

Disparate Outcomes from US Domestic Migration

Ben Klemens *

December 10, 2019

Abstract

Economic models of migration, domestic and international, typically begin with the assumption that a moving household's primary goal is to attain higher income than it would earn by staying. This article uses administrative records for almost all people earning formal market income in the U.S., 2001–2015, totaling about 1.7 billion household observations with 82 million long-distance moves, to develop a detailed match between movers and comparable stayers and thus a comparison of movers' income changes relative to stayers. In aggregate, movers see about a median 1% gain in income after moving relative to the counterfactual of staying, with wide variance. Even a decade later, about two out of five households have lower income relative to staying, with an overall median relative income gain of about 6%. Pecuniary benefits are not evenly distributed: movers leaving school and younger single households without children are likely to see higher income relative to staying, but other movers, most notably single parents, are roughly half as likely to see a relative income gain. The overall story is a bifurcated population of movers. Roughly half move to higher income relative to staying, and the rest do not, indicating for whom the hypothesis of income maximization is difficult to support, and where future research about the many motives for moving may focus.

1 Introduction

Moving large distances can be a challenging and costly endeavor for any household. Yet 4.7% of households in the United States choose to uproot themselves and start anew at a location more than 80km away. But is it for money that they pack their boxes, or something else?

With some exceptions, the Economics literature on the subject of migration, both domestic and international, builds on the rationale that moving households do so primarily for higher or less risky incomes, or lower costs [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. Even familial ties are often characterized as primarily useful for reducing costs and raising expected income [12, 13]. Anthropologists and

*Thanks to Randy Capps, Adam Cole, Robin Fisher, Gray Kimbrough, Elizabeth Landau, Ithai Lurie, Elizabeth Maggie Penn, and Joshua Tauberer, and the compilers of the data bank, Raj Chetty, John Friedman, Emmanuel Saez, Danny Yagan, and their counterparts at IRS.

sociologists often explore the non-pecuniary motives of migration [14, 15], but typically in narrative and not empirical form. Doing such analysis on population subgroups would be difficult, as empirical studies conducted to date have largely relied on relatively small samples, such as the results in [10] based on 124 moves or the tabulations of 7,823 moves in [6], or pooled aggregates that lose many individual characteristics [4].

To more fully explore the outcomes and subgroups of long-distance movers, this article uses the universe of people in the United States with formal market income, 2001–2015, to run pseudo-experiments comparing changes in income after a household moves to those for matched households who do not move. The dataset, developed in [16], includes 1,748,802,270 usable observations, with 82,711,474 moves over 80 km.

Aggregate outcomes split the population roughly evenly, with 55.8% of movers seeing a change in income higher than a comparable stayer sees, rising to 63.5% 10 years after the pre-move year.

One could characterize moving as a choice between two bets, one to draw future income given staying and another to draw future income given moving, and then divide the population into groups where the bet has a good payoff and groups where it does not. Section 3.1 shows that relative income gains are primarily driven by movers with certain characteristics, some previously discussed in papers such as [17] or [18].

As the characteristics of a mover deviate from the archetype of the young, childless, early-career single man, movers' incomes fare relatively worse. Households moving while leaving school do especially well, seeing a median income change 22.6% larger than the counterfactual median income change if the household stayed. In aggregate, the median change in income of movers not leaving school are 0.4% above the counterfactual median change given staying, and only 4.6% above a decade later. For movers not leaving school and over 45 years of age, their median income change is below the median change given staying. Among those under 45 making under \$100,000 a year and who are not leaving school, 75.1% of single men without children see incomes better than comparable stayers a decade after the move, but only 56.7% of single women with child dependents do.

In the framing of the choice of moving as a choice between two bets, the evaluation also requires knowing the level of risk for the bets. Section 3.2 shows that the likelihood that a mover's income rises is greater than the likelihood that a stayer's income rises, and the same is true for the likelihood that a mover's income falls. This effect persists: movers' incomes are more dynamic even 14 years after the move.

The roughly half of the population that has high move likelihood and high expectations of income gain from moving relative to staying have been well covered by the Economic modeling literature. That leaves the other half.

That movers are choosing a bet that is worse in volatility and often a worse payoff implies that their decision process is either based on an incorrect evaluation process, or is simply not primarily about maximizing income. But there are myriad other dimensions worthy of adding to a migration decision model:

moving to a better school district, moving to be closer to relatives, leaving rural areas for the amenities of the city, leaving the city for the open space of rural areas, or even simple variety seeking. Retirees are sometimes given special treatment in the literature[19], sometimes described as moving to save costs and sometimes for amenities or family. Leaving residence at a college is also a common motivation for moving; the college-educated population also gets special treatment in the literature [18, 20] and in this paper. Quantifying the prevalence of these considerations as inputs to the move decision is reserved for future work, but this article details the populations most worth studying to evaluate these motives.

Utility for tax policy This article is a component of a larger study on the demographics underlying models of the inputs to tax revenue calculations, and how they evolve over time. The administration of tax policy is impossible without these models, and the quality of tax policy-related estimates are an increasing function of the quality of the underlying models. The comparison of movers to stayers is one view on a large set of statistics on changes in inputs to the tax forms, most of which will not be discussed in this short article. For example, if a homeowner's primary asset is the home itself, the relationship between move rates and homeownership rates are vital to estimating wealth. Health insurance appears in several points of the tax calculation, and is very likely to change for movers—especially those whose income shifts significantly—making the tabulations directly applicable to imputations and projections of health insurance. Not all retirees stop working, and the categorization into all-else-equal cells detailed in Section 2.2 implicitly includes a tabulation of retiree incomes necessary for any retiree-specific estimates. Whether home ownership and local tax payments enter into the tax calculation depend partly on marital status, and the all-else-equal categorization also provides a detailed evolution of the joint distribution of these elements. Migration is a good predictor of simultaneous changes in these elements, so projections and imputations that treat movers and stayers identically will underestimate the correlations in changes. All of these tabulations are necessary components of the analysis underlying this article, although the results presented in this short paper focus on only one component, income.

