

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 79, No. 1 (January, 2011), 211–251

THE EFFECT OF EXPECTED INCOME ON INDIVIDUAL MIGRATION DECISIONS

JOHN KENNAN

University of Wisconsin, Madison, WI 53706, U.S.A. and NBER

JAMES R. WALKER

University of Wisconsin, Madison, WI 53706, U.S.A. and NBER

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

THE EFFECT OF EXPECTED INCOME ON INDIVIDUAL MIGRATION DECISIONS

BY JOHN KENNAN AND JAMES R. WALKER¹

This paper develops a tractable econometric model of optimal migration, focusing on expected income as the main economic influence on migration. The model improves on previous work in two respects: it covers optimal sequences of location decisions (rather than a single once-for-all choice) and it allows for many alternative location choices. The model is estimated using panel data from the National Longitudinal Survey of Youth on white males with a high-school education. Our main conclusion is that interstate migration decisions are influenced to a substantial extent by income prospects. The results suggest that the link between income and migration decisions is driven both by geographic differences in mean wages and by a tendency to move in search of a better locational match when the income realization in the current location is unfavorable.

KEYWORDS: Migration, dynamic discrete choice models, job search, human capital.

1. INTRODUCTION

THERE IS AN EXTENSIVE LITERATURE on migration.² Most of this work describes patterns in the data: for example, younger and more educated people are more likely to move; repeat and especially return migration account for a large part of the observed migration flows. Although informal theories explaining these patterns are plentiful, fully specified behavioral models of migration decisions are scarce, and these models generally consider each migration event in isolation, without attempting to explain why most migration decisions are subsequently reversed through onward or return migration.

This paper develops a model of optimal sequences of migration decisions, focusing on expected income as the main economic influence on migration. The model is estimated using panel data from the National Longitudinal Survey of Youth (NLSY) on white males with a high-school education. We emphasize that migration decisions are reversible and that many alternative locations must be considered. Indeed (as we show in Section 2), repeat migration is a prominent feature of the data, and in many cases people choose to return to

¹The National Science Foundation and the NICHD provided research support. We thank Taisuke Otsu for outstanding research assistance. We also thank the editor and referees for very detailed constructive criticism of several earlier versions of the paper. We are grateful to Joe Altonji, Kate Antonovics, Peter Arcidiacono, Gadi Barlevy, Philip Haile, Bruce Hansen, Igal Henedel, Yannis Ioannides, Mike Keane, Derek Neal, John Pencavel, Karl Scholz, Robert Shimer, Chris Taber, Marcelo Veracierto, Ken Wolpin, Jim Ziliak, and seminar and conference participants at the Chicago Federal Reserve Bank, Carnegie–Mellon, Duke, Iowa, IZA, Ohio State, Penn State, Rochester, SITE, the Upjohn Institute, Virginia, Wisconsin, and Yale for helpful comments.

²See Greenwood (1997) and Lucas (1997) for surveys.

a location that they had previously chosen to leave, even though many unexplored alternative locations are available. A dynamic model is clearly necessary to understand this behavior.

Structural dynamic models of migration over many locations have not been estimated before, presumably because the required computations have not been feasible.³ A structural representation of the decision process is of interest for the usual reasons: we are ultimately interested in quantifying responses to income shocks or policy interventions not seen in the data, such as local labor demand shocks or changes in welfare benefits. Our basic empirical question is the extent to which people move for the purpose of improving their income prospects. Work by [Keane and Wolpin \(1997\)](#) and by [Neal \(1999\)](#) presumes that individuals make surprisingly sophisticated calculations regarding schooling and occupational choices. Given the magnitude of geographical wage differentials, and given the findings of [Topel \(1986\)](#) and [Blanchard and Katz \(1992\)](#) regarding the responsiveness of migration flows to local labor market conditions, one might expect to find that income differentials play an important role in migration decisions.

We model individual decisions to migrate as a job search problem. A worker can draw a wage only by visiting a location, thereby incurring a moving cost. Locations are distinguished by known differences in wage distributions and amenity values. We also allow for a location match component of preferences that is revealed to the individual for each location that is visited.

