Federal Reserve Bank of Minneapolis
Research Department

Understanding the Long-Run Decline in Interstate Migration*

Greg Kaplan and Sam Schulhofer-Wohl

Working Paper 697

April 2012

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 show that micro data rule out many popular explanations for this decline, including aging of the population, the rise of two-earner households, other compositional changes, regional changes, and the rise in real incomes. We argue instead that the fall in migration is due to a decline in the geographic specificity of occupations and an increase in workers' ability to learn about other locations before moving there, through both information technology and inexpensive travel. We develop a theory to formalize these ideas and show that a plausibly calibrated version is consistent with cross-sectional and time-series patterns of interstate migration, occupations, and incomes.

Keywords: Interstate migration; Labor mobility; Gross flows; Information technology; Learning
JEL: D83, J11, J24, J61, R12, R23

---

*Kaplan: University of Pennsylvania; gkaplan@sas.upenn.edu. Schulhofer-Wohl: Federal Reserve Bank of Minneapolis; wohls@minneapolisfed.org. We thank Diego Amador and Chloe Booth for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

In the early 1990s, about 3 percent of Americans moved between states each year. Today, that rate has fallen in half. Many policymakers have worried that the decline in migration heralds a less-flexible economy where workers do not move to places with good jobs. In this paper, we argue that the worry is misplaced: Low migration is *good* news, because it means either that workers have better information about their opportunities or that workers do not need to move to obtain good jobs. We begin by using micro data to rule out many popular explanations for the decline in migration. For example, the decline is not due to aging of the population, changes in the number of two-earner households, or the rise in real incomes. We are left with two explanations that the data cannot rule out. The first explanation 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. The second explanation 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. In either case, the fall in migration does not demand a policy response. 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.

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 because most of the data are collected in March.[1] The net rate comes from the Census Bureau's annual state population estimates (U.S. Census Bureau, 1999, 2009).[2] 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 are much smaller than the overall decline

[1]See section 2 for details on the CPS data.

[2]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 $t+1$.

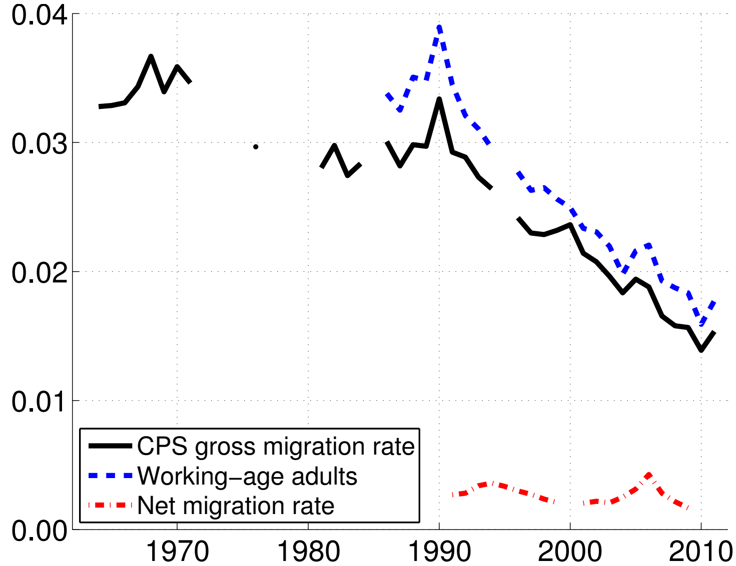

Figure 1: Gross and net interstate migration.
Source: Authors' calculations from Current Population Survey (CPS) micro data and Census Bureau population estimates.

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.[3] These patterns suggest that to understand the decline in migration, we must look for factors that affect gross flows rather than net flows; factors that vary over long time horizons rather than at business cycle frequencies; and factors 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.

[3]We define this sample in section 2.