It is of course the primary input to the tax calculation, and this article argues that models of tax revenue would benefit from separately treating the moving subpopulation as having a more volatile income than the bulk of the population. A model where movers are solely maximizing income and minimizing costs would produce very different predictions for the effect of a policy shift than would a model accommodating a bifurcated population where up to half of movers are not pecuniarily motivated.

2 Methods

This section gives a brief overview of the key terms and caveats; full detail can be found in the methods appendix.

2.1 Definitions

The unit of analysis is the household, not the individual.

All information is taken from informational returns, typically based on addresses during the given tax year, and income tax forms filed by households, typically in the first four months of the next year. Where possible, addresses are based on informational statements sent by employers and other parties based on addresses in the year the transaction transpired. The data set includes households who gain almost any type of formal market income, have a mortgage, paid tuition, or filed a personal income tax return (IRS Form 1040).

Educational attainment and race are often pointed to as causes of differences in migration outcomes, but they are not available via tax forms. However, whether a person is currently paying tuition to a school is available, and will be shown to be of significance below.

Let R be the reference year, before any potential move. A *move* is a change of address over over 80km, measured from the closest interior point to the centroid of the U.S. Postal Service Zone Improvement Plan (ZIP) code in year R to the interior centroid of the ZIP code in year $R + 1$. Although other statuses will be measured at time $R + 2$, $R + 3$, \dots , the classification as a mover is strictly based on a change between years R and $R + 1$. The goal of the pseudo-experiments is to observe the outcome from the single event of the case group moving and the control group staying. Therefore, households do not change status, even if households tagged as stayers choose to move later, or households tagged as movers return home in year $R + 2$.

Pseudo-experiments below will measure changes in characteristics for movers versus stayers in year $R + N$, $N = 1, 2, \dots, 14$. Each uses all available data for that year, including 2001–2002, 2002–2003, \dots for $R + 1$, but only 2001–2015 for $R + 14$. Therefore, year-to-year comparisons are not a panel, as some households will have entries in some spans but not others. Results presented below will usually begin with pre- and post-move characteristics at year $R + 2$, and end at year $R + 10$, so that all results are based on post-move data unlikely to include pre-move information for part of the year, and at least five years of data.

Because returns are based on tax data, definitions match tax law more than definitions often used in the migration literature. For example, a household on a tax form is one earner or two married earners plus a set of dependents, even though real-world households show diversity beyond this form. We do not know if a household owns or rents their domicile, but we do know whether they have a mortgage. We know if a household has dependents under 18 years old, but do not know if the household heads are their parents.

Each ZIP code is also given a population density in high, medium, and low categories, cost of housing as a percentage of income (a proxy for cost of living),

and unemployment rate.

Students paying tuition receive a 1098-T reporting their tuition paid. Define *leaving school* as having a 1098-T in year R , but not having one in year $R + N$.

2.2 *All-else-equal* cells

This article uses a pseudo-experimental method that begins by putting households into cells where all members of a cell are identical in period R adjusted gross income category, sex if single, marital status, age category, number of dependents under 18, UI reciprocity, retirement status, local tax category, housing tenure, tertiary enrollment status, and location population density category, cost of living category, and unemployment rate category. For a given cell, some statistics describing decisions and outcomes for movers versus stayers will be calculated. Statistics for each cell will be built, then aggregates reported for the full population.

For relative change in income, the core statistic, herein Δ , is the aggregate percent change in income for movers minus the aggregate percent change in income for stayers in the same cell. For example, a zero Δ at $R + 2$ for a given cell indicates that movers with the list of characteristics associated with the cell see an average post-move income change two years after the reference year R that is equal to the average income change for stayers with exactly the same characteristics two years after year R . Statistics describing the distribution of the per-cell values of Δ as a whole and for subsets of cells will be presented below.

In this data set, as in the literature generally, income is log-normally distributed, a distribution which standard central limit theorems show can be generated via a series of random multiplicative changes in income. The Δ statistic reflects this by measuring the multiplicative change rather than alternatives such as nominal dollar amounts.

The risk of an income rise for movers in a given cell is the count of movers whose income rose divided by the number of movers. The risk ratio is that risk divided by the analogous risk for stayers. The next step is to aggregate risk ratios for each all-else-equal cell to a population-level statistic, via the Cochran-Mantel-Haenszel (CMH) statistic [21, 22, 23], an aggregation method commonly used in epidemiological studies. See the methods section for the full definition and a brief discussion of its characteristics.

Dividing populations into cells obviates the need to simulate *ceteris paribus* experiments using regressions or other models built using a subjectively selected set of covariates. There is still the problem of reporting a 15-dimensional distribution in a comprehensible manner, which this article primarily approaches by listing aggregates and some descriptive statistics about their distributions, sometimes for the full population and sometimes for subsets excluding certain groups that show distinct patterns.

The size of the data set puts it beyond the utility of statistical methods for hypothesis testing. The data set is not a sample, so any uncertainty depends on non-statistical factors such as whether the data gathering process detailed

in the methods appendix was reliable and whether differences are large enough to be policy significant. Even if the data set were a sample, with over a billion observations even a 0.1% subset gives CMH statistics with a χ^2 distribution of over a million degrees of freedom.

3 Results

Following the metaphor of moving or staying as a choice of bet, this section will first present the expected payoffs for the bet, then the risk of gain or loss. Both will be traced for a decade from the reference year, and presented for the full population and several subgroups.