The decision problem is too complicated to be solved analytically, so we use a discrete approximation that can be solved numerically, following [Rust \(1994\)](#). The model is sparsely parameterized. In addition to expected income, migration decisions are influenced by moving costs (including a fixed cost, a reduced cost of moving to a previous location, and a cost that depends on distance), by differences in climate, and by differences in location size (measured by the population in each location). We also allow for a bias in favor of the home location (measured as the state of residence at age 14). Age is included as a state variable, entering through the moving cost, with the idea that if the simplest human capital explanation of the relationship between age and migration rates is correct, there should be no need to include a moving cost that increases with age.

³[Holt \(1996\)](#) estimated a dynamic discrete choice model of migration, but his framework modeled the move–stay decision and not the location-specific flows. Similarly, [Tunali \(2000\)](#) gave a detailed econometric analysis of the move–stay decision using microdata for Turkey, but his model does not distinguish between alternative destinations. [Dahl \(2002\)](#) allowed for many alternative destinations (the set of states in the United States), but he considered only a single lifetime migration decision. [Gallin \(2004\)](#) modeled net migration in a given location as a response to expected future wages in that location, but he did not model the individual decision problem. [Gemici \(2008\)](#) extended our framework and considered family migration decisions, but defined locations as census regions.

Our main substantive conclusion is that interstate migration decisions are indeed influenced to a substantial extent by income prospects. There is evidence of a response to geographic differences in mean wages, as well as a tendency to move in search of a better locational match when the income realization in the current location is unfavorable.

More generally, this paper demonstrates that a fully specified econometric model of optimal dynamic migration decisions is feasible and that it is capable of matching the main features of the data, including repeat and return migration. Although this paper focuses on the relationship between income prospects and migration decisions at the start of the life cycle, suitably modified versions of the model can potentially be applied to a range of issues, such as the migration effects of interstate differences in welfare benefits, the effects of joint career concerns on household migration decisions, and the effects on retirement migration of interstate differences in tax laws.⁴

2. MIGRATION DYNAMICS

The need for a dynamic analysis of migration is illustrated in Table I, which summarizes 10-year interstate migration histories for the cross-section sample of the NLSY, beginning at age 18. Two features of the data are noteworthy. First, a large fraction of the flow of migrants involves people who have already moved at least once. Second, a large fraction of these repeat moves involves people returning to their original location. Simple models of isolated move–stay decisions cannot address these features of the data. In particular, a model of return migration is incomplete unless it includes the decision to leave the

TABLE I
INTERSTATE MIGRATION, NLSY 1979–1994^a

	Less Than High School	High School	Some College	College	Total
Number of people	322	919	758	685	2,684
Movers (age 18–27)	80	223	224	341	868
Movers (%)	24.8%	24.3%	29.6%	49.8%	32.3%
Moves per mover	2.10	1.95	1.90	2.02	1.98
Repeat moves (% of all moves)	52.4%	48.7%	47.4%	50.5%	49.5%
<i>Return migration</i> (% of all moves)					
Home	32.7%	33.1%	29.1%	23.2%	28.1%
Not home	15.5%	7.1%	6.8%	8.6%	8.4%
Movers who return home	61.3%	56.5%	51.3%	42.8%	50.2%

^aThe sample includes respondents from the cross-section sample of the NLSY79 who were continuously interviewed from ages 18 to 28 and who never served in the military. The home location is the state of residence at age 14.

⁴See, for example, Kennan and Walker (2010) and Gemici (2008).

initial location as well as the decision to return. Moreover, unless the model allows for many alternative locations, it cannot give a complete analysis of return migration. For example, a repeat move in a two-location model is necessarily a return move, and this misses the point that people frequently decide to return to a location that they had previously decided to leave, even though many alternative locations are available.

3. AN OPTIMAL SEARCH MODEL OF MIGRATION

We model migration as an optimal search process. The basic assumption is that wages are local prices of individual skill bundles. We assume that individuals know the wage in their current location, but to determine the wage in another location, it is necessary to move there, at some cost. This reflects the idea that people may be more productive in some locations than in others, depending on working conditions, residential conditions, local amenities, and so forth. Although some information about these things can of course be collected from a distance, we view the whole package as an experience good.

