An Analysis of U.S. Domestic Migration via Subset-stable Measures of Administrative Data

Ben Klemens*

December 31, 2020

Abstract

How does the likelihood of moving across U.S. regions vary with changes in household characteristics, and how does the risk of a change in status vary given a move? Statistics aimed at these questions are calculated for households who earned formal market income in the U.S., 2001–2015, totaling about 1.7 billion observations with 82.7 million long-distance moves, and covering statuses such as income, education and employment status, local and federal tax payments, number of children, retirement or marital status. The key theoretical result of this article shows that the Cochran-Mantel-Haenszel statistic is the unique aggregate risk ratio within a broad class that has the "subset stability" property: If a statistic has value s_1 for one subset and s_2 for another, then the statistic for the union of the two sets is between s_1 and s_2. A sequence of pseudo-experiments generate a wealth of tests regarding the relationship between moving and a broad range of household characteristics, for the full population and salient subsets, with some focus on the characteristics of the 44.2 % of movers who see negative income returns relative to the counterfactual of staying.

JEL categories: J61, C14, H24, D19
Keywords: migration, administrative records, demographic analysis, tax

1 Introduction

This article estimates the relationship between the choice to move within the United States and a set of statuses including characteristics such as marital status, number of children, employment status, income, and local tax payments. The data set includes the full universe of earners in the formal U.S. economy, totaling 1,748,802,270 observations. Of these, an average of 4.7 % of households move over 80km from any two-year period to the next, giving 82,711,474 moves used in the analysis.

*Thanks to David Bridgeland, Randy Capps, Adam Cole, Aaron Schumacher, Bethany DeSalvo, Robin Fisher, Chung Kim, Gray Kimbrough, Elizabeth Landau, Ithai Lurie, Nick Turner, Elizabeth Maggie Penn, Joshua Tauberer, and the compilers of the data bank, Raj Chetty, John Friedman, Emmanuel Saez, Danny Yagan, and their counterparts at IRS.

The article focuses on two questions: all else equal, how do certain characteristics change the likelihood of moving over 80km (50 miles), and given such a move, what is the relative likelihood that the characteristics of a household change? A pseudo-experimental method is developed to answer these questions.

The method is intended as a drop-in replacement for regression methods, to generate statistics that answer questions of the form *holding all other factors equal, how does an outcome change given a change in one characteristic?* A typical pseudo-experiment begins by dividing the population into cells that match on thirteen characteristics but differ on one, then calculates statistics to express, for example, the ratio of likelihood of moving for observations given a mortgage to the likelihood of moving for the all-else-equal observations who do not have a mortgage. These pseudo-experiments will be re-run up to fourteen years after the move.

Without the need for such delicate model building, this article and the extensive tables in the appendix test hundreds of hypotheses using fully controlled experiments.

Raw risk ratios (also known as relative risks) and typical estimates via regression models sometimes lead to awkward conclusions regarding subsets of the data. If we know a household makes more than $50k/year then its risk of relative income gain after moving is lower than the risk for the general population, but it is also lower if we know a household does not make more than $50k/year. That is, *any* income information for an observation lowers its associated risk ratio, and this is a difficult paradox to design policy around. The ideal aggregate risk ratio would be *subset stable*, meaning that if subgroup A has some statistic RR_A, and the set of observations outside of subgroup A has statistic $RR_{\neg A}$, then for the full data set the statistic is bounded between RR_A and $RR_{\neg A}$.

Two novel results, Lemma 1 and Theorem 1, state that the Cochran-Mantel-Haenszel (CMH) statistic [Mantel and Haenszel, 1959], an aggregate risk ratio measure originally developed in the literature on controlled medical experiments, is the *unique* subset stable aggregate risk ratio within a very broad class of candidates.

Smaller surveys, typically analyzed via generalized linear models, require great care in model design, meaning that only a handful of hypotheses can be tested at a time. The pseudo-experimental method dispenses with many such issues, including selection of interaction terms, nonlinearities, and causal confounders when causality is appropriately directed; in the language of causal modeling, a pseudo-experiment is a *do* operation over all observable data.[1] For situations where a single direction of causality is reasonable, such as the ef-

[1]Nonetheless, these are pseudo-experiments. A true random experiment would allocate moves or characteristics via an external source, as close as possible to dropping a random household in a new location. But such an inquiry is not useful in a free society where such helicopter drops or allocation of characteristics are extremely infrequent. Military moves due to redeployment perhaps approximate a truly random allocation [Burke and Miller, 2017], but individuals who chose to be in military families have unobservable characteristics systematically different from those who do not. Whether these results would apply to families where one member is randomly drafted into the military is unknown, and perhaps not worth studying in a situation where there is no such draft.

fect of many of the pre-move statuses on moving, the results of these pseudo-experiments can be read as hypothesis tests, providing a controlled measure of the effect of appropriate characteristics on moving, or of moving on appropriate characteristics. The traditional null hypothesis of no effect is always rejected, because every pre-status will have some statistically significant effect on move propensity, and moves will have an effect on every post-status. An appendix to this article gives further detail on the pseudo-experimental method.

Summary of migration results No effort will be made to provide a grand unified theory of migration, but the pseudo-experimental method and data set allow estimation of a wealth of details about the relationship between household characteristics and moving, for the overall population and—via the use of subset-stable statistics—diverse subgroups. This section lists 22 headline results.

In fact, the results show that there may be no grand unifying theory to be had. The economics literature focuses heavily on moving for pecuniary reasons, but a claim that this is the primary motivator for moving does not hold up for a very large portion of movers. Instead, we must look at the many individual situations like retirement, marriage, or leaving school to cover more of the population.

Although some of these results appear as stylized facts in the literature, comprehensive administrative records allow their quantification without survey error, provide new information for situations where the opposite stylized fact also appears in the literature, make possible results about features like local tax payments available only via administrative records, allows tests a decade after the reference year, and allow the calculation of statistics regarding small subsets of small subsets like those earning under $12,000/year and exiting a marriage during a move (0.1 %, which would be about 100 observations in a survey with 100,000 respondents, but here is 1,929,910 observations).

1. The analysis finds abundant evidence that money is not a key consideration in the move decision for a very large percent of movers: 41.8 % of movers are no better post-move in terms of adjusted gross income (AGI) relative to the counterfactual of staying, coarsely-measured local tax category, or cost of living; are not retiring; and are retaining the same mortgage and unemployment status. Aggregate income outcomes initially 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 % ten years after the pre-move year.

2. Looking only at those whose cost of living does not change (measured via median rent as a percentage of income in an area) does not have a qualitatively significant effect. The U.S. government's General Services Administration specifies locality pay adjustments to the standard payment schedules, which provides a well-researched measure of cross-regional differences. As of 2017, the difference between lowest and highest locality adjustment is 15.3%, setting aside the outlier of the San Francisco area

(22.2%).[2] A household wishing to raise or lower its costs by 15% could readily do so without moving over 80km, so we expect to have cost-cutting and cost-upgrading households in both the cross-region movers and within-region stayers.

3. Every major change in household status (except a fall in the number of dependents under 18) raises the likelihood of moving.

4. This unique data set allows us to trace households up to fourteen years past the reference year, and the relative dynamism persists for most characteristics. That is, over a decade after the move, movers remain a distinct population.

5. Regarding the debate over whether movers seek lower taxes or more amenities, holding relevant characteristics constant reveals that local taxes are more likely to go up for movers relative to stayers than to go down for movers relative to stayers. This effect also persists for several years after the move.

6. Using the microdata reveals that there is a great deal of churn in statuses after a move. For example, holding all else equal (including mortgage status, children, presence of unemployment income, school status, . . .), moving inflates the chance of marrying by 37.6 % but it raises the chance of exiting marriage by 37.2 %. Using only aggregate totals would obscure such bidirectional changes and the dynamism of the moving population.

7. The likelihood of moving holding all else equal is U-shaped in income, with those in the highest and lowest AGI bands more likely to move than those in the middle bands.

8. 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. But those supranormal returns from moving are tied only to that single initial post-school move, not steadily increasing post-move gains.

9. The effect of schooling is mostly due to exits from part-time education or graduate school; full-time undergraduates see a median income gain only 8.3 % larger than the counterfactual non-moving income gain.

10. Those in tertiary education are decidedly more mobile than the general population, especially those leaving graduate school. Although 6.3 % of all observations are in school, 18.3 % of moving observations were in school.

11. Concerning the all-else-equal likelihood of moving forward in school enrollment during a move, the likelihood is lower for those earning under $12,000/year (2010 dollars) than those earning more.

[2] `https://www.federalpay.org/gs/locality`, accessed December 2020.

12. Moving single heads of households with dependents under 18 see worse income outcomes from moving than singles without children or married households. 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 55.4 % of single men with child dependents do.

13. The risk of a mover's income falling relative to the counterfactual of staying is higher than risk of a mover's income rising relative to the counterfactual of staying. This holds across many subgroups, and persists for several years before the relative risks of rise or fall converge.

14. The all-else-equal likelihood of moving for middle-income earners is lower relative to others.

15. Holding all else equal, most notably housing tenure and dependent count, a married household is 1.18 times as likely to move than an otherwise identical unmarried household.

16. Concerning the all-else-equal likelihood of exiting marriage during a move, the likelihood is greater for those earning under $12,000/year (2010 dollars) than those earning more.

17. There is also a great deal of churn between location characteristics, as movers are about as likely to move to higher cost-of-living areas as they are to move to lower [Walker, 2017]. The same holds for unemployment rate categories.

18. By definition, retiring and retiree movers are not moving primarily for work, meaning they must be attracted to social network or amenities. However, holding all else equal, retiree movers are more likely to see a lower future income stream than non-moving retirees, which advises that they may be seeking cost-cutting amenities more than lifestyle-enhancing amenities.

19. By raw percentages and looking only at the unmarried, women move more than men; controlling for other statuses, including housing tenure, men are more likely to move than women. Similarly, the married move less often than singles by raw percentages, but controlling for all other characteristics are more likely to move.

20. By raw percentages and looking only at the unmarried, women move more than men; controlling for other statuses, including housing tenure, men are more likely to move than women.

21. Those who exited unemployment insurance in the all-else-equal pseudo-experiments are 1.49 times as likely to have moved than those that did not move, suggesting that moving may be a popular strategy for finding a job among the 6.6 % of the observations that received some unemployment compensation on average over all years.

22. Even controlling for all relevant factors in the data set, federal taxes for movers are still more volatile than for non-movers.

