

NBER WORKING PAPER SERIES

UNDERSTANDING THE LONG-RUN DECLINE IN INTERSTATE MIGRATION

Greg Kaplan
Sam Schulhofer-Wohl

Working Paper 18507
<http://www.nber.org/papers/w18507>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2012

We thank Diego Amador and Chloe Booth for excellent research assistance and Joan Gieseke for editorial assistance. We also thank Esteban Rossi-Hansberg and numerous seminar and conference participants for helpful comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by Greg Kaplan and Sam Schulhofer-Wohl. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Understanding the Long-Run Decline in Interstate Migration
Greg Kaplan and Sam Schulhofer-Wohl
NBER Working Paper No. 18507
November 2012
JEL No. D83,J11,J24,J61,R12,R23

ABSTRACT

We analyze the secular decline in interstate migration in the United States between 1991 and 2011. Gross flows of people across states are about 10 times larger than net flows, yet have declined by around 50 percent over the past 20 years. We argue that the fall in migration is due to a decline in the geographic specificity of returns to occupations, together with an increase in workers' ability to learn about other locations before moving there, through information technology and inexpensive travel. These explanations find support in micro data on the distribution of earnings and occupations across space and on rates of repeat migration. Other explanations, including compositional changes, regional changes, and the rise in real incomes, do not fit the data. We develop a model to formalize the geographic-specificity and information mechanisms and show that a calibrated version is consistent with cross-sectional and time-series patterns of migration, occupations, and incomes. Our mechanisms can explain at least one-third and possibly all of the decline in gross migration since 1991.

Greg Kaplan
Department of Economics
Princeton University
Fisher Hall
Princeton, NJ 08544
and NBER
gkaplan@princeton.edu

Sam Schulhofer-Wohl
Research Department
Federal Reserve Bank of Minneapolis
90 Hennepin Ave.
Minneapolis MN 55480-0291
wohls@minneapolisfed.org

1. Introduction

In the early 1990s, about 3 percent of Americans moved between states each year. Today, that rate has fallen by half. We show that micro data rule out many popular explanations for this change, such as aging of the population or changes in the number of two-earner households. But the data do support two novel theories. The first theory is that labor markets around the country have become more similar in the returns they offer to particular skills, so workers need not move to a particular place to maximize the return on their idiosyncratic abilities. The second theory is that better information — due to both information technology and falling travel costs — has made locations less of an experience good, reducing the need for young people to experiment with living in different places. We build a model that makes these ideas precise and show that a plausibly calibrated version is consistent with cross-sectional and time-series patterns.

Many policymakers have worried that the decline in migration heralds a less-flexible economy where workers cannot move to places with good jobs. In such an economy, the labor market might adjust more slowly to shocks, potentially prolonging recessions and reducing growth. Low migration has thus been proposed as an explanation for the slow recovery from the 2007–’08 financial crisis (see, for example, Batini et al., 2010). But the causes of decreased migration that we identify suggest that the economy may not be less flexible after all. Rather, low migration means that workers either do not need to move to obtain good jobs or have better information about their opportunities. In either case, the appropriate policy response may differ from the appropriate response to a decrease in workers’ ability to move. Thus, understanding the causes of the decline in gross migration is an important goal for economists.

Figure 1 shows gross and net interstate migration rates over the past half-century. The gross rate — the fraction of U.S. residents at least 1 year old who lived in a different state one year ago — comes from the Annual Social and Economic Supplement to the Current Population Survey, commonly known as the March CPS. The net rate comes from the Census Bureau’s annual state population estimates (U.S. Census Bureau, 1999, 2009a). Several key patterns are immediately apparent. First, net flows are an order of magnitude smaller than gross flows. Second, while the gross flows exhibit some cyclical fluctuations, these fluctuations

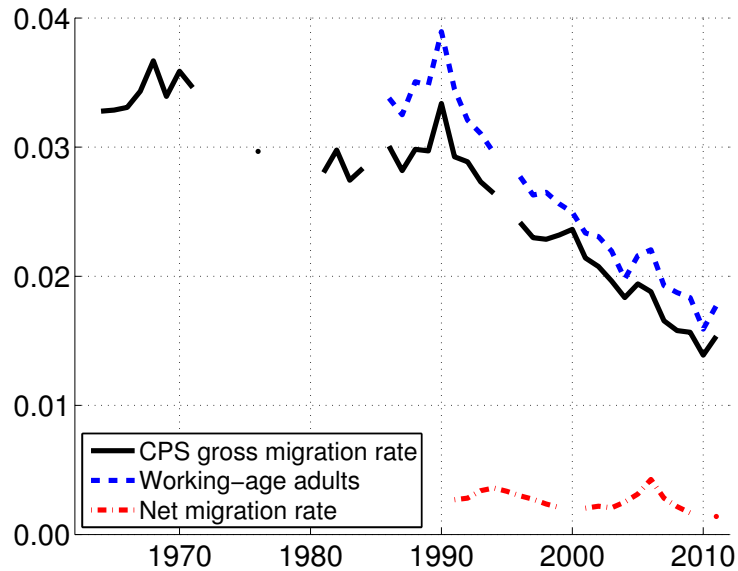


Figure 1: Gross and net interstate migration.

Source: Authors' calculations from Current Population Survey (CPS) micro data and Census Bureau population estimates. The numerator of the net migration rate is one-half of the sum of absolute values of inflows minus outflows in each state. (This number is the minimum number of moves that would have to be prevented to set net migration to zero in every state.) The denominator of the rate between years t and $t + 1$ is the U.S. population at t minus deaths between t and $t + 1$.

are much smaller than the overall decline over the past 20 years. Third, the trend in gross flows is virtually identical when we restrict the analysis to a sample of working-age adults in civilian households. These patterns suggest that to understand the decline in migration, we must look for factors that affect gross flows rather than net flows; that vary over long time horizons rather than at business cycle frequencies; and that affect working-age people, rather than only people making life cycle-related transitions such as retiring or moving for college.

Two additional patterns guide our focus on information and on workers. Figure 2(a) shows that even among recent immigrants to the United States, the fraction who move between states after arriving has fallen over time. This decline is broadly consistent with both theories we propose. Improved information may make immigrants better able to choose a good initial destination. Alternatively, if immigrants choose their initial destinations based on family or ethnic ties (e.g., MacDonald and MacDonald, 1964), then later move to places where their idiosyncratic job matches are better, a decline in interstate migration by new immigrants is consistent with the hypothesis that locations have become more similar in the jobs they offer, so that there is less reason for immigrants to change their initial locations.

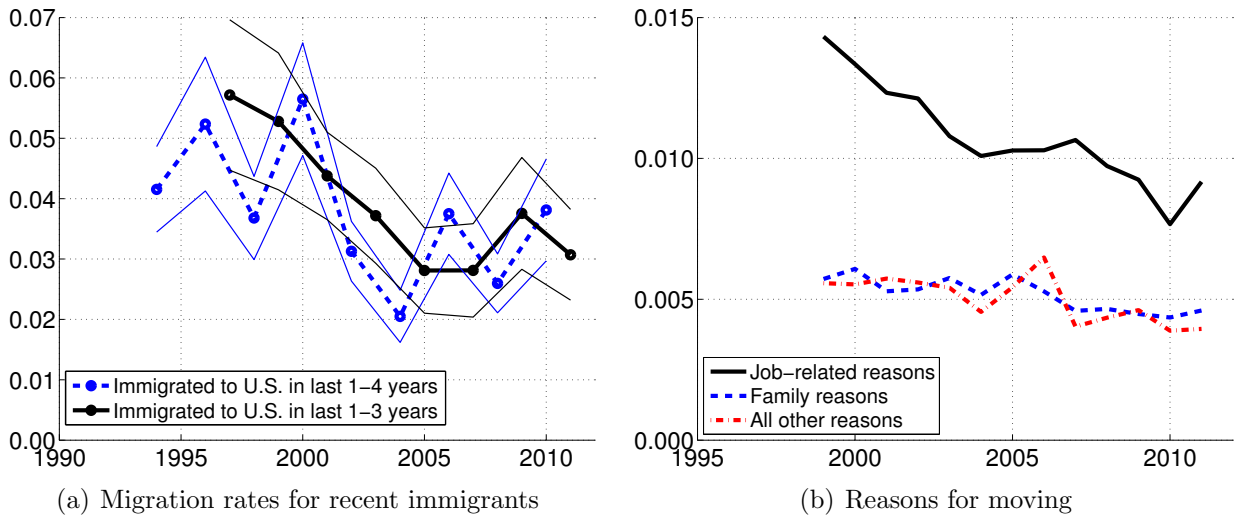


Figure 2: Key patterns in migration rates.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Estimates shown for all years when variables are available. In figure 2(a), sample is further restricted to individuals with non-imputed data on number of years in the United States, estimates are standardized to the mean age distribution for new immigrants over the years shown, and thin lines show 1-standard-error confidence band around point estimates.

Figure 2(b) examines the dimensions of information that may matter by showing the fraction of Americans who say they moved between states for various reasons. Job-related reasons — primarily moves for new jobs or job transfers — have declined sharply, while other types of moves have declined more slowly. Of course, the reasons people give in a survey may not be their true reasons for moving. However, when a survey respondent says she moved for a new job, we think it is highly likely that she changed jobs around the time of the move — even if other factors, such as local amenities, motivated the desire to search for a job in a new location. Thus, to understand why migration is falling, we need to understand why people have become less likely to make moves that happen around the same time as job changes.

The decline in job-related moves suggests that the potential improvements in job opportunities from moving are smaller than in the past. However, any decline in the impact of moving on job opportunities cannot come simply from convergence of mean incomes across states: Such a change would reduce net migration, not gross migration. Rather, there must be a change in the importance of the match between a particular worker and a particular location. In our model, workers choose between two locations and two occupations. Each

worker has different skills in the two occupations, and each occupation is more productive in one location. Changes in this occupation-location premium, which we call the geographic specificity of occupations, have no effect on net flows but do change gross flows by reducing workers' need to sort into the places where their particular skills are most productive.

A decrease in the geographic specificity of occupations cannot be the whole story, however. If locations offer more similar jobs, workers will be less likely to move for work but more likely to move for amenity-related reasons, because a smaller difference in amenities is now required to overcome the difference in earnings. But as figure 2(b) shows, amenity-related moves have not risen. Thus, some other factor must also be at play.

In our model, this other factor is information. The two locations in our model differ in both the job opportunities and the local amenities they offer. Based on evidence that most workers who move for job-related reasons do so with a new job in hand, we assume that workers can search remotely for a job and know the distribution of job opportunities in remote locations. However, we assume that amenities are an experience good: Workers do not know how much they will like the sun in California until they live there. If workers in one location are sufficiently uncertain about amenities in the other location, they may move simply to acquire information. We call this a move for experimentation purposes. If the new location proves to be good, an experimenting worker will stay there; if not, she may return to her original location. We model an increase in information as an increase in the precision of workers' priors about the amenities in each location. Tighter priors have two opposing effects on migration. First is a "news" effect: Some workers discover that they would prefer a different location and decide to move. Second is an experimentation effect: Some workers who would have made experimental moves no longer do so because the tighter prior reduces their need to acquire information. Because people who move for experimentation purposes often dislike the new location and return to the origin, the reduction in experimental moves is larger than the increase in news-driven moves, and more information leads to lower migration overall.

We provide direct empirical evidence for both of our theories. To support the hypothesis that migration has fallen because job opportunities have become more similar across locations, we show that occupations have become more evenly spread across the country. Fur-

ther, this change appears to result from a decrease in the dispersion of productivity rather than a change in the supply of workers willing to take jobs in particular places, because we also find that the variance across states of the average income for a given occupation has fallen. (If, instead, workers increasingly desired jobs in unproductive places — due for example to an exogenous decrease in mobility — the dispersion of incomes would rise.) In addition, we show that migration responds to the geographic specificity of occupations: On average, workers move to states where their particular occupations are better paid.

Recent advances in information technology and decreases in travel costs clearly make it easier for workers to learn about faraway places.¹ The hypothesis that increases in information have reduced the need to migrate also has a testable prediction: Rates of repeat migration should have declined because migrants will be more likely to be satisfied with their destinations. We turn to panel data to test this prediction and find that repeat migration indeed has declined, although the estimates are imprecise.

We use a calibrated model to demonstrate that our theories not only are consistent with the data but also can explain a substantial portion of the decline in gross migration. The decrease in geographic specificity of occupations explains one-third of the fall in migration since 1991. An increase in information can explain as much as all of the remaining decrease.

Our work is related to a substantial literature. Molloy, Smith, and Wozniak (2011) survey research on internal migration in the United States and describe important patterns in the decline in interstate migration, finding, as we do, that compositional changes cannot explain much of the decline. Our analysis of compositional changes extends theirs by considering more fine-grained measures of some variables and by formally calculating counterfactual migration rates that hold composition fixed.²

¹In principle, lower travel costs might raise migration by making it easier to move. However, much of the cost of moving is a time cost — the migrant must find a new home, pack and unpack belongings, and find local services such as doctors and schools — that lower airfares cannot offset. Moreover, if lower travel costs should have increased migration, the observed decline is simply a larger puzzle; our mechanisms then explain less of the decrease relative to the appropriate counterfactual, leaving more room for other explanations.

²One factor that has received much attention but that we do not consider here is fluctuations in the housing market. The trend we document is a secular decline in migration over at least 20 years, during which house prices and homeownership rose and then fell. If house prices and homeownership are important determinants of gross migration, it is difficult to explain why the decline in migration was monotonic while the housing market fluctuated sharply. In addition, Molloy, Smith, and Wozniak (2011) show that the decline in house prices since the mid-2000s plays at most a small role in the drop in migration over that period.

Theories of migration, such as the classic models by Harris and Todaro (1970) and Roback (1982), generally focus on net flows. In related empirical work, Ganong and Shoag (2012) analyze the relationship between income convergence, net flows, and housing regulation, while Partridge et al. (2012) study the response of net migration to demand shocks.

Kennan and Walker (2011) structurally estimate a model in which workers choose locations to maximize their expected lifetime income. Differences in expected income across locations imply that the model features both gross and net flows. The price of studying net flows is that Kennan and Walker (2011) must allow workers to choose among many locations, which means the model has many state variables and must be highly simplified along many dimensions to remain tractable. By studying only gross flows, we can reduce our model to two locations and add realism along other important dimensions, such as utility from amenities, learning, and geographic specificity of skills. Bayer and Juessen (2012) similarly study gross flows in a two-location model but focus on how the autocorrelation of income affects selection into migration; their model does not include amenities or learning, and they investigate cross-sectional patterns rather than the change in migration over time. Coen-Pirani (2010) builds a model to explain gross and net flows but does not analyze the decline in gross flows.

Our analysis is also connected to the literatures on agglomeration effects, city growth rates, and the concentration of industries in particular regions. To our knowledge, the decrease in the geographic specificity of occupations has not been described previously in the economic literature; it is distinct, for example, from the shift of aggregate employment toward less-dense areas described as “deconcentration” by Carlino and Chatterjee (2002). Our finding that workers move to states where their occupations are better paid is similar to Borjas, Bronars, and Trejo’s (1992) conclusion that high-skill workers tend to move to places with higher returns to skill, although they focus on one-dimensional measures of skill such as education or aptitude test scores rather than on a multitude of occupations. The New Economic Geography literature, starting with Krugman (1991), studies how transportation costs and local economies of scale lead workers and firms to concentrate in one location. In the typical model, these effects largely result in net flows: workers move toward the more populated region. Empirically, Crozet (2004) tests the ability of a New Economic Geography model to explain labor migration. More broadly, changes in agglomeration effects could help

explain the changes in the geographic dispersion of wages that we take as exogenous; Duranton and Puga (2004) review a variety of mechanisms that generate agglomeration effects and, hence, might affect the geographic dispersion of wages. The literatures on city population growth, industry concentration, and concentration of skilled workers in particular cities also essentially analyze net rather than gross flows. Again, though, the theoretical mechanisms proposed in these literatures, such as learning through interaction with other workers (Glaeser, 1999), linkages between human capital and entrepreneurship (Berry and Glaeser, 2005; Glaeser, Ponzetto, and Tobio, forthcoming), or technological diffusion and knowledge spillovers (Desmet and Rossi-Hansberg, 2009), could help explain changes in the geographic dispersion of wages.

The paper proceeds as follows. In section 2, we describe the CPS data and compare migration rates in the CPS and other datasets. In section 3, we review a litany of demographic and economic theories of falling migration and show that they are incompatible with the data. Section 4 presents direct evidence for the key mechanisms in our model: We show that the returns to working in particular occupations have become less geographically dispersed and that repeat migration rates have declined, and we review evidence for falling costs of learning about distant locations. Section 5 lays out our model of information and migration. Section 6 calibrates the model, examines its success in fitting the data, and quantifies how much of the decline in migration our mechanisms can explain. Section 7 concludes.

2. Data

We focus our analysis on working-age adults in civilian households in the March CPS from 1991 to 2011.³ (We start the analysis in 1991 because, as shown in figure 1, the CPS migration rate spikes in 1990, but the cause of this spike is unclear and we do not want it to unduly influence our results.) We define a civilian household as one where no household member is in the military; excluding military households is important because military households move frequently and the military has become smaller (Pingle, 2007). We define a working-age adult as one who is no more than 55 years old and either (a) has a bachelor's degree and is at least 23 years old, or (b) does not have a bachelor's degree, is not currently enrolled in

³Our data omit 1995 because the CPS did not measure one-year migration that year.

school, and is at least 19 years old.⁴ Thus, we concentrate on people who have completed their education but are not yet approaching retirement. From 1996 onward, we follow Kaplan and Schulhofer-Wohl (2012) and exclude observations with imputed migration data so that changes in CPS imputation procedures do not produce spurious fluctuations in the migration rate. (The imputation rate before 1996 is negligible.) We obtain most of the data from the Integrated Public Use Microdata Series (King et al., 2010) but identify imputed observations with the imputation flags on the original public-use files from the Bureau of Labor Statistics.

The CPS measures migration with retrospective questions: Did the respondent live in the same home one year ago, and if not, where did he or she live? We drop respondents who did not live in the United States one year ago so that fluctuations in immigration do not affect our results. Since we are interested in how internal migration affects the labor market, we ideally would measure migration between distinct labor markets, such as the commuting zones defined by the Bureau of Labor Statistics (Tolbert and Sizer, 1996). However, we cannot identify migrants' origin and destination commuting zones because commuting zones are groups of counties and origin counties are not available in the CPS public-use files. Instead, we examine migration between states. In most parts of the country, states are large enough that labor markets do not cross state borders. Of course, by looking at interstate migration, we miss some migration between distinct labor markets within a state and include some migration that does not entail changing labor markets, such as when a worker in Manhattan moves to a New Jersey suburb. We show below that our results are robust to controlling for the latter bias by excluding two sets of states where the problem is particularly severe — New York, New Jersey, and Connecticut, as well as Maryland, Virginia, and Washington, D.C. — from the data.