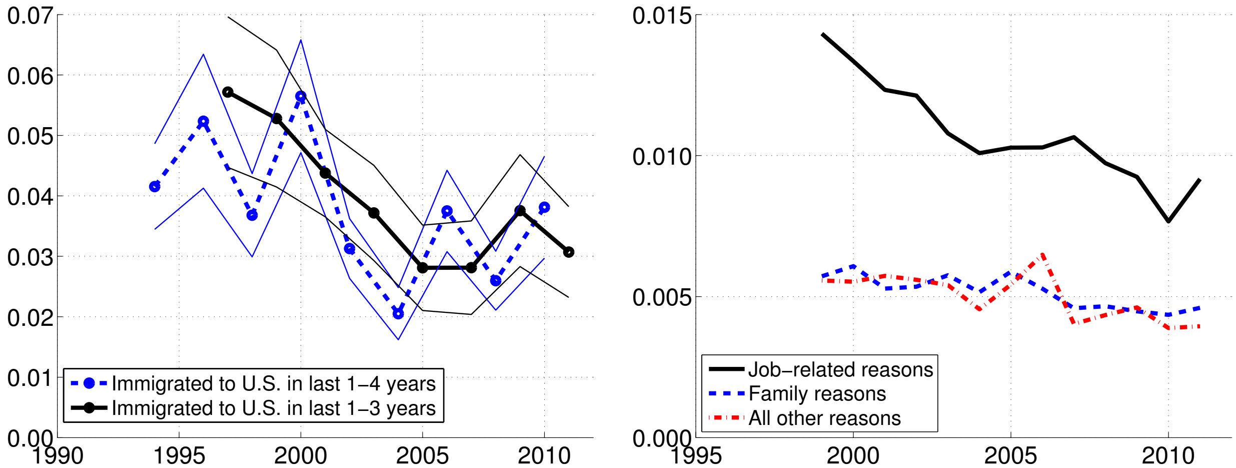

(a) Migration rates for recent immigrants

(b) Reasons for moving

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 be important 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 he or she moved for a new job, we think it is highly likely that the respondent changed jobs around the time of the move — even if other factors, such as local amenities, may have 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 of the two locations. 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.

The two locations in our model — call them "here" and "there" — differ in both the job opportunities they offer and the local amenities they provide. Workers thus potentially have two motives for changing locations: First, they will move if moving raises their expected earnings or expected utility from amenities. Second, if workers "here" are sufficiently uncertain about the job opportunities and amenities available "there," they may move "there" simply to acquire information, especially if they are young and have many years to benefit from additional information. Changes in the similarity of the two locations affect the first motive for migration, while changes in information affect the second motive.

Increased similarity of locations Locations may differ in the types of jobs they offer, and different jobs may be more or less suitable for particular workers. Suppose that locations become more similar in the types of jobs they offer. Then workers will be less likely to move to get a suitable job, but more likely to move for amenity-related reasons, because a smaller difference in amenities is now required to overcome the difference in the quality of the job match. But such a pattern contradicts the evidence: As figure 2(b) shows, amenity-related moves have not risen. Likewise, if locations become more similar in terms of amenities, job-related moves will rise and amenity-related moves will fall. That pattern, too, contradicts the data. An increase in the similarity of places can be part of the story, but it cannot be the whole story because no type of move has risen.

Increased information Based on evidence that most workers who move for job-related reasons do so with a new job already in hand, and the ease with which individuals can search for jobs remotely, 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. We model an increase in information as an increase in the precision of workers' prior beliefs about the utility they will derive from local amenities in each location. This increase in information decreases migration by reducing workers' motivation to move in order to acquire information; the effect is particularly strong among the young, consistent with empirical evidence that the decline in migration is strongest among the young.

Our two theories are difficult to test empirically because it is difficult to measure the counterfactual match quality a worker would have if he lived in a different place, or to measure the information people have about potential destinations before they move. However, we provide indirect evidence that supports the assumptions underlying each of the theories. To support the hypothesis that job opportunities have become more similar across locations, we show that occupations have become more evenly spread across the country. Further, this change likely results 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 show that the variance across states of the average income for a given occupation has fallen. (If, instead, workers increasingly desired to hold jobs in unproductive places — due for example to an exogenous decrease in mobility — the dispersion of incomes would rise.) To support the hypothesis that people have increasing amounts of information about non–labor market features of distant locations, we appeal to changes in the costs of obtaining such information due to advances in information technology and decreases in travel costs. We believe further measurements of the factors driving our theories are a fruitful area for future research. In addition, we show in this paper that other theories are inconsistent with the data, while a quantitative model embodying our theories can match many important features of the data both across workers and over time.

Our work is related to a substantial empirical and theoretical literature. Molloy, Smith, and Wozniak (2011) survey existing research on internal migration in the United States and empirically 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.[4] Theories of migration, such as the classic models by Harris and Todaro (1970) and Roback (1982), generally focus on net flows. Kennan and Walker (2011) structurally estimate a model of individual migration decisions in which workers choose locations to maximize their expected lifetime income. Differences in expected income across locations imply that the model features both gross flows and net flows. The price of studying net flows is that Kennan and Walker (2011) must allow workers to choose among a large number of locations, which means the model has many state variables and must be highly simplified along many dimensions to remain tractable. Our key theoretical insight is that if we set aside net flows and study only gross flows, the model needs to have only two locations, "here" and "there." This simplification allows us to add realism along other important dimensions. For example, workers in our model care about amenities as well as income; workers gradually learn about other locations, instead of having full information; skills are geographically specific; and workers living in one location can search remotely for jobs elsewhere, allowing us to study the effect of technological changes that make remote search less costly.

The paper proceeds as follows. In section 2, we describe the CPS data in detail 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 theory: We show empirically that the returns to working in particular occupations have become less geographically dispersed in recent decades and review evidence for falling costs of learning about distant locations. Section 5 lays out our model of information and migration. Section 6 concludes.

[4] One factor that has received much attention in the literature but that we do not consider here is fluctuations in the housing market and homeownership rates. 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 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) review an extensive literature and present their own estimates, showing that the decline in house prices since the mid-2000s plays at best a small role in the drop in migration over that period.