Utility for tax policy This article is a component of a larger study on the population characteristics underlying models of the inputs to tax revenue calculations, and how they evolve over time. 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 not discussed in this article.

For example, if a home owner'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, making the tabulations directly applicable to imputations and projections of health insurance.

Income 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, who have more volatile income than the bulk of the population. Even holding income and other relevant factors constant, there is still more volatility in federal tax revenue from movers than from stayers.

Economic models frequently focus only on pecuniary gains, but this may not be ideal given the results here where up to half of movers do not show evidence of being pecuniarily motivated.

Outline Section 2 reviews the literature on mereological methods, the use of administrative records in migration research, and the household characteristics used in this article.

Section 3 briefly describes the statuses used in this article. An appendix answers detailed questions about their precise construction.

Section 4 describes the pseudo-experimental case-control method, including the procedure of defining all-else-equal cells, how one could calculate risk ratios for each, and how the CMH statistic is used as an aggregate risk ratio. The use of the CMH statistic is motivated by the novel uniqueness result in Theorem 1.

Section 5 presents results. But the pseudo-experimental method allows the efficient construction of controlled experiments via relatively simple tabulations, so the results section organizes its many results via a sequence of distinct narratives, beginning with population-level estimates regarding income, taxes, and individual characteristics, continuing to subsets that have a large influence on the top-level measures or address certain questions in the literature. Readers working on certain subjects should be able to focus on the relevant segments without difficulties.

There are four appendices to this article.

In a study of this magnitude, there are endless minor details of data handling to be dealt with. Appendix 1 gives full detail on the construction of the characteristics briefly described in Section 3. Detailed questions about the construction of variables used in the paper should be directed to this appendix.

Appendix 2 discusses the CMH statistic further, including notes on its characteristics and proofs of the results. It includes a section entitled *Why not just run a 13-dimensional regression?* going into far greater detail on that question for readers who are not specialists in methodology.

Appendix 3 provides additional results using this unique data set, including more about the relative risk of income change for some subgroups, unemployment, mortgage, age and retirement, children, federal tax, geographic characteristics, and low income movers.

Appendix 4 gives per-year tables of basic percentages and CMH statistics. The results may be useful to researchers interested in certain characteristics. Given that distribution-based statistical methods are not appropriate for large-scale administrative data, the per-year tables allow evaluation of the reliability of estimates. Over the course of 2001–2015, few of the statistics stray more than a few percent from the statistic calculated using all years.

2 Literature

This appendix covers the literature on the methodological question of subset stability, followed by a review of the methods used to generate statistics regarding migration comparable to the microdata approach used here. It will then review the literature regarding each of the variables that will be used in the analysis of this article.

2.1 The CMH statistic and subset stability

Perhaps the concern in the literature closest to (but not identical to) subset stability is Simpson's Paradox [Yule, 1903, Simpson, 1951], in which redividing a contingency table between two variables along a third can change the correlation between the first two from positive to negative, or vice versa.

A search for "Simpson's paradox" in any academic search engine will return a stream of articles by authors who ran into the paradox while analyzing their data. When Simpson was writing in the 1950s, it was infeasible to have a regression of more than a handful of variables; in the present day, regressions on a dozen or more parameters are easy to find in the economics literature. Lerman [2017] gives several examples, warning that larger data sets require greater caution regarding subsets.

The paradox indicates that elements which are not a direct cause of the outcome are included in the contingency table or regression [Arah, 2008, Hernán et al., 2011, Tunaru, 2001, Yarnold, 1996]. The solution, to develop a causal graph and then build the regression around that graph, can be impracticable in a situation like migration where almost anything is a direct cause for moving in some cases, or in cases where researchers want to control for a large number of variables. Without this careful development of a causal framework, generalized linear models are susceptible to large changes in parameter estimates given seemingly small changes in the selection of regressors.

For a small number of parameters and an appropriate set of assumptions, regression techniques can generate parameters that are comparable to CMH statistics [Barros and Hirakata, 2003, Wacholder, 1986, Rogers and Swaminathan, 1993, Hidalgo and López-Pina, 2004], but whether this comparability extends to a modern high-dimensional setting is an open question.

2.2 Administrative records and government surveys

Estimates of cross-region migration rates in the literature are typically below about five percent, so micro-level analyses need as large a population as possible, preferably without sampling. But full records are not always available, and studies have often relied on a variety of surveys. DeWaard et al. [2018] offers a thorough overview of the relative benefits of the Federal Reserve Bank of New York and Equifax's Consumer Credit Panel (CCP), the U.S. Census Bureau's American Community Survey (ACS), and Current Population Survey (CPS). The authors mention data compiled by the Internal Revenue Service's Statistics of Income division, which is used by the U.S. Census Bureau (along with other sources) to generate its estimates of cross-county migration, and is used in this article.

There are a few extant examples of broad descriptive surveys using government data sets. Hernández-Murillo et al. [2011] use the U.S. Census Bureau's Survey of Income and Program Participation (SIPP), a panel study which follows respondents should they move within the United States. Their study, covering 7,823 moves, primarily consists of tabulations indicating broad differences in migration rates and outcomes for various characteristics. The tabulations make no attempt to adjust for correlations across factors, but nonetheless find that the educated move more often than the uneducated, women move slightly more often than men, the never married move more often than other statuses, losing or changing jobs lowers move propensity, and of course, homeowners move less often than others. This article finds that using controlled experiments reverses some of these results: men move more often than women and the married move more often than singles, *ceteris paribus*.

Ihrke and Faber [2012] present tabulations based on the CPS, which is a 'rooftop survey' that does not follow households as they move, but does ask survey participants migration-related questions. The most common responses to the question of why a household moved are housing-related, followed by family-related, then employment-related.

2.3 Characteristics

The literature typically focuses on a set of characteristics supported both by theory and our intuition about what might enter into a household's decision to move. This section discusses those featured in this article.

Income Making the decision to uproot one's social circle, employment, schools, and overall lifestyle requires balancing non-pecuniary factors including personal and familial characteristics, and pecuniary factors, including income, unemployment insurance, and tax burden [Stark and Bloom, 1985].

With some exceptions, the economics literature on the subject of migration, both domestic and international, focuses heavily on pecuniary factors: higher or less risky incomes, or lower costs [Sjaastad, 1962, Greenwood and Sweetland, 1972, Stark and Bloom, 1985, Borjas, 1998, Jackman and Savouri, 1992, Krieg, 1997, Detang-Dessendre et al., 2004, McKinnish, 2005, 2007, Hernández-Murillo et al., 2011, Kennan and Walker, 2011, Young et al., 2016, Nunn et al., 2018, Fee et al., 2019]. Even familial ties are often characterized as primarily useful for reducing costs and raising expected income [Chau, 1997, Palloni et al., 2001, Garip, 2008].

Studies of only wealthy households such as Young et al. [2016] do find evidence that that portion of the population moves for higher income or lower taxes. Kleven et al. [2020] surveys primarily cross-country movements of millionaires and finds comparable results. Other wealthier households move for "lifestyle" reasons [Benson and O'Reilly, 2015, O'Reilly, 2016]. But only 2.6 % of observations in this data have annual income over $216,300 (in 2010 dollars), and there is no reason that results or intuition about that subpopulation would apply to the other 97.4 %.

Mortgage status In the middle of the period studied here, beginning in 2006, there was a collapse in home values, leaving many homeowners with mortgage debt larger than the value of their home (an "underwater" mortgage). Using regional differences in the impact of the housing market collapse, Molloy et al. [2011] find no evidence that homeowners with underwater mortgages show different migration patterns than those who have positive equity. This advises that it is reasonable to ignore the relationship between home value and mortgage debt in considering the decision to move.

The appendix shows statistics broken down per year, allowing the reader to check for patterns in moves given mortgage status and vice versa.

Local conditions Since Tiebout [1956], there has been a thread of literature testing whether people move toward better amenities or away from higher taxes. These studies typically use region-wide measures of taxes and benefits, and have produced mixed results about whether local taxes are or are not an important determinant of migration [Cebula, 1974, Herzog and Schlottmann, 1986, Clark and Hunter, 1992, Pack, 1973, Preuhs, 1999, Fox et al., 1989, Nelson and Wyzan, 1989, Edwards, 2018].

Localities with higher taxes and higher density typically provide more amenities [Vedder, 1990].

Roback [1982] finds that, among the measures she considers, the best correlate to quality of life is population density itself. She finds no significant correlation to disamenities like crime and pollution, while Dyck et al. [2011]

finds a strong quality of life decrease from these density-related disamenities. There is some debate about whether amenities attract density [Rapoport, 2018] or density attracts amenities [Duncan et al., 2012].

Retirement income People who have retired from work by definition have less need to concern themselves with an area's labor market and so are more likely to move for nonpecuniary purposes [Clark and Hunter, 1992, Brown, 2008]. Young et al. [2016] find a "Florida effect" that their estimates have a far higher dependence on moves to Florida than moves to any other state, perhaps due to pull factors like good weather and beaches, an established infrastructure for retirees and relatively low local taxes.[3] Preliminary tests with this data set using Florida as a separate characteristic did not find an effect noticeable enough to retain it, which would be the case if a small number of millionaires are especially enamored of Florida but after controlling for household characteristics such as age, income, or retirement status, the great majority of households are not.

Unemployment insurance and work status Although unemployment insurance is not typically studied separately, there are mixed results about movement toward or away from public support [Nunn et al., 2018]. McKinnish [2005] used a cross-border approach to find a relatively large sensitivity to Aid to Families with Dependent Children (AFDC). However, McKinnish [2007] finds the results to be somewhat unstable and dependent on specification, such as whether "welfare mothers" are treated separately from the childless and men.

Education This article will bolster and add nuance to the findings that leaving residence at a college or finding post-school work is a common motivation for moving and that educated movers show supranormal returns relative to stayers [Gurak and Kritz, 2000, Quinn and Rubb, 2005, Faggian and McCann, 2009].

3 Data

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 and its variants).

The unit of observation is the household, not a simple head count, under the presumption that a long-distance move is a collective decision among the household members [Tcha, 1995]. In all cases, the decision to uproot one's life and move to a new locale is treated as a binary choice, without further analysis of the subsequent choice of distance and direction.

