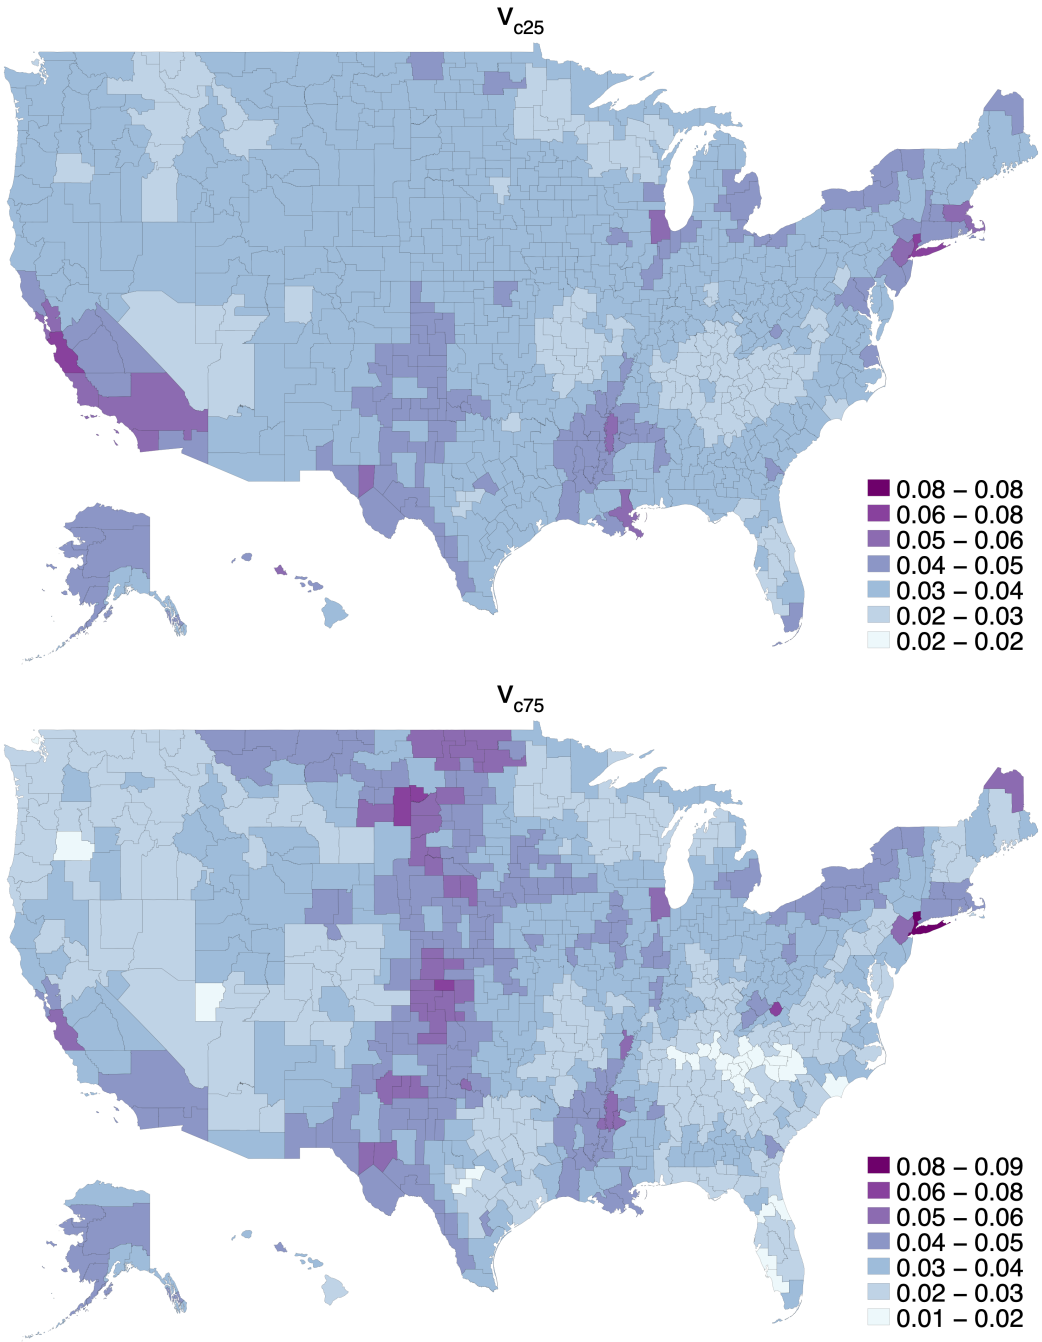
Counterfactual indirect effects of 24 years of exposure to each CZ on child family income rank (in percentiles) in adulthood defined as the difference between total effects from Estimator A4 and exposure effects from Estimator A1, relative to the national mean, estimated using different exposure durations for children who moved between CZs at different ages during childhood. Estimation as in Table 2 (N=21,848,889), but replacing parent rank p_i with latent parent rank $\hat{q}_i$ for Estimator A4. Estimates displayed only if the confidence interval of the effect of 24 years of exposure is contained within 50 percentile ranks of the national mean (all CZs in Table 2, Column 1-2). Overseas territories CZs omitted. Top: effect for children at the 25th percentile of parent family income. Bottom: effect for children at the 75th percentile of parent family income.

Figure 25: Parents' Revealed Preference, Top 50 CZs by Population



Parents' revealed preference valuation for each CZ, estimated as the first eigenvector of the row-normalized migration matrix for the population of movers (N=21,848,889). The vertical axes plot the rank of each CZ among the top 50 CZs by population in the 2000 decennial census, each of which has a population greater than 1,170,000 residents. Top: effect for children below median parent family income. Bottom: effect for children above median parent family income.

Figure 26: Parents' Revealed Preference by Commuting Zone



Parents' revealed preference for CZs based on moves between CZs. The revealed preference measure is the first eigenvector of the row-normalized migration matrix for the population of movers, and represents the result of foot-voting for destination CZs by moving to them, weighted by the rank of the origin CZs, iterated to convergence. Higher values of v_{cp} are preferred. The mobility matrix is obtained from the population of movers used to estimate exposure effects (N=21,848,889). Top: revealed preference for parents at the 25th percentile of parent family income. Bottom: revealed preference for parents at the 75th percentile of parent family income. Overseas territories CZs omitted.