2. Data

We focus our analysis on working-age adults in civilian households in the CPS from 1991 to 2011.[5] (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 for accurately measuring trends in civilian migration 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 school, and is at least 19 years old.[6] Thus, we concentrate on people who have completed their educations but are not yet approaching retirement. From 1996 onward, we follow Kaplan and Schulhofer-Wohl (forthcoming) 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 available 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,

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

[6]The 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.

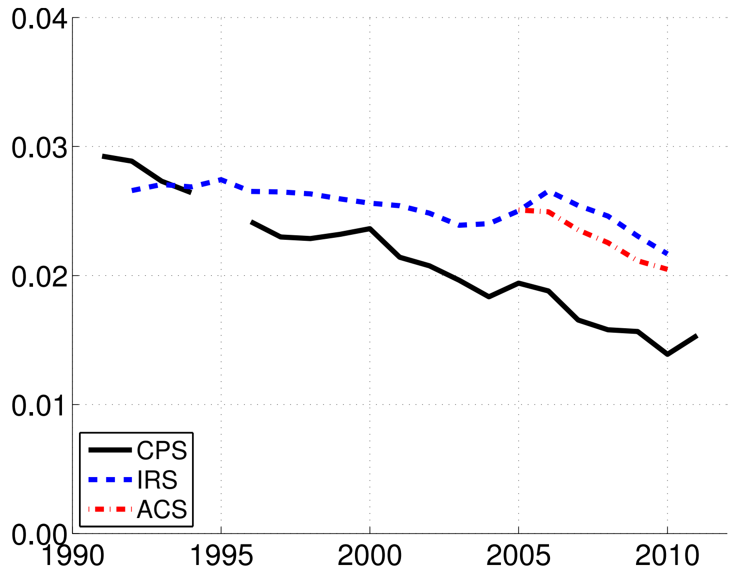

Figure 3: Gross migration measured from different data sources.
Source: Authors' calculations from Current Population Survey (CPS) and American Community Survey (ACS) micro data and Internal Revenue Service (IRS) state-to-state migration tabulations.

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 New York, New Jersey, and Connecticut from the data.

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 interstate migration rates appears in other data as well. Figure 3 compares migration rates in the CPS, in American Community Survey (ACS) micro data (Ruggles et al., 2010), and in Internal Revenue Service (IRS) data. The ACS migration rate parallels the CPS rate from 2005 to 2010 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 migration rate. 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 owe 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).[7]

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.[8]

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 gain a sense of what dimensions are important to model and to show that a variety of common beliefs about the decline 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

[7] Based 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.

[8] From 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, this assumption is not strictly correct, but the available sampling information in the public-use files does not allow us to easily relax it.)

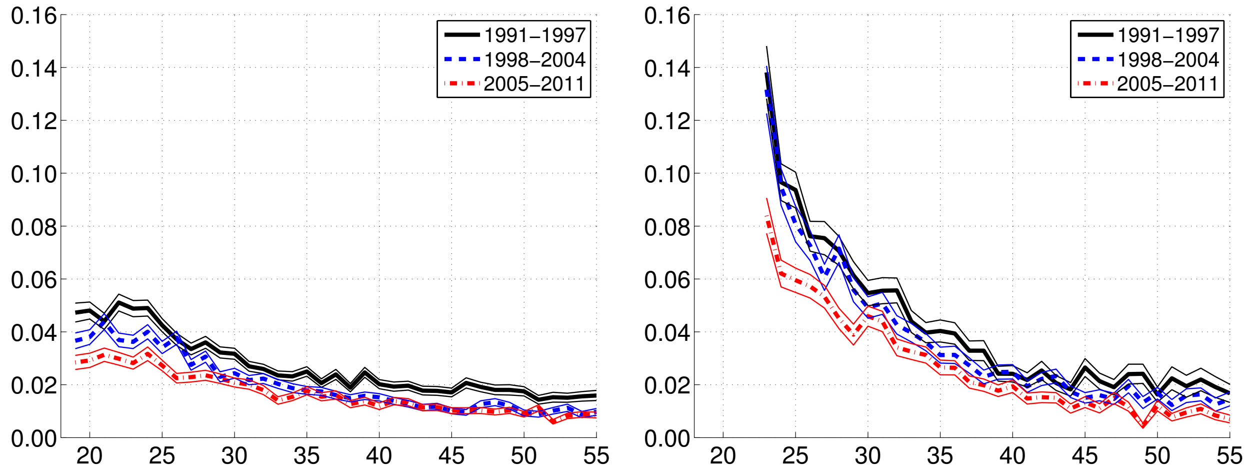

(a) No bachelor's degree

(b) Bachelor's degree

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.

larger for the young — a fact we will ask our calibrated model to reproduce. Because the age profiles are quite different for the two education groups, we will produce two calibrations of our model, one for each group.

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, 2010; Mincer, 1978) who cannot move because their partners cannot move.