The model aims to describe the migration decisions of young workers in a stationary environment. The wage offer in each location may be interpreted as the best offer available in that location.⁵ Although there are transient fluctuations in wages, the only chance of getting a permanent wage gain is to move to a new location. One interpretation is that wage differentials across locations equalize amenity differences, but a stationary equilibrium with heterogeneous worker preferences and skills still requires migration to redistribute workers from where they happen to be born to their equilibrium location. Alternatively, it may be that wage differentials are slow to adjust to location-specific shocks, because gradual adjustment is less costly for workers and employers.⁶ In that case, our model can be viewed as an approximation in which workers take current wage levels as an estimate of the wages they will face for the foreseeable future. In any case, the model is intended to describe the partial equilibrium response of labor supply to wage differences across locations; from the worker's point of view, the source of these differences is immaterial, provided that they are permanent. A complete equilibrium analysis would of course be much more difficult, but our model can be viewed as a building-block toward such an analysis.

⁵This means that we are treating local match effects as relatively unimportant: search within the current location quickly reveals the best available match.

⁶Blanchard and Katz (1992), using average hourly earnings of production workers in manufacturing, by state, from the Bureau of Labor Statistics (BLS) establishment survey, described a pattern of "strong but quite gradual convergence of state relative wages over the last 40 years." For example, using a univariate AR(4) model with annual data, they found that the half-life of a unit shock to the relative wage is more than 10 years. Similar findings were reported by Barro and Sala-i Martin (1991) and by Topel (1986).

Suppose there are J locations, and individual i 's income y_{ij} in location j is a random variable with a known distribution. Migration decisions are made so as to maximize the expected discounted value of lifetime utility. In general, the level of assets is an important state variable for this problem, but we focus on a special case in which assets do not affect migration decisions: we assume that the marginal utility of income is constant, and that individuals can borrow and lend without restriction at a given interest rate. Then expected utility maximization reduces to maximization of expected lifetime income, net of moving costs, with the understanding that the value of amenities is included in income, and that both amenity values and moving costs are measured in consumption units. This is a natural benchmark model, although of course it imposes strong assumptions.

There is little hope of solving this expected income maximization problem analytically. In particular, the Gittins index solution of the multiarmed bandit problem cannot be applied because there is a cost of moving.⁷ But by using a discrete approximation of the wage and preference distributions, we can compute the value function and the optimal decision rule by standard dynamic programming methods, following Rust (1994).

3.1. The Value Function

Let x be the state vector (which includes wage and preference information, current location, and age, as discussed below). The utility flow for someone who chooses location j is specified as $u(x, j) + \zeta_j$, where ζ_j is a random variable that is assumed to be independent and identically distributed (i.i.d.) across locations and across periods, and independent of the state vector. Let $p(x'|x, j)$ be the transition probability from state x to state x' if location j is chosen. The decision problem can be written in recursive form as

$$V(x, \zeta) = \max_j (v(x, j) + \zeta_j),$$

where

$$v(x, j) = u(x, j) + \beta \sum_{x'} p(x'|x, j) \bar{v}(x')$$

and

$$\bar{v}(x) = E_{\zeta} V(x, \zeta),$$

and where β is the discount factor and E_{ζ} denotes the expectation with respect to the distribution of the J -vector ζ with components ζ_j . We assume that

⁷See Banks and Sundaram (1994) for an analysis of the Gittins index in the presence of moving costs.

ζ_j is drawn from the type I extreme value distribution. In this case, following McFadden (1974) and Rust (1987), we have

$$\exp(\bar{v}(x)) = \exp(\bar{\gamma}) \sum_{k=1}^J \exp(v(x, k)),$$

where $\bar{\gamma}$ is the Euler constant. Let $\rho(x, j)$ be the probability of choosing location j , when the state is x . Then

$$\rho(x, j) = \exp(\bar{\gamma} + v(x, j) - \bar{v}(x)).$$

We compute v by value function iteration, assuming a finite horizon T . We include age as a state variable, with $v \equiv 0$ at age $T + 1$, so that successive iterations yield the value functions for a person who is getting younger and younger.

4. EMPIRICAL IMPLEMENTATION

A serious limitation of the discrete dynamic programming method is that the number of states is typically large, even if the decision problem is relatively simple. Our model, with J locations and n points of support for the wage distribution, has $J(n+1)^J$ states for each person at each age. Ideally, locations would be defined as local labor markets, but we obviously cannot let J be the number of labor markets; for example, there are over 3,100 counties in the United States. Indeed, even if J is the number of States, the model is computationally infeasible,⁸ but by restricting the information available to each individual, an approximate version of the model can be estimated; this is explained below.