Let R be the reference year, before any potential move. Household data elements include a US Postal Service Zone Improvement Plan (ZIP) code, which is translated to a US Census Bureau ZIP Code Tabulation Area (ZCTA), as

[3] It also features exposure to climate change; McLeman and Hunter [2010] discuss the resulting out-migration.

detailed in the appendix. A *move* is a change of address over 80km, measured from the closest interior point to the centroid of the ZCTA in year R to the interior centroid of the ZCTA in year $R + 1$.[4] Although other statuses will be measured at time $R+2$, $R+3$, ..., the classification as a mover is strictly based on a change between years R and $R + 1$. Most of the statistics in this article then compare statuses in year R to year $R + 2$, the earliest year when we can be confident that the records refer to post-move income.

At $R + 2$ and later years, the question the pseudo-experiments ask is "what is the outcome from the *single event* of the case group moving while the control group stays?" For example, if movers from some area in year one return home in year two, we would like our measure to indicate whether the movers do or do not behave like stayers. If households in the group who chose to not move between years R and $R + 1$ choose to move later, then we can expect the distinction between the movers and counterfactual stayers to dissipate over time, and the effect of the single event of moving between R and $R+1$ to weaken; this is what our measure will do.[5]

Each pseudo-experiment uses all available data for that year, including 2001–2002, 2002–2003, ... for $R+1$, but only 2001–2015 for $R+14$. Period-to-period comparisons are not quite a panel, as some households will have entries in some spans but not others. Results presented below will usually begin with post-move characteristics at year $R + 2$, so that all results are based on post-move data unlikely to include pre-move information for part of the year, and end at year $R + 10$, to ensure results using at least five years of data.

Here are summaries of each variable, how it is generated, and key caveats. An appendix provides complete detail.

- Income is intended to match the adjusted gross income (AGI) line on Form 1040.

  All incomes are adjusted to 2010 dollars, and incomes are sometimes discretized into rough ventiles, which induces a roughly log scale for the bin cutoffs. Bin zero includes households making up to \$5,640 per year, in 2010 dollars; for reference, this is equivalent to about 15 hours/week at the federal minimum wage. For a small family, official poverty lines fall around \$12k, the bottom of category 4. Category 12 includes incomes from about \$43k up to exactly \$50k. Category 17 includes incomes from about \$82k up to exactly \$100k. Category 20 is incomes over about \$216k, up to millions.

[4] As discussed in the appendix, 80km is also the definition of a move used by the U.S. Internal Revenue Service. These are not small moves: the IRS Statistics of Income division estimates \$3.5 billion in moving expenses claimed by those moving over 80km in 2016.

[5] Alternatives to the strict adherence to a controlled pseudo-experiment, instead relying on household history, create more difficulties than they resolve. Classifying movers by their full pattern of moves is error-prone (what if a household moves twice in the same year?), and requires arbitrary decisions about how to treat different series. Is a mover who moves in years 1, 2, and 4 comparable to one who moves in years 1, 3, and 4? Throwing out moving households after they move again creates a sample that answers the question "what is the outcome from moving once and never moving again relative to the counterfactual of never moving?", but this is is a biased measure of any activity among the full population. Specific questions about chain migrants versus once-in-a-lifetime movers is reserved for future research.

- Marital status is based on filing status on Form 1040, meaning that those observations for whom only informational returns are available may not have a marital status assigned.
- Sex is taken from Social Security Administration records, and is used here only for households not marked as married.
- Age is divided into ages 19–24, 25–34, 35–44, 45–54, 55–64, 65+. For married couples, the mean age is used.
- The number of dependents under 18 (referred to as *kids* when space is limited) is divided into categories of 0, 1, 2, and three or more.
- Unemployment insurance (UI) recipiency is based on the presence of Form 1099-G, which is sent if a recipient receives one week or more of UI benefits.
- Retirement status is one iff the household is over 55 years old, has individual retirement account (IRA) or Social Security income, and all other income is less than that IRA or Social Security income.
- Local tax category is from Form 1040, Schedule A, and includes actual state and local income or sales tax, property tax, and other smaller local taxes. However, about half of filers do not file Schedule A. This indicates that their local taxes are likely below their standard deduction level, a value derived from their filing status and number of dependents. The three categories are therefore local taxes observed or assumed to be below standard deduction, local taxes between 100% and 125% of standard deduction, and over 125%. Although this is a coarse categorization, movers are more likely to see a change in local tax category as defined here than change in marital, retirement, UI, and other infrequently-changing statuses.
- Housing tenure is based on the presence of a mortgage statement.
- School status is based on the presence of a tuition statement, Form 1098-T. Households are divided into not in school, part-time, full-time undergrad, full-time graduate categories. Define *leaving school* as having a 1098-T in year R, but not having one in year $R + N$.
- For each ZCTA, the US Census Bureau provides population densities; median rent as a percentage of income as of 2011, used here as a proxy for cost of living; and unemployment rate as of 2011. Each is divided into three categories that give a roughly even split. Although these measures all change over time, it is uncommon for an area to shift from, say, the lowest tercile of costs of living to the highest tercile.

Other missing statuses Administrative records covering the full population do not include race, ethnicity, veteran status, languages spoken, extended family information, reliable personal wealth measures, or visa status for international immigrants. There is little question that these characteristics affect the choice to move, but they are set aside for the duration of this article.

By defining *amenities* as all location-specific considerations which are not economic (including familial and social ties), we get the tautology that movers are pushed or pulled by economic or amenity considerations. There is no question that amenities affect the move decision [Greenwood and Sweetland, 1972], but they are not the strength of the administrative records, and so are set aside for the duration of this article (apart from the limited proxies are used as above). The focus of the results will be on the characteristics of movers.

4 The method: *Ceteris paribus* pseudo-experiments

This section develops statistics that address the question *as a single variable such as home ownership or income changes, how does the likelihood of migration change?* and the converse, *comparing movers versus stayers, how does the likelihood of being in, entering, or exiting a status change?*

The goal is a set of statistics regarding one characteristic that control for the others. This is the typical interpretation of coefficients in generalized linear models, and an understanding of parameter interpretation in that context could be applied to the Cochran-Mantel-Haenszel statistics used throughout the article. The key difference is that *controlling for* or *holding all else equal* in the context of regressions is done via adding controlled-for variables to the regression, while the CMH statistic follows the tradition of controlled experiments, where the population is simply separated into groups using the controls.

4.1 *All-else-equal* cells

When considering the relationship between UI and migration, for example, divide the population into cells where each cell has a single, fixed value of all other pre-move characteristics: year $\times$ age category $\times$ mortgage status $\times$ AGI category $\times \dots$. Any given cell affords a single case-control experiment measuring the outcome of UI given migration status (or vice versa) with all else held equal. These will be referred to as *ceteris paribus* (*c.p.*) cells. The full cross of all of the above categories, move/stay, and year generates 45,417 cells with at least fifty movers.

Table 1 shows the schematic breakdown of the constituents of a cell with all held constant but unemployment insurance and move/stay status.

	Stay in UI	Exit UI	Stay out of UI	Enter UI
Moved	A	B	E	F
Stayed	C	D	G	H

Table 1: A reference for one *ceteris paribus* cell.

The exit rate rate out of UI for movers is $B/(A+B)$, excluding those who began out of UI and therefore can not exit it; for stayers, the exit rate is $D/(C+$

D); so the *ceteris paribus* ratio of the likelihood of exiting UI given a move versus staying is the ratio of these two, which can be written as

$$RR = \frac{\frac{B}{A+B}}{\frac{D}{C+D}} = \frac{B(C+D)}{D(A+B)}. \tag{1}$$

Consider a mover and a doppelgänger identical in all other characteristics, representing the counterfactual that the mover chose to stay. Then the risk ratio, also known as the relative risk, gives the estimate of the change in likelihood of exit from UI for a mover relative to the counterfactual of staying.

Conversely, the relative risk of moving given exit out of UI $\frac{A(B+D)}{B(A+C)}$. One could also write down the risk ratio of moving given starting in UI versus not, or of being on UI post-move given moving versus staying, and so on.

Unlike UI recipiency, income is a continuous variable, so pre- and post-comparisons can typically be done with a simple percent change. For a given cell, let Δ be the percent change for movers minus the percent change for stayers, giving the all-else-equal gain or loss from moving.

4.2 Aggregation

The above process is for one *c.p.* cell, and the process could be repeated to generate a separate measures for every other. Now we need a means of aggregating these measures.

For Δ, a weighted mean is adequate. Some other measures to describe the variation in Δ will also be used below.

We might initially presume that a simple weighted average across cells would work, but Mantel and Haenszel [1959] find that a modified weighted average has better statistical properties.[6]

Normalise the four counts in Equation 1 to a set of probabilities such that $A+B+C+D=1$. Attach a weight w_i to cell i proportional to the count of all households in that cell, so $\Sigma_i w_i = 1$. Then the proposed measure for entry into UI given move/stay status, now known as the Cochran-Mantel-Haenszel (CMH)

[6] The more common version of the CMH statistic is an odds ratio, not a risk ratio. *Odds* is calculated by the ratio of count of occurrence of an event over count of non-occurrence; *risk* is the ratio of the same occurrence count over the full count of the population [Ranganathan et al., 2015]. An odds ratio or risk ratio is the ratio of two so-defined odds or risks.

This article relies on the risk ratio. Colloquial references to the *chance*, *likelihood*, and typically even *odds* of an event refer to the risk, not the odds as defined here. The odds ratio is symmetric, giving equal odds to the chance of moving among retirees versus non-retirees, and the odds of retiring among movers versus stayers. The risk ratio gives distinct values for the two, which can better advise causal inquiries.

For relatively unlikely events, such as a health condition in a typical medical study, the odds ratio approximates the risk ratio, but as the likelihood of the event grows, the odds overestimates the risk to the point of being almost unusable for discussing the chance that an event will occur [Schmidt and Kohlmann, 2008].

statistic [Pearce, 2004, Rothman, 2012], is

$$\frac{\Sigma_i w_i B_i (C_i + D_i)}{\Sigma_i w_i D_i (A_i + B_i)}.$$

This reduces to the risk ratio when there is only one cell, behaves like the simple average of risk ratios in many cases, and gracefully handles situations where D is zero in some cell, giving an infinite risk ratio for that cell.

Let $CMH(\text{enter UI}|\text{move})$ indicate this aggregate risk ratio for entering UI for movers relative to comparable stayers. More generally, $CMH(\alpha|\beta)$ will indicate the CMH aggregate risk ratio of α for observations with characteristic β versus all-else-equal observations without characteristic β.