3.1 Relative change in income

Figure 1 displays the distribution of Δ s, the expected difference in change in income for movers versus stayers. The darker, more peaked curve in Figure 1 is the distribution of Δ at year $R + 2$. As time progresses, the distribution gets wider and starts to skew upward, as per the lighter curve showing the distribution at year $R + 10$. In earlier periods, the distribution for the overall population is largely balanced, with a median within a few percent of zero. But a not-small portion of the distribution stretches out to a 50% larger change for movers or for stayers or beyond. For some subgroups, the distribution is not so symmetric. The right-hand histogram is the distributions of Δ s for movers leaving school from $R + 2$ to $R + 10$, which have most of their density above zero. Outcomes for this group will be discussed in detail below.

Table 2 shows the characteristics of Δ for several subsets of the population in years $R + 2$ and $R + 10$, for cells with a positive number of movers. The table gives some summary statistics to describe the group and the distribution, including the percent of the moving population in the given group, and for $R + 2$ and $R + 10$ the percent of movers with Δ above zero and the median of the Δ distribution. The mean behaves qualitatively like the median (notably having the same sign) for almost all subsets, but its exact value is sensitive to assumptions about outliers and the treatment of special cases, and so is not reported.

The top line shows the population aggregate, where Δ is positive but not overwhelmingly so: 55.8% of movers have higher income relative to stayers in year $R + 2$, and 63.5% in year $R + 10$.

A person who takes a 10% pay cut to live in a city with a 20% lower cost of living is arguably not doing worse. The second row of the table shows statistics for only those for whom the cost of living category in the post-reference year, as measured via median gross rent as a percent of income, is less than or equal to the category in year R . As a reminder, households who move to a lower-rent ZIP code less than 80km away are classed as stayers, so this subset includes a large number of both movers and stayers. This subset's outcomes are qualitatively similar to those of the full population. Pre-move cost of living is still taken into




Figure 1: Density of Δ , income change for movers relative to comparable stayers. Left: all households. Right: only those leaving school. The taller, darker curve is at year $R + 2$, then the curves for $R + 5$, $R + 9$, and $R + 10$ are increasingly lighter and flatter. The medians appear below the main histogram. The $\Delta = 0$ line is shown for reference. Plot is smoothed using a $\text{Uniform}[-.05, 0.05]$ kernel.

account when generating all-else-equal cells, but the post-move cost of living makes a small enough qualitative difference that it will be set aside for the remainder of the article.

Leaving school. Households who move after leaving school have excellent initial outcomes relative to stayers. Figure 1 presents the histograms of Δ for this group. As per Figure 2, the median Δ is 22.6% at time $R + 2$, with 78.5% showing a higher Δ than comparable stayers.

In the generation of all-else-equal cells, students are divided into part-time (4.3% of movers), undergraduate (1.2%), and graduate students (1.2%). The results printed in Table 2 are based on the distribution of Δ s including all three types of all-else-equal-cell. Dividing into subgroups, the median Δ at $R + 2$ for half-time students leaving school is 24.3%, for full-time undergraduates, 8.3%, and for full-time graduate students, 25.9%.

The median Δ for households leaving school at $R + 10$ is 107.9% of its value at $R + 2$, compared with a growth of 603.2% over the same period for non-college movers. The relative lack of growth in this statistic may indicate that the difference in outcomes for movers with a tertiary education versus those without may mostly be due to the single post-school move.

Given that the outcomes from post-school movers are so distinct, the remainder of this section excludes school-leaving movers and stayers, although

Subgroup	% of movers	$R + 2$		$R + 10$	
		%pos	median	%pos	median
All	100 %	55.8 %	0.01	63.5 %	0.06
All, no cost-of-living drop	71.0 %	54.2 %	0.01	62.5 %	0.06
Leaving school, all	6.7 %	78.5 %	0.23	72.0 %	0.24
All others	93.3 %	53.0 %	0.00	61.9 %	0.05
<hr/> Without school leavers					
Retiring	0.7 %	32.5 %	-0.08	34.8 %	-0.06
Retired	1.5 %	40.2 %	-0.03	45.0 %	-0.02
<hr/> Also without retirees					
19 ≤ age < 25	20.1 %	55.2 %	0.02	69.3 %	0.12
25 ≤ age < 35	24.6 %	61.5 %	0.03	73.5 %	0.14
35 ≤ age < 45	15.0 %	56.1 %	0.01	63.2 %	0.05
45 ≤ age < 55	11.7 %	44.5 %	-0.01	53.2 %	0.01
55 ≤ age < 65	8.8 %	36.6 %	-0.03	43.1 %	-0.02
65 ≤ age	10.9 %	54.8 %	0.00	46.7 %	-0.00
<hr/> Also without age ≥ 55					
\$0 < AGI < \$22,500	17.6 %	58.7 %	0.02	71.7 %	0.12
\$22,500 ≤ AGI < \$50k	13.6 %	53.9 %	0.01	73.5 %	0.12
\$50k ≤ AGI < \$100k	8.5 %	54.8 %	0.01	70.5 %	0.09
\$100k ≤ AGI	3.0 %	39.5 %	-0.02	58.0 %	0.03
<hr/> Also without AGI > \$100k					
Single men, no children	18.9 %	62.9 %	0.04	75.1 %	0.13
Single women, no children	13.3 %	60.6 %	0.03	71.7 %	0.13
Single men, 1+ children	3.2 %	51.6 %	0.00	55.4 %	0.03
Single women, 1+ children	4.8 %	44.9 %	-0.01	56.7 %	0.03
Married, no children	5.5 %	63.6 %	0.04	74.0 %	0.15
Married, 1+ children	9.9 %	61.7 %	0.03	72.3 %	0.12

Table 2: Statistics describing the percent change in income for movers minus percent change in income for comparable stayers, two and ten years after the reference year. The percent of the distribution over zero and the median for each subgroup are listed. Some groups showing distinct patterns—households leaving school, retirees, households with age over 45, and high earners—are incrementally excluded.

those whose education entirely precedes or succeeds their move will remain.