It is unlikely that the long-run decline we describe is a mechanical result of under-sampling people who move to newly built homes. First, the CPS sample frame is designed to capture new construction (U.S. Census Bureau, 2006, chap. 3).⁵ In addition, if a bias

⁴The CPS measures current school enrollment only for people ages 16 to 24. We treat all people over age 24 as not currently enrolled in school.

⁵The CPS technical documentation (U.S. Census Bureau, 2006, chaps. 3, 15) does report that the sample may miss some newly built group quarters. However, excluding group quarters residents from the sample does not change the gross interstate migration rate by more than 0.01 percentage point in any year.

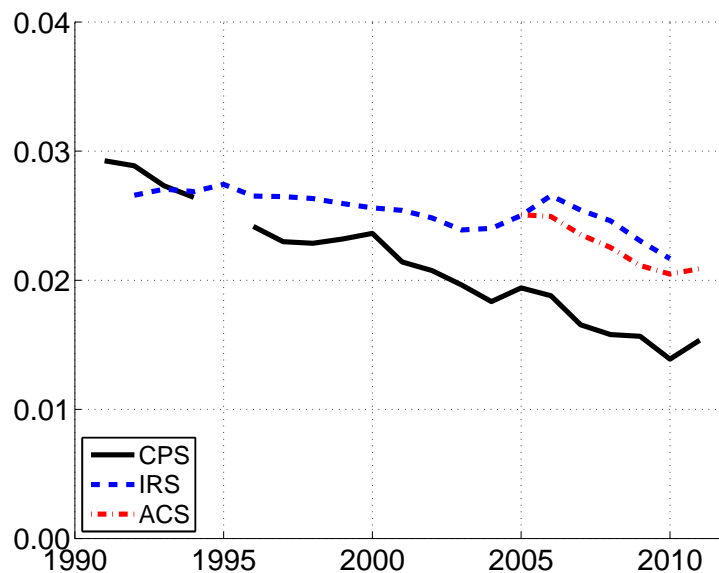


Figure 3: Gross migration measured from different data sources.

Source: Authors' calculations from Current Population Survey (CPS) and American Community Survey (ACS) micro data and Internal Revenue Service (IRS) state-to-state migration tabulations.

associated with new construction were the main driver of changes in the CPS migration rate, the rate would have fallen sharply during the housing boom of the mid-2000s and risen during the housing bust; it did not.

We use the CPS data because they cover many years and contain a myriad of covariates that allow us to test hypotheses about the decline in migration. However, the decline in measured annual interstate migration rates appears in other data as well.⁶ Figure 3 compares migration rates in the CPS, in micro data from the Census Bureau's American Community Survey (ACS) (Ruggles et al., 2010), and in Internal Revenue Service (IRS) data. The ACS migration rate parallels the CPS rate from 2005 to 2011 but is about one-half of a percentage point higher in each year, likely because the ACS pursues nonrespondents more intensively (Koerber, 2007). We do not examine earlier ACS data because, before 2005, the ACS was a pilot project and occasional changes in survey procedures may have affected the estimated

⁶It is also possible to measure migration over horizons longer than a year, for example by asking whether individuals lived in a different state five years ago or were born in a different state than the one where they live now. Molloy, Smith, and Wozniak (2011) show that these long-term migration rates have fallen much less than annual migration rates. However, long-term migration rates respond to fundamentally different factors than annual rates because they ignore some return migration (someone who moves away and returns within a two-year period will not count as a migrant in the five-year measure) and integrate individuals' behavior over many years.

migration rate.⁷ The IRS data cover more years; they, too, show a decline, albeit smaller than in the CPS. However, the IRS data are not a perfect measure of migration: They cover only people with incomes high enough to file taxes, track mailing addresses rather than home addresses, and can be distorted by changes in household formation and in the time of year when people file their returns (Internal Revenue Service, 2008).⁸

Because the CPS is a very large sample — more than 200,000 individuals in 2011 — the standard errors of our estimates are typically minuscule, on the order of one-tenth of a percentage point, and we omit them from most of the graphs in the next section in the interest of legibility. However, we show standard errors when their magnitude is meaningful.⁹

⁷The ACS initially mails a survey form to sample housing units, then tries to telephone those who do not return the mailed form, and finally sends field representatives to personally interview a subsample of those who are not reached by mail or phone (U.S. Census Bureau, 2009b, chap. 7). Estimated migration rates are lower for mail and telephone respondents than for in-person respondents, probably because the mail and telephone surveys are less likely to reach recent movers (Koerber, 2007). (Koerber, 2007, finds that migration rates vary little by response mode in the CPS because the CPS does not use mail surveys and is more likely than the ACS to have up-to-date phone numbers for respondents.) Thus, any variation in the percentage of respondents interviewed in person could translate into variation in the estimated migration rate. From 2000, when the ACS became a large national demonstration project, through 2005, the first year of full implementation, survey procedures changed at least four times in ways that could have affected the rate of in-person interviews and hence the migration rate. First, in 2002, budget constraints caused the ACS to conduct no surveys in July and to skip telephone and in-person follow-ups in June (Garrett and Williams, 2006). Second, in mid-2002, the rate of telephone interviews increased because the ACS obtained more telephone numbers from the decennial census and was able to contact more housing units by phone (Garrett and Williams, 2006). Third, in 2004, budget constraints caused the ACS to skip telephone and in-person follow-ups in January (U.S. Census Bureau, 2012). Fourth, through 2004, all housing units that were not reached by telephone or mail had a one-in-three chance of inclusion in the personal-interview subsample (U.S. Census Bureau, 2004), but since 2005, the probability of a personal interview has varied by census tract and by whether the housing unit has an address where mail can be delivered (U.S. Census Bureau, 2005, 2009b). The pre-2005 public-use data do not indicate the response mode, and the data since 2005 do not distinguish between telephone and in-person interviews, so we cannot adjust the data to account for these effects.

⁸Based on IRS news releases reporting the number of returns filed each week during the filing season from 1996 to 2010, the median filing date of individual income tax returns appears to be shifting earlier by about one day every two years. However, the news release data are too imprecise to allow us to measure the second derivative of the median filing date, which is what determines the timing bias, if any, in the IRS data.

⁹From 2005 onward, we calculate standard errors using the person-level replicate weights provided by the Census Bureau that account for the design of the CPS sample. Before 2005, replicate weights are not available, so we calculate standard errors by assuming that the survey weights are inversely proportional to the probability of sampling and that the sample is clustered by households. Clustering on households and replicate weights give virtually identical standard errors for the interstate migration rate in 2005 and later years. We do not follow Davern et al.'s (2006) method of clustering on geographic areas because it gives larger standard errors than the replicate weights, likely because clustering on geography is too conservative when analyzing a variable such as interstate migration that is not highly correlated across neighboring households. When we combine estimates for multiple years, we calculate separate point estimates for each year, take the unweighted average across years, and calculate the standard error assuming the estimates in different years are independent. (Because the CPS is a rotating panel of addresses, this assumption is not strictly correct, but the available sampling information in the public-use files does not allow us to easily relax it.)

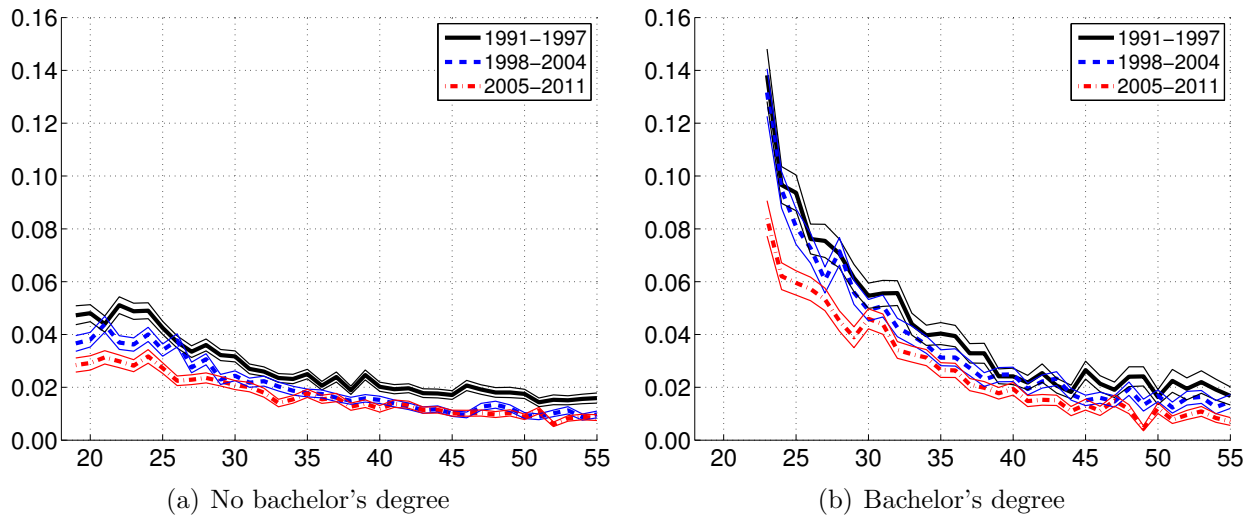


Figure 4: Age profile of interstate migration.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Thin lines show 1-standard-error confidence bands around point estimates.

One limitation of the CPS is that it does not follow migrants over time and thus does not let us see how earnings and other outcomes change when a particular person moves, or even to see how likely a migrant is to migrate again. When we examine repeat migration, we turn to panel data from the Survey of Income and Program Participation. However, for most purposes, panel data are not ideal for measuring migration because results will depend on the survey's success rate in tracking respondents who move, and this success rate could change over time independent of any changes in actual behavior. Thus, we focus most of our analysis on the CPS, which is not subject to attrition bias precisely because it is cross-sectional.

3. Patterns of Migration: Theories and Data

This section describes demographic and economic patterns in migration over the past two decades. We use these patterns both to learn what dimensions are important to model and to show that various common beliefs about the fall in migration do not match the data.

A. Life cycle patterns and composition effects

Figure 4 shows the age profile of migration rates separately for college graduates and nongraduates in our sample of working-age adults. Migration rates decline sharply with age, but this decline is steeper for college graduates, who migrate much more than nongraduates

up to about age 40. Since 1991, the migration rate has fallen at all ages. Thus, although the population is aging and older people migrate less, the aggregate decline in migration cannot be due solely to population aging; the aggregate rate would have fallen even if the age distribution had remained the same. Importantly, however, the decline in migration is larger for the young — a fact we will ask our calibrated model to reproduce.

Figure 5(a) quantifies the importance of population aging by calculating what the interstate migration rate would have been in each year if the age distribution had not changed after 1991. The effect is tiny: Holding the age distribution fixed, the migration rate would have been 0.1 percentage point higher in 2011. We find similar results when we adjust for changes in the distribution of education, marital status, or number of labor force participants in the household: Figures 5(b), 5(c), and 5(d) show that migration rates have fallen at all education levels, for people of all marital statuses, and for both single-earner and multiple-earner households. Further, figure 5(a) shows almost unnoticeable effects on gross migration of holding the population distributions of these variables fixed at the 1991 distribution. Thus, although the demographics of the U.S. population have changed in many ways since 1991, these changes have no power for explaining the decline in interstate migration. In particular, the findings on marital status and number of earners demonstrate that the fall in migration is not due to changes in the number of “tied stayers” (Gemici, 2011; Guler, Guvenen, and Violante, 2012; Mincer, 1978) who cannot move because their partners cannot move.

B. Occupation and industry effects

Over the past several decades, the service sector has expanded while manufacturing has declined. If workers’ mobility rates differed across industries, this sectoral shift could produce a decline in migration. However, figure 6(a) shows that service-industry workers have approximately the same mobility as workers in other industries and that mobility has declined in parallel for workers in all industries. Further, figure 5(a) shows that when we hold the industry distribution fixed at the 1991 distribution, the migration trend does not change. Thus, the rise of the service sector seems unlikely to explain the decline in migration.

Another hypothesis is that new communications technologies reduce migration by allowing some workers to do their jobs from anywhere in the country, instead of having to

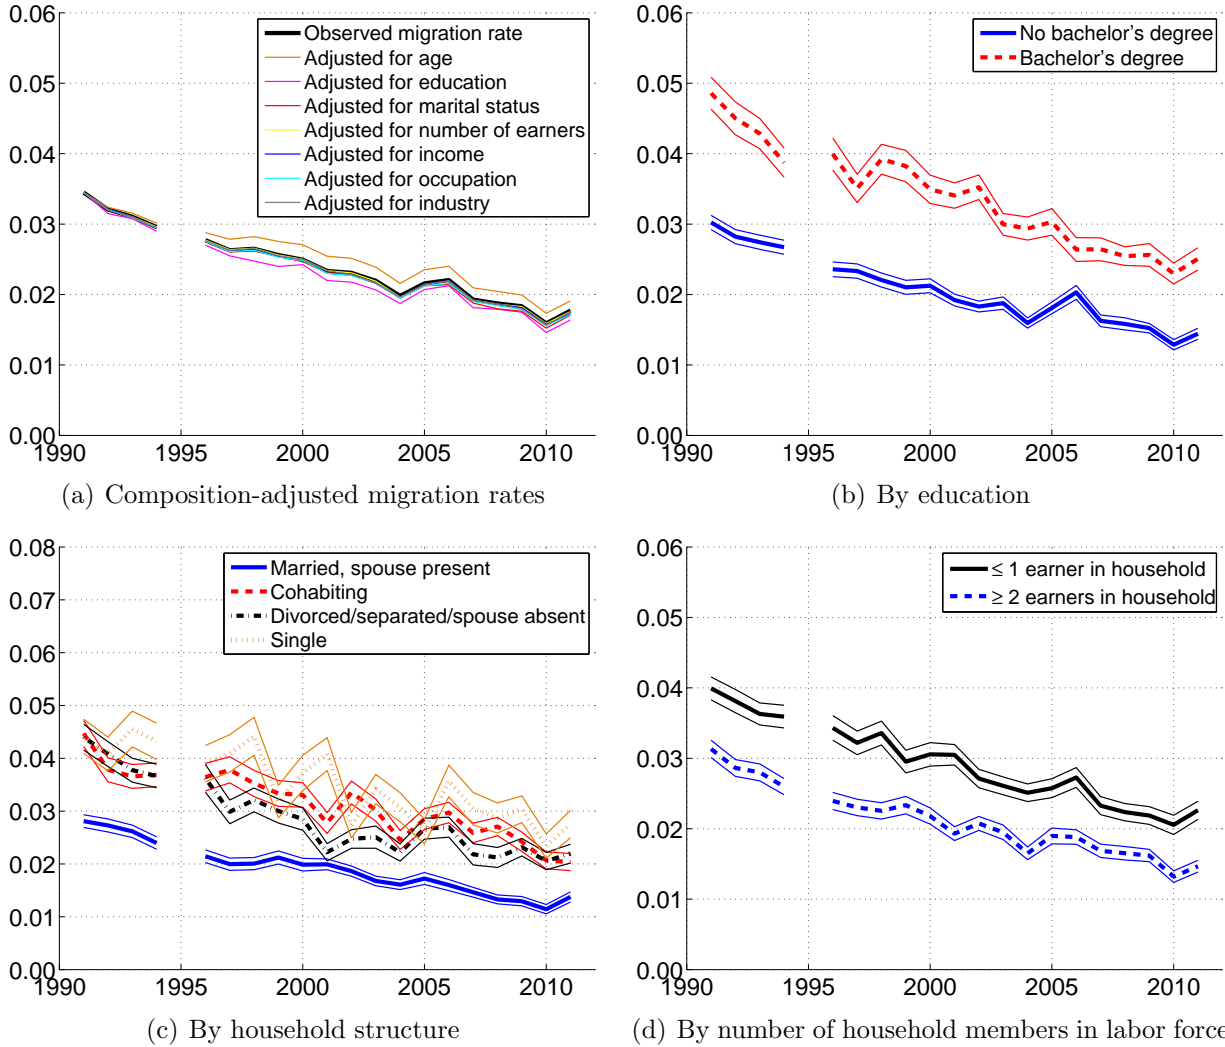


Figure 5: Time series of interstate migration by population subgroups.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Composition-adjusted rates hold the following variables constant at their 1991 distribution: respondent's age (single years), respondent's education (single years), respondent's marital status (four categories shown in figure 5(c)), number of labor force participants in respondent's household (two categories shown in figure 5(d)), and real income per capita of respondent's household (20 equal-population bins in 1991). Thin lines in figures 5(b), 5(c), and 5(d) show 1-standard-error confidence bands around point estimates.

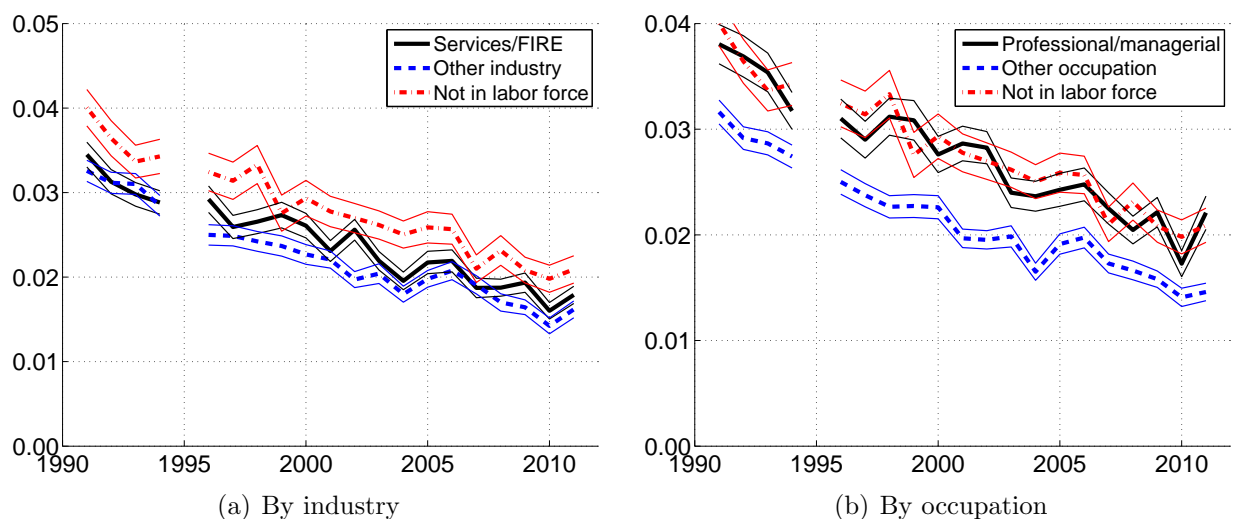


Figure 6: Migration rate by industry and occupation.