Subset stability One could easily imagine other means of aggregating the per-cell risk ratios beyond the CMH statistic, and it is reasonable to include a stability conditions in the search:

Definition 1 *An aggregate statistic over a set of cells is* subset stable *iff the statistic calculated using the aggregate data is bounded between the smallest and largest per-cell statistic.*

Setting aside risk ratios for a moment, for Δ, where the aggregate statistic is the weighted mean of the two separately-calculated Δs, subset stability is easily satisfied.

But finding failures of subset stability for the simple risk ratio is not at all difficult. Table 17 in the appendix will show risk ratios for AGI moving up and for AGI moving down, plus three binary splits of both. Of these six splits, two fail subset stability.

An appendix also uses Logit regressions to estimate parameters describing move propensity given some characteristics, and finds that the parameters also fail subset stability.

That is, for measures regarding the effects of changes regarding discrete categories, subset stability is not at all given. But:

Lemma 1 *The CMH statistic satisfies subset stability.*

The proof is relegated to an appendix.

In fact, the CMH statistic is the only option, within a broad range of candidates. Define *well-behaved* to mean any function where all derivatives (first, second, third, ... in all variables) are defined and finite when all variables $(w_1, w_2, \ldots, A_1, A_2, \ldots, B_1, B_2, \ldots)$ are within the open interval $(0, 1)$. This is a broad class including reasonable candidates such as the weighted average over the risk ratios for every cell, sums of logs, and far more exotic options. We exclude from the well-behaved set those functions with components that will evaluate to zero, infinity, or $0/0$ in the very common case where one cell (out of possibly thousands) has a zero value in some configuration. This condition eliminates the maximum, minimum, and expressions consisting of weighted products with forms like $\prod_i A_i^{w_i}$ or $\prod_i (C_i + D_i)^{w_i}$.

As a final stipulation, the set of cells is unordered, so a well-behaved aggregation function should be symmetric: For any two cells c_1 and c_2, we should have $f(c_1, c_2) = f(c_2, c_1)$.

Theorem 1 *Within the class of well-behaved functions, the CMH statistic is the unique aggregation function that satisfies subset stability.*

The appendix provides a proof, and further nuance regarding the CMH statistic, including discussion of how estimates via regression models are not subset stable and show instabilites the CMH does not, how small cells are handled, and a lemma measuring the gap between the simple risk ratio and the CMH statistic and how this advises the bias given missing data.

Mantel and Haenszel [1959] demonstrate that the CMH statistic provides an unbiased and consistent aggregation of per-cell odds ratios to the full population under basic distributional assumptions, and demonstrate how several alternative forms do not have those desirable conditions under the same distributional assumptions. That the CMH statistic is the only aggregation that satisfies subset stability may help to explain aspects of the statistical stability they discuss.[7]

5 Results

Section 2.1 touched on how the results of regressions, where parameters for all characteristics are estimated simultaneously, are sensitive to details of the causal framework developed to select input characteristics. As a result, regression-based tests can be difficult to design and typically come in small batches. The pseudo-experimental method developed to this point, in which estimates are done for one test statistic at a time, is not sensitive to misspecification of the underlying causal framework linking control variables and can not show the instabilities regression-based methods show. We can therefore generate a wealth of tests of hypotheses regarding the interaction of household characteristics and moving.

This section will begin with some overall descriptive statistics about migrants, including some basic statistics about who moves and some per-status details. It then organizes the many tests of claims regarding the interaction of household characteristics and moving into a series of somewhat separate narratives, beginning broadly and exploiting subset stability to ask questions about subsets with certain characteristics.

[7]Other statistical issues with CMH statistics are not relevant in this context. In medical studies, when subjects are selected *ex ante* and split into *ceteris paribus* cells based on observed covariates, there is bias in the measure of odds or risk ratio to the extent that those controlled covariates correlate to the outcome [Costanza, 1995, Deeks, 1998, Rose and van der Laan, 2009]. But that is not the situation in typical administrative record or commercial data sets, with a defined universe of observations with no subject selection. Multiple testing issues in mereological methods [Dixon and Simon, 1992] are not a consideration for descriptive studies, or can be adjusted via methods such as Bonferroni corrections.

The hypothesis that households move for income gain is the norm in in the economics literature, but Section 5.2 rejects the hypothesis for about 45% of the population. The narrative in Section 5.3 elaborates on this at a micro-level, finding that income is more likely to go down relative to the counterfactual stayer than up relative to the counterfactual stayer.

After the pecuniary results, the CMH and other statistics will be used to measure how household characteristics in the population and various subgroups relate to the risk of moving.

Because this unique data set is difficult to access, an appendix provides many additional details regarding the relative risk of income change for some subgroups, unemployment, mortgage, age and retirement, children, federal tax, and geographic characteristics.

5.1 Descriptive overview

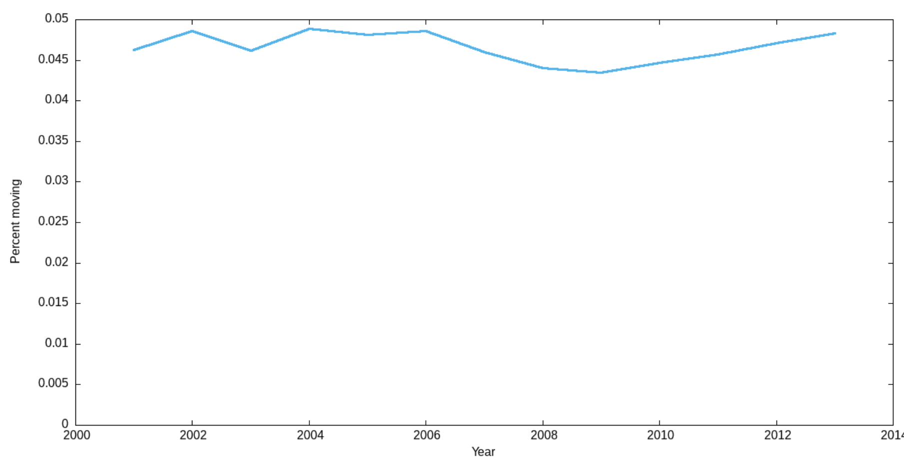

Figure 1: The rate of moving over 80km, by households, from 2001 to 2014.

Male, single	5.1 %	0 children	4.9 %
Female, single	6.1 %	1 child	4.4 %
Unmarried	5.5 %	2 children	4.2 %
Married	3.7 %	3+ children	4.6 %
No mortgage	5.7 %	No Local tax	4.0 %
Mortgage	3.0 %	L.tax < SD	5.8 %
No UI	4.7 %	SD< L.tax < 1.25 SD	3.1 %
UI	4.6 %	1.25 SD< L.tax	2.8 %
Not retired	4.7 %	$18 <$ age <25	9.6 %
Retired	2.7 %	$25 \leq$ age <35	7.2 %
Not in school	4.2 %	$35 \leq$ age < 45	4.1 %
Part-time school	8.8 %	$45 \leq$ age < 45	3.1 %
FT undergrad school	6.9 %	$55 \leq$ age < 65	3.0 %
FT graduate school	11.1 %	$65 \leq$ age	2.6 %

Table 2: Simple percent in each pre-status moving. UI=unemployment insurance; FT=full time; SD=standard deduction.

To start the overview broadly, Figure 1 shows the overall move rate, an average over all years of 4.7 %.

Unless otherwise stated, all statistics in this section are an aggregate of the entire period, 2002–2015, and presented changes are based on post-statuses in year $R + 2$. Aggregate risk ratios are calculated by generating *ceteris paribus* cells as described in Section 4.1. For example, to estimate risk ratios regarding income and migration, each cell compares only observations who are identical in year and reference period marital status, sex if unmarried, unemployment insurance status, retirement status, housing tenure, number of children, local tax category, age category, tertiary education enrollment type, and regional density, unemployment rate, and cost of living categories.

Table 2 shows the unadjusted percent of observations with each pre-move status who moved.

Male, single	107.0 %	0 children	122.4 %
Female, single	93.5 %	1 child	86.9 %
Unmarried	84.5 %	2 children	86.8 %
Married	118.4 %	3+ children	95.3 %
No mortgage	136.4 %	No Local tax	114.4 %
Mortgage	73.3 %	L.tax < SD	115.5 %
Not on UI	101.2 %	SD< L.tax < 1.25 SD	91.2 %
On UI	98.8 %	1.25 SD< L.tax	80.1 %
Not retired	95.0 %	18< age <25	151.9 %
Retired	105.2 %	25≤ age <35	153.6 %
Not in school	84.1 %	35 ≤ age < 45	95.9 %
Part-time school	103.5 %	45 ≤ age < 55	75.3 %
FT undergrad school	106.4 %	55 ≤ age < 65	74.6 %
FT graduate school	167.3 %	65 ≤ age	50.6 %

Table 3: For various pre-move statuses, $CMH(\text{move}|\text{pre-status})$. For example, single men are 1.07 times as likely as otherwise comparable single women to move, and those attending college part time are 1.03 times as likely to move as comparable households in any other schooling status. For binary statuses, the two risk ratios are reciprocals. UI=unemployment insurance; FT=full time; SD=standard deduction.

Table 3 shows the aggregate relative risk of moving given the pre-move status listed. For example, a household with one dependent under 18 is 86.9 % as likely to move as a comparable household with zero, two, or more children. A comparable breakdown by AGI category appears in Figure 6.

	Married	Mortgage	UI	Retired	School
Enter status|Moved	137.6 %	120.6 %	106.1 %	124.5 %	137.7 %
Exit status|Moved	137.2 %	235.6 %	112.2 %	–	116.6 %

	Kid	AGI	Local tax	Fed tax
Enter status|Moved	103.2 %	112.9 %	124.0 %	117.6 %
Exit status|Moved	99.8 %	119.8 %	143.2 %	128.9 %

Table 4: $CMH(\cdot|\cdot)$ statistics. For dependent, AGI, education, and tax measures, *entry* means rising in category, *exit* means moving to a lower category.

Changes in status Table 4 gives an overview of the risk ratios of likelihood of shifting up/down in status between movers and stayers. For example, movers are 1.21 times as likely to enter into a mortgage from not having one than comparable households who stayed within an 80km radius, and are 2.36 times as likely to shift from having a mortgage to not.

	Married	Mortgage	UI	Retired	School
Moved|Enter status	141.0 %	122.9 %	106.1 %	125.2 %	139.7 %
Moved|Exit status	139.2 %	271.0 %	148.7 %	–	147.1 %

	Kid	AGI	Local tax	Fed tax
Moved|Enter status	103.6 %	145.2 %	128.9 %	136.0 %
Moved|Exit status	99.7 %	148.7 %	149.6 %	152.2 %

Table 5: $CMH(\cdot|\cdot)$ statistics. For dependent, AGI, education status, and tax measures, *entry* means rising in category, *exit* means moving to a lower category.