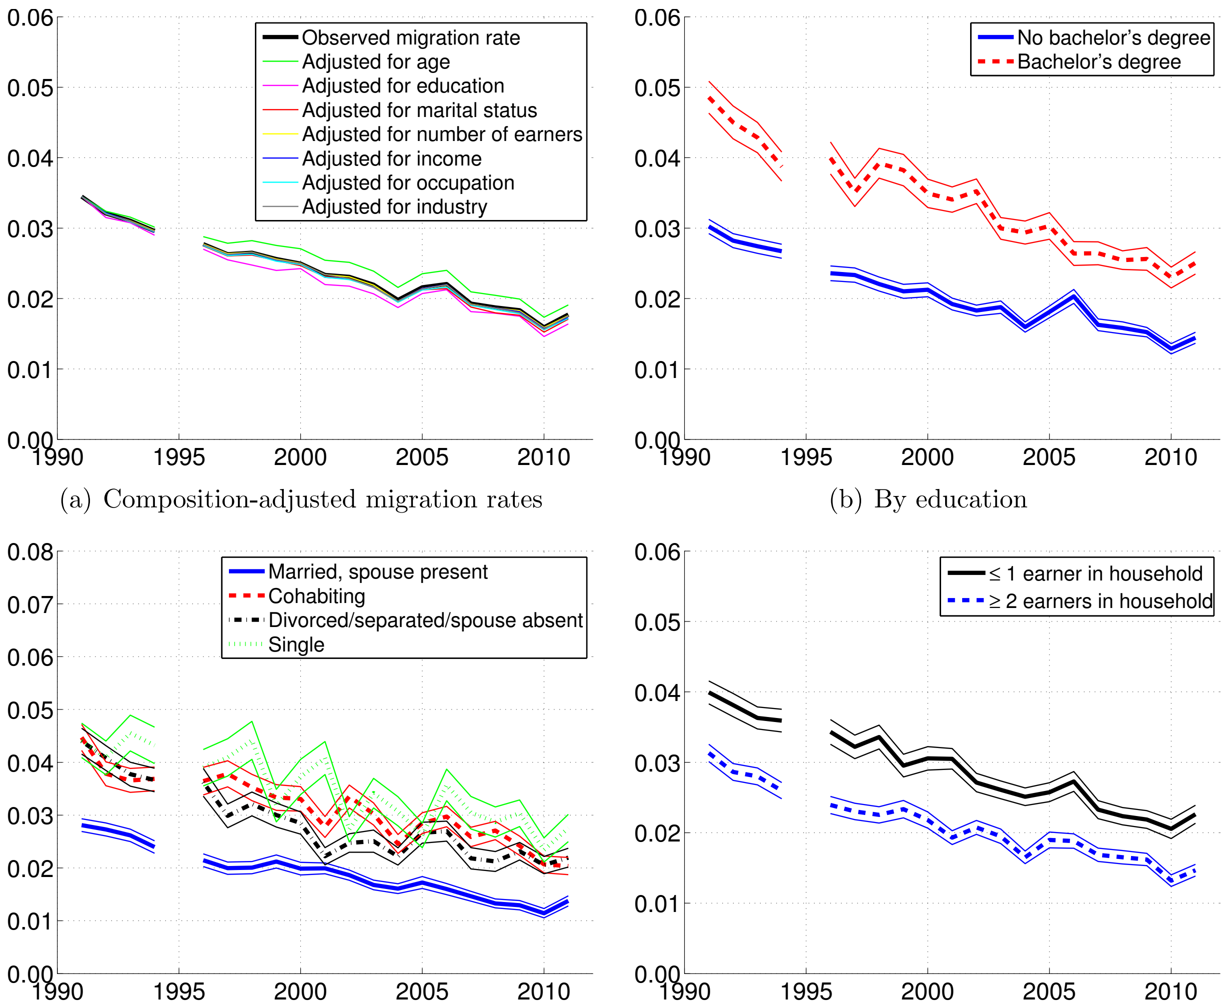

(a) Composition-adjusted migration rates

(b) By education

(c) By household structure

(d) By number of household members in labor force

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.