4.1. *A Limited History Approximation*

To reduce the state space to a reasonable size, it seems natural in our context to use an approximation that takes advantage of the timing of migration decisions. We have assumed that information on the value of human capital in alternative locations is permanent, so if a location has been visited previously, the wage in that location is known. This means that the number of possible states increases geometrically with the number of locations. In practice, however, the number of people seen in many distinct locations is small. Thus by restricting the information set to include only wages seen in recent locations, it is possible to drastically shrink the state space while retaining most of the information actually seen in the data. Specifically, we suppose that the number

⁸And it will remain so: for example, if there are 50 locations, and the wage distribution has 5 support points, then the number of dynamic programming states is

40,414,063,873,238,203,032,156,980,022,826,814,668,800.

of wage observations cannot exceed M , with $M < J$, so that it is not possible to be fully informed about wages at all locations. Then if the distributions of location match wage and preference components in each of J locations have n points of support, the number of states for someone seen in M locations is $J(Jn^2)^M$, the number of possible M -period histories describing the locations visited most recently, and the wage and preference components found there. For example, if $J = 50$, $n = 3$, and $M = 2$, the number of states at each age is 10,125,000, which is manageable.

This approximation reduces the number of states in the most obvious way: we simply delete most of them.⁹ Someone who has “too much” wage information in the big state space is reassigned to a less-informed state. Individuals make the same calculations as before when deciding what to do next, and the econometrician uses the same procedure to recover the parameters governing the individual’s decisions. There is just a shorter list of states, so people with different histories may be in different states in the big model, but they are considered to be in the same state in the reduced model. In particular, people who have the same recent history are in the same state, even if their previous histories were different.

Decision problems with large state spaces can alternatively be analyzed by computing the value function at a finite set of points, and interpolating the function for points outside this set, as suggested by [Keane and Wolpin \(1994\)](#).¹⁰ In our context this would not be feasible without some simplification of the state space, because of the spatial structure of the states. Since each location has its own unique characteristics, interpolation can be done only within locations, and this means that the set of points used to anchor the interpolation must include several alternative realizations of the location match components for each location; allowing for n alternatives yields a set of n^J points, which is too big when $J = 50$ (even if n is small). On the other hand, it is worth noting that our limited history approximation works only because we have discretized the state space. If the location match components are drawn from continuous distributions, the state space is still infinite, even when the history is limited (although interpolation methods could be used in that case).

⁹Note that it is not enough to keep track of the best wage found so far: the payoff shocks may favor a location that has previously been abandoned, and it is necessary to know the wage at that location so as to decide whether to go back there (even if it is known that there is a higher wage at another location).

¹⁰For example, this method was used by [Erdem and Keane \(1996\)](#) to analyze the demand for liquid laundry detergent, and by [Crawford and Shum \(2005\)](#) to analyze the demand for pharmaceuticals. In these applications, the agents in the model do not know the flow payoffs from the various available choices until they have tried them, just as our agents do not know the location match components until they have visited the location.

4.2. *Wages*

The wage of individual i in location j at age a in year t is specified as

$$w_{ij}(a) = \mu_j + v_{ij} + G(X_i, a, t) + \eta_i + \varepsilon_{ij}(a),$$

where μ_j is the mean wage in location j , v is a permanent location match effect, $G(X, a, t)$ represents a (linear) time effect and the effects of observed individual characteristics, η is an individual effect that is fixed across locations, and ε is a transient effect. We assume that η , v , and ε are independent random variables that are identically distributed across individuals and locations. We also assume that the realizations of η and v are seen by the individual.¹¹

The relationship between wages and migration decisions is governed by the difference between the quality of the match in the current location, measured by $\mu_j + v_{ij}$, and the prospect of obtaining a better match in another location k , measured by $\mu_k + v_{ik}$. The other components of wages have no bearing on migration decisions, since they are added to the wage in the same way no matter what decisions are made. The individual knows the realization of the match quality in the current location and in the previous location (if there is one), but the prospects in other locations are random. Migration decisions are made by comparing the expected continuation value of staying, given the current match quality, with the expected continuation values associated with moving.