Table 5 shows the complementary risk ratios for moving given changes in status relative to those who do not change status. Just as moving households show more dynamism, households who are undergoing changes in these statuses are more likely to move. The high rate of both entry and exit from every status advocates for the use of microdata in studying migration: aggregate measures of status post-move can be expected to understate the amount of change.

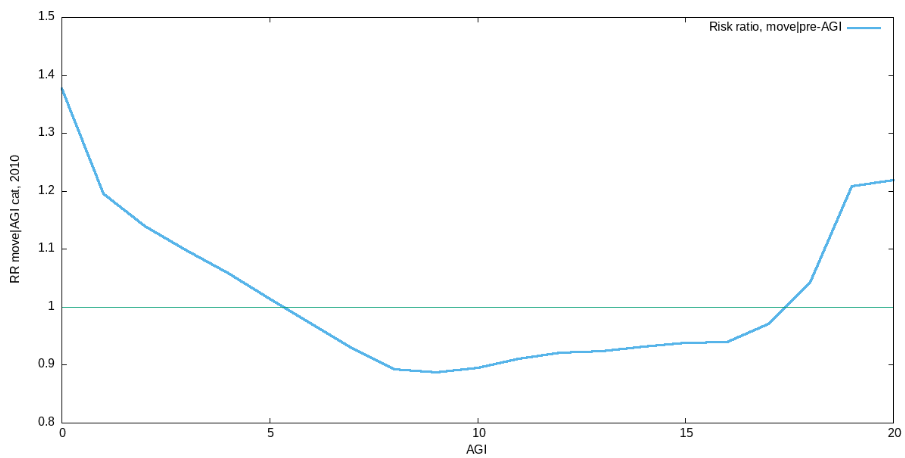

Figure 6: For each income band B, $CMH(\text{move}|\text{band } B)$.

What income categories are moving? Figure 6 shows the all-else-equal likelihood of moving for a given AGI range relative to all other movers, using the same *all else equal* pseudo-experiments. The all-else-equal likelihood of moving for middle-income earners is lower relative to others.

A non-controlled tally of the percent of each AGI category moving will appear in Figure 10.

5.2 Relative change in income

Matching movers to stayers who are indistinguishable from them in the data set, we can read the stayers' incomes as the counterfactual income movers would see if they had stayed. Although one could imagine exceptions, for the great majority of the 82 million movers in the data set it is reasonable to assume that a move in the near future does not determine pre-move age, choice of income, the choice to have a child, and several other statuses. In this case, the controlled pseudo-experiments test hypotheses of the form *what is the effect of moving on counterfactual income*, for salient subsets of the population. This section will run several dozen such hypothesis tests.

Let Δ be the expected difference in percent change in income for movers versus comparable stayers. Figure 7 displays the distribution of Δs for each all-else-equal cell. The darker, more peaked curve in Figure 7 is the distribution of Δs at year $R+2$. As time progresses, the distribution gets wider and starts to lean slightly toward positive values, 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 % or larger change for movers than for stayers, or vice versa.

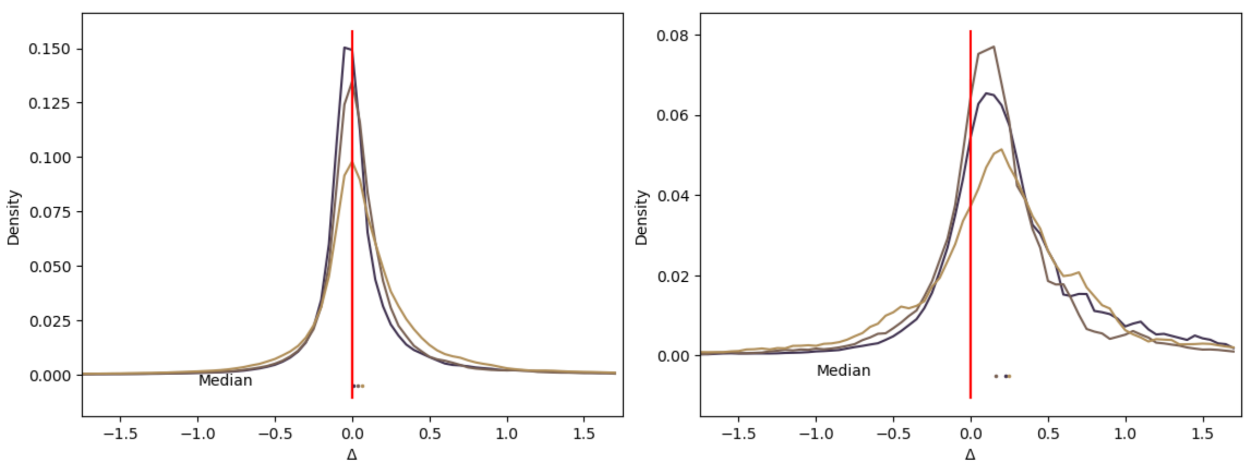

Figure 7: Density of Δ, income change for movers relative to comparable stayers. Left: all households. Right: only those leaving school. The taller, darker curve is at year $R+2$, then the curves for $R+5$, and $R+10$ are increasingly lighter and flatter. The medians appear below the main histogram. The $\Delta = 0$ line is shown for reference. Plots are smoothed using a Uniform$[-.05, 0.05]$ kernel.

For some subgroups, the distribution is not so symmetric. The right-hand histogram in Figure 7 is the distributions of Δs 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.

Showing plots for every subgroup would be overwhelming, so Table 8 shows summary statistics describing 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 summary statistics for the left plot in Figure 7 covering the full population, where 55.8 % of movers have higher income relative to comparable stayers in year $R + 2$, rising to 63.5 % in year $R + 10$.

A person who takes a 10% pay cut to live in an area with a 20% lower cost of living is arguably not doing financially 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 greater than or equal to the category in year R. As a reminder, the set of stayers includes households who moved within their area to a ZCTA lower on the cost of living measure. This subset's outcomes are qualitatively similar to those of the full population.

		$R+2$		$R+10$	
Subgroup	% of movers	%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
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
Also without retirees					
$19 \leq$ age < 25	20.1 %	55.2 %	0.02	69.3 %	0.12
$25 \leq$ age < 35	24.6 %	61.5 %	0.03	73.5 %	0.14
$35 \leq$ age < 45	15.0 %	56.1 %	0.01	63.2 %	0.05
$45 \leq$ age < 55	11.7 %	44.5 %	−0.01	53.2 %	0.01
$55 \leq$ age < 65	8.8 %	36.6 %	−0.03	43.1 %	−0.02
$65 \leq$ age	10.9 %	54.8 %	0.00	46.7 %	−0.00
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
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
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
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

Table 8: 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. See text for further discussion.

Leaving school. Households who move after leaving school have excellent initial outcomes relative to stayers. Figure 7 presents the histograms of Δ for this group, and Figure 8 gives the summary statistics for the histograms: 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 8 are based on the distribution of Δs including the three separate types of all-else-equal-cell. Dividing into subgroups, the median Δ at $R+2$ for part-time students leaving school is 24.3 %, for full-time undergraduates, 8.3 %, and for full-time graduate students, 25.9 %.

But after the initial post-move gain, Δ does not continue a pattern of standout growth. The median Δ for households leaving school at $R+10$ is 1.08 times as large as its value at $R+2$. Movers not leaving school see a much smaller initial median Δ, but the median grows to 6.03 times as large over the same period. That is, the supranormal gains the educated see from moving seem to be due to the single post-school move.

The result that those who are leaving school see such large returns offers a caveat that papers focusing on the 2.4 % of movers leaving school [Card, 2001, Detang-Dessendre et al., 2004, Quinn and Rubb, 2005, Faggian and McCann, 2009, Fee et al., 2019] may give an inaccurate impression of the experience of the remainder of migrants.

Given that the outcomes from post-school movers are so distinct, the remainder of this section excludes school-leaving movers and stayers, although those whose education entirely precedes or succeeds their move will remain.

Retiring By the definition used in this article, the majority of post-retirement income is from Social Security or IRA payments, which are not contingent on location, and are an increasing function of lifetime earnings. But retirees who move see worse pecuniary outcomes from moving than comparable retirees who do not, with only 32.5 % seeing positive income gains over the counterfactual of staying.

Defining amenities as non-work considerations, retirees out of the work force are amenity-based movers [Sander and Bell, 2013, Greenwood, 1981]. One may have an image of "a growing and increasingly affluent retirement-age population [who] has responded to the environmental amenities associated with warmer climates" by moving to chase the sun [Greenwood, 1981], but the data set shows, in the balance, a cohort of moving retirees whose lifetime earnings have fallen behind. 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. If retiring or retired movers are typically not as affluent relative to comparable staying retirees, movers may be seeking amenities such as lower rent or familial support over lifestyle upgrades [Fournier et al., 1988].

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. 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, though 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 $R+2$, 39.5 % of movers in this set see a change in income better than the change comparable stayers see.

Households making over $100k, 3.0 % 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.

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, 44.9 % of single women with child dependents initially have a Δ greater than zero, versus 61.7 % of married households with child dependents. At $R+10$ single men with no dependents have a median Δ that is 10.3 % higher than single men with child dependents. Defining amenities as considerations outside of expanding income, these non-married households with dependents under 18 are seeking amenities, which may include familial support.

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

Conclusion to this narrative This narrative presented a breakdown of the population using top-line characteristics to find those subgroups with supra- or sub-normal returns to moving relative to comparable stayers. Minor changes in the order of the narrative, such as looking at pre-move income category before age category, did not change the qualitative results. Summing the percentages, the groups showing an aggregate relative return below zero comprise 44.6 % of the population.

5.3 Pecuniary characteristics

Even within one *ceteris paribus* cell, if ten movers lose income and one gains significantly, while all stayers see a small gain, Δ may be positive but the risk ratio for the chance of an income loss given moving relative to staying will be negative. If our goal is to evaluate hypotheses about macro-level outcomes, Δ is more appropriate; if our goal is to test claims that individuals are motivated to move by higher income, as asserted by the many articles cited in the introduction, then the risk of seeing an income gain of any sort may be more relevant.