Retiring By the measures here, retirees who move see worse pecuniary outcomes from moving than retirees who do not. The relative lack of pecuniary progress supports claims in the literature that retired or retiring movers are more likely to be seeking amenities or family than income.

The distribution of Δ leans low especially for households moving while retiring, but is also low for the subset of the population already in retirement pre-move.

Because the data and the literature show retirees to be a distinct demographic, those who are in retirement post-move will also be excluded for the subsets below.

Age Even having removed households leaving school or retired or retiring, the returns from moving are higher for the younger, and steadily decrease with age. For example, movers between 25 and 35 have a median Δ in year $R + 10$ of 13.8%, while movers between 55 and 65 have a median Δ in year $R + 10$ of -1.5%.

As we continue to search for non-school-leavers who saw relative gains from moving, those over 45 will be excluded for the next subsets. For the remaining population, all-else-equal cells will still be constructed using age in the pre-move year.

Pre-move income Generally, the distribution of Δ for lower-income earners is more favorable than for higher-income earners. The pattern may be partly a mechanical consequence of a reversion to the mean exacerbated by the higher volatility of movers' incomes.

The small number of moving households with an income over \$100,000/year (in 2010 dollars) have a distribution of Δ lower in median and percent over zero, both at $R + 2$ and $R + 10$. In year $R + 2$, 60.5% of movers in this set see a change in income worse than the change comparable stayers see.

Households making over \$100k, 2.8% of all movers, will also be excluded from the subpopulations to follow, although pre-move income will continue to be used for generating all-else-equal cells for the remaining population.

Household composition Leaving only households under 45 earning under \$100,000/year who are not leaving school, single parents now stand out as having relatively worse outcomes by the measures here. For example, 55.1% of single women with child dependents initially have a Δ less than zero, and at $R + 10$ single men with no dependents have a median Δ that is 10.3% higher than single men with child dependents.

The others, including singles with no children and married households, have roughly uniform outcomes.




Figure 3: Risk ratio of upward (lower curve) and downward (upper curve) shift in wage category, movers versus stayers. At $R + 2$, movers are 106.4% more likely than stayers to see an upward shift in income and 122.3% more likely to see a downward shift, but by $R + 10$ the upward and downward risks are roughly equal.

3.2 Risk of gain or loss

The distributions of Δ , such as those in Figure 2, show a large spread which could be due to dynamic shifts in incomes for both movers and stayers. But the CMH statistics presented in this section demonstrate that the large Δ s seen above are more due to the dynamism of movers.

It is no surprise that movers' income changes are generally greater than that of stayers, but this section adds two additional observations: the risk of change is biased toward income loss, and persists in most cases more than a decade into the future.

The upper curve in Figure 3 shows the Cochran-Mantel-Haenszel aggregate ratio of the risk of a drop in income for movers to the risk of a drop for stayers. The horizontal axis is time from the reference year R , with CMH statistics presented for year $R + 2, \dots, R + 14$. The lower curve shows the aggregate risk ratio for a rise in income for movers relative to the odds that stayers see a rise.

For reference, the risk ratio of one is also plotted. That the curves remain above one throughout shows that the one pseudo-treatment of moving from time R to $R + 1$ has an effect that persists more than a decade later.

Returning to the metaphor of future income as a bet, the bet given moving is in aggregate higher risk in both the up and down directions, with a lean downward. This pattern holds for the population as a whole and for subsets of the population (not reported here due to their qualitative similarity).

4 Conclusion

For some, especially people making a move while leaving school, the intuitive story of moving for better prospects and higher paying work fits well. The 54.4% of movers who are leaving college or are non-retirees under 45 making less than \$100k and not single with dependents eventually see an above-median relative gain in aggregate.

But the more the household drifts from the archetype of the unfettered individual—young, no high-paying income stream, single with no children—the worse the imaginary bet of moving becomes. The 44.6% of movers who are not leaving school, but are over 45, make over \$100k/year, or are single with dependents see at best a break-even median income change initially and at $R+10$ a median income smaller than comparable stayers or only a few percent higher.

Some in all subgroups see relative gains and losses, but looking at the overall population, a comparable story results: most of the population sees a relative loss immediately post-move, and those who see a relative gain by year $R+10$ are still only 63.5% of movers.

The results are intended to open the question of why such a large portion of the population chooses to uproot their lives, changing homes, jobs, schools, and local social networks, in return for a future median return of a few percent or even a loss. Households who see a drop in income may have chosen to move based on an intention to expand income, but were overoptimistic, or simply saw bad luck, or may have had anomalously bad expectations from staying compared to others who match on all 14 available characteristics. There is correlation in the data here between moving and many life transitions, including entering or exiting marriage, adding a child to the household, or retirement. A large percentage of movers change from high-, medium-, or low-density areas to an area with a different density. People often prefer to be near family, even for nonpecuniary reasons. If movers have adventurous personalities, the risks of moving compared to staying may be appealing in their own right.

Do these motivations apply to international moves? On a human scale, all of the above stories are without regard to borders. But the stakes are higher in costs and effort, and in changes in income, so we expect the relative weight of factors to be different. The differences in scale can be an obstacle to testing hypotheses regarding motivations, because if two countries have income differences large enough that a migrant is guaranteed to see a rise in income, it is mechanically impossible to reject the hypothesis that all movers are motivated by income.