Source: Authors’ calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Thin lines show 1-standard-error confidence bands around point estimates. Professional/managerial occupations include all managerial and professional specialty occupations, and service/FIRE industries include all service industries and finance, insurance, and real estate as coded by King et al. (2010).

live in the city where their employer has its operations. These changes affect some occupations much more than others. But figure 6(b) shows that the migration rate for professional and managerial workers — who may have the most opportunities to work remotely — has declined only slightly more than the migration rate for workers in other occupations. Thus, we must seek an explanation for decreased migration that applies to all workers, not just those who can do their jobs over the Internet.

C. Income effects

The recent recession notwithstanding, the United States has grown wealthier since 1991. If living in one place for a long time is a normal good, the rise in incomes could cause a fall in migration. Figure 7 tests this hypothesis by estimating the migration rate as a function of real household income per capita, controlling for age.¹⁰ Controlling for age is important

¹⁰We obtain the graph by estimating a partially linear model in which migration depends linearly on a full set of age indicator variables and nonparametrically on income: $migration = \mathbf{x}\boldsymbol{\beta} + f(income) + \epsilon$, where \mathbf{x} represents the age indicators. We estimate the linear part using the method of Yatchew (1997) with tenth-order differencing, then estimate the nonparametric part with local linear regression. We normalize the age effects to have mean zero.

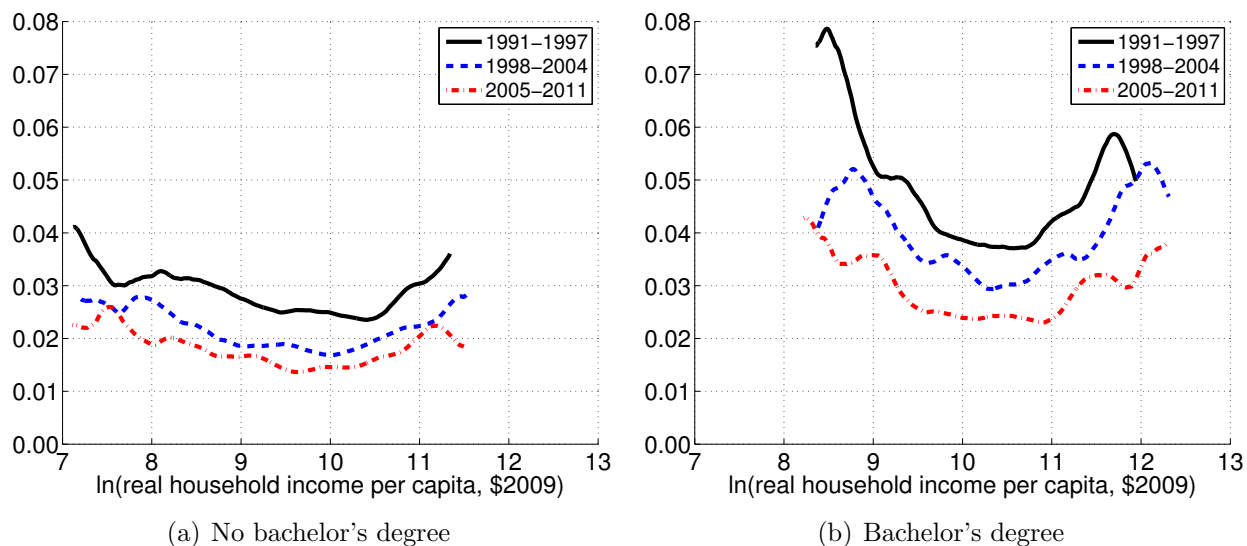


Figure 7: Migration rate by real household income, controlling for age.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Lines are local linear regression estimates of the migration rate as a function of income, controlling for age indicator variables in a partially linear model using the method of Yatchew (1997). Estimates use bandwidth 0.15 and Epanechnikov kernel. The graph is truncated at the 1st and 99th percentiles of the income distribution, conditional on education, in each period.

because young people tend to have lower incomes and migrate more. Even after we control for age, migration is indeed higher at the low end of the income distribution. However, migration also ticks up at the high end of the distribution, so if income gains were concentrated among the already well off, the overall rise in incomes would not necessarily reduce migration. In addition, the figure shows that migration rates fell uniformly across the income distribution, and figure 5(a) shows that holding the real income distribution constant would not change the overall migration rate. Thus, rising real incomes do not explain the fall in migration.

D. Regional effects

Throughout U.S. history, high migration rates have been associated with large flows from one part of the country to another, most prominently in the Great Migration of African Americans out of the South. Is the recent decline in migration merely the result of a change in flows into or out of one part of the country? The net migration rate shown in figure 1 suggests not: Even if all net interstate migration were eliminated, gross flows would barely change. Figure 8 examines this question another way by disaggregating the gross migration

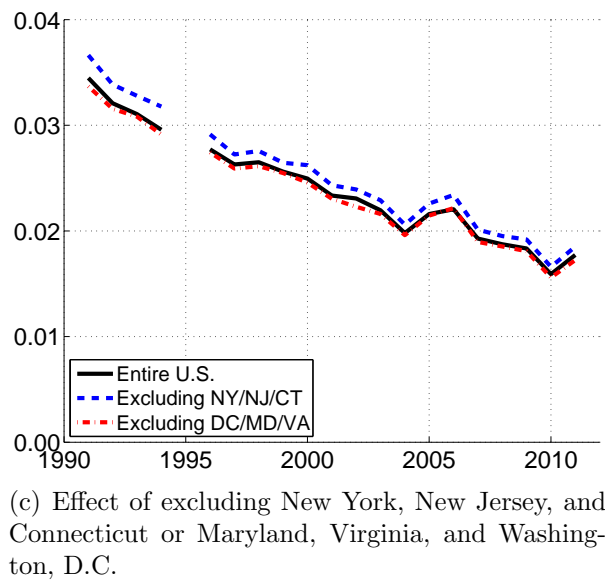
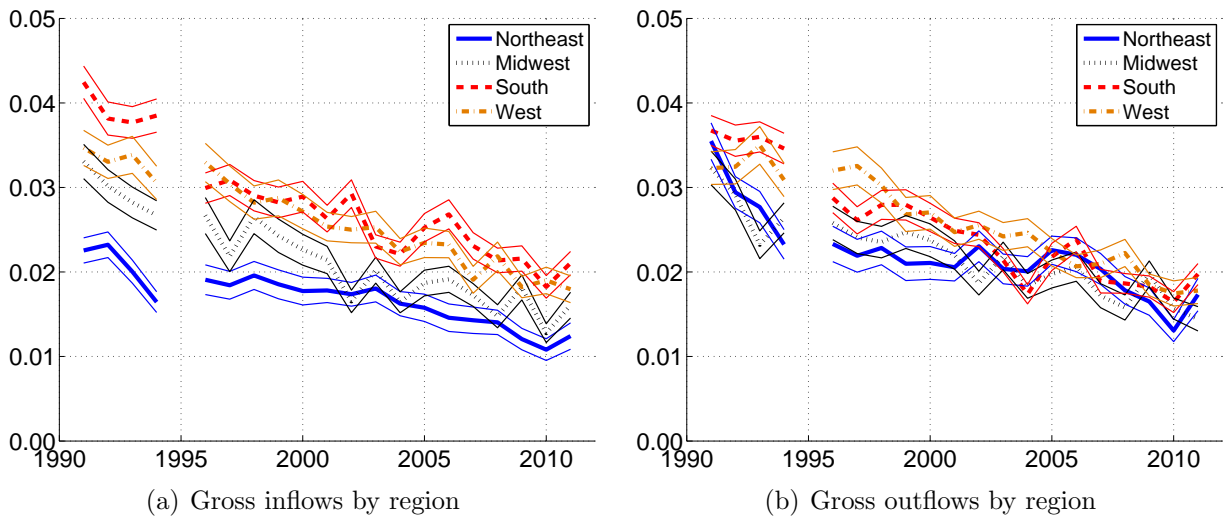


Figure 8: Regional patterns in migration.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Thin lines in figures 8(a) and 8(b) show 1-standard-error confidence bands around point estimates.

rate by region. Figure 8(a) shows each region's gross inflow rate: the fraction of people in the region who lived in a different state (whether in the region or outside it) one year ago. Similarly, figure 8(b) shows each region's gross outflow rate: the fraction of people who lived in the region one year ago and have since moved to a different state (whether in the region or outside it). The graphs show that both immigration and outmigration have fallen substantially in all regions. Thus, the driving force in the decline in migration cannot be a simple change in Americans' desire or ability to move to or from one particular part of the country.

Another possibility is that in some parts of the country, interstate migration is a poor proxy for migration between labor markets. If migration from cities to suburbs has fallen over time, and if interstate migration captures some urban-suburban moves, we could mistakenly conclude that moves between labor markets have fallen when in fact they have not. We conjecture that this problem is likely to be most severe in the New York metropolitan area, which extends to large parts of New Jersey and Connecticut, and the Washington, D.C., metropolitan area, which extends into Maryland and Virginia. However, figure 8(c) shows that the decline in migration is actually larger when we exclude all respondents who live in New York, New Jersey, and Connecticut in the survey year, and is virtually identical when we exclude respondents in Maryland, Virginia, and Washington, D.C.¹¹

E. Job search

Improvements in information technology and reductions in travel costs potentially change workers' information about both job opportunities and local amenities. One possibility is that increased information about faraway jobs reduces the number of workers who move to a distant location simply in order to search for a job there. However, the data on workers' reasons for moving suggest that this mechanism is not at work. Figure 9 disaggregates the reasons for moving that we showed earlier in figure 2(b). The decline in job-related moves comes mainly from a decline in people who move for a new job or job transfer; there has been little change in the fraction of people who move because they lost a job or to look for work.

¹¹We exclude the entirety of the states because the boundaries of the metropolitan areas have changed over time.

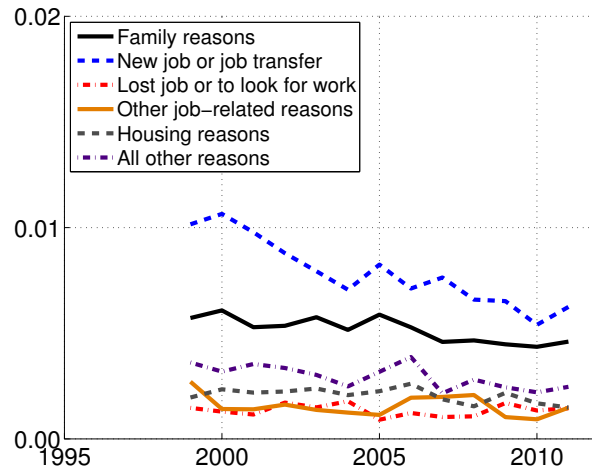


Figure 9: Detailed reasons for moving.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Estimates shown for all years when variables are available.

Although self-reported reasons for choices must be interpreted with caution, these findings are the opposite of what we would expect if better information had made it easier for workers to search remotely without moving first: In that case, we would see an increase in moves for new jobs and a decrease in the number of people who move to look for work. Nonetheless, because moves for a new job or job transfer are much more common than moves to look for work, the ability to search for jobs in remote locations appears to be an important component of migration decisions.

4. Direct Evidence for Our Mechanisms

Our theory relies on two mechanisms to generate a decrease in migration: an increase in the similarity of job opportunities in different parts of the country and a decrease in the cost of learning about amenities in faraway locations. This section presents direct evidence for each mechanism.

A. Increases in the similarity of job opportunities

We test whether job opportunities around the country have become more similar and whether geographic differences in job opportunities are related to migration by examining both prices and quantities. First, we show that the dispersion of incomes across states and metropolitan areas within occupations has fallen. In other words, the earnings of workers

in a given occupation have become more similar across space. This convergence in the *price* of workers in various occupations might result from a change in either demand (e.g., an increase in the productivity of certain occupations in places where those occupations used to be unproductive) or supply (e.g., workers moving from places with low productivity to places with high productivity until marginal productivity is equated across space). A change in demand would reduce migration in the model we present later in the paper; a change in supply is merely a consequence of migration and would not itself cause migration to fall. To distinguish between demand and supply effects, we examine the distribution of the number of workers in each occupation around the country. If productivity in particular occupations becomes less geographically specific, occupations will become less geographically segregated — that is, the distribution of occupations in each state will become more similar to the national average. By contrast, if workers move to places where their occupations are more productive, each location will become more specialized and occupations will become more geographically segregated. We find that occupations and industries have become less geographically segregated across states and metropolitan areas, supporting the view that occupations’ productivity levels have become less geographically specific. The change in productivity dispersion, as measured by the change in income dispersion, will be a key input to our model below. We then connect productivity dispersion to migration by showing that, on average, a migrant’s occupation brings higher pay in the destination state than in the origin state. Thus, migrants tend to move toward states where their occupations earn higher pay, a key mechanism in our model.

The dispersion of incomes within occupations

We study the geographic specificity of occupations’ income levels by estimating a statistical model in which incomes depend on a state-occupation interaction and characterizing the variance of the state-occupation interaction. Our model is

$$\ln y_{iost} = a_{st} + b_{ot} + \mathbf{x}'_{iost} \boldsymbol{\beta}_t + \xi_{ost} + \epsilon_{iost}, \quad (1)$$

where y_{iost} is the wage, salary, and self-employment income of worker i in occupation o , state s , and year t ; a_{st} is a state-year fixed effect; b_{ot} is an occupation-year fixed effect; \mathbf{x}_{iost} is a

vector of controls, including sex, dummy variables for single year of education, and a quartic polynomial in potential experience; and ξ_{ost} is the state-occupation interaction of interest. We model ξ_{ost} as a normally distributed random effect with mean 0 and variance $\sigma_{\xi,t}^2$ and estimate the variance $\sigma_{\xi,t}^2$ for each year.¹² When $\sigma_{\xi,t}^2$ is smaller, the variance of incomes across states within an occupation is smaller, after controlling for individual demographics \mathbf{x}_{iost} and factors b_{st} that affect all occupations in a state. Thus, if occupations' productivity becomes less geographically specific, $\sigma_{\xi,t}^2$ will fall.

We use single-digit occupations as listed in appendix A1 (except for military, unemployed, and not in the labor force) to keep the number of parameters manageable.¹³ Although detailed occupation coding in Census Bureau datasets has changed over time, these changes should have had little impact on how workers are classified at the one-digit level, so we think it is unlikely that our results are driven by changes in occupation coding.

We estimate the model year by year in data from the CPS, the decennial census, and the ACS.¹⁴ Figure 10(a) shows moments of the posterior distribution of $\sigma_{\xi,t}^2$ for each year. A clear decline from 1970 to 2000 can be seen in the decennial census data. The downward trend also appears in the CPS, although the CPS estimates are volatile from year to year. In the ACS, the posterior mean has a slight upward trend, but the posterior density is too dispersed to demonstrate whether the true trend is upward or downward.

¹²Treating ξ_{ost} as a fixed effect and then calculating the variance of the estimated fixed effects would produce upward-biased estimates of $\sigma_{\xi,t}^2$ because some of the variance in the estimated fixed effects would come from sampling error. This bias would depend on the sample size, and the size of the sample changes over time, so estimates based on fixed effects would not be comparable over time and could not be used to determine the trend in $\sigma_{\xi,t}^2$.

¹³Nonetheless, the large samples and large number of parameters make maximum likelihood estimators converge very slowly, so we estimate the posterior distribution of the parameters by Markov chain Monte Carlo using algorithm 16 of Chib (2001). For these estimates only, we do not use survey weights because we cannot derive the weighted likelihood without detailed information on sampling and weighting procedures (see, e.g., Gelman, 2007). The demographic controls \mathbf{x}_{iost} should absorb most of the subpopulation heterogeneity that survey weights are intended to account for. We use standard uninformative priors for the fixed effects a_{st} and b_{ot} and the coefficients β_t and standard weakly informative priors for the variances of ξ_{ost} and ϵ_{iost} . The prior has a larger effect on the posterior when the sample is smaller. We choose a prior for $\sigma_{\xi,t}^2$ centered on 0.001, which is smaller than any of the posterior means. Thus, the prior could produce a downward trend in the posterior estimates of $\sigma_{\xi,t}^2$ only if the samples were getting smaller over time; in reality, the sample sizes increase over time, which means that our results are, if anything, an underestimate of the decrease in income dispersion.

¹⁴We use the 1 percent form 1 and form 2 state samples from the 1970 census, the 5 percent samples from the 1980 and 1990 censuses, the 5 percent and 1 percent samples from the 2000 census, and annual samples from the ACS. We obtain all census and ACS data from Ruggles et al. (2010).

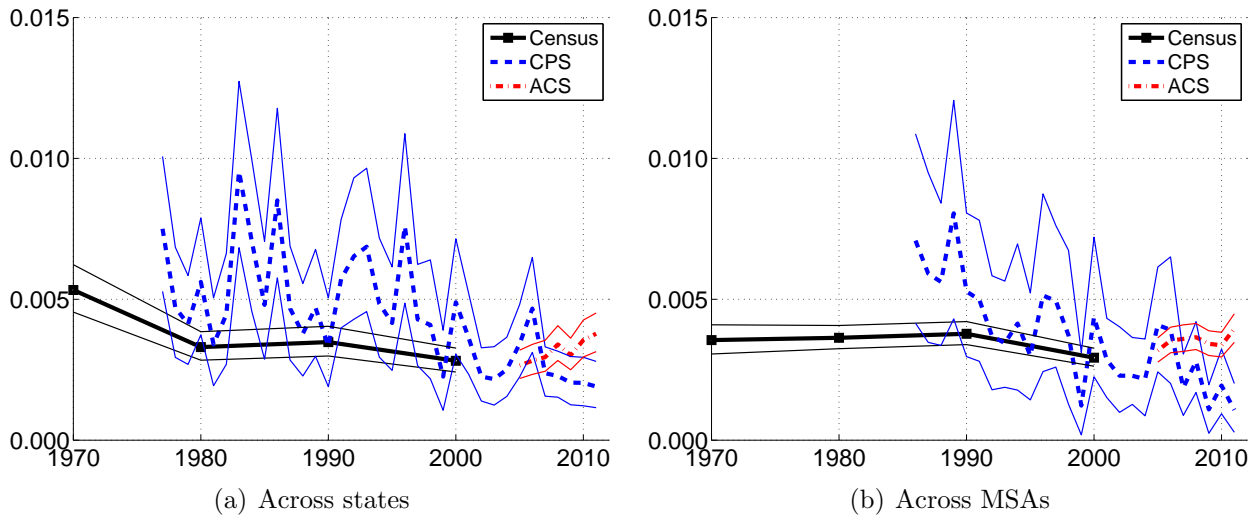


Figure 10: Variance of location-occupation interaction in income.