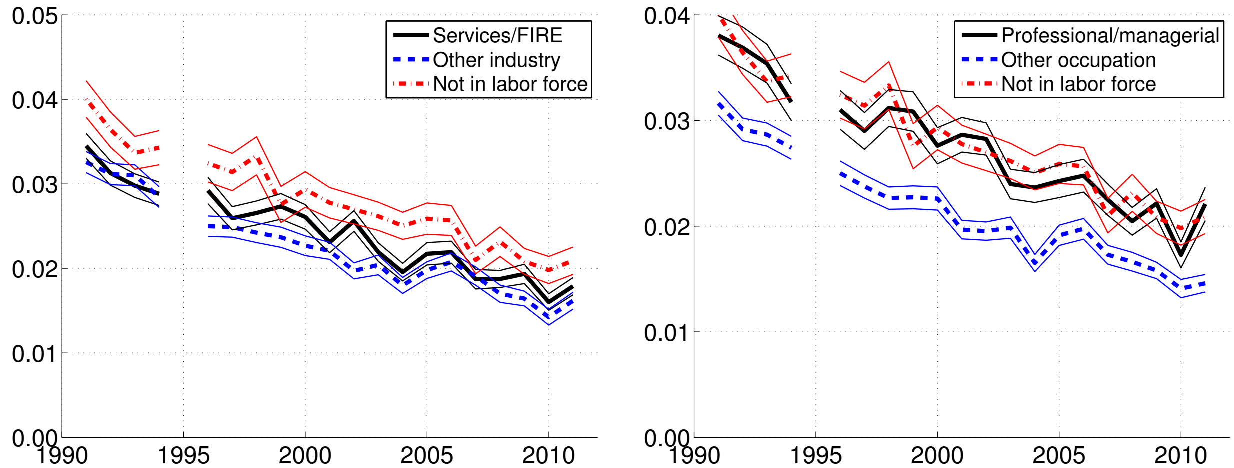

(a) By industry (b) By occupation

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).

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 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

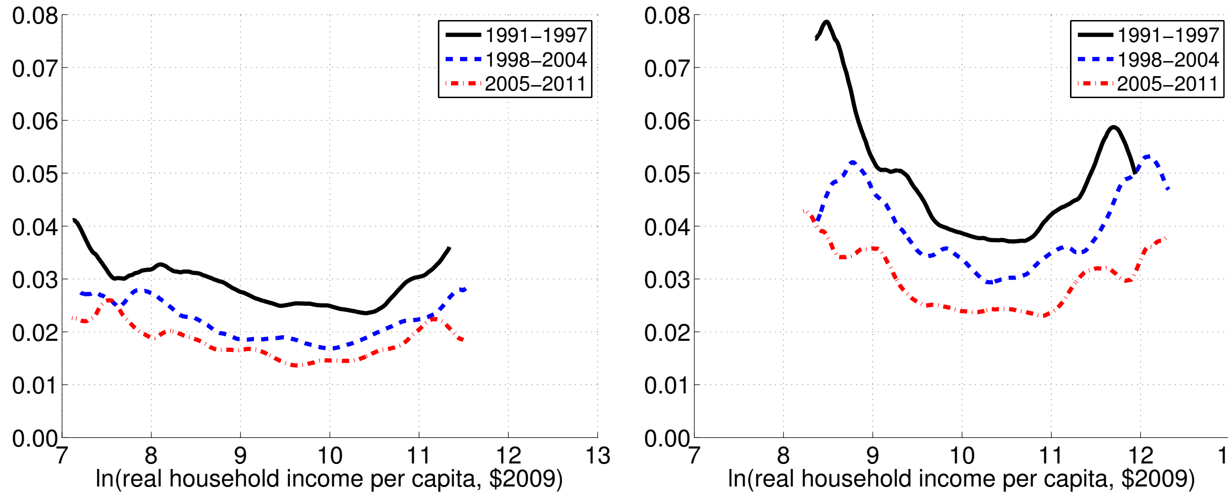

(a) No bachelor's degree

(b) Bachelor's degree

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.

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.[9] Controlling for age is important 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,

[9] We 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: $\mathit{migration} = \mathbf{x}\boldsymbol{\beta} + f(\mathit{income}) + \epsilon$, 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.

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 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 inmigration 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. 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.[10]