Modeling implications The administration of tax policy is impossible without accurate underlying models of the population, and this project provides more detailed inputs. Among the full U.S. population, there is a subpopulation of movers whose characteristics are more volatile. Their income, benefits, housing tenure, education credits, and health insurance statuses are likely to change simultaneously, while for the staying population, changes are less likely and less likely to be simultaneous. As per the discussion in the introduction,

this has application to many of the aspects of the imputation and projection of tax revenue under the current regime and potential changes. Rates of co-change can be vital to some estimates, and a model that separately describes movers and stayers will gain accuracy relative to a model with a distribution giving an equal pattern of changes to the full population.

The results here show that a migration model whose primary focus is on a goal of income maximization is implicitly a model whose focus is on young, low-to middle-income households, probably leaving school, excluding singles with dependents. Such a model is likely to find statistical significance to the claim that households seek higher income—people still prefer more money to less—but the data here shows that its efficacy will likely reach its upper limit at explaining a little over half of moves. A model of a more diverse population would require a wider variety considerations. The results here advise which groups merit greater focus, and may be useful for calibration of more fine-grained models, especially those estimating longer-term outcomes. Analyses with less comprehensive data sets require some structural assumptions, and the results suggest key features those models could include, such as school leavers (but not necessarily past graduates) and single parents.

References

- [1] Borjas, G. Immigration and welfare magnets. Tech. Rep. (1998). URL <https://doi.org/10.3386/w6813>.
- [2] Fee, K., Wardrip, K. & Nelson, L. Opportunity occupations revisited: Exploring employment for sub-baccalaureate workers across metro areas and over time. Tech. Rep., Philadelphia Federal Reserve (2019).
- [3] Greenwood, M. J. & Sweetland, D. The determinants of migration between standard metropolitan statistical areas. *Demography* **9**, 665 (1972).
- [4] Jackman, R. & Savouri, S. Regional migration in Britain: An analysis of gross flows using NHS central register data. *The Economic Journal* **102**, 1433 (1992).
- [5] Stark, O. & Bloom, D. E. The new economics of labor migration. *The American Economic Review* **75**, 173–178 (1985).
- [6] Hernández-Murillo, R., Ott, L. S., Owyang, M. T. & Whalen, D. Patterns of interstate migration in the United States from the Survey of Income and Program Participation **93**, 169–85 (2011).
- [7] Nunn, R., Kawano, L. & Klemens, B. Unemployment insurance and worker mobility. Tech. Rep., Urban-Brookings Tax Policy Center (2018).
- [8] McKinnish, T. Importing the poor. *Journal of Human Resources* **40**, 57–76 (2005).

- [9] McKinnish, T. Welfare-induced migration at state borders: New evidence from micro-data. *Journal of Public Economics* **91**, 437–450 (2007).
- [10] Kennan, J. & Walker, J. R. The effect of expected income on individual migration decisions. *Econometrica* **79**, 211–251 (2011).
- [11] Young, C., Varner, C., Lurie, I. Z. & Prisinzano, R. Millionaire migration and taxation of the elite. *American Sociological Review* **81**, 421–446 (2016).
- [12] Palloni, A., Massey, D. S., Ceballos, M., Espinosa, K. & Spittel, M. Social capital and international migration: A test using information on family networks. *American Journal of Sociology* **106**, 1262–1298 (2001).
- [13] Garip, F. Social capital and migration: How do similar resources lead to divergent outcomes? *Demography* **45**, 591–617 (2008).
- [14] Benson, M. & O’Reilly, K. From lifestyle migration to lifestyle immigration: Categories, concepts and ways of thinking. *Migration Studies* **4**, 20–37 (2015).
- [15] O’Reilly, K. *Lifestyle Migration* (Routledge, 2016).
- [16] Chetty, R., Friedman, J. N., Saez, E. & Yagan, D. The SOI databank: A case study in leveraging administrative data in support of evidence-based policymaking. *Statistical Journal of the IAOS* **34**, 99–103 (2018).
- [17] Clark, D. E. & Hunter, W. J. The impact of economic opportunity, amenities and fiscal factors on age-specific migration rates. *Journal of Regional Science* **32**, 349–365 (1992).
- [18] Quinn, M. A. & Rubb, S. The importance of education-occupation matching in migration decisions. *Demography* **42**, 153–167 (2005).
- [19] Brown, D. *Rural Retirement Migration* (Springer, Dordrecht, 2008).
- [20] Gurak, D. T. & Kritz, M. M. The interstate migration of U.S. immigrants: Individual and contextual determinants. *Social Forces* **78**, 1017–1039 (2000).
- [21] Pearce, N. Effect measures in prevalence studies. *Environmental Health Perspectives* **112**, 1047–1050 (2004).
- [22] Rothman, K. J. *Modern Epidemiology* (LWW, 2012).
- [23] Mantel, N. & Haenszel, W. Statistical aspects of the analysis of data from retrospective studies of disease. *Journal of the National Cancer Institute* (1959).
- [24] Kaplan, G. & Schulhofer-Wohl, S. Interstate migration has fallen less than you think: Consequences of hot deck imputation in the current population survey. *Demography* **49**, 1061–1074 (2012).