Source: Authors' calculations from Current Population Survey (CPS) micro data, 1977–2011; American Community Survey (ACS) micro data, 2005–2011; and decennial census micro data (1970, 1980, 1990, and 2000). Thick lines are posterior means; thin lines show 2.5th and 97.5th percentiles of posterior distribution. Sample restricted to employed civilians ages 16 and over.

Figure 10(b) shows that we also find a decline in $\sigma_{\xi,t}^2$ when we define geographic locations by metropolitan statistical areas (MSAs) rather than states, but the decline is now steeper in the CPS and less steep in the decennial census. Using MSAs instead of states poses challenges because some people do not live in MSAs, not all MSAs are identified in public-use datasets, and MSA boundaries change over time. Also, MSAs can be identified in the CPS only starting in 1986. In the figure, we limit the sample for each dataset (Census, CPS, or ACS) to respondents who live in MSAs that are identified in every year for that dataset. We do not adjust for changes in the counties that each MSA includes because the public-use data do not always show the respondent's county and because expansions of MSAs often reflect expansions of the geographic area that forms a single labor market.¹⁵

We interpret the downward trend in geographic dispersion of incomes as reflecting convergence across states in productivity within occupations. In this view, a worker with given skills faces more similar wages across states now than he did several decades ago. However, another possible interpretation is that unobserved differences across states in worker

¹⁵The MSA analysis use the 1 percent form 1 and form 2 metro samples instead of the state samples from the 1970 census and drop the 1 percent sample from the 2000 census because it identifies relatively few MSAs.

quality may have gotten smaller, which would mean that a worker with given skills does not necessarily face more similar wages across states now than in the past. Econometrically distinguishing differences across states in unobservable worker quality from differences across states in skill returns is a complicated task, and we do not attempt it here.¹⁶ But we believe our implicit assumption that workers with identical education, living in the same state but educated in different states, face identical wages is a reasonable one. This assumption is common in the literature (see, e.g., Dahl, 2002), and even papers that study the effect of local school quality on wages, such as Card and Krueger (1992), acknowledge large differences across space in the returns to education for workers with identical quality.

The dispersion of quantities of workers

In figure 11, we compute the Theil information-theory index of segregation (Theil and Finizza, 1971) for various categories of occupations and industries. The Theil index is commonly used to measure racial segregation, and Reardon and Firebaugh (2002) show that it has many desirable properties that other indices lack, especially when measuring segregation of more than two groups. The Theil index compares the distribution of occupations or industries in each location to the distribution for the nation as a whole. The index ranges from 0 to 1, with higher values indicating more segregation; the index is 1 when each occupation or industry is found in only one location, and 0 when each location has the same distribution of occupations or industries as the nation as a whole.¹⁷ As with our analysis of income

¹⁶For example, re-estimating (1) with state-specific returns to education would not allow us to rule out differences in worker quality because state-specific returns to education could reflect either differences in the wages for identical workers or differences in the quality of education.

¹⁷The formula for the index is

$$H = \frac{1}{E} \sum_{s=1}^S \frac{N_s}{N} (E - E_s),$$

where s indexes states, N_s is the number of workers in state s , N is the national number of workers, and E and E_s are the national and state entropy indices. The entropy indices are defined by

$$E = - \sum_{j=1}^J \pi_j \ln \pi_j, \quad E_s = - \sum_{j=1}^J \pi_{js} \ln \pi_{js},$$

where j indexes groups (occupations or industries), π_j is the fraction of U.S. workers who are in group j , and π_{js} is the fraction of state s 's workers who are in group j . We compute standard errors for the index by the delta method, treating the number of workers in each state as nonstochastic and assuming that when zero workers are observed in a state-group cell, that cell contains zero workers in the population.

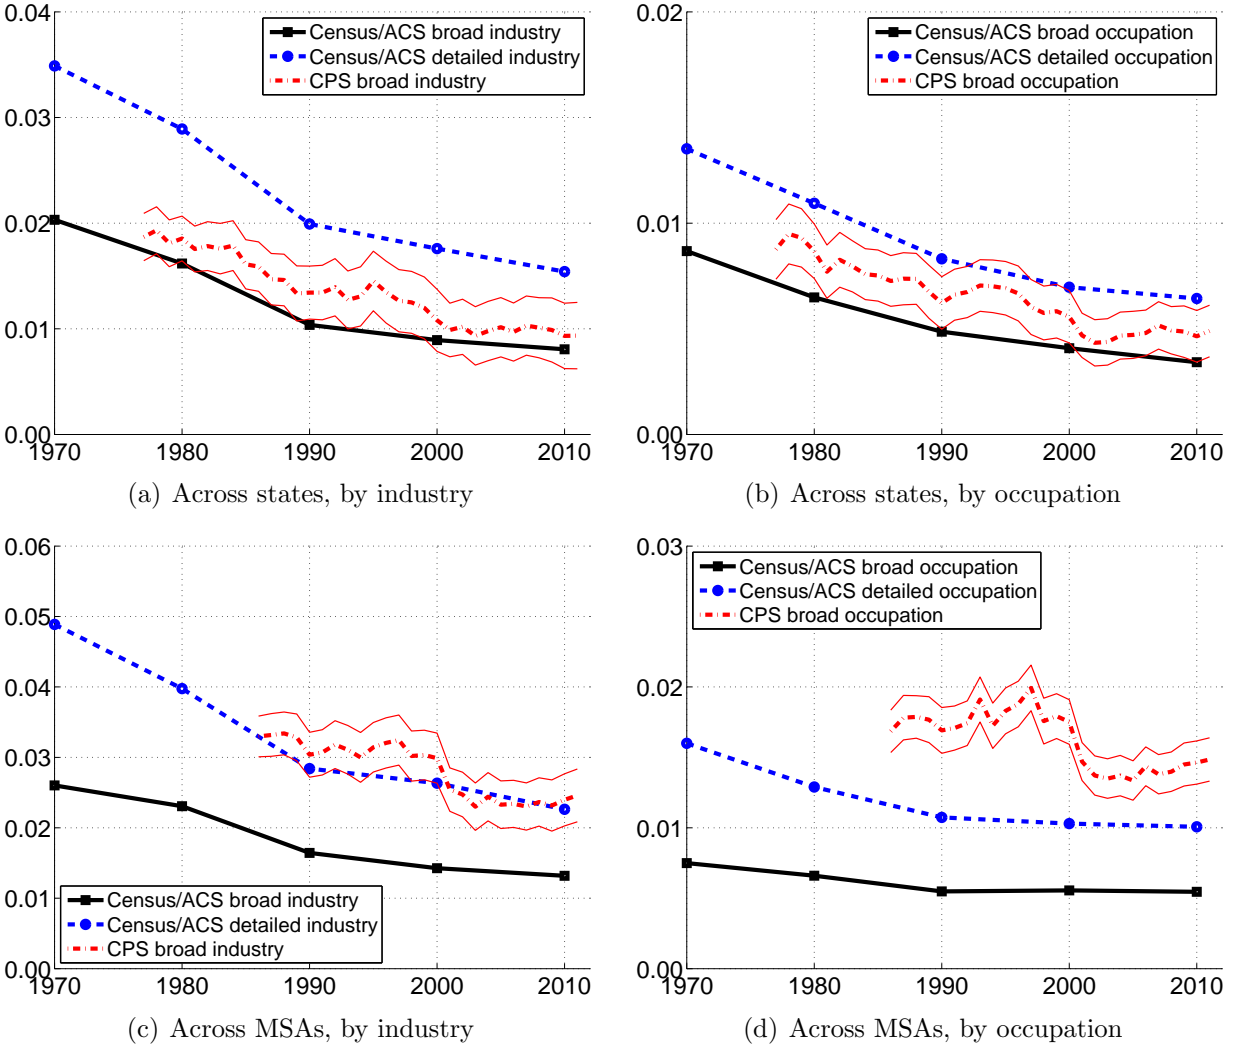


Figure 11: Theil indices of segregation of industries and occupations.

Source: Authors' calculations from Current Population Survey (CPS) micro data, 1981–2011; American Community Survey (ACS) micro data, combined 2006–2010 sample; and decennial census micro data (1970, 1980, 1990, and 2000). Detailed industry and occupation categories are three-digit codes for the IND1990 and OCC1990 variables in Ruggles et al. (2010). Broad industry and occupation categories are one-digit codes listed in appendix A1. Sample restricted to employed civilians ages 16 and over. Weighted by number of workers. Figures 11(c) and 11(d) include only MSAs that can be identified in every year for a given dataset. Thin lines show 1-standard-error confidence bands around CPS point estimates. Standard errors of census and ACS estimates are too small to be visible.

dispersion, we focus on states as locations but check that our results are robust to using MSAs instead.

Figure 11 shows that, over time, states’ and metropolitan areas’ distributions of workers across industries and occupations have become more similar. This pattern holds whether we look at single-digit industries and occupations or more detailed categories. (To examine more detailed categories, we must go to decennial census and multiyear ACS data,¹⁸ because the CPS contains too few observations to reliably estimate, say, the number of workers in hardware stores in Vermont in 2010.) In results not shown here, we found that the decline is not solely due to the shift from manufacturing to services; an index of segregation of detailed industries within manufacturing also falls over the past four decades.

Migration and the geographic specificity of occupations

We examine whether migrants move to states where their occupations bring higher pay by testing whether the state-occupation interaction ξ_{ost} in equation (1) is larger in the migrant’s destination state than in the migrant’s origin state. For each year, we regress log income on state fixed effects, occupation fixed effects, and the demographic controls, then estimate ξ_{ost} by the average of the residuals from this regression in occupation o , state s , and year t . We use the estimates of ξ_{ost} to calculate the difference in earnings between migrants’ origin and destination states, conditional on the migrants’ occupations. Specifically, for a migrant i who moved from state s to state s' and is currently working in occupation o , we define

$$\Delta_{it} = \hat{\xi}_{o,s',t} - \hat{\xi}_{ost}, \quad (2)$$

where $\hat{\xi}_{ost}$ is our estimate of the state-occupation interaction ξ_{ost} . The quantity Δ_{it} is the difference between i ’s predicted income from equation (1) at the destination state and i ’s predicted income at the origin state, holding constant i ’s occupation and demographics and controlling for differences in incomes that affect all occupations in a given state. That is,

¹⁸We use the same census samples as for the income dispersion analysis. For the ACS, we use the 2006–2010 combined dataset, which is equivalent to a 5 percent sample. Despite the large size of these samples, some occupations and industries are not observed in all years. We combine industries or occupations that are not observed in all years into an “all other” category so that the groups over which the index is calculated are constant over time.

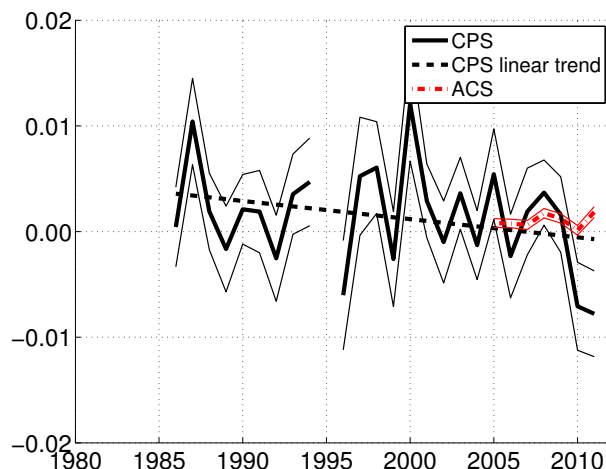


Figure 12: Mean within-occupation income gains from moving.

Source: Authors' calculations from Current Population Survey (CPS) and American Community Survey (ACS) micro data. Samples restricted to interstate migrants who are working-age adults in civilian households and report an occupation. CPS sample restricted to those with non-imputed migration data. ACS sample restricted to those not living in group quarters. Estimates shown for all years when variables are available. Thin lines show 1-standard-error confidence bands around point estimates.

Δ_{it} represents i 's within-occupation income gains from moving, net of any difference in the average income across all occupations between the origin and destination states. If migrants move toward states where their occupations are better paid, we expect Δ_{it} to have a positive mean.

Figure 12 shows the mean of Δ_{it} in each year in the CPS and ACS. The CPS estimates begin in 1986 because that is the earliest year when we can identify working-age adults. We define occupations by one-digit codes, we estimate ξ_{ost} using data on all workers, not just migrants, and we estimate the mean of Δ_{it} using data on all interstate migrants, even if they are not currently employed. The results show that, on average, Δ_{it} is positive: Migrants move toward states where their occupations are higher paid. In the CPS, the estimates are volatile and imprecisely estimated but appear to trend downward, consistent with the hypothesis that the gaps in occupations pay between states have gotten smaller. From 1991 to 1997, the mean of Δ_{it} in the CPS is 0.0013, which means that the average migrant's occupation pays 0.13 percent more in the destination state than in the origin state. By the 2005–'11 period, that average falls to zero. The ACS estimates are more precise owing to the larger sample size; the average gap in the ACS from 2005 to 2011 is 0.10 percent, and we can reject at the

5 percent significance level the null hypothesis that the average in the ACS is zero.

B. Decreases in the cost of information

If people have better information about distant locations, they will be less likely to move somewhere only to find it unsatisfactory and move again soon afterward. Thus, a decrease in the cost of information about faraway places should reduce the rate of repeat and return migration. We cannot measure repeat and return migration in the cross-sectional CPS, so we turn to panel data from the U.S. Census Bureau’s Survey of Income and Program Participation (SIPP) as well as data on place of birth and residence five years ago from the decennial census. Each of these data sources has drawbacks, but given the need for repeated measures of migration, the SIPP and the decennial census are the best available datasets for our purpose.

As we discuss in section 2, panel survey data are not ideal for measuring migration because results can depend on the survey’s procedure for finding migrants at their new locations, and any changes in measured migration might result from changes in survey procedures. However, the SIPP makes significant efforts to locate respondents who move (see U.S. Census Bureau, 2009c, chap. 2) and starts with a large sample, about 50,000 households in recent years, so that the population at risk of repeat migration is large enough to obtain reasonably precise estimates. (Overall interstate migration rates in the SIPP are similar in magnitude to those in the CPS, although the downward trend is less pronounced.) The SIPP consists of a series of independent panels that started in various years and were followed for several years each. Respondents are interviewed every four months. For each panel, we calculate repeat and return annual interstate migration rates in the first two years that the panel was followed. Specifically, for a panel first interviewed in year t , the repeat migration rate is the probability of living in a different state at the seventh interview (year $t + 2$, 24 months after the first interview) than at the fourth interview (year $t + 1$, 12 months after the first interview), conditional on making an interstate move between the first and fourth interviews (years t and $t + 1$). The return migration rate is the probability of returning to the year- t state at $t + 2$, conditional on making an interstate move between t and $t + 1$. These definitions ignore moves that happen within a single year — even though the SIPP measures such moves

— so that we are measuring annual rates that can be compared with the annual rates in the CPS. In some panels, the SIPP public-use data files combine certain small states to protect respondents’ anonymity. We use these combinations of states in all panels to ensure that our results are not driven by changes in the state coding.

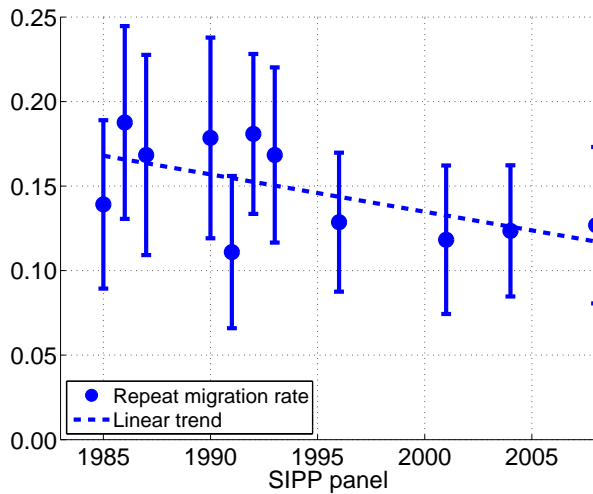
Because the census is collected cross-sectionally, it is not affected by panel attrition. Long-form census questionnaires ask respondents where they were born and where they lived five years ago. We calculate the repeat migration rate as the probability of living in a different state at the time of the census than five years earlier, conditional on making an interstate move between birth and five years ago. The return migration rate is the probability of currently living in the birth state, conditional on not living in the birth state five years ago. These rates are not directly comparable to the one-year migration rates among workers that are the focus of our paper, because they include some moves by children and some moves in the distant past, and ignore some moves at frequencies higher than five years.¹⁹ Nonetheless, these rates are useful indicators because decreases in high-frequency return migration by workers should lead, all else equal, to decreases in the return migration rate that we measure in the census.

To make the SIPP and census data as comparable as possible to our results from the CPS, we limit the SIPP and census samples to people who are working-age adults in civilian households. Additional details on sample selection and the calculation of confidence intervals are in appendix A2.

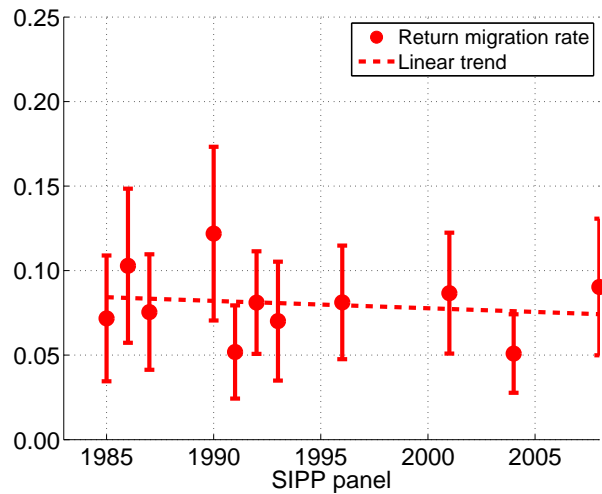
Figure 13 shows the results. Repeat and return migration rates are high: The SIPP data show that someone who leaves a state in one year has roughly a 7 percent chance of returning the next year and a similar chance of moving to a third state the next year. But these rates have fallen over time. In particular, the annual repeat migration rate in the SIPP appears to have fallen by about 5 percentage points in the past two decades. The census data also show high but declining repeat and return migration.