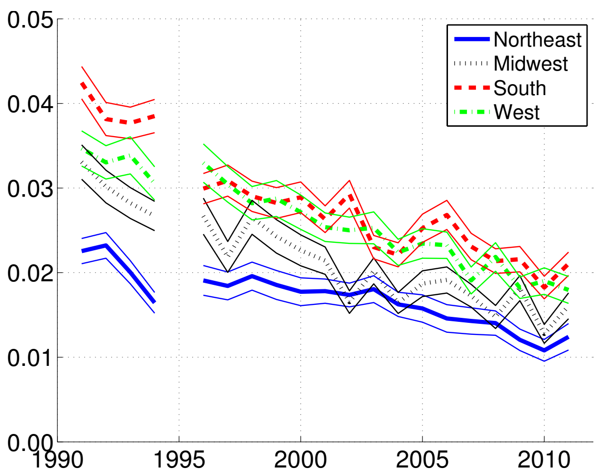

(a) Gross inflows by region

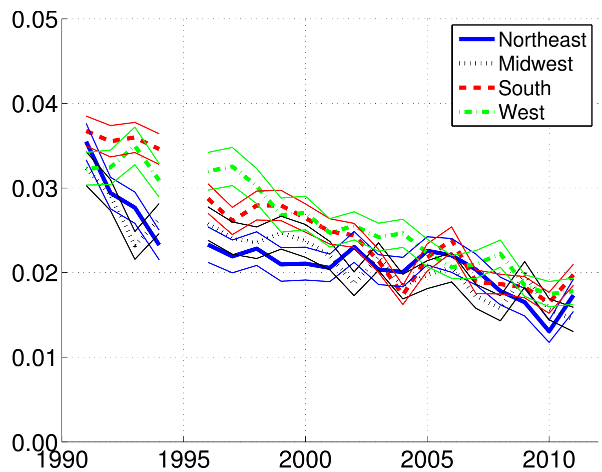

(b) Gross outflows by region

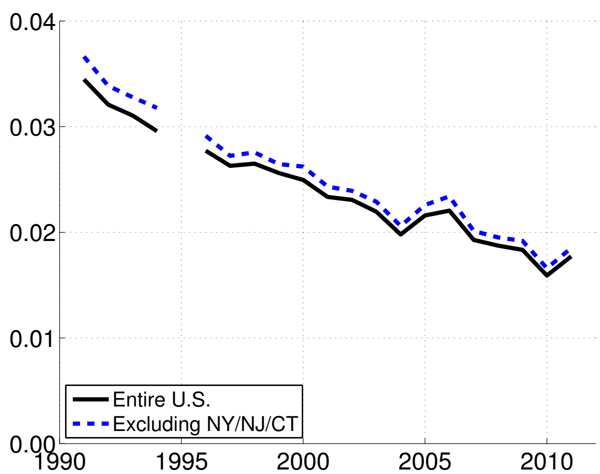

(c) Effect of excluding New York, New Jersey, and Connecticut

Figure 8: Regional patterns in migration.

Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults.

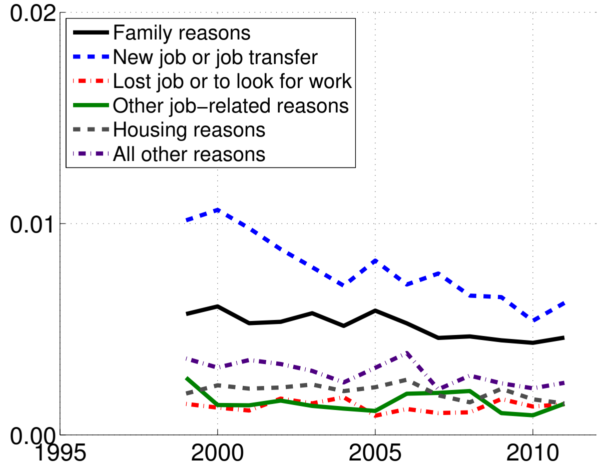

Figure 9: Detailed reasons for moving.
Source: Authors' calculations from Current Population Survey (CPS) micro data. Sample restricted to working-age adults. Estimates shown for all years when variables are available.

E. Job search

Improvements in information technology 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. Although self-reported reasons for choices must be interpreted with caution, these findings are the opposite of what we would expect if information technology 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.

[10] We exclude the entirety of the three states because the boundaries of the metropolitan area have changed over time.

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. Decreases in the dispersion of job opportunities

We test whether job opportunities around the country have become more similar by examining both quantities and prices. First, we show that occupations and industries have become less geographically segregated around the country. That is, the distribution of occupations and industries among workers in each state has become more similar to the national average. This change in the *quantity* of workers in various occupations and industries 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., an increase in workers' willingness to do their occupations in places where those occupations have low productivity). 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 a decrease in migration and would not itself cause migration to fall. To distinguish between demand and supply effects, we examine the dispersion of incomes across states within occupations. If productivity in particular occupations becomes less geographically specific, income dispersion within occupations will fall; if, on the other hand, workers become willing to do their occupations in unproductive places, income dispersion within occupations will rise. We find that income dispersion within occupations has fallen, supporting the view that occupations' productivity levels have become less geographically specific.

The dispersion of quantities of workers

In figure 10, 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

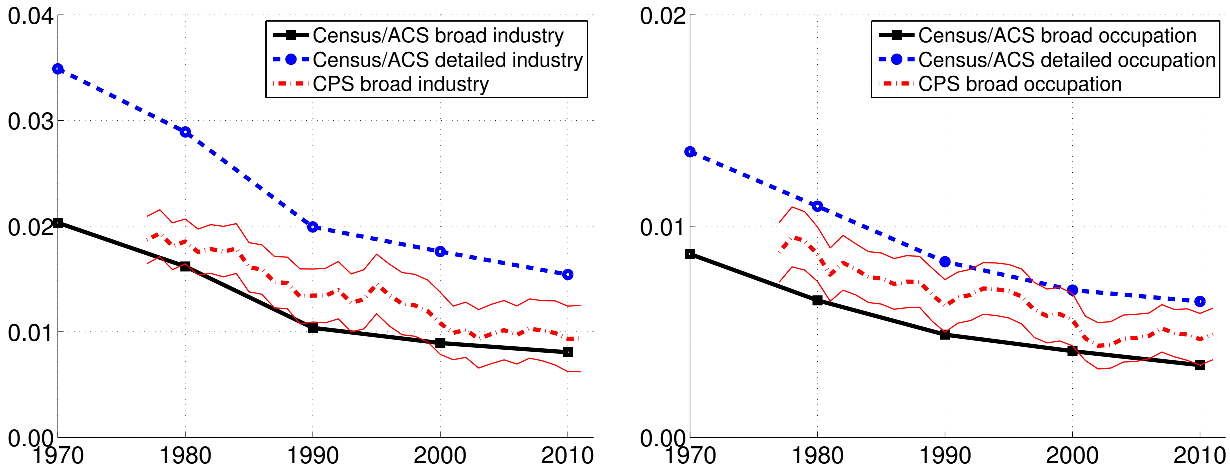

(a) By industry (b) By occupation

Figure 10: 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. 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.