- [25] Tcha, M. Altruism, household size and migration. *Economics Letters* **49**, 441–445 (1995).
- [26] Fisher, R., Gee, G. & Looney, A. Sex couples after Windsor: Characteristics of married tax filers in 2013 and 2014. Tech. Rep. 108, U.S. Treasury Office of Tax Analysis (2016).
- [27] Feenberg, D. & Coutts, E. An introduction to the TAXSIM model. *Journal of Policy Analysis and Management* **12**, 189 (1993).
- [28] Poot, J., Alimi, O., Cameron, M. & Maré, D. The gravity model of migration: The successful comeback of an ageing superstar in regional science. *Investigaciones Regionales* **2016**, 63–86 (2016).
- [29] Molloy, R., Smith, C. L. & Wozniak, A. Internal migration in the United States. *Journal of Economic Perspectives* **25**, 173–196 (2011).
- [30] Grubestic, T. H. & Matisziw, T. C. On the use of ZIP codes and ZIP code tabulation areas (ZCTAs) for the spatial analysis of epidemiological data. *International Journal of Health Geographics* **5** (2006).
- [31] Krieger, N. *et al.* ZIP code caveat: Bias due to spatiotemporal mismatches between ZIP codes and US Census-defined geographic areas—the Public Health Disparities Geocoding Project. *American Journal of Public Health* **92**, 1100–1102 (2002).
- [32] Ratcliffe, M., Burd, C., Holder, K. & Fields, A. Defining rural at the U.S. Census Bureau. Tech. Rep. ACSGEO-1, U.S. Census Bureau (2016).
- [33] Forsyth, A. Defining suburbs. *Journal of Planning Literature* **27**, 270–281 (2012).

5 Methods Appendix

This section describes the sources of data, the precise meaning of each measure, and how missing or unusual elements are handled. Once the data set is prepared as described here, the calculation of statistics as reported in the main article are straightforward.

The core data set was compiled and cleaned by Raj Chetty, John Friedman, and Danny Yagan in cooperation with the U.S. Treasury [16]. The universe of observations is the set of people who filed or received one or more of the following forms:

- The individual tax return: Form 1040, 1040A, or 1040EZ.
- Any wages, salaries, or tips: Form W-2.
- Social security payments: Form 1099-SSA.
- Mortgage interest: Form 1098.

- Retirement benefits: Form 1099-R.
- Unemployment insurance (UI) benefits: Form 1099-G.
- Non-employee (typically contractor) payments, including payments from the contracting party or third parties: Forms 1099, 1099-MISC, and 1099-K.
- Tuition statement: Form 1098-T

The data bank includes one observation for every person for whom any of the above forms exist for any year, 2001–2015. The deceased typically have no income or informational returns, and so disappear from the data set without explicitly being handled. Those international movers who file a return with no ZIP code also disappear from the data set.

The process of producing a single sequence of data for each individual is described in [16]. The tax and income records are matched to Social Security Administration records using Social Security Numbers (SSNs) to determine birth dates and sex.

The U.S. government is happy to collect taxes from noncitizens, who are given an Individual Taxpayer Identification Number (ITIN) in lieu of a Social Security number, and are included in this data set. The application for an ITIN (Form W-7) asks for sex and date of birth.

The analysis is based on the full data set, without sampling.

Missing data is not imputed. Many of the variables, including having a mortgage, UI payments, and Social Security payments, are based on the presence or absence of a form from an administrative agency. The lack of presence of a form is always coded as a zero, and determining which zeros are in error is a question of understanding problems in form processing which are unlikely to be amenable to typical imputation methods aimed at data missing at random. Methods to impute missing survey data have been known to create anomalies in the completed data. For example, the U.S. Census Bureau completes incomplete or missing survey responses on many major surveys using a method that simulates reusing punched cards that have been run through an overheating card reader (“hot deck imputation”), which [24] show to be the cause of a large part of the reported decline in interstate migration during the early 2000s.

The pseudo-experimental method requires dividing all statuses into categories, including income and tax payments. Income is divided into twenty categories, but local taxes have inherent measurement limitations that make fine-grained categories difficult, as discussed further below.

The migration decision may better be modeled as being made by a full household, rather than by an individual [25]. The administrative data is a correct representation of the full household in cases where there is one earner or two married earners plus dependents with no substantial income of their own, but there are myriad other types of household. Nonetheless, the word *household* will be used as the unit of discussion below to stress that the statistics reported are not by head count.¹

¹Perhaps *taxpaying unit* is more accurate, but is deeply dehumanizing.

5.1 Per-status details

Marital status is defined by filing status: married filing jointly or married filing separately versus filing as single or head of household.

Filers who mark themselves as married filing separately are taken as separate entities. There are 28,647,209 returns in the data where both spouses filed a separate return, of which 26.1% list a spouse in a different ZIP code. Of those where one spouse moved, 40.1% of the other spouses did not move. This gives evidence that many people choose to file separately because they are maintaining separate households.

All income is summed for both spouses filing jointly, and if either has a mortgage payment or wage, Social Security or UI income, the household will be listed with that status.

The data set includes each spouse in a married-filing-jointly household, meaning that a simple sum would double-count the households, total income, and taxes. The solution is to simply count each individual as half a filer, then total as usual.

For total counts, and income and tax dollar totals, a household is a single unit. When a person enters or exits a marriage, the interpretation is normally correct, because counts are always based only on the household as described in the pre-status state. However, the categorization into AGI category as used in the CMH calculations has the anomaly that if two filers in a low-AGI bin marry but do not change jobs, those filers find themselves in a higher-AGI bin post-marriage, which gives the misleading impression that both found higher pay. To solve this problem in cases where post-AGI bin is a consideration, cells with pre-period singles and post-period marrieds use a post-AGI based on half of the household's post-income, and cells with pre-period married households who become post-period single households are compiled using pre-AGI based on half of the household's income. Most situations in this article do not use post-move AGI and this correction is not applied.

The status of primary filer has no legal meaning and is not used. In practice, [26] find that in cases of married couples, more than 95% of returns have a male primary filer and female secondary filer, making it a noisy measure of sex.