Several important caveats apply to these findings. First, the estimates from the SIPP are very imprecise. The reason is that respondents are part of the sample used to estimate

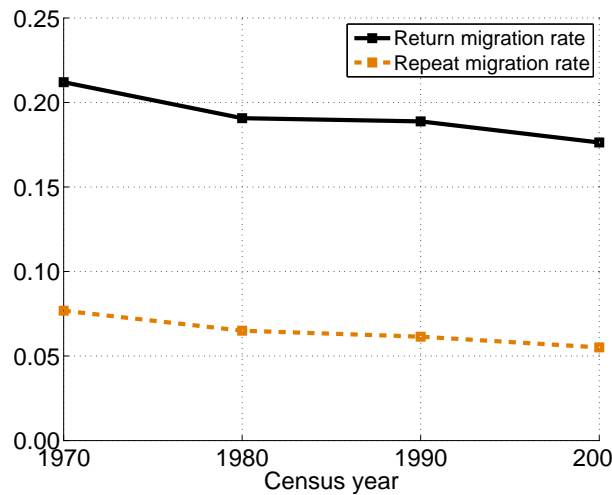
¹⁹For example, a person who is born in Minnesota, moves to Wisconsin at age 1, returns to Minnesota at age 29, and is a census respondent at age 30 will count as a return migrant; a person who lives in Minnesota up to age 27, moves to Wisconsin at age 28, returns to Minnesota at age 29, and is a census respondent at age 30 will not count as a return migrant.



(a) Repeat migration in the SIPP



(b) Return migration in the SIPP



(c) Repeat and return migration in the census

Figure 13: Repeat and return interstate migration rates.

Source: Authors' calculations from Survey of Income and Program Participation (SIPP) and decennial census micro data. Working-age adults in civilian households; additional sample selection details in appendix A2. Dashed lines in figures 13(a) and 13(c) are predicted values from a linear regression of the repeat or return migration rate on the initial year of the SIPP panel. Vertical bars in 13(a) and 13(c) show 95 percent confidence intervals for point estimates. Confidence intervals for census estimates are too small to be visible.

repeat or return migration only if they migrate in the initial year; because interstate migration is rare to begin with, this sample is small even though the overall SIPP sample is large. The downward trends in the SIPP repeat and return migration rates are not statistically significant at conventional significance levels. Second, the changes we observe in the SIPP could theoretically be due to changes in procedures for following respondents who move, although we are not aware of any such changes. Third, in the census, the timing of the changes in moving behavior is unclear because the estimated rates are a function of migration over the entire life cycle, and the observed changes could result from changes in moving rates among families with children rather than from changes in moving rates among workers.

In addition, a decline in repeat and return migration is not the strongest possible evidence for our hypothesis about migration, because other changes that reduce migration in general could mechanically reduce repeat and return migration as well. A stronger prediction of our theory is that the fraction of moves in a year that are repeat or return moves should have declined. Unfortunately, we cannot estimate this fraction with any precision because the estimated migration rates in the SIPP are noisy and the census does not give us the annual migration rates we would need.

Still, there are reasons beyond the decline in repeat and return migration to believe that people have more information about distant locations than in the past. The past several decades have seen dramatic changes in several technologies and markets that help people to gather this information. Most obvious, of course, is the development of the Internet, which allows people to inexpensively and rapidly learn about life in other cities. But other changes have also sharply reduced the cost of information. Following the breakup of AT&T in 1984, competition in the market for long-distance telephone calls rose, prices fell by 50 percent in seven years, and demand for long-distance services doubled (Taylor and Taylor, 1993). Thus, even before the Internet, the cost of learning about distant places by picking up the telephone was decreasing.

Travel costs are also an important influence on the cost of gathering information. A person who wants to learn whether she will like the weather in California can best do so by going to California on vacation. After the United States deregulated the airline industry in 1978, airfares fell significantly (though the exact size of the decrease is difficult to calculate)

and airlines offered more flights to more destinations (Borenstein and Rose, 2008). A decrease in the cost of air travel reduces the cost of gathering information both by reducing actual outlays on travel and, for workers who substitute to air travel from other modes of transportation, by reducing the time required to reach the destination.

C. Summary

The data show that numerous hypotheses about the decline in migration do not match the facts. Migration has not fallen because of any of the following:

- Changes in the distribution of age, education, marital status, number of labor force participants per household, or real household incomes.
- Changes affecting only people of particular ages, people at particular education levels, people in particular kinds of households, people at particular income levels, or people in particular occupations or industries.
- Changes affecting only one region of the country.
- Changes affecting the relative desirability of some regions compared with other regions in the same way for all people.

In contrast, the data do support two other mechanisms: a decrease in the geographic specificity of occupations and a decrease in the cost of information.

5. A Model of Information, Specialization, and Gross Migration

Guided by our empirical findings, we construct a model in which broad-based changes in information technology and the structure of labor markets impel all workers to migrate less. Our model contains five features that make it suitable for our purposes.

First, our model features only two symmetric locations, which can be thought of as “here” and “there.” By formulating a model with only two locations, each of which contains half of the population, rather than multiple locations with different populations, we limit our ability to make inferences about net population movements to or from particular locations. However, since it is gross migration rates and not net migration rates that have changed, this modeling choice does not come at any cost. Instead, it imparts some important benefits. Unlike existing models (e.g., Davis, Fisher, and Veracierto, 2010; Kennan and Walker, 2011)

that have multiple locations, our model is simple enough to allow for the inclusion of richer environmental features that are at the heart of theories of the decline in migration.

Second, agents in our model can be employed in one of two distinct occupations or be nonemployed. Each occupation commands a higher wage in one of the two locations. One could think, for example, of banking in New York and acting in Los Angeles. Individuals in our model have occupation-specific skills that evolve stochastically, so that there is heterogeneity across households in their comparative advantage at working in an occupation, and thus their labor market incentives for living in each location.

Third, locations in our model are an experience good. This means that individuals have imperfect information about the non-labor market (amenity) values that they derive from living in each location. Only by living in a location do individuals learn about their preferences for living there.

Fourth, the labor market in our model is frictional, in the sense that individuals must search for employment opportunities. Moreover, living in one location does not preclude an individual from searching for a job in the other location. The possibility of remote search is important because it allows us to capture the notion that even if the fundamental reason for a move is a change in amenity-related preferences, the move may not take place until the individual finds a job opportunity in the desired location.

Finally, our model has a life cycle element, since we showed in section 3A that the likelihood of migration varies greatly with age.

A. Environment

Demographics and preferences Individuals (which we will also refer to as households or agents) live for T periods, $t = 1, \dots, T$. In each period, they live in one of two locations, $j \in \{a, b\}$, and either work in one of two occupations, $k \in \{A, B\}$, or are nonemployed, $k = u$. They choose locations, occupations, and job search strategies to maximize expected discounted utility:

$$\mathbb{E} \sum_{t=1}^T \beta^{t-1} (y_t + u_t),$$

where y_t is income and u_t is utility derived from non-labor market features of the location where the individual lives at age t . We let $n_t^j \in [0, t]$ denote the number of periods that the

individual has lived in location j , up to and including period t . Note that $n_t^a = t - n_t^b$, since in any period an individual who is not in location a must be in location b , and vice versa.

Information and amenities Agents' preferences for local amenities, $v = (v^a, v^b)'$, are fixed over time. However, individuals do not know these preferences and must learn them over time, through living in the two locations. Each period, an individual who lives in location j_t receives non-labor market utility $u_t = v^{j_t} + \epsilon_t$, that is, the sum of his underlying unknown preference for the location and an i.i.d. random preference shock. The individual observes only u_t and must use this information to update his belief about v^{j_t} . We denote the initial prior mean and precision of beliefs by m_0^j and τ_0 . We assume that the ϵ shocks and the v values are normally distributed with precisions τ_ϵ and τ_v , respectively:

$$\epsilon \sim N\left(0, \frac{1}{\tau_\epsilon^2}\right), \quad v^j \sim N\left(0, \frac{1}{\tau_v^2}\right).$$

We assume that v^a and v^b are independent so that a strong preference for living in either location imparts no information about the *absolute* preference for living in the other location.

Labor markets Labor markets are arranged according to an island structure in each location, in the spirit of Lucas and Prescott (1974). There are two islands, one on which production takes place and one populated by nonemployed households who receive a nonemployment benefit. To find the production island, nonemployed households are required to search. On the production island, there is a competitive labor market for each occupation. Technology is constant returns to scale in skills, and labor is the only input for production. Thus, the wage rate per unit of skill equals the marginal product of skills in each occupation; we take this marginal product as exogenous.

An individual at age t is characterized by his skills in each of the two occupations, $s_t = (s_t^A, s_t^B)'$, which evolve according to an exogenous Markov process normalized so that $E[e^{s_t}] = 1 \forall t$. These skills are revealed to the individual at the beginning of period t . (That is, while our model features learning about amenity preferences, there is perfect information about skills. We abstract from learning about skills because observed occupational mobility is much higher than long-distance geographic mobility — implying that people will learn

faster about their skills than about their geographic preferences — and because explaining occupational mobility within geographic locations is not our goal.) We also assume that skills have a deterministic life cycle component ψ_t to capture the evolution of average wages over the life cycle.

We denote the price of skills for occupation k in location j by p_k^j and assume that

$$p_B^a = p_A^b < p_A^a = p_B^b.$$

This specification encodes two assumptions. First, islands and occupations are symmetric. Second, there is a geography-occupation interaction in the price of skills: An occupation commands a higher price per unit of skill when it is performed in the location where it has a comparative advantage. We normalize $p_B^a = p_A^b = 1$ and define $\theta = p_A^a = p_B^b$. Thus, $\theta > 1$ is the wage premium for working in a matched location and occupation.

Incomes depend on skills s_t , the life cycle component ψ_t , and a time cost of moving. Specifically, income is

$$y_t(j_t, k_t, s_t) = \begin{cases} (1 - \kappa \mathbf{1}_{\text{migrate}}) \psi_t e^{s_{k_t}} p_k^j & \text{employed agent} \\ (1 - \kappa \mathbf{1}_{\text{migrate}}) q & \text{nonemployed agent,} \end{cases}$$

where κ is the time cost of moving, q is the nonemployment benefit, and a worker who migrates between periods t and $t + 1$ loses a fraction κ of his work time during period t .

Both nonemployed and employed workers can choose to search for the production island in either location (but not both), regardless of where they are currently located. Search is costless and equally efficient for employed and nonemployed agents. Searches succeed with probability λ .²⁰ In addition, employed workers randomly lose their jobs and are forced to move to the nonemployment island with probability δ .

Timing An individual enters period t with the following relevant information: the location where he resides (j_t); his current island, that is, production or nonemployment (i_t); his skills in

²⁰Although we could introduce search costs and differences in search efficiency, we do not need to do so to explain the decrease in migration.

the two occupations at the end of the previous period (s_{t-1}); his beliefs about his preferences for living in the two locations, conditional on information at the end of period $t - 1$ (m_{t-1}); and the number of periods he has lived in location a (n_{t-1}^a). Recall that $n_{t-1}^b = t - 1 - n_{t-1}^a$.

Within each period, the timing of events is as follows:

1. The individual's skills in each of the two occupations, s_t , are realized.
2. The individual receives his non-labor market utility, $u_t = v^{j_t} + \epsilon_t$.
3. The individual updates the number of periods he has lived in location a , n_t^a , and his beliefs about his utility from living in location j_t , $m_t^{j_t}$. He does not update his beliefs about utility from living in the other location because he has no new information about that parameter. A formal description of the learning problem and formulas for updating beliefs can be found in appendix A3.
4. If the individual is on the production island ($i_t = 1$), he chooses his occupation for the current period k_t . He may also choose to quit to the nonemployment island.
5. The individual works (if employed) and receives his earnings or nonemployment benefit.
6. After working, an employed worker may randomly lose his job with probability δ .
7. The individual decides whether to search for the production island in either location, and the results of search are realized.
8. Conditional on the outcome of search, the individual makes his migration decision, that is, he chooses his location for $t + 1$. This consists of the choice of a location-island pair (j_{t+1}, i_{t+1}) . Migrants pay a moving cost κ , proportional to income. There is no cost for switching occupations.

Given this timing, we can express the worker's maximization problem recursively. The associated Bellman equations are given in appendix A4.

B. Incentives to migrate

Why might an individual in this model decide to migrate? First, consider a shock to skills in one or both occupations. When $\theta > 1$, so that each location has a comparative advantage in one of the two occupations, a shock to an individual's relative abilities in the two occupations changes his relative earnings potential in the two locations. If θ is large enough that the effect on earnings dominates any difference in the locations' perceived amenity values,

this shock to the worker’s skills will lead him to migrate. Second, consider a low realization of non–labor market utility u_t . Such a realization causes an individual to revise downward his beliefs about his underlying preference for the current location, and hence to revise upward his belief about his relative preference for living in the other location. If this change in beliefs is big enough to overcome any difference in potential earnings across the two locations, then the individual will choose to migrate. The likelihood of such a move depends on both the tightness of individuals’ prior beliefs about their preferences for each location τ_0 and the information content of the signals that they obtain through living in a location, τ_ϵ .

The two reasons for migration just described can be considered the fundamental reasons for moving in the model, since it is the exogenous shocks to either skills or beliefs that change individuals’ relative desire to live in the two locations. However, because of the frictional labor market, the proximate cause of migration may differ from the fundamental cause. Consider an individual who desires to migrate because he has received a series of bad draws for his amenity-related utility in the current location. Knowing that he desires to live in the other location, this individual may search for a job there, yet may move only once he finds a job. The proximate reason for this individual’s migration is the outcome of search — a job offer in the remote location. Hence, if asked in a survey about his reason for migrating, he may well answer that it was to take a new job. However, the fundamental reason for migrating was actually the shock to his beliefs about his non–labor market preferences.

Finally, the model generates one additional type of migration, which we refer to as experimentation. Consider a worker in location a who believes that he prefers the amenities in a (i.e., $m_t^a > m_t^b$) but is quite uncertain about his beliefs regarding location b (i.e., has a small precision τ_t^b). This worker may migrate to location b , even though in expectation the amenities there are worse, simply because the information gained from the move is valuable. However, a move made for reasons of experimentation may lead to return migration if, once the worker learns more about b , he becomes relatively certain that he prefers location a . Increases in initial information will reduce both the initial experimental moves and the subsequent return migration.

6. Quantifying our proposed mechanisms' effect on migration

Above, we showed *qualitatively* that our proposed explanations for the secular decline of gross interstate migration fit the facts, while many other theories do not. This section shows that our explanations succeed *quantitatively*: Given the size of the observed fall in the geographic specificity of returns to skills, and for plausible improvements in information, the decline in migration that our model predicts is consistent with what has been observed in the data.

Our quantitative exercise compares steady states of the model under different parameter values. We start by fitting the model to cross-sectional data from the period 1991–'97.²¹ Our parameterized model for this period closely fits the salient features of the migration and labor market data. We then reduce the geographic concentration of returns to skills and increase the available amount of information, compute the model's new steady state, and compare it with data from 2005–'11.

We emphasize that the purpose of the quantitative analysis is not to attempt a full structural estimation of a life cycle model of migration, occupation choice, and labor market flows, as in, for example, Kennan and Walker (2011). Rather, our goal is to provide some confidence that the two mechanisms we propose as the source of the decline in migration not only are qualitatively consistent with the evidence, but also generate the right quantitative drop in migration rates.

A key element in our quantitative exercise is the size of the change in our two proposed mechanisms that we feed into the model. For the decline in the geographic specificity of skill prices, we can measure the change directly because there is a one-to-one mapping between the wage premium for working in a matched location-occupation in the model, θ_t , and the measured state-occupation interaction in earnings regressions, $\sigma_{\xi,t}^2$.²² We measure θ in the same years as the other data; thus, our comparison of steady states assumes that the economy will quickly converge to a new steady state after a change in θ . Section 6C below analyzes transition dynamics in our model. We find that under reasonable assumptions, convergence

²¹We choose 1991–'97, rather than just 1991, as our initial condition because pooling several years of data gives us sharper estimates of the empirical moments that we want our model to match.

²²If we ran the earnings regressions from section 4A in simulated data from the model and included skills as one of the controls, we would obtain $\sigma_{\xi,t}^2 = (\theta_t - 1)^2/4$. Thus, $\theta_t = 1 + 2\sigma_{\xi,t}$.

is indeed rapid. In addition, our results are not particularly sensitive to this assumption because, if convergence were slow, the solution would be to calibrate the model to lagged changes in θ , such as the change from 1970 to 1990, that are similar in magnitude to the contemporaneous changes.

For the increase in information, there is no direct analogue in the data. The increase in information that our mechanism emphasizes is an increase in the precision of initial beliefs about local amenities, τ_0 , because this precision reflects what individuals know about a place without living there. It is difficult to say exactly how much the technological changes we describe above have increased this precision. However, we think it is unlikely that technologies such as the Internet and low-cost travel could give people more information than what they could learn by actually living in a location. Therefore, it is natural to measure the improvement in information as a fraction of the precision of the annual signals, τ_ϵ . We experiment with improvements in initial information that are equivalent to a range of fractions of τ_ϵ . If we had a precise estimate of the change over time in the return or repeat migration rate, we could use this estimate to discipline the increase in information. However, as figure 13 shows, we cannot estimate this rate precisely even in the largest available suitable dataset. Although we do not directly target the change in return or repeat migration, we do show below that our model generates a decline in repeat migration that is consistent with the empirical evidence.

Our quantitative analysis features an important asymmetry: We change the geographic dispersion of returns to skills but not the geographic dispersion of the value of amenities. We make this choice because a decrease in the geographic dispersion of amenity values would not fit the data. In particular, the model predicts that job-related moves will rise when amenity values become more similar, but in the data, job-related moves fell.

In section 3, we showed that the age profile of migration is very different for college-educated and non-college-educated workers. We therefore perform all of the quantitative analysis separately for these two education groups. We do not have a theory of how well people are matched with their initial locations, so we drop from our simulations the first model period — when many agents move because of initial conditions — and match the age profile starting with migration between the second and third model periods.