industries in each state 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 state, and 0 when each state has the same distribution of occupations or industries as the nation as a whole.[11]

Figure 10 shows that, over time, states' distributions of workers across industries have become more similar. This pattern holds whether we look at single-digit industries or more

[11] The 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 indexes 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 j'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.

detailed categories. (To examine more detailed categories, we must go to decennial census and multiyear ACS data,[12] because the CPS contains too few observations to reliably estimate, say, the number of workers in hardware stores in Vermont in 2010.) The same pattern holds when we look at occupations instead of industries. In addition, 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.

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

$$\ln y_{iost} = a_{st} + b_{ot} + \mathbf{x}'_{iost}\boldsymbol{\beta}_t + \theta_{ost} + \epsilon_{iost},$$

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 θ_{ost} is the state-occupation interaction of interest. We model θ_{ost} as a normally distributed random effect with mean 0 and variance $\sigma^2_{\theta,t}$ and estimate the variance $\sigma^2_{\theta,t}$ for each year.[13] When $\sigma^2_{\theta,t}$ 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

[12] 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, the 5 percent and 1 percent samples from the 2000 census, and the 2006–2010 combined ACS, which is equivalent to a 5 percent sample. We obtain all census and ACS data from Ruggles et al. (2010). 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.

[13] Treating θ_{ost} as a fixed effect and then calculating the variance of the estimated fixed effects would produce upward-biased estimates of $\sigma^2_{\theta,t}$ 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 $\sigma^2_{\theta,t}$.

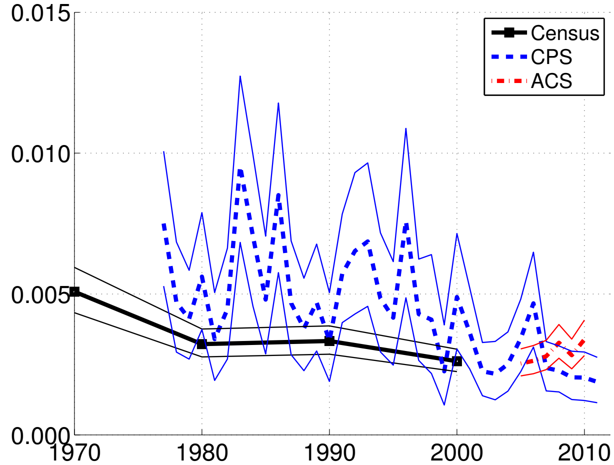

Figure 11: Variance of state-occupation interaction in income.
Source: Authors' calculations from Current Population Survey (CPS) micro data, 1977–2011; American Community Survey (ACS) micro data, 2005–2010; and decennial census micro data (1970, 1980, 1990 and 2000). Thick lines are posterior means; thin lines show 2.5th and 97.5th percentiles of posterior distribution.

less geographically specific, $\sigma^2_{\theta,t}$ will fall.[14]

Figure 11 shows the posterior distribution of σ^2_{ost} for each year. There is a clear decline from 1970 to 2000 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 be sure whether the true trend is upward or downward.

B. Decreases in the cost of information

The past several decades have seen dramatic changes in the ease with which people can gather information about distant locations. Most obvious, of course, is the development of the Internet, which allows people to inexpensively and rapidly gather information on life

[14]We use single-digit occupations to keep the number of parameters manageable. Nonetheless, 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 procedures. 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 $\boldsymbol{\beta}_t$ and standard weakly informative priors for the variances of θ_{ost} and ϵ_{iost}. The prior has a larger effect on the posterior when the sample is smaller. We choose a prior for σ^2_{ost} centered on 0.001, which is smaller than any of the posterior means. Thus, the downward trend over time in the posterior estimates of σ^2_{ost} could arise 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.

in other cities. But other changes in technology and markets 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, this 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 unemployed. 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. This feature of our model also allows us to explore the possibility that changes in the cost of searching for jobs remotely have contributed to the decline in migration.

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 geography Individuals (which we will also refer to as households or agents) live for T periods, $t = 1 \ldots T$. 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 unemployed, $k = u$. An individual in period t is characterized by his skills in each of the two occupations, $s_t = \left(s_t^A, s_t^B\right)'$, which evolve according to an exogenous Markov process, and his preferences for the non–labor market features of the two locations, $v = \left(v^a, v^b\right)'$, which are fixed over time. 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 (although, as we show below, it may impart some information about the *relative* preference). 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 Each individual's skills in the two occupations, s_t, are revealed to the individual at the beginning of period t. However, an individual's preferences for the non–labor market features of each location, v, are not known to the individual and must be learned over time, through living in the two locations. Each period, an individual receives non–labor market utility from living in location j_t, given by $u_t = v^{j_t} + \epsilon_t$, i.e., 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:

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

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 unemployed households. To find the production island, unemployed 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 reflects the marginal product of skills in each occupation. We thus specify the price of skills directly.