During this period, same-sex couples were not legally recognized by the IRS. Nonetheless, about 0.1% of couples are same sex, with roughly equal male and female pairs.²

Children are those people listed as dependents whose age is less than or equal to 18. The status categories are 0, 1, 2, and 3 or more children. Dependents are frequently re-listed on multiple returns, often because separated parents both list the same child. This is incorrect for tax purposes, but the multiple listing

²Because some U.S. Census Bureau surveys had a policy during this period of editing same sex couples by changing the gender of one partner, the administrative tax data may be the only source of migration data on same-sex couples; such study is reserved for future work. Any couple that marks itself as married on the tax forms is counted as married in this article, regardless of genders.

may be a better description of the household than the tax-conformant rule that each dependent be claimed by only one household.

Sex is derived from matching Social Security Administration records with the tax records. The join is very effective, assigning a sex to nearly all records. [26] discuss misclassification of sex in the SSA records themselves, and find errors in the SSA records that are small enough to be ignorable for the purposes of this article.

Households with a married status (filing jointly or separately) are treated as a third category in the tabulations with regard to sex, and will be ignored in the statistics reported by sex.

Age is based on Social Security records and is divided into these categories:

Category	1	2	3	4
Ages	19–24	25–44	45–64	65+

Those with a pre-move age under 19 are entirely excluded from the data set, including those moving out of their homes to go to college. It is perhaps beneficial to exclude them, because the question of why people who move to college chose to move answers itself, and this makes their decision process distinct from the decision process of the great majority of the population.

For married couples, the mean of the filer's and spouse's age is used.

In school Students are in school if they receive a tuition statement, Form 1098-T. The form includes information about whether the student is in school part-time, tuition is toward full-time undergraduate, or full-time graduate studies. Students who have no tuition, or whose tuition is paid by an employer or military program, receive no 1098-T. Because the 1098-T first appeared shortly before the period under study, there is no way to determine from the data set what previous education any given adult has received.

Adjusted gross income (AGI) is income from all sources, minus a short list of expenses including student loan interest and health savings account deductions. Dollar amounts are inflation-adjusted to 2010 dollars. AGI is divided into a zero bin, then 19 additional cut points that are roughly equidistant on a log scale, beginning with cut points at \$1, \$5,640, \$9,005, and on up to \$148,676 and \$216,033+.

Unemployment insurance (UI) This is coded as one if there is any positive payment on form 1099-G—even for a single week of income—else zero.

Local taxes On Schedule A of the 1040, depending on the year, filers could deduct state and local income taxes, real estate taxes, a sales tax allowance,

Status	2001 deduction	2015 deduction
Single	\$4,550	\$6,300
Married filing jointly	\$7,600	\$12,600
Married filing separately	\$3,800	\$6,300
Head of household	\$6,650	\$9,250

Table 4: The range of standard deductions, for use in dividing local taxes into categories.

personal property taxes, new motor vehicle taxes, and some other taxes. The local tax measure in this article is the sum of these.

More than half of the population did not file a Schedule A. One can expect (but not guarantee) that a person who received taxable income would claim any local taxes over the standard deduction. Based on this assumption, local taxes are divided into four categories:

- Category 0: Households with zero income and no Schedule A, or a Schedule A with a zero report for local taxes.
- Category 1: Other households who did not file a Schedule A. These taxpayers are likely to have paid at least \$1 in those local taxes listed on Schedule A.
- Category 2: Households with local taxes between the standard deduction and 1.25 times the standard deduction.
- Category 3: Households with local taxes above 1.25 times the standard deduction.

The cutoff of 1.25 times the standard deduction splits the deduction-taking population roughly in half, but only 33.0% of observations filed a Schedule A claiming nonzero local taxes.

The standard deduction rises every year. To give a sense of the range, Table 4 shows the deduction for each filing status at the beginning and end of the period.

This is a coarse measure chosen based on available data, but nonetheless, 19.6% of moving observations did change categories as given here. This makes change in local tax category more common than some other changes; for example, 0.9% of movers enter retirement.

By using the exact amount of local taxes paid, there is no need to consider nominal region-wide tax rates, or to rely on simulations such as those provided in the current version of [27].

Federal tax Federal tax is unused in this paper, but is included in the data set. The measure is total taxes before credits. Credits such as Earned Income Tax Credit and the child tax credit are federal programs that are tied to income and

therefore listed on the tax form, so using the pre-credit is the desired measure for policy question. Prepayment such as withholding is not relevant.

Tax is divided into nine bins, with a zero bin, then cut points at $\log_{10}(\text{tax})=1.5, 2.5, 3, 3.5, 4, 4.5, 5$, and up.

Retirement People can begin receiving Social Security income (SSI) between ages 62 and 70, and can continue to work while receiving benefits, subject to an income limitation.

Setting aside death, leaving social security is very rare—and even if a recipient dies, his or her spouse continues to receive benefits. This article assumes that once persons first appear on Social Security, they stay on.

One can make pre-retirement withdrawals from an individual retirement account (IRA) under certain circumstances or by paying a penalty. After age 59.5 there is no penalty. The compromise used in this data set is to count an IRA withdrawal as potentially toward retirement income if the filer or filer's spouse is over 59.

Our concern is how the decision to move is affected by the life event of retirement, rather than the beginning of a new SSI or IRA income stream. A household is coded as in retired status iff it is marked as having retirement income as above, and has wages smaller than Social Security and IRA income.

Geography: moves All of the administrative records include a residence address with a Zone Improvement Plan (ZIP) code from the U.S. Postal Service. Where possible, addresses are taken from the informational returns (W-2, 1099, ...) that are sent early the next year after the given tax year, often using an address registered the year before, with the address on the tax return filed as late as April of the next year used as a last resort.