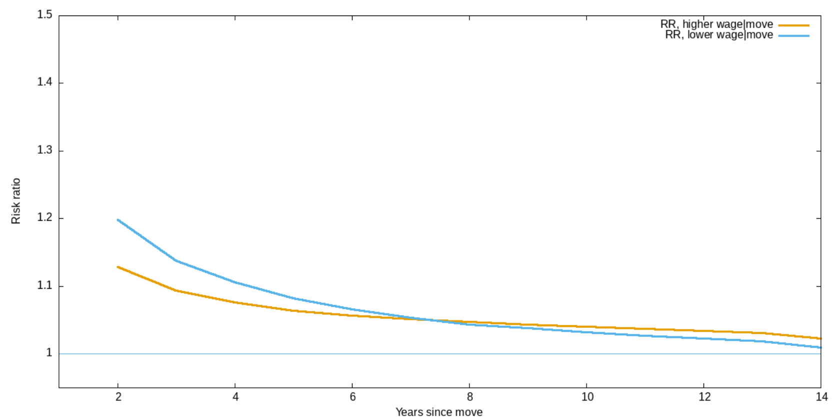

Figure 9: Relative risk of downward (initially upper curve) and upward (lower curve) shift in wage category, movers versus stayers.

Figure 9 shows the evolution of the aggregate risk ratios for relative rise and relative fall in income. The figure shows that, for the first few years, CMH(income drop|moved)>CMH(income drop|stayed), and in later years the two are not different with policy significance. That is, by income, movers are more likely to be relatively worse off after moving than relatively better off. At $R+2$, CMH(income rose|moved) =112.9 % and CMH(income fell|moved) =119.8 %.

The relatively greater chance of lower relative outcomes than higher persists for several years until they are roughly equal.

Both aggregate risk ratios are always greater than one, indicating a larger risk of change relative to stayers. This would not be visible in the aggregate Δs, where ten households rising in income and ten falling in income may look like no change at all. Although in later years these statistics may be close enough to one to not have policy significance, the greater dynamism of migrants relative to their corresponding all-else-equal stayers persists relatively far into the future. Even seven years post-move, both the risk of rising and falling in income relative to comparable stayers are over 105%. Except where noted, every plot of this type will be on the same scale, with the relative risk ratios of changing status given a move below 0.95 (indicating stayers more dynamic than movers) unused and omitted.

See the appendix for some sub-category breakdowns of the relative risks of income rise or fall.

How much income moves? Figure 10 shows the un-controlled percent of people moving in each band and percent of dollars moving. The total is over all two-year periods. The top bin is omitted.[8]

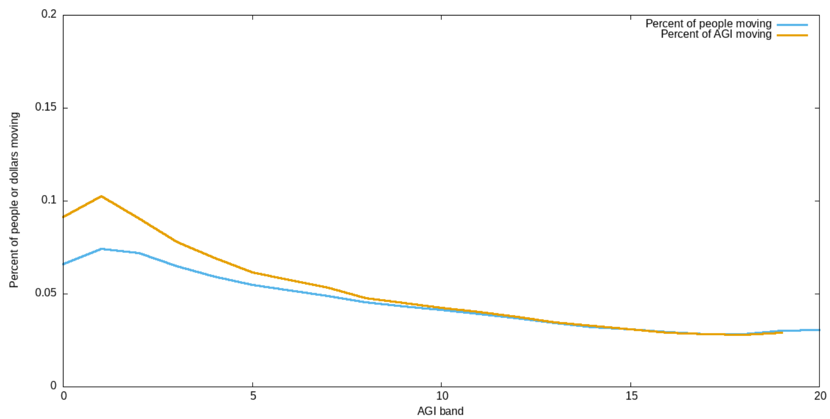

Figure 10: For the population in each pre-AGI band, what percent of household observations moved at time $R+2$ (blue lower line), and what percent of AGI at time $R+2$ is earned by the movers (orange upper line)? If 4% of households moved and 4% of AGI moved, this indicates post-move income parity between movers and stayers. If a larger percentage of AGI moved than households, we conclude that the moving households in the given band are seeing accordingly larger growth in their income relative to stayers.

For incomes lower than about AGI category 6—about $20,000/year in 2010 dollars—a relatively larger percentage of AGI goes to movers in the post-period than stayers. The difference fades, and in higher income categories about the same percent of observations and AGI move. Again, this may be explainable via a regression to the mean, in which movers have a more volatile income than stayers. At higher incomes, the volatility does not change the average, but because income can not fall below zero, the average income for the low-income categories is a rising function of income volatility.

[8]The top AGI bin goes from $216,000 to infinity, making comparison of pre- and post-move income impossible or useless.

5.4 Taxes

One could expect higher local taxes at a destination to move migration in a few directions:

- Down, because households respond to the disincentive of higher taxes.
- Indifferent, because it is difficult to predict future local taxes, based on the amount of future income that will be taxed, state deductions and exemptions, and including property taxes—or because local taxes are just not something movers think about.
- Up, because higher taxes are often correlated to more local amenities, especially when moving from lower to higher density areas.

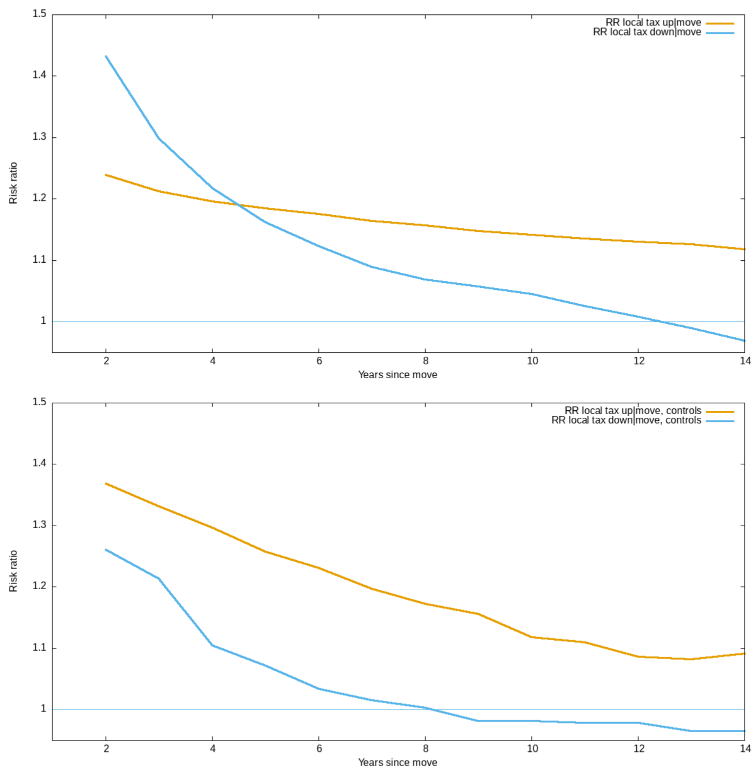

Figure 11: Blue curves which fall below 1: CMH(upward shift in local tax category|moving). orange curves which are always above 1: CMH(downward shift in local tax category|moving). Top: universe. Bottom: subset with no change in other pecuniary categories and density.

The data set shows evidence that could support all of these. The first plot of Table 11 shows the pseudo-experiments tracking the likelihood of a rise or fall in local taxes relative to stayers. The CMH(drop in taxes|moving) is initially large but falls over the years, which by itself might be consistent with a large set of movers who want lower taxes, but have difficulty predicting their taxes four years out.

	up|move	down|move	% of pop
Local tax	124.0 %	143.2 %	100 %
Local tax|density	128.7 %	144.4 %	90.9 %
Local tax|AGI, UI, tenure, retirement	136.9 %	127.7 %	30.8 %
Local tax|AGI, UI, tenure, retirement, density	136.9 %	126.1 %	29.0 %

Table 12: At $R+2$, CMH(rise in local taxes|moving) and CMH(fall in local taxes|moving), conditional on holding certain characteristics constant from pre- to post-move.

However, most statistics throughout this article control only for pre-status, meaning that two people who start with a mortgage at time R could be compared to each other even if only one leaves the mortgage at time $R+2$, thus almost guaranteeing a change in local tax situation. Table 12 shows the relative risk ratios of changing status for various subgroups at time $R+2$. For example, there was no change in any of UI, retirement, housing tenure, or AGI category statuses for 30.8 % of observations. For this subgroup, the risk ratios for changes upward or downward were CMH(rise in local taxes|moving)=136.9 % and CMH(fall in local taxes|moving)=127.7 %.

Higher-tax jurisdictions can also be higher-amenity jurisdictions, and as per the literature in Section 2.3, cities typically have both higher amenities and higher taxes. Adding unchanged urban density category to the list of statuses held constant brings the relative risk ratios to 136.9 % up and 126.1 % down. After all other controls are taken into account, a small number of movers change density categories, meaning that the CMH statistics not controlling for density are comparable.

The second plot of Figure 11 shows the results of the pseudo-experiments over time. As per Table 12, holding pecuniary characteristics constant shifts the aggregate risk ratios for moving down in local tax bin to below the aggregate risk ratios for moving up in local taxes relative to stayers. When considering the subset with other relevant changes held constant, the relative risk of moving to higher taxes is noticeably higher than the relative risk of moving to lower taxes.

Federal tax Even controlling for changes in relevant characteristics, federal tax payments show a perhaps surprising amount of variation, given that households file the same forms regardless of location. Movers pay relatively more in federal taxes than stayers.

5.5 Other statuses

Because movers are more likely to see a fall in income than a rise, all else equal, other factors must be an important part of the decision process. This section considers several. Where the direction of causality is clear, these statistics can be read as controlled hypothesis tests measuring outcome effects.

5.5.1 Tertiary Education

Representation of households paying tuition among movers is 2.91 times that of their prevalence in the general population.

Movers are 1.38 times more likely than comparable stayers to advance in their education (no education to some, part-time education to full-time, full-time undergraduate to full-time graduate school). This is consistent over time, indicating that what difference in educational entry between movers and stayers happens at the move, followed by largely at-parity education entry by movers and stayers post-move.

Initially, movers are 1.17 times the likelihood of exiting education (reversing the ordering above) than comparable stayers. After a few years, both movers and stayers who were in school before moving have likely left, meaning the relative risk of exiting given moving versus staying converges to approximately one.

	No school	Part-time	Undergrad	Grad school
% of each status moving	4.2 %	8.8 %	6.9 %	11.1 %
CMH(moving|pre-status)	84.1 %	103.5 %	106.4 %	167.3 %
CMH(pre-status|moving)	97.7 %	102.5 %	106.8 %	171.2 %

Table 13: Some statistics for observations in the given education categories.