Table 1: Calibrated parameters

Parameter	Description	Non-college	College
q	value of nonemployment*	2.10	0.42
δ	separation rate	0.084	0.127
ρ	skills process: autoregressive coefficient	0.931	0.862
σ_0^2	skills process: initial variance	0.304	0.036
σ_η^2	skills process: innovation variance	0.138	0.161
κ	moving cost	0.52	0.24
τ_0	amenities: precision of initial prior beliefs	10.8	8.4
τ_ϵ	amenities: precision of preference shocks	7.9	21.0

*As multiple of average earnings conditional on working in the first model period.

A. Initial steady state: 1991–1997

Parameterization

The model period is annual. We fix the annual discount factor, β , at 0.96 and the arrival probability of a job offer, λ , at 0.5. Given our available data, these parameters are difficult to separately identify from the remaining model parameters and have little impact on our ultimate findings.²³ Also, we set the dispersion in amenity values across locations, $\frac{1}{\tau_v}$, equal to 1. This choice is essentially a normalization. Migration decisions are determined by beliefs about amenity values, not by the true values, so the dispersion of true values influences decisions only by influencing the dispersion of beliefs.

The wage premium for working in a matched location-occupation, θ , is set to 1.15, consistent with the average measured state-occupation interaction in earnings in the CPS from 1991 to 1997. This value implies that an identical worker earns 15 percent higher wages in the matched location-occupation than the unmatched one.

Ten parameters remain: the cost of moving, κ ; the parameters describing the learning process for amenity values (τ_0, τ_ϵ); the value of nonemployment, q ; the separation rate, δ ; and the parameters describing the stochastic process for skills (a quadratic in t for the deterministic component ψ_t , and an AR(1) process with initial variance σ_0^2 , innovation variance σ_η^2 , and

²³In particular, the arrival probability of a job offer is difficult to separately identify from the value of nonemployment and the separation rate without data on flows in and out of employment. Since our focus is on understanding migration, not labor market flows, we fix λ exogenously.

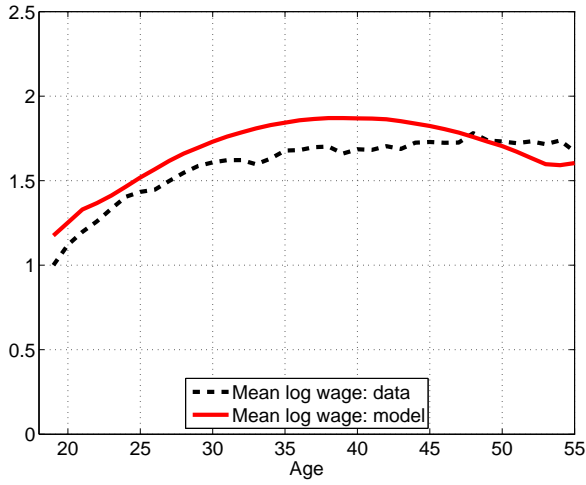
autoregressive parameter ρ for the stochastic component s_t). We choose these parameters to match the age profiles of migration, employment, mean log earnings, and the variance of log earnings, as well as two scalar moments: the autocorrelation of log earnings and the average log earnings difference between migrants and non-migrants.²⁴ Although we cannot formally prove that these moments identify the parameters, below we provide an intuitive argument for why the parameters are identified, alongside our discussion of the model fit. Table 1 shows the calibrated parameter values.

Model fit

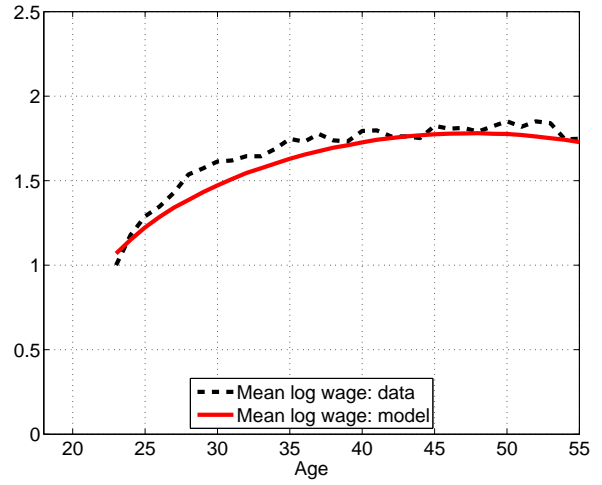
Figure 14 shows how the model fits the age profiles of labor market moments. The graphs in the left column of figure 14 refer to the non-college sample, while those in the right column refer to the college sample. Figures 14(a) and 14(b) show the fit for mean log earnings conditional on working. Conditional on migration and occupation decisions, which affect mean earnings through the geographic specificity parameter, θ , the age profile of mean log earnings pins down the quadratic age profile for ψ_t . Earnings grow over the life cycle both because of this deterministic age component and because workers move toward locations that match their occupations as they age. (The fraction of workers in a matched location-occupation rises from 62 percent to a maximum of 70 percent for the non-college group and from 60 percent to 68 percent for the college group.)

Figures 14(c) and 14(d) show the model fit for the variance of log earnings. Together with the autocorrelation of log earnings, this profile pins down the parameters that govern the stochastic process for skills $(\rho, \sigma_0^2, \sigma_\epsilon^2)$. The variance among the young determines the initial variance σ_0^2 . Conditional on ρ , the variance among the old determines the innovation variance σ_ϵ^2 , because with $\rho < 1$ the variance among the old will be that of the stationary distribution

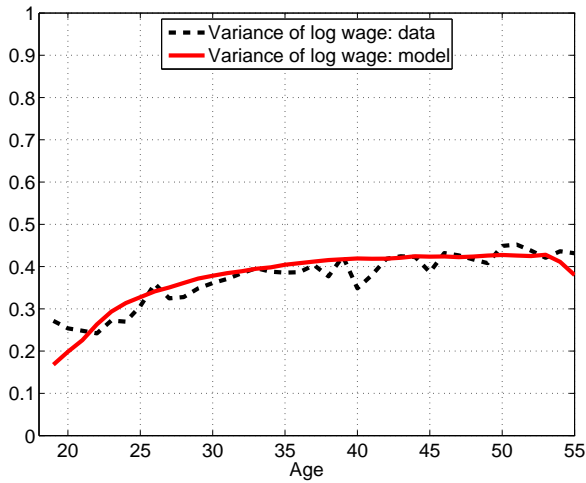
²⁴The earnings variable for the age profiles and the migrant–non-migrant earnings difference is usual weekly earnings. The CPS measures this variable at the time of the survey — that is, after any migration — in contrast to other income variables that are for the previous year and may include income before or after migration. We calculate the migrant–non-migrant earnings difference as the coefficient on a migration indicator in a regression of log earnings on the migration indicator and age indicators. For the autocorrelation of log earnings, we cannot use cross-sectional CPS data, so we turn to the Panel Study of Income Dynamics and calculate the autocorrelation of residuals from a regression of log labor income on experience, experience squared, and year indicators. (We make the PSID samples comparable to our CPS samples by including only people ages 23 to 55. We use PSID data for 1968 to 1997.) When we ask the model to match the earnings difference and autocorrelation, we run the same regressions in the model as in the data.



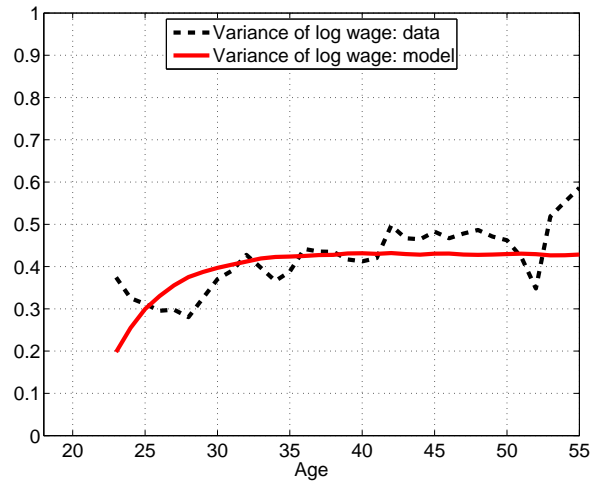
(a) Mean log wage, non-college sample



(b) Mean log wage, college sample



(c) Variance of log wage, non-college sample



(d) Variance of log wage, college sample



(e) Employment rate, non-college sample



(f) Employment rate, college sample

Figure 14: Model fit, labor market moments.

Table 2: Scalar targeted moments

Moment	Non-college		College	
	Data	Model	Data	Model
Autocorrelation of log earnings	0.767	0.822	0.780	0.811
Migrant vs. non-migrant log earnings difference	-0.094	-0.070	-0.001	-0.022

of the stochastic process. Finally, the curvature of the age profile and the autocorrelation of earnings shown in table 2 determine the autoregressive parameter, ρ . Of course, these parameters are determined jointly with all the other parameters of the model since, for example, more migration and high levels of θ also serve to increase the variance of log earnings.

Figures 14(e) and 14(f) show the fit for the employment rate. The average level of employment over the life cycle pins down the separation rate, δ , because this parameter determines the outflow from employment and hence the steady-state employment rate. The slope and curvature of this age profile pin down the value of nonemployment, q , given migration choices as well as the stochastic process for skills and the mean wage profile ψ_t . This is because the employment rate at a given age depends on where q falls in the distribution of potential earnings, which in turn depends on age through ψ_t .²⁵

Figure 15 shows the age profile of migration in the model. The model closely captures the overall level of migration as well as the way that migration varies with age for the two groups. Our calibrated model yields an average migration rate for the non-college sample of 2.56 percent, compared with 2.69 percent in the 1991–’97 data; and an average migration rate for the college sample of 4.40 percent, compared with 4.16 percent in the 1991–’97 data.²⁶ The age profile of migration can be thought of as encompassing three moments: the level of migration for the old, the level of migration for the young, and the curvature of migration between young and old ages. The stochastic process for skills and the location-occupation match premium θ together pin down the level of migration for the old, because old agents

²⁵Our model cannot match the decrease in employment with age for the college sample because the model does not include any of the mechanisms that could produce such a decline, such as marriage and child-rearing, differential taxation of secondary earners, and disability. We think it is unlikely that including these mechanisms would significantly change the model’s implications for migration.

²⁶We compute these averages weighting all ages equally.

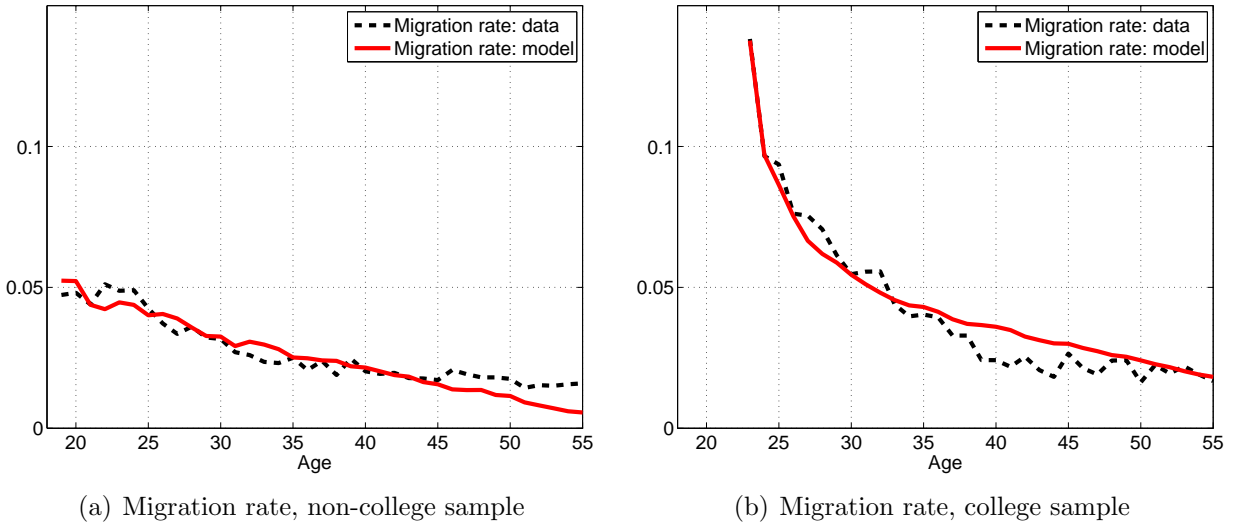


Figure 15: Model fit, migration.

migrate primarily for labor-market-related reasons. (Old agents have largely completed the learning process and have little uncertainty about amenities. In addition, they have higher wages and thus respond more to the location-occupation premium.) Meanwhile, migration rates among the young depend primarily on τ_0 , which describes the tightness of individuals' initial priors about the amenity value of each location. A tighter prior leads to lower migration for experimentation reasons in the first half of the working life. Finally, the speed with which individuals learn about their preferences, and thus the rate at which migration for experimentation reasons slows with age — the curvature of the age profile — is determined by the precision of the signals, τ_ϵ .

It may seem natural to conjecture that the age profile of migration should also pin down the cost of moving, κ . But while κ affects migration rates, other parameters can adjust to offset changes in κ . Instead, the moving cost in our model is primarily pinned down by the mean difference in log earnings between migrants and non-migrants, which we report in table 2. Two components determine the size and sign of this difference. First, there is the average increase in earnings that an individual receives by moving, compared with the counterfactual earnings that individual would have if he did not move. Such wage gains arise primarily when individuals move to a matched location-occupation; thus, this component is determined by θ . In the model, this component is strongly positive and yields an earnings

difference that is larger than in the data. Second, and offsetting these wage gains, the model generates negative selection into migration because the cost of moving is a time penalty; higher-wage workers pay higher moving costs when measured in units of income. Because the cost of migration rises with income, lower-income workers are more likely to move. This negative selection reduces the cross-sectional average earnings of migrants and increases the average earnings of non-migrants. The larger is κ , the more negative selection occurs. Hence, the moving cost is pinned down by the difference between the wage gains from moving that θ alone would imply and the observed earnings gap between migrants and non-migrants.

In the calibrated model, the annual occupational switching rate conditional on being employed is approximately 6.3 percent for the non-college sample and 13.3 percent for the college sample. Although Kambourov and Manovskii (2008) empirically document a similar overall level for annual mobility between one-digit occupations, they report that workers with more education have lower occupational mobility, contrary to the results in our simulations. A likely explanation for the discrepancy is that our model does not feature occupation-specific human capital, which Kambourov and Manovskii (2008) argue reduces migration among highly educated workers. We think it is unlikely that introducing a complication such as occupation-specific human capital would have a meaningful impact on migration rates in our model. In both the data and the model, occupation switches tend to take place early in the life cycle, but migration among the young in the model is mostly driven by amenity preferences and experimentation, not by occupational switches.

B. Change in gross migration

Decline in geographic specificity of returns to skills

We measure the effect on migration of the fall in geographic specificity of skills by changing θ from 1.15 to 1.10 — consistent with the average measured state-occupation interaction in earnings in 2005 to 2011 — but leaving all other parameters the same and simulating the model's new steady state. Figures 16(a) and 16(b) show the results for the non-college and college samples, respectively. The solid red line in each figure shows the age profile of migration in the baseline model. The dashed blue line shows the age profile of migration in the model with a lower geographic specificity of returns to skills. On average over all

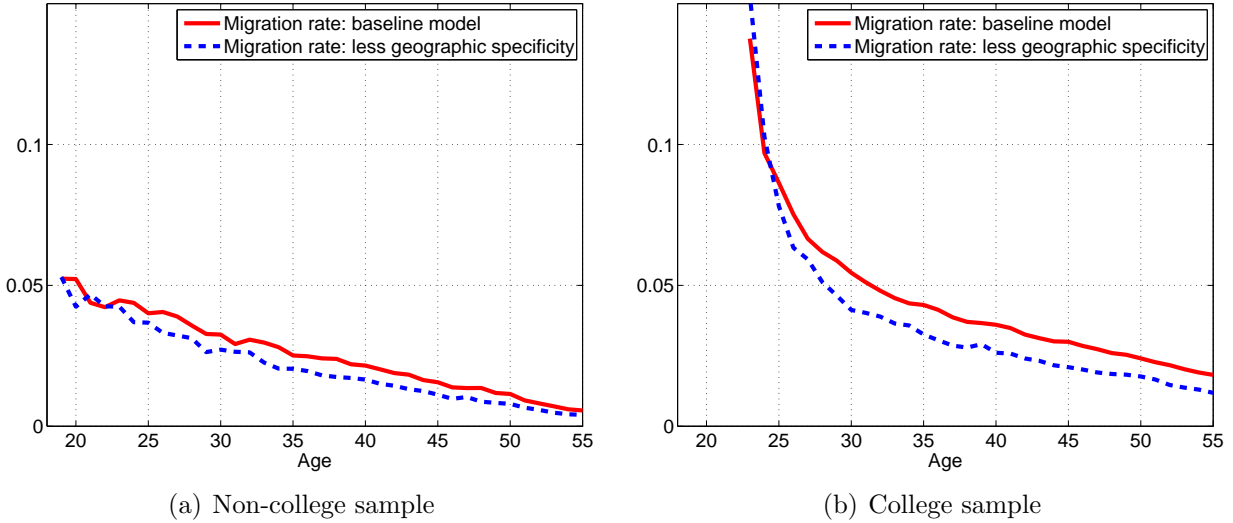


Figure 16: Effects of reduced geographic specificity on migration.

ages, the decline in the geographic specificity of skills reduces the migration rate from 2.56 percent to 2.14 percent for the non-college sample and from 4.40 percent to 3.64 percent for the college sample. These changes are equivalent to 42 percent of the observed drop in migration for the non-college group and 52 percent of the observed drop for the college group from the 1991–’97 baseline to the 2005–’11 period. Thus, the increased similarity of returns to occupational skills across space can account for about half of the total observed decline in migration between 1991–’97 and 2005–’11.

It is not surprising that the decline in the wage premium for working in a matched location-occupation has a larger effect for the college sample and that the resulting declines in migration are concentrated among older workers. (In fact, for the college sample, this mechanism has essentially no impact on migration in the first five years in the labor market.) Higher-wage (i.e., older and college-educated) workers respond more to changes in θ because these workers have more to gain in absolute terms from a wage premium that generates a multiplicative increase in wages for being in a matched location-occupation. In addition, older workers have more incentive to move locations to be better matched to their particular skills because they are less likely to want to switch occupations before their careers end.