The U.S. Census Bureau tabulates individuals at residential addresses of various types into ZIP Code Tabulation Areas (ZCTAs) which largely match ZIP codes. But a large number of tax forms list post office boxes or business addresses large enough to merit their own ZIP code, so another crosswalk from the U.S. Health Resources and Services Administration imputes each ZIP code to its closest ZCTA-related ZIP code.³ After this process, a small fraction of a percent of households are not mapped to a ZCTA, and are thus excluded from the analysis.

The ZIP codes are used for two types of calculation: determining whether a move occurs and location characteristics.

Although one could treat distance travelled as an output to the process, as is commonly used in gravity models such as [28], this article restricts itself to the binary question of whether a person moved or did not.

For any Census-defined areas, the Census gazetteer includes the latitude and longitude of an *interior point*, which is the closest point to the centroid within the area. U.S. law offers a tax credit for long-distance moves, which are roughly defined as a change in workplace of over 80km (50 miles). Following this, a

³<https://www.udsmapper.org/zcta-crosswalk.cfm>

move is defined as changing to a ZCTA whose interior point is more than 80km from the interior point of the previous ZCTA.⁴

The literature often defines a long-distance move as a crossing of state boundaries [6], but many metro area boundaries or commuter rail lines cross state lines (Boston, Chicago, the District of Columbia, Kansas City, New York, Philadelphia, Pittsburgh, St Louis, ...). [29] report additional problems with using a more structured (and therefore more complicated) definition of moving based on commute-based statistical areas (CBSAs), consolidated statistical areas (CSAs), and states.

There is some disparity between ZIP codes and ZCTAs, which easily creates issues for fine-grained analysis [30, 31]. But this analysis here is based on the characteristics of large areas and moves over 80km which are unlikely to suffer neighborhood or block-level issues.

Geography: characteristics An area is given three characteristics, intended to represent the health of the local job market, cost of living, and population density. The focus of the article is on individual statuses, so there may not be benefit in burdening the model with several typically correlated measures to describe each class in detail.

The only source for ZIP-level economic characteristics is the U.S. Census Bureau's American Community Survey, using five-year aggregates, which first became available in the middle of the time span studied here. Unemployment rate and housing cost will be fixed at the 2011 value, the closest available year to the middle of the range of years used here. Therefore, rapid changes in these categories in the ten years prior to 2011 and four years after, for example where a market swings from low-ranked cost of living to high-ranked cost of living, are not accommodated.

Each individual is already marked as receiving unemployment insurance or not, but if a person is considering a job change, the overall state of the market may have an impact on location choice. Unemployment rate is divided into three categories, chosen to put about a third of ZIP codes in each category, with cut points between categories at 5.6% and 9.4%.

The second local characteristic, median gross rent as a percentage of household income, is used to represent cost of living in an area. Again, there are three categories with an even number of ZIP codes in each, with cutoffs at 27.3% of income spent on rent and 34.6%.

The third characteristic, population density, is defined as the number of people living in a ZIP code as of the 2010 Census divided by square kilometers of land, and is divided into urban, suburban, and rural categories.

The density data was generated using tract-level data from the Census Bureau's Geography division, aggregated to ZIP code using a crosswalk file pro-

⁴The calculation is approximate, using the straight-line distance $\sqrt{111(lat_1 - lat_2) + 85(lon_1 - lon_2)}$. This formula is a compromise between precision and speed of calculation.

	Income does not rise	Income rises
Stayed	A	C
Moved	B	D

Table 5: A reference for the risk ratio calculation.

vided by the U.S. Department of Housing and Urban Development.⁵

The Census’s block-level definition of a rural area assigned 19.3% of the 2010 population to rural areas [32], and using a cutoff density of 55 households per km² over the entire ZIP code comes close to this percentage.⁶ There is limited consensus about how to define suburbs [33], and the Census Bureau does not offer a definition. A cutoff of 500 households per km² gave an even split of the population between middle- and high-density areas, and will be used in this article. This density is more typical of outer suburbs.

5.2 The Cochran-Mantel-Haenszel statistic

Figure 5 shows an all-else-equal cell for reference. For a rise in income, the risk ratio is the likelihood of a mover’s income rising between period R and $R + 2$, $D/(B + D)$, divided by the likelihood of a stayer’s income rising, $C/(A + C)$:

$$\frac{C(B + D)}{D(A + C)}.$$

For a set of cells where cell i has total population $n_i \equiv A_i + B_i + C_i + D_i$, the CMH aggregate is

$$\frac{\sum_i C_i(B_i + D_i)/n_i}{\sum_i D_i(A_i + C_i)/n_i}.$$

A full analysis of the CMH statistic is beyond the scope of this article, but it has some properties that make it appealing for the aggregation of statistics from a large number of cells. For example, if one cell for single men has risk ratio RR_m , and a cell with otherwise identical characteristics for single women has risk ratio RR_f , then we are guaranteed that the CMH statistic calculated by ignoring sex and merging the two cells is bounded between these two values. If a cell has twice as much weight as another, it is given twice as much weight in the numerator and denominator of the CMH statistic. A simple average of risk ratios could be significantly shifted by one cell with an extreme risk ratio (possibly even infinite), but summing numerators and denominators separately allows the cell’s information to be added to the bulk of other cells’ information while mitigating the instability of small-cell estimates.

⁵https://www.huduser.gov/portal/datasets/usps_crosswalk.html

⁶The Census definition is based on the block level, while this article uses the ZIP code level. A block is by definition a populated area, but most ZIP codes include some uninhabited areas, from parks to freeways to entirely commercial areas, making the Census’s urban cutoff of 386 people per km² inapplicable at the ZIP level.

5.3 Data availability statement

All data is held on servers at the U.S. Internal Revenue Service's Statistics of Income division. Access is available under the conditions of 26 U.S. Code §6103.