Table 13 gives a breakdown by enrollment status. By simple percent, part-time and full-time undergraduate students are more likely to move over 80km than the non-school population. The aggregate risk ratios for these statuses relative to all others tell a similar story, that part-time and full-time undergraduate students are more likely to move than those not in the given status, holding all else equal, and movers are more likely to be in one of these statuses.

But those in full-time graduate school at reference year R show even more of a propensity to move: CMH(moving|FT graduate)= 167.3 %, and CMH(FT graduate|moving)= 171.2 %.

5.5.2 Sex and marital status

By the raw percentages, 5.1 % of men move and 6.1 % of women do. But a test controlling for all other factors finds the reverse: CMH(move|male)=107.0 %.

Perhaps the main confounder that causes this reversal is mortgage status: 80.4 % of observations of single mortgage holders are men. Knowing that households with a mortgage move less often, it is no surprise that single male-headed households by simple count would move less than single female-headed households.

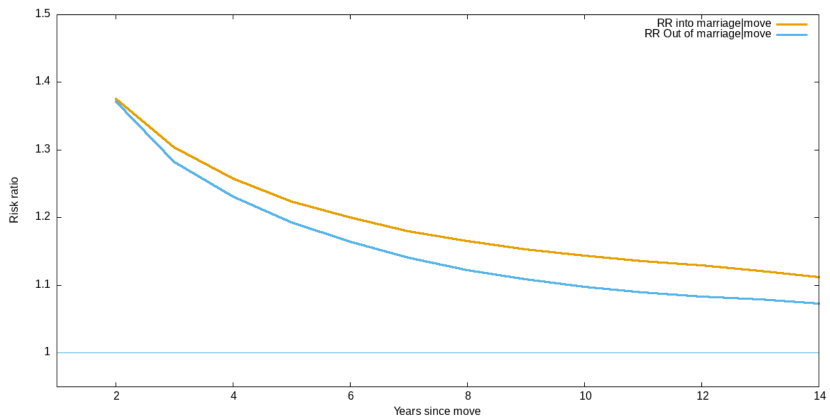

Figure 14: Upper curve: CMH(enter marriage|moving). Lower curve: CMH(exit marriage|moving).

Regarding marriage, churn is again apparent in Figure 14: CMH(entering marriage|moving)= 137.6 %, but in the other direction, CMH(exiting marriage|moving)= 137.2 %. These differences from stayers fall with time but are still apparent 14 years after the move.

By simple percents, 39.2 % of household observations are married, but only 30.8 % of moving observations are married.

But CMH(move|married)=118.4 %—all else equal, a married household is *more* likely to move than an unmarried household. In the same manner as the male/female statistical reversal above, the primary means by which the *all else equal* clause bears weight seems to be housing tenure, as 65.5 % of married observations carry a mortgage versus 23.8 % of single observations.

Splitting the population into those earning less than \$12k and those earning more, CMH(enter marriage|moving) shows little difference: 137.5 % for those earning under \$12k and 137.7 % for those earning more. But CMH(exit marriage|moving) for those earning under \$12k is 140.4 %, versus 132.1 % for those earning more.

6 Conclusion

The pseudo-experimental method presented in this article strips away many potentially arbitrary decisions that could skew results. Given its stability assurances and estimation of one parameter at a time, we need not make arbitrary decisions about what variables and interactions to include in the model. There are no parametric assumptions in this article. The main modeling assumption is that observations that are indistinguishable in the data set, save for the two characteristics under consideration, are the best comparison to each other. When aggregating *ceteris paribus* cells, the Cochran-Mantel-Haenszel statistic has a nontrivial form, but there are still no decisions that a researcher has to make, because Theorem 1 tells us it is the only reasonable option.

After stripping away as many places as possible where arbitrary decisions could influence our estimates, tests of claims regarding migration are easy to construct. The primary starting point in the sequence of tests was the starting point of the economics literature, that movers are seeking higher incomes. Hyatt et al. [2018] advises researchers to "exercise caution when explaining changes in overall migration using arguments specific to economic migration." The archetype of the unfettered individual—young, no high-paying pre-move income stream, single with no children—is often the focus of study (e.g., Detang-Dessendre et al. [2004], Kennan and Walker [2011]), but the more the household drifts from this archetype, the worse moving as an imaginary bet 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 relative to comparable stayers initially. Even ten years after the reference year, only 63.5 % of movers see a relative gain.

It is easy to find migration models, or a migration component of a larger model, which assumes a goal of income maximization by movers, but that implicitly assumes the modeled agents act like the unfettered archetype. Used for hypothesis testing, such a model may find statistical significance to the claim that households seek higher income—all else equal, people do prefer more money to less. But when used for prediction, the model will have less efficacy for the nearly half of movers who do not meet the archetype.

In aggregate, the statistics confirm that it is impossible to present a single unifying theory of why households move, but it instead takes diverse stories to add up to the migration we observe. This motivates the exploratory method of this article and its focus on describing subsets of the population, rather than testing a single isolated hypothesis about the move decision. Most of the life transitions in this article are infrequent, but in aggregate, only 12.8 % of movers experience none of the following: having a child, making more money, taking on a mortgage, changing to a different urban/suburban/rural density area, changing marital status, retiring, or getting off of unemployment insurance, compared to 34.1 % of those who do not move over 80km.

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.

References[9]

Onyebuchi A Arah. The role of causal reasoning in understanding Simpson's paradox, Lord's paradox, and the suppression effect: Covariate selection in the analysis of observational studies. *Emerging Themes in Epidemiology*, 5 (1), February 2008. doi: 10.1186/1742-7622-5-5.

Aluísio JD Barros and Vânia N Hirakata. Alternatives for logistic regression in cross-sectional studies: an empirical comparison of models that directly estimate the prevalence ratio. *BMC Medical Research Methodology*, 3(1), October 2003. doi: 10.1186/1471-2288-3-21.

Michaela Benson and Karen O'Reilly. From lifestyle migration to lifestyle inmigration: Categories, concepts and ways of thinking. *Migration Studies*, 4(1): 20–37, October 2015. doi: 10.1093/migration/mnv015.

George Borjas. Immigration and welfare magnets. Technical report, November 1998.

David Brown. *Rural Retirement Migration*. Springer, Dordrecht, 2008.

Jeremy Burke and Amalia R. Miller. The effects of job relocation on spousal careers: Evidence from military change of station moves. *Economic Inquiry*, 56(2):1261–1277, December 2017. doi: 10.1111/ecin.12529.

David Card. Estimating the return to schooling: Progress on some persistent econometric problems. *Econometrica*, 69(5):1127–1160, September 2001. doi: 10.1111/1468-0262.00237.

Richard J Cebula. Interstate migration and the Tiebout hypothesis: An analysis according to race, sex and age. *Journal of the American Statistical Association*, 69(348):876–879, December 1974. doi: 10.1080/01621459.1974.10480221.

Nancy H. Chau. The pattern of migration with variable migration cost. *Journal of Regional Science*, 37(1):35–54, February 1997. doi: 10.1111/0022-4146.00042.

Raj Chetty, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. Where is the land of opportunity? The geography of intergenerational mobility in the United States. Technical report, National Bureau of Economic Research, January 2014.

David E Clark and William J Hunter. The impact of economic opportunity, amenities and fiscal factors on age-specific migration rates. *Journal of Regional Science*, 32(3):349–365, August 1992. doi: 10.1111/j.1467-9787.1992.tb00191.x.

[9] This is the reference list for both the article and its appendices.

M.C. Costanza. Matching. *Preventive Medicine*, 24(5):425–433, September 1995. doi: 10.1006/pmed.1995.1069.

J Deeks. When can odds ratios mislead? Odds ratios should be used only in case-control studies and logistic regression analyses. *BMJ*, 316 (7136):989–91, March 1998.

Cecile Detang-Dessendre, Carine Drapier, and Hubert Jayet. The impact of migration on wages: Empirical evidence from French youth. *Journal of Regional Science*, 44(4):661–691, November 2004. doi: 10.1111/j.0022-4146.2004.00353.x.

Jack DeWaard, Janna E. Johnson, and Stephan D. Whitaker. Internal migration in the United States: A comparative assessment of the utility of the Consumer Credit Panel, March 2018.

Dennis O. Dixon and Richard Simon. Bayesian subset analysis in a colorectal cancer clinical trial. *Statistics in Medicine*, 11(1):13–22, 1992. doi: 10.1002/sim.4780110104.

Dustin T. Duncan, Jared Aldstadt, John Whalen, Kellee White, Marcia C. Castro, and David R. Williams. Space, race, and poverty: Spatial inequalities in walkable neighborhood amenities? *Demographic Research*, 26:409–448, May 2012. doi: 10.4054/demres.2012.26.17.

Delfien Van Dyck, Greet Cardon, Benedicte Deforche, and Ilse De Bourdeaudhuij. Do adults like living in high-walkable neighborhoods? Associations of walkability parameters with neighborhood satisfaction and possible mediators. *Health & Place*, 17(4):971–977, July 2011. doi: 10.1016/j.healthplace.2011.04.001.

Chris Edwards. Tax reform and interstate migration. Technical Report 84, Cato Institute, September 2018.

Alessandra Faggian and Philip McCann. Universities, agglomerations and graduate human capital mobility. *Tijdschrift voor Economische en Sociale Geografie*, 100(2):210–223, April 2009. doi: 10.1111/j.1467-9663.2009.00530.x.

Kyle Fee, Keith Wardrip, and Lisa Nelson. Opportunity occupations revisited: Exploring employment for sub-baccalaureate workers across metro areas and over time. Technical report, Philadelphia Federal Reserve, April 2019.

Daniel Feenberg and Elisabeth Coutts. An introduction to the TAXSIM model. *Journal of Policy Analysis and Management*, 12(1):189, 1993. doi: 10.2307/3325474.

Robin Fisher, Geof Gee, and Adam Looney. Sex couples after Windsor: Characteristics of married tax filers in 2013 and 2014. Technical Report 108, U.S. Treasury Office of Tax Analysis, August 2016.

Ann Forsyth. Defining suburbs. *Journal of Planning Literature*, 27(3):270–281, June 2012. doi: 10.1177/0885412212448101.

Gary M. Fournier, David W. Rasmussen, and William J. Serow. Elderly migration: For sun and money. *Population Research and Policy Review*, 7(2): 189–199, May 1988. doi: 10.1007/bf00125466.