In the quantitative exercise, we feed into the model an exogenous change in skill prices and allow agents’ occupation and location decisions to respond endogenously. We can check

whether our model features a quantitatively realistic endogenous response to the change in skill prices by comparing the implied change in the distribution of workers across space with the change observed in the data. In the baseline model, the Theil index of occupational segregation is 0.092 for the non-college group and 0.083 for the college group. These indices fall to 0.060 in both groups in the experiment with lower geographic specificity θ .²⁷ The level of the Theil index of occupational segregation depends on the number of occupations and hence is quite different in the simulated model (with two occupations) than in the data (where we consider many more occupations). To adjust for this difference in levels, we compare the percent change in the Theil index in the model and the data. In the model, the change in geographic specificity θ reduces the Theil index by 35 percent for the non-college sample and by 28 percent for the college sample. By comparison, figure 11(b) shows that in the data, from 1990 to 2010, the Theil index for broad occupations fell by 30 percent and the index for detailed occupations fell by 23 percent. Since our calibration strategy did not target changes in the Theil index, this close mapping between data and model should be taken as evidence that our relatively simple model of migration and occupation decisions does a relatively good job of capturing workers' endogenous responses to changes in the geographic specificity of skills.

Increase in information

Figure 17 shows the decline in migration that the model generates when we give individuals more information about their preferences for living in different locations. The solid red line in each figure is the age profile of migration in the baseline model. The other four lines show the age profile of migration when individuals are better informed about their preferences. As discussed above, we model the information change as an increase in the initial precision of beliefs about preferences for amenities, τ_0 , and we set as an upper bound for this increase the improvement in precision that would be obtained by living in a location for one year, τ_ϵ . The figures also report the corresponding reductions in migration if we increase the amount of available information by smaller amounts: one-quarter, one-half, and three-quarters of a year's worth of signals. Overall, these relatively modest increases in information

²⁷Results are similar in the experiment below with more information as well as lower geographic specificity. Adding information alone slightly reduces the index.

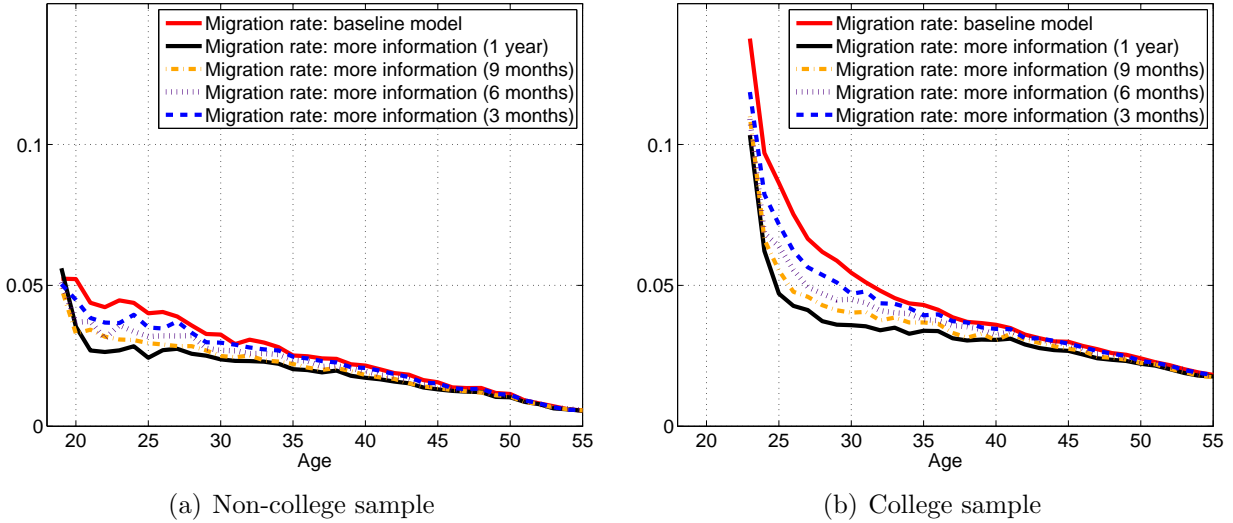


Figure 17: Effects of increased information on migration.

generate large declines in migration. For the non-college sample, the overall decline ranges from a 7 percent reduction for three months' worth of information to a 24 percent reduction for one year's worth of information. For the college sample, the overall decline ranges from an 8 percent reduction for three months' worth of information to a 25 percent reduction for one year's worth of information. Expressed as a fraction of the observed decline in migration, this mechanism generates between 19% and 60% of the decline in the data for the non-college sample, and between 25% and 74% for the college sample.

Increased availability of information mainly affects young workers. Figure 17 reveals that as retirement approaches, this mechanism has almost no impact on migration rates. This finding is intuitive: By the time individuals reach the second half of their working lives, most of the initial uncertainty about their preferences has been resolved, so changes in the precision of initial beliefs have little effect on migration rates at older ages.

We can check whether the informational changes we feed into the model are reasonable by examining their effect on the repeat migration rate. (In the model, repeat and return migration are identical because there are only two locations, so we compare repeat migration in the model with repeat migration in the data.) In the baseline model, the one-year repeat migration rate is 1.0 percent for the non-college group and 5.4 percent for the college group. These rates fall to 0.8 percent and 4.4 percent, respectively, in the experiment with six

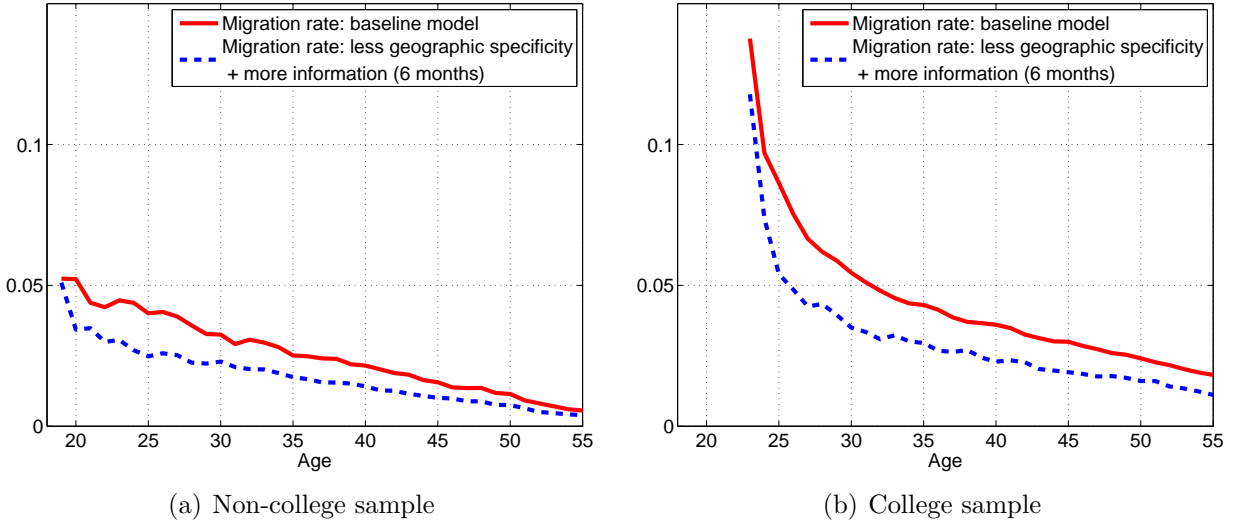


Figure 18: Combined effects of geographic specificity and information on migration.

additional months of information.²⁸ The levels of repeat migration in the model are much lower than in the data, most likely because the model does not include heterogeneity in moving costs κ and thus cannot reproduce a phenomenon that we think is likely to be important in the real world: Migrants are those with lower moving costs and thus are disproportionately likely to migrate again. However, adding six months of information reduces repeat migration by a factor of about one-fourth to one-fifth in the model, a change that is broadly consistent with the observed decline in figure 13.

Combined effect of the two mechanisms

The previous two experiments showed that a reduction in the geographic specificity of returns to skills is quantitatively consistent with the observed decline in migration at older ages, while an increase in the availability of information is quantitatively consistent with the observed decline in migration at younger ages. Since the data show that migration has declined at all ages, it is natural to conjecture that the combined effect of our two proposed mechanisms can quantitatively account for the overall observed reduction in migration. Figure 18 shows the results of an experiment in which we increase information by the equivalent of

²⁸The changes are proportionately larger or smaller when we add more or less information. Changing θ alone has almost no effect on the repeat migration rate.

Table 3: Average migration rates in the data and the model.

	Non-college	College
Data, 1991–’97	2.69%	4.16%
Model, baseline	2.56%	4.40%
Model, less geographic specificity*	2.14%	3.64%
Model, +6 months’ information [†]	2.21%	3.76%
Model, both mechanisms* [†]	1.76%	3.02%
Data, 2005–’11	1.69%	2.69%

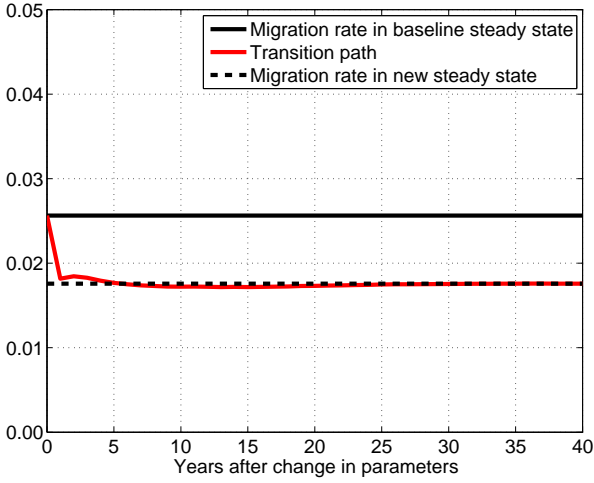
*Reduce θ from 1.15 to 1.10. [†]Increase τ_0 by $\tau_\epsilon/2$.

six months’ worth of signals and reduce the wage premium for being in a matched location-occupation by the amount observed in the data. As shown in table 3, the model generates an overall reduction in migration of 0.8 percentage point for the non-college sample, which is equivalent to 80 percent of the corresponding decline in the data; and a reduction of 1.38 percentage points for the college sample, which is equivalent to 93 percent of the corresponding decline in the data. Additional information would further reduce the migration rate. Since we are somewhat agnostic on the actual size of the increase in information that took place over the period from 1991 to 2011, and since the change in θ is somewhat imprecisely estimated, we conclude that our combined mechanisms can account for at least one-third and possibly all of the observed reduction in gross interstate migration in the United States.

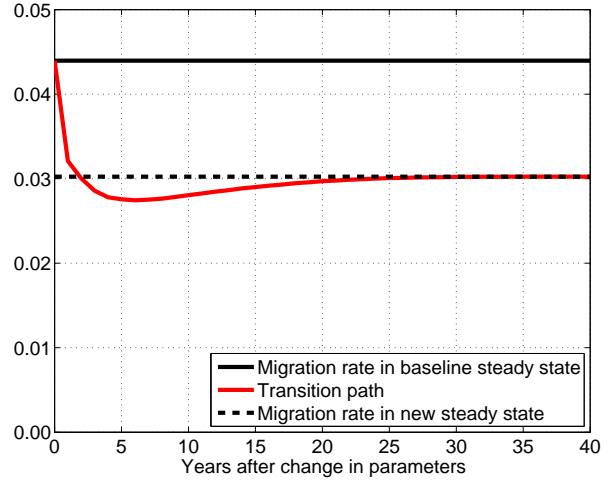
C. Transition dynamics

We claimed above that it was appropriate to measure θ in the same years as the migration rate and to compare steady states with different parameters because convergence to the steady state is rapid in our model. Here, we explore the model’s transition dynamics to provide evidence for this claim.

Figure 19 shows how the average migration rate in the model evolves after we change the geographic specificity of occupations and the amount of information agents have. In the experiment, agents begin in the steady state corresponding to our baseline calibration. Then we impose two one-time, unexpected, permanent parameter changes: a decrease in θ from 1.15 to 1.10 and an increase in the precision of every agent’s beliefs corresponding to six



(a) Non-college sample



(b) College sample

Figure 19: Transition dynamics of the average migration rate after a change in geographic specificity and information.

months of information.

As the figure shows, the migration rate reaches the level of the new steady state within a few years. This rapid convergence occurs because the only potentially slow-moving endogenous state variable in our model is the precision of agents' beliefs, but our experiment directly shocks this precision.

Convergence in our model would be slower if we assumed that the new parameters applied only to cohorts newly reaching working age and not to agents who were of working age before the parameter change. Such an assumption would be quite strong, however. For example, it would mean assuming that workers who were 20 years old when the geographic specificity of occupations changed would ignore that change for the rest of their lives.

In any event, our model's overall ability to explain a large fraction of the decline in migration is not sensitive to what we assume about whether parameter changes affect new cohorts or all cohorts. This timing assumption has little effect on transition dynamics related to information, because changes in information mainly affect young agents, and after a few years, all young agents are from cohorts that reached working age after the parameter change. Thus, the timing assumption mainly has an impact on transition dynamics related to a change in θ , the geographic specificity of occupations. In additional experiments, we found

that the migration rate reaches the new steady-state value after 20 to 30 years when the new parameters apply only to newborn cohorts. Therefore, if we assume new parameters affect only new cohorts, we should measure the change in θ over a period 20 to 30 years earlier than the period when we measure a decline in migration. Figure 10(a) shows that the decline in the geographic specificity of occupations across states from 1970 to 1990 was as large as the decline from 1991 to 2011.²⁹ As a result, our comparison of steady states based on 1991 and 2011 parameters gives quantitatively similar results to what we would obtain if we calibrated the model to the change in θ between 1970 and 1990, simulated transitions assuming that the new parameters affect only newborn cohorts, and compared the simulated migration rates in 1991 and 2011.

7. Conclusion

We argue that interstate migration is falling in the United States because of a combination of two factors: a reduction in the geographic specificity of returns to different types of skills and an increase in workers' information about how much they will enjoy living in alternative locations. Micro data reject numerous alternative explanations but do support our two hypotheses. We build a model of migration that makes these hypotheses precise. In the model, workers choose locations on the basis of both income and local amenities, search for jobs both locally and remotely, and gradually learn about the amenities in different locations. The calibrated model provides a good fit to the data and shows that our mechanisms can account for at least one-third and possibly all of the decline in interstate migration over the past two decades.

Our empirical analysis reveals a novel fact about U.S. labor markets: Returns to occupations have become less geographically specific over time. While our analysis takes this change as exogenous and studies its implications for migration, looking into the causes of the decrease in geographic specificity would be a valuable subject for future research. Understanding these causes is important for determining what policies, if any, are appropriate

²⁹In addition, improvements in information have continued over many decades, and it seems reasonable to assume that the changes from 1970 to 1990 (related to airline and telephone deregulation) were similar in magnitude to those between 1991 and 2011. But such an assumption is not necessary because, as noted, the model converges rapidly after a change in information even if that change affects only new cohorts.

in response to the decline in migration. It would also be valuable to know whether similar changes in the geography of work have happened in other parts of the world, and if so, what the effect on migration has been. (Molloy, Smith, and Wozniak, 2011, show that Canada and most European countries have not experienced a secular decline in migration, although a variety of factors such as European economic integration may cause the international experience to differ from the U.S. one.)

The decline in interstate migration is not the only recent change in gross worker flows in the United States: Davis, Faberman, and Haltiwanger (2012) show that job creation and destruction rates have also fallen in the past two decades, while Kambourov and Manovskii (2008) show that occupational and industry mobility increased between 1968 and 1997. Because job-to-job flows and occupational and industry changes are much more common than long-distance migration, we did not ask our model to explain these phenomena. However, future research could examine whether there is a connection between changes in interstate migration and changes in other gross flows, and if so, whether that connection helps explain the fraction of the migration decrease that our mechanisms leave unexplained.

Appendix

A1. One-digit industry and occupation categories

We use the following one-digit industry categories in figures 5(a), 11(a), and 11(c):

1. Unknown
2. Agriculture/forestry/fishing
3. Mining/construction
4. Durable manufacturing
5. Nondurable manufacturing
6. Transportation/utilities
7. Trade
8. Finance/insurance/real estate
9. Services
10. Government
11. Not in labor force

We use the following one-digit occupation categories in figures 5(a), 10, 11(b), and 11(d):

1. Executive, administrative, and managerial
2. Professional specialty
3. Technicians and related support
4. Sales
5. Administrative support
6. Service
7. Farming, forestry, and fishing
8. Precision production, craft, and repair
9. Operators, fabricators, and laborers
10. Military
11. Unemployed not classified
12. Not in labor force

A2. Details on SIPP and census data on repeat and return migration

SIPP. We use data from the 1985, 1986, 1987, 1988, 1990, 1991, 1992, 1993, 1996, 2001, 2004, and 2008 SIPP panels. We omit the 1984 panel because it uses different groupings of small states that would cause us to lose information and the 1989 panel because it was followed for less than two full years.

The SIPP interviews respondents every four months. Thus, the wave 1 interview tells us the respondent's location at baseline, the wave 4 interview tells us the location in year $t + 1$, and the wave 7 interview tells us the location in year $t + 2$. We restrict the sample to respondents who (a) had data collected in all of waves 1, 4, and 7 and (b) as of wave 4 were working-age adults (defined as in our CPS sample) and lived in households where no one was a member of the military. Restriction (a) implies that we exclude anyone who moves and is not followed by the SIPP. Restriction (b) makes our sample as comparable as possible to the working-age adults sample from the CPS (for which we measure migration between $t - 1$ and t among people who are working-age adults in non-military households at t).

We construct a combined state code for respondents in Maine and Vermont, and a second combined state code for respondents in Alaska, Idaho, Iowa, Montana, North Dakota, South Dakota, and Wyoming, because states within these groups cannot be distinguished in some years in the public-use data files.

We weight the data by the SIPP panel weights. For panels where the Census Bureau constructed several panel weights, we use the panel weight that corresponds to the first two years of the survey. For pre-2001 panels, we use the half-sample and stratum codes provided with the public-use files to account for the survey sampling design when we estimate confidence intervals. Starting with the 2001 panel, we use the panel replicate weights to estimate confidence intervals.

Census. We use the 1 percent form 1 and form 2 state samples from the 1970 census, the 5 percent samples from the 1980 and 1990 censuses, and the 5 percent and 1 percent samples from the 2000 census, all obtained from Ruggles et al. (2010). We restrict the sample to respondents who were born in the United States, lived in the United States five years earlier, do not have imputed birth states or locations five years earlier, and are working-age adults

in non-military households at the time of the census. We use the stratum and cluster codes provided by Ruggles et al. (2010) to account for the census public-use sample design when we estimate confidence intervals.