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

$$\theta\left(j_t, k_t\right) = \begin{cases} p_J^j & \text{if } j_t = k_t \\ 1 & \text{if } j_t \neq k_t \end{cases}.$$

Thus, $\theta \geq 1$ is the wage premium for working in a matched location and occupation. Individuals' skills in each of the two occupations evolve according to independent Markov processes. These processes are normalized so that $E\left[e^{s_t}\right] = 1\ \forall t$. In addition, we assume that skills have a deterministic life cycle component ψ_t, in order to accurately capture the evolution of average wages over the life cycle. Wages of employed workers are hence given by

$$w_t\left(j_t, k_t, s_t\right) = \theta\left(j_t, k_t\right)\psi_t e^{s_t}.$$

Individuals who are unemployed receive a constant benefit b. We define $w_t\left(j_t, k_t = u, s_t\right) \equiv b$.

Both unemployed and employed workers can choose to search for the production island in either location, regardless of where they are currently located. However, search is costly. Moreover, search is more efficient for unemployed agents than for employed agents, and for individuals searching in their current location than for individuals searching remotely.

There are thus three arrival probabilities (i.e., probabilities of finding the production island, conditional on searching for it): $\lambda_{u,0}$, for unemployed households searching in their current location; $\lambda_{u,1}$, for unemployed households searching in the remote location; and $\lambda_{e,1}$, for employed households searching in the remote location. There are three associated costs of searching: $(c_{u,0}, c_{u,1}, c_{e,1})$. Because the labor market is competitive on the production island, there is no reason for an employed worker to search in his current location — such an individual has already found the production island and would not gain from searching for it again.

Preferences Individuals have risk-neutral preferences and make search, occupation, and migration decisions to maximize the discounted sum of expected utility:

$$\mathbf{E}\sum_{t=1}^{T}\beta^{t-1}\left[u_t + w_t - \kappa\mathbf{1}_{j_t\neq j_{t+1}} - c\mathbf{1}_{\text{search}}\right].$$

u_t is the utility derived from the non–labor market features of the location j_t where the individual lives in period t. This non–labor market utility comprises the individual's underlying (unknown) preference for that location, v^{j_t}, plus the random utility shock, ϵ_t. w_t is the wage or unemployment benefit received in period t. κ is the cost of moving, paid only if the individual lives in two different locations in t and $t+1$. The final term reflects the cost associated with searching for the production island. This cost may differ depending on the individual's current employment status and the location in which he searches, but to simplify the notation, the above expression does not reflect these differences in search costs. There are no constraints on borrowing. Given the timing described below, this maximization problem can be expressed recursively. The associated Bellman equations are given in the Appendix.

Timing An individual enters period t with the following relevant information: the location where he resides (j_t); his current island, i.e., production or unemployment (i_t); his skills in 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^a_{t-1}). Recall that $n^b_{t-1} = t - 1 - n^a_{t-1}$.

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^a_t, and his beliefs about his utility from living in location j_t, $m^{j_t}_t$. 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 formulae for updating beliefs can be found in the Appendix.
4. If the individual is on the production island ($i_t = 1$), then he chooses his occupation for the current period k_t. He may also choose to quit to the unemployment island.
5. Individuals work. Employed individuals are paid their earnings, $w_t\left(j_t, k_t, s_t\right)$. Unemployed individuals receive their benefit b.
6. Individuals decide whether to search for the production island in either location, paying the appropriate cost if they do so, and the results of search are realized.
7. Conditional on the outcome of search, the individual makes his migration decision, i.e., he chooses his location for $t + 1$. This consists of the choice of a location-island pair (j_{t+1}, i_{t+1}). There is a moving cost κ to change locations. There is no cost for switching occupations.

B. Incentives to migrate

Why might an individual in this model decide to migrate? First, consider a shock to skills in one or both of the two occupations, $\left(s^a_t, s^b_t\right)$. 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 so 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. A move that was made for reasons of experimentation may lead to return (or onward) migration if realized utility in the new location turns out to be lower than the worker expected.

6. Conclusion

We argue that interstate migration is falling in the United States due to a combination of two factors: (i) a reduction in the geographic specificity of returns to different types of skills, and (ii) 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. In ongoing work, we plan to calibrate our model and quantitatively assess the fall in gross migration when workers get more information about amenities and when skills become less geographically specific.

Appendix

A1. One-digit industry and occupation categories

We use the following one-digit industry categories in figures 5(a) and 10(a):

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) and 10(b):

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.

References

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.

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.

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.

Gemici, Ahu, 2011, "Family Migration and Labor Market Outcomes." Manuscript, New York University.

Guler, Bulent, Fatih Guvenen, and Giovanni L. Violante, 2010, "Joint-Search Theory: New Opportunities and New Frictions." Manuscript, New York University.

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.

Kaplan, Greg, and Sam Schulhofer-Wohl, forthcoming, "Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey." *Demography*.

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/`.

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.

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, 2009, "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`.

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