William F Fox, Henry W Herzog, and Alan M Schlottman. Metropolitan fiscal structure and migration. *Journal of Regional Science*, 29(4):523–536, November 1989. doi: 10.1111/j.1467-9787.1989.tb01242.x.

Filiz Garip. Social capital and migration: How do similar resources lead to divergent outcomes? *Demography*, 45(3):591–617, 2008. doi: 10.1353/dem.0.0016.

William H Greene. *Econometric Analysis*. Macmillan Publishing, second edition, 1993.

Michael J Greenwood. *Migration and EconomicGrowth in the United States: National, Regional, and Metropolitan Perspectives*. Academic Press, 1981.

Michael J Greenwood and Douglas Sweetland. The determinants of migration between standard metropolitan statistical areas. *Demography*, 9(4):665, November 1972. doi: 10.2307/2060673.

Tony H Grubesic and Timothy C Matisziw. 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(58), December 2006. doi: 10.1186/1476-072X-5-58.

Douglas T Gurak and Mary M Kritz. The interstate migration of U.S. immigrants: Individual and contextual determinants. *Social Forces*, 78(3):1017–1039, 2000. doi: 10.2307/3005940.

Deanna B. Haunsperger and Donald G. Saari. The lack of consistency for statistical decision procedures. *The American Statistician*, 45(3):252, August 1991. doi: 10.2307/2684305.

Miguel A Hernán, David Clayton, and Niels Keiding. The Simpson's paradox unraveled. *International Journal of Epidemiology*, 40(3):780–785, March 2011. doi: 10.1093/ije/dyr041.

Rubén Hernández-Murillo, Lesli S. Ott, Michael T. Owyang, and Denise Whalen. Patterns of interstate migration in the United States from the Survey of Income and Program Participation. *Review*, 93(3):169–85, May/June 2011. doi: 10.20955/r.93.169-186.

Henry W. Herzog and Alan M Schlottmann. State and local tax deductibility and metropolitan migration. *National Tax Journal*, 39(2):189–200, 1986.

M. Dolores Hidalgo and José Antonio López-Pina. Differential item functioning detection and effect size: A comparison between logistic regression and Mantel-Haenszel procedures. *Educational and Psychological Measurement*, 64 (6):903–915, December 2004. doi: 10.1177/0013164403261769.

Henry Hyatt, Erika McEntarfer, Ken Ueda, and Alexandria Zhang. Interstate migration and employer-to-employer transitions in the United States: New evidence from administrative records data. *Demography*, 55(6):2161–2180, October 2018. doi: 10.1007/s13524-018-0720-5.

David K Ihrke and Carol S Faber. Geographical mobility: 2005 to 2010. Technical report, U.S. Census Bureau, December 2012.

Richard Jackman and Savvas Savouri. Regional migration in Britain: An analysis of gross flows using NHS central register data. *The Economic Journal*, 102(415):1433, November 1992. doi: 10.2307/2234799.

Greg Kaplan and Sam Schulhofer-Wohl. Interstate migration has fallen less than you think: Consequences of hot deck imputation in the current population survey. *Demography*, 49(3):1061–1074, 2012.

John Kennan and James R Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251, 2011. doi: 10.3982/ecta4657.

Henrik Kleven, Camille Landais, Mathilde Muñoz, and Stefanie Stantcheva. Taxation and migration: Evidence and policy implications. *Journal of Economic Perspectives*, 34(2):119–142, May 2020. doi: 10.1257/jep.34.2.119.

Randall G. Krieg. Occupational change, employer change, internal migration, and earnings. *Regional Science and Urban Economics*, 27(1):1–15, February 1997. ISSN 01660462. doi: 10.1016/S0166-0462(96)02142-4.

Nancy Krieger, Pamela Waterman, Jarvis Chen, Mah-Jabeen Soobader, S. V. Subramanian, and Rosa Carson. 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(7):1100–1102, 2002.

Kristina Lerman. Computational social scientist beware: Simpson's paradox in behavioral data. *Journal of Computational Social Science*, 1(1):49–58, November 2017. doi: 10.1007/s42001-017-0007-4.

Nathan Mantel and William Haenszel. Statistical aspects of the analysis of data from retrospective studies of disease. *Journal of the National Cancer Institute*, April 1959. doi: 10.1093/jnci/22.4.719.

Terra McKinnish. Importing the poor. *Journal of Human Resources*, XL(1): 57–76, 2005. doi: 10.3368/jhr.xl.1.57.

Terra McKinnish. Welfare-induced migration at state borders: New evidence from micro-data. *Journal of Public Economics*, 91(3-4):437–450, April 2007. doi: 10.1016/j.jpubeco.2006.09.002.

Robert A. McLeman and Lori M. Hunter. Migration in the context of vulnerability and adaptation to climate change: Insights from analogues. *Wiley Interdisciplinary Reviews: Climate Change*, 1(3):450–461, May 2010. doi: 10.1002/wcc.51.

Raven Molloy, Christopher L Smith, and Abigail Wozniak. Internal migration in the United States. *Journal of Economic Perspectives*, 25(3):173–196, August 2011. doi: 10.1257/jep.25.3.173.

Michael A. Nelson and Michael L. Wyzan. Public policy, local labor demand, and migration in Sweden, 1979–84. *Journal of Regional Science*, 29(2):247–264, May 1989. doi: 10.1111/j.1467-9787.1989.tb01235.x.

Ryan Nunn, Laura Kawano, and Ben Klemens. Unemployment insurance and worker mobility. Technical report, Urban-Brookings Tax Policy Center, February 2018.

Karen O'Reilly. *Lifestyle Migration*. Routledge, May 2016. doi: 10.4324/9781315592398.

Janet Rothenberg Pack. Determinants of migration to central cities. *Journal of Regional Science*, 13(2):249–260, August 1973. doi: 10.1111/j.1467-9787.1973.tb00399.x.

Alberto Palloni, Douglas S. Massey, Miguel Ceballos, Kristin Espinosa, and Michael Spittel. Social capital and international migration: A test using information on family networks. *American Journal of Sociology*, 106(5):1262–1298, March 2001. doi: 10.1086/320817.

Blaise Pascal. On the arithmetic triangle. In David Eugene Smith, editor, *A Source Book in Mathematics*, chapter 10, pages 67–84. McGraw-Hill Book Co, 1929.

Neil Pearce. Effect measures in prevalence studies. *Environmental Health Perspectives*, 112(10):1047–1050, July 2004. doi: 10.1289/ehp.6927.

Jacques Poot, Omoniyi Alimi, Michael P. Cameron, and David C. Maré. The gravity model of migration: The successful comeback of an ageing superstar in regional science. *Investigaciones Regionales*, 2016(36–Specialissue):63–86, 1 2016.

Robert R Preuhs. State policy components of interstate migration in the United States. *Political Research Quarterly*, 52(3):527–547, 1999.

Michael A. Quinn and Stephen Rubb. The importance of education-occupation matching in migration decisions. *Demography*, 42(1):153–167, 2005. doi: 10.1353/dem.2005.0008.

Priya Ranganathan, Rakesh Aggarwal, and CS Pramesh. Common pitfalls in statistical analysis: Odds versus risk. *Perspectives in Clinical Research*, 6(4): 222, 2015. doi: 10.4103/2229-3485.167092.

Jordan Rapoport. The faster growth of larger, less crowded locations. *Economic Review*, 103(4):5–38, Fourth Quarter 2018.

Michael Ratcliffe, Charlynn Burd, Kelly Holder, and Alison Fields. Defining rural at the U.S. Census Bureau. Technical Report ACSGEO-1, U.S. Census Bureau, December 2016.

Jennifer Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6):1257–1278, December 1982. doi: 10.1086/261120.

H. Jane Rogers and Hariharan Swaminathan. A comparison of logistic regression and Mantel-Haenszel procedures for detecting differential item functioning. *Applied Psychological Measurement*, 17(2):105–116, June 1993. doi: 10.1177/014662169301700201.

Sherri Rose and Mark J. van der Laan. Why match? Investigating matched case-control study designs with causal effect estimation. *The International Journal of Biostatistics*, 5(1), January 2009. doi: 10.2202/1557-4679.1127.

Kenneth J. Rothman. *Modern Epidemiology*. LWW, December 2012.

Nikola Sander and Martin Bell. Migration and retirement in the life course: An event history approach. *Journal of Population Research*, 31(1):1–27, December 2013. doi: 10.1007/s12546-013-9121-1.

Carsten Oliver Schmidt and Thomas Kohlmann. When to use the odds ratio or the relative risk? *International Journal of Public Health*, 53(3):165–167, June 2008. doi: 10.1007/s00038-008-7068-3.

E. H. Simpson. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2):238–241, July 1951. doi: 10.2307/2984065.

Larry A. Sjaastad. The costs and returns of human migration. *Journal of Political Economy*, 70(5, Part 2):80–93, October 1962. doi: 10.1086/258726.

Oded Stark and David E. Bloom. The new economics of labor migration. *The American Economic Review*, 75(2):173–178, 1985. doi: 10.2307/1805591.

MoonJoong Tcha. Altruism, household size and migration. *Economics Letters*, 49(4):441–445, October 1995. doi: 10.1016/0165-1765(95)00702-h.

Charles M Tiebout. A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424, 1956.

Radu Tunaru. Models of association versus causal models for contingency tables. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 50(3): 257–269, 2001.

Richard Vedder. Tiebout, taxes, and economic growth. *Cato Journal*, 10(1): 91–108, 1990.

Roman Vershynin. *High-Dimensional Probability*. September 2018.

Sholom Wacholder. Binomial regression in GLIM: Estimating risk ratios and risk differences. *American Journal of Epidemiology*, 123(1):174–184, January 1986. doi: 10.1093/oxfordjournals.aje.a114212.

Kyle E. Walker. The shifting destinations of metropolitan migrants in the u.s., 2005-2011. *Growth and Change*, 48(4):532–551, January 2017. doi: 10.1111/grow.12187.

Paul R. Yarnold. Characterizing and circumventing Simpson's paradox for ordered bivariate data. *Educational and Psychological Measurement*, 56(3):430–442, June 1996. doi: 10.1177/0013164496056003005.

Cristobal Young, Charles Varner, Ithai Z. Lurie, and Richard Prisinzano. Millionaire migration and taxation of the elite. *American Sociological Review*, 81 (3):421–446, May 2016. doi: 10.1177/0003122416639625.

G. Undy Yule. Notes on the theory of association of attributes in statistics. *Biometrika*, 2(2):121–134, 1903. doi: 10.1093/biomet/2.2.121.