A3. Updating formulas for beliefs

This appendix derives the updating formulas for beliefs about amenities. Because the prior and the signal are both normally distributed, the posterior after any number of signals will also be a normal distribution and can be completely described by its mean and variance. Let $m_{t,n}$ and $\frac{1}{\tau_{t,n}^2}$ be the mean and variance of the posterior at date t after n signals. Using Bayes' theorem and the definitions of normal densities, we have the following relationship between the kernels of the posterior, signal, and prior after one signal:

$$\exp\left(-\frac{1}{2}\tau_{t,1}^2(v - m_{t,1})\right) \propto \exp\left(-\frac{1}{2}\tau_\epsilon^2(u_t - v)\right) \exp\left(-\frac{1}{2}\tau_0^2(v - m_0)\right),$$

which implies

$$\tau_{t,1}^2 = \tau_0^2 + \tau_\epsilon^2, \quad m_{t,1} = \frac{\tau_0^2 m_0 + \tau_\epsilon^2 u_t}{\tau_0^2 + \tau_\epsilon^2}.$$

Repeating the same analysis given priors with mean $m_{t,n-1}$ and precision $\tau_{t,n-1}^2$, we arrive at the following general updating formulas for the moments of the belief distribution:

$$\begin{aligned} \tau_{t,n}^2 &= \tau_{t-1,n-1}^2 + \tau_\epsilon^2 = \tau_0^2 + n\tau_\epsilon^2, \\ m_{t,n} &= \frac{[\tau_0^2 + (n-1)\tau_\epsilon^2]m_{t-1,n-1} + \tau_\epsilon^2 u_t}{\tau_0^2 + n\tau_\epsilon^2}. \end{aligned}$$

Because there are only two locations, we only need to keep track of n_t^a , the number of periods lived in location a up to time t . The number of periods lived in location b is then given by $t - n_t^a$, and the precision of beliefs can be expressed as a function of n_t^a . The updating formulas conditional on the location j_t where the agent lives in period t are thus:

$j_t = a$	$j_t = b$
$n_t^a = n_{t-1}^a + 1$	$n_t^a = n_{t-1}^a$
$m_t^a = \frac{(\tau_0^2 + n_{t-1}^a \tau_\epsilon^2) m_{t-1}^a + \tau_\epsilon^2 u_t}{\tau_0^2 + n_t^a \tau_\epsilon^2}$	$m_t^a = m_{t-1}^a$
$m_t^b = m_{t-1}^b$	$m_t^b = \frac{[\tau_0^2 + (t-1 - n_{t-1}^a) \tau_\epsilon^2] m_{t-1}^b + \tau_\epsilon^2 u_t}{\tau_0^2 + (t - n_t^a) \tau_\epsilon^2}$

The conditional distribution of the time $t + 1$ signal, u_{t+1} , given information available at the end of period t , is normal with mean and variance given by

$$\begin{aligned} \mathbb{E} \left[u_{t+1}^{j_{t+1}} | j_{t+1}, n_t^{j_{t+1}}, m_t^{j_{t+1}} \right] &= m_t^{j_{t+1}}, \\ \text{Var} \left[u_{t+1}^{j_{t+1}} | j_{t+1}, n_t^{j_{t+1}}, m_t^{j_{t+1}} \right] &= \frac{1}{\tau_0^2 + n_t^{j_{t+1}} \tau_\epsilon^2} + \frac{1}{\tau_\epsilon^2}, \end{aligned}$$

where we have used the fact that the time $t + 1$ location decision is known at the end of period t .

A4. Bellman equations

This appendix describes the Bellman equations associated with the decision problem in the model of section 5. We consider the expected present value of an individual in period t just before making his location and island choice for $t + 1$, (j_{t+1}, i_{t+1}) . The state variables at this point are $x_t = (j_t, s_t, m_t, n_t, o_t)$, where j_t is the current location, s_t is the vector of skills in the two occupations, m_t is the vector of beliefs about preferences over the two locations, n_t is the number of periods lived in location a , and $o_t = (o_t^a, o_t^b)'$ are indicator variables denoting whether the individual has an offer to work in each location. Because employed workers are free to choose either occupation, the current occupation is not relevant when deciding the future location once we condition on o_t ; hence, k_t is not a state variable. Furthermore, because individuals are always free to quit to unemployment, an individual will never choose a pair $(j, 0)$ over $(j, 1)$ if he has an offer at location j . Nonetheless, even if j_{t+1} is the only choice at a given state, it is convenient to define the value functions in terms of location-island pairs. We will denote this choice as $y_t \equiv (j_{t+1}, i_{t+1})$. Consequently, let $J_t(x_t, y_t)$ be the expected present value of an individual in period t who has state variables x_t and chooses the location-island pair y_t .

Agents make two other decisions in each period: occupation and search choices. It is

useful to define a beginning-of-period value function. Let $V_t(j_t, i_t, s_t, m_t, n_t^a) \equiv V_t(\Omega_t)$ be the expected present value of an individual who begins period t on island i_t at location j_t with state variables (s_t, m_t, n_t^a) , and who makes optimal occupation, search, and migration choices from then onward.

The choice-specific value functions are then given by

$$\begin{aligned} J_t(x_t, y_t) &= \sum_{s_{t+1} u_{t+1}} \int [u_{t+1} + \beta V_{t+1}(\Omega_{t+1})] dF(u_{t+1} | m_t, j_{t+1}) \Pr(s_{t+1} | s_t) - \kappa_t(x_t, y_t) \\ &= \sum_{s'} \int_{u'} [u' + \beta V'(\Omega')] dF(u' | m, n, j') \Pr(s' | s) - \kappa(x, y), \end{aligned}$$

where in the second line we have used primes to denote $t + 1$ variables. The migration cost $\kappa_t(x_t, y_t)$ for individuals who migrate is $\kappa p_k^j \psi_t e^{s_t}$ for those who worked in period t and κq for those who were unemployed, where $p_A^a = p_B^b = \theta$ and $p_A^b = p_B^a = 1$.

Using the conditional distribution of u_{t+1} derived above, we then have

$$J(x, y) = m^{j'} + \beta \sum_{s'} \int_{u'} V'(\Omega') dF(u' | m, n, j') \Pr(s' | s) - \kappa(x, y),$$

where u' shows up inside the integral in the $m^{j'}$ component of Ω' . This holds because $m' = E[u' | j', n, m]$.

We now derive the value function V . Define

$$L_t(x_t) = \max_{y_t} J_t(x_t, y_t).$$

The available choices y_t are determined by the offers o_t . Expanding $y_t = (j_{t+1}, i_{t+1})$ and making use of the fact that $J(x, (j, 1)) \geq J(x, (j, 0))$ always, we have

$$L_t(x_t) = \max \{ J_t(x_t, a, o_t^a), J_t(x_t, b, o_t^b) \}.$$

For ease of notation, we suppress the rest of the state space x_t and the time subscript because these choices are made within one time period, and we denote this value function as $L(o^a, o^b)$.

Consider now the search decision. Let ζ denote the search decision, with $\zeta = 0$ being

no search, $\zeta = 1$ being search in the opposite location, and $\zeta = 2$ being search in the current location (which is relevant only for unemployed households). Let $H_e(\zeta)$ and $H_u(\zeta)$ be the expected present values of employed and unemployed agents who choose search strategy ζ . For an agent at location a , these search-specific value functions are

$$\begin{aligned}
H_e(0) &= L(1, 0) \\
H_e(1) &= \lambda_e L(1, 1) + (1 - \lambda_e) H_e(0) - c_e \\
H_u(0) &= L(0, 0) \\
H_u(1) &= \lambda_{u,1} L(0, 1) + (1 - \lambda_{u,1}) H_u(0) - c_{u,1} \\
H_u(2) &= \lambda_{u,2} L(1, 0) + (1 - \lambda_{u,2}) H_u(0) - c_{u,2},
\end{aligned}$$

where again we have suppressed the dependence on the state variables (j, s, m, n) . These functions are analogously defined for agents in location b .

Finally, let $K(j, k)$ be the expected present value of an agent in location j who works in occupation k . These occupation-specific value functions are

$$K(j, k) = \begin{cases} p_k^j \psi e^{s^j} + (1 - \delta) \max_{q \leq 1} H_e(q) + \delta \max_q H_u(q) & \text{if employed } (k \in \{A, B\}) \\ q + \max_q H_u(q) & \text{if unemployed } (k = u). \end{cases}$$

The beginning-of-period value function V can thus be written as

$$V_t(j_t, i_t, s_t, m_t, n_t^a) = \begin{cases} \max_k K(k, j_t, i_t, s_t, m_t, n_t^a) & \text{if } i_t = 1 \\ K(u, j_t, i_t, s_t, m_t, n_t^a) & \text{if } i_t = 0, \end{cases}$$

where the second line reflects the fact that if $i_t = 0$, the individual is unemployed and so $k_t = u$.

References

Bayer, Christian, and Falko Juessen, 2012, “On the Dynamics of Interstate Migration: Migration Costs and Self-Selection.” *Review of Economic Dynamics* 15(3), 377–401.

- Batini, Nicoletta, Oya Celasun, Thomas Dowling, Marcello Estevão, Geoffrey Keim, Martin Sommer, and Evridiki Tsounta, 2010, “United States: Selected Issues Paper.” International Monetary Fund Country Report No. 10/248. Available at <http://www.imf.org/external/pubs/ft/scr/2010/cr10248.pdf>.
- Berry, Christopher R., and Edward L. Glaeser, 2005, “The Divergence of Human Capital Levels Across Cities.” *Papers in Regional Science* 84(3), 407–444.
- Borenstein, Severin, and Nancy L. Rose, 2008, “How Airline Markets Work . . . Or Do They? Regulatory Reform in the Airline Industry.” Forthcoming in *Economic Regulation and Its Reform: What Have We Learned?* Chicago: University of Chicago Press. Available at <http://www.nber.org/chapters/c12570.pdf>.
- Borjas, George J., Stephen G. Bronars, and Stephen J. Trejo, 1992, “Self-Selection and Internal Migration in the United States.” *Journal of Urban Economics* 32(2), 159–185.
- Card, David, and Alan B. Krueger, 1992, “Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States.” *Journal of Political Economy* 100(1), 1–40.
- Carlino, Gerald, and Satyajit Chatterjee, 2002, “Employment Deconcentration: A New Perspective on America’s Postwar Urban Evolution.” *Journal of Regional Science* 42(3), 445–475.
- Chib, Siddhartha, 2001, “Markov Chain Monte Carlo Methods: Computation and Inference.” In *Handbook of Econometrics*, vol. 5, ed. James J. Heckman and Edward Leamer, 3569–3649. Amsterdam: Elsevier.
- Coen-Pirani, Daniele, 2010, “Understanding Gross Worker Flows Across U.S. States.” *Journal of Monetary Economics* 57(7), 769–784.
- Crozet, Matthieu, 2004, “Do Migrants Follow Market Potentials? An Estimation of a New Economic Geography Model.” *Journal of Economic Geography* 4(4), 439–458.

- Dahl, Gordon B., 2002, "Mobility and the Return to Education: Testing a Roy Model with Multiple Markets." *Econometrica* 70(6), 2367–2420.
- Davern, Michael, Arthur Jones Jr., James Lepkowski, Gestur Davidson, and Lynn A. Blewett, 2006, "Unstable Inferences? An Examination of Complex Survey Sample Design Adjustments Using the Current Population Survey for Health Services Research." *Inquiry* 43(3), 283–297.
- Davis, Morris, Jonas Fisher, and Marcelo Veracierto, 2010, "The Role of Housing in Labor Reallocation." Federal Reserve Bank of Chicago Working Paper 2010-18.
- Davis, Steven J., R. Jason Faberman, and John Haltiwanger, 2012, "Labor Market Flows in the Cross Section and Over Time." *Journal of Monetary Economics* 59(1), 1–18.
- Desmet, Klaus, and Esteban Rossi-Hansberg, 2009, "Spatial Growth and Industry Age." *Journal of Economic Theory* 144(6), 2477–2502.
- Duranton, Gilles, and Diego Puga, 2004, "Micro-foundations of Urban Agglomeration Economies." In *Handbook of Regional and Urban Economics*, vol. 4, ed. J.V. Henderson and J.-F. Thisse, 2063–2117. Amsterdam: Elsevier.
- Ganong, Peter, and Daniel Shoag, 2012, "Why Has Regional Convergence in the U.S. Stopped?" Harvard Kennedy School Research Working Paper 12-028.
- Garrett, B. Dale, and Andre L. Williams, 2006, "Meeting 21st Century Demographic Data Needs — Implementing the American Community Survey: Report 6: The 2001–2002 Operational Feasibility Report of the American Community Survey." Available at http://www.census.gov/acs/www/Downloads/library/2004/2004_Garrett_01.pdf.
- Gelman, Andrew, 2007, "Struggles with Survey Weighting and Regression Modeling." *Statistical Science* 22(2), 153–164.
- Gemici, Ahu, 2011, "Family Migration and Labor Market Outcomes." Manuscript, New York University.
- Glaeser, Edward L., 1999, "Learning in Cities." *Journal of Urban Economics* 46(2), 254–277.

- Glaeser, Edward L., Giacomo A.M. Ponzetto, and Kristina Tobio, forthcoming, “Cities, Skills, and Regional Change.” *Regional Studies*.
- Guler, Bulent, Fatih Guvenen, and Giovanni L. Violante, 2012, “Joint-Search Theory: New Opportunities and New Frictions.” *Journal of Monetary Economics* 59(4), 352–369.
- Harris, John R., and Michael P. Todaro, 1970, “Migration, Unemployment and Development: A Two-Sector Analysis.” *American Economic Review* 60(1), 126–142.
- Internal Revenue Service, 2008, “Supplemental Documentation for Migration Data Products.” Available at <http://www.irs.gov/pub/irs-soi/countmiguser0708.doc>.
- Kambourov, Gueorgui, and Iourii Manovskii, 2008, “Rising Occupational and Industry Mobility in the United States: 1968–97.” *International Economic Review* 49(1), 41–79.
- Kaplan, Greg, and Sam Schulhofer-Wohl, 2012, “Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey.” *Demography* 49(3), 1061–1074.
- Kennan, John, and James R. Walker, 2011, “The Effect of Expected Income on Individual Migration Decisions.” *Econometrica* 79(1), 211–251.
- King, Miriam, Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick, 2010, Integrated Public Use Microdata Series, Current Population Survey: Version 3.0. [Machine-readable database]. Minneapolis: University of Minnesota.
- Koerber, Kin, 2007, “Comparison of ACS and ASEC Data on Geographic Mobility: 2004.” Manuscript, U.S. Census Bureau. Available at http://www.census.gov/acs/www/library/by_year/2007/.
- Krugman, Paul, 1991, “Increasing Returns and Economic Geography.” *Journal of Political Economy* 99(3), 483–499.
- Lucas, Robert E., Jr., and Edward C. Prescott, 1974, “Equilibrium Search and Unemployment.” *Journal of Economic Theory* 7(2), 188–209.

- MacDonald, John S., and Leatrice D. MacDonald, 1964, "Chain Migration, Ethnic Neighborhood Formation, and Social Networks." *Milbank Memorial Fund Quarterly* 42(1), 82–97.
- Mincer, Jacob, 1978, "Family Migration Decisions." *Journal of Political Economy* 86(5), 749–773.
- Molloy, Raven, Christopher L. Smith, and Abigail Wozniak, 2011, "Internal Migration in the United States." *Journal of Economic Perspectives* 25(3), 173–196.
- Panel Study of Income Dynamics, 2012, Public use dataset produced and distributed by the Institute for Social Research, Survey Research Center, University of Michigan, Ann Arbor, Mich.
- Partridge, Mark D., Dan S. Rickman, M. Rose Olfert, and Kamar Ali, 2012, "Dwindling U.S. Internal Migration: Evidence of Spatial Equilibrium or Structural Shifts in Local Labor Markets?" *Regional Science and Urban Economics* 42(1–2), 375–388.
- Pingle, Jonathan F., 2007, "A Note on Measuring Internal Migration in the United States." *Economics Letters* 94(1), 38–42.
- Reardon, Sean F., and Glenn Firebaugh, 2002, "Measures of Multigroup Segregation." *Sociological Methodology* 32(1), 33–67.
- Roback, Jennifer, 1982, "Wages, Rents, and the Quality of Life." *Journal of Political Economy* 90(6), 1257–1278.
- Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek, 2010, Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]. Minneapolis: University of Minnesota.
- Taylor, William E., and Lester D. Taylor, 1993, "Postdivestiture Long-Distance Competition in the United States." *American Economic Review* 83(2) (Papers and Proceedings), 185–190.

- Theil, Henri, and Anthony J. Finizza, 1971, "A Note on the Measurement of Racial Integration of Schools by Means of Informational Concepts." *Journal of Mathematical Sociology* 1(2), 187–193.
- Tolbert, Charles M., and Molly Sizer, 1996, "U.S. Commuting Zones and Labor Market Areas: A 1990 Update." U.S. Department of Agriculture Economic Research Service Staff Paper No. AGES-9614.
- U.S. Census Bureau, 1999, "State Population Estimates and Demographic Components of Population Change: Annual Time Series, April 1, 1990 to July 1, 1999." Available at <http://www.census.gov/popest/data/state/totals/1990s/tables/ST-99-07.txt>.
- U.S. Census Bureau, 2004, "Accuracy of the Data (2004)." Available at http://www.census.gov/acs/www/Downloads/data_documentation/Accuracy/accuracy2004.pdf.
- U.S. Census Bureau, 2005, "Accuracy of the Data (2005)." Available at http://www.census.gov/acs/www/Downloads/data_documentation/Accuracy/accuracy2005.pdf.
- U.S. Census Bureau, 2006, "Design and Methodology: Current Population Survey." Technical Paper 66. Available at <http://www.census.gov/prod/2006pubs/tp-66.pdf>.
- U.S. Census Bureau, 2009a, "Annual Population Estimates, Estimated Components of Resident Population Change, and Rates of the Components of Resident Population Change for the United States, States, and Puerto Rico: April 1, 2000 to July 1, 2009." Available at http://www.census.gov/popest/data/national/totals/2009/files/NST_EST2009_ALLDATA.csv.
- U.S. Census Bureau, 2009b, "Design and Methodology: American Community Survey." Available at http://www.census.gov/acs/www/Downloads/survey_methodology/acs_design_methodology.pdf.
- U.S. Census Bureau, 2009c, "Survey of Income and Program Participation Users' Guide," third edition (initially published 2001, revised 2009). Available at <http://www.census.gov/sipp/usrguide.html>.

U.S. Census Bureau, 2012, American Community Survey, “Response Rates — Data,” December 17, accessed Jan. 17, 2013. Available at https://www.census.gov/acs/www/methodology/response_rates_data/.

Yatchew, A., 1997, “An Elementary Estimator of the Partially Linear Model.” *Economics Letters* 57, 135–143.