4.3. *State Variables and Flow Payoffs*

Let $\ell = (\ell^0, \ell^1, \dots, \ell^{M-1})$ be an M vector containing the sequence of recent locations (beginning with the current location), and let ω be an M vector recording wage and utility information at these locations. The state vector x consists of ℓ , ω , and age. The flow payoff for someone whose "home" location is h is specified as

$$\tilde{u}_h(x, j) = u_h(x, j) + \zeta_j,$$

where

$$u_h(x, j) = \alpha_0 w(\ell^0, \omega) + \sum_{k=1}^K \alpha_k Y_k(\ell^0) + \alpha^H \chi(\ell^0 = h) \\ + \xi(\ell^0, \omega) - \Delta_\tau(x, j).$$

¹¹An interesting extension of the model would allow for learning, by relaxing the assumption that agents know the realizations of η and v . In particular, such an extension might help explain return migration, because moving reveals information about the wage components. Pessino (1991) analyzed a two-period Bayesian learning model along these lines and applied it to migration data for Peru.

Here the first term refers to wage income in the current location. This is augmented by the nonpecuniary variables $Y_k(\ell^0)$, representing amenity values. The parameter α^{fi} is a premium that allows each individual to have a preference for their native location (χ_A denotes an indicator meaning that A is true). The flow payoff in each location has a random permanent component ξ ; the realization of this component is learned only when the location is visited. This location match component of preferences is analogous to the match component of wages (v), except that ξ can only be inferred from observed migration choices, whereas both migration choices and wages are informative about v . The cost of moving from ℓ^0 to ℓ^j for a person of type τ is represented by $\Delta_\tau(x, j)$. The unexplained part of the utility flow, ζ_j , may be viewed as either a preference shock or a shock to the cost of moving, with no way to distinguish between the two.

4.4. Moving Costs

Let $D(\ell^0, j)$ be the distance from the current location to location j and let $\mathbb{A}(\ell^0)$ be the set of locations adjacent to ℓ^0 (where states are adjacent if they share a border). The moving cost is specified as

$$\begin{aligned} \Delta_\tau(x, j) = & (\gamma_{0\tau} + \gamma_1 D(\ell^0, j) - \gamma_2 \chi(j \in \mathbb{A}(\ell^0)) \\ & - \gamma_3 \chi(j = \ell^1) + \gamma_4 a - \gamma_5 n_j) \chi(j \neq \ell^0). \end{aligned}$$

We allow for unobserved heterogeneity in the cost of moving: there are several types, indexed by τ , with differing values of the intercept γ_0 . In particular, there may be a “stayer” type, meaning that there may be people who regard the cost of moving as prohibitive, in all states. The moving cost is an affine function of distance (which we measure as the great circle distance between population centroids). Moves to an adjacent location may be less costly (because it is possible to change states while remaining in the same general area). A move to a previous location may also be less costly, relative to moving to a new location. In addition, the cost of moving is allowed to depend on age, a . Finally, we allow for the possibility that it is cheaper to move to a large location, as measured by population size n_j . It has long been recognized that location size matters in migration models (see, e.g., [Schultz \(1982\)](#)). California and Wyoming cannot reasonably be regarded as just two alternative places, to be treated symmetrically as origin and destination locations. For example, a person who moves to be close to a friend or relative is more likely to have friends or relatives in California than in Wyoming. One way to model this in our framework is to allow for more than one draw from the distribution of payoff shocks in each location.¹² Alternatively, location size may affect moving costs; for example,

¹²Suppose that the number of draws per location is an affine function of the number of people already in that location and that migration decisions are controlled by the maximal draw for

friends or relatives might help reduce the cost of the move. In practice, both versions give similar results.

4.5. Transition Probabilities

The state vector can be written as $x = (\tilde{x}, a)$, where $\tilde{x} = (\ell^0, \ell^1, x_v^0, x_v^1, x_\xi^0, x_\xi^1)$ and where x_v^0 indexes the realization of the location match component of wages in the current location, and similarly for the other components. The transition probabilities are

$$p(x' | x, j) = \begin{cases} 1, & \text{if } j = \ell^0, \tilde{x}' = \tilde{x}, a' = a + 1, \\ 1, & \text{if } j = \ell^1, \tilde{x}' = (\ell^1, \ell^0, x_v^1, x_v^0, x_\xi^1, x_\xi^0), a' = a + 1, \\ \frac{1}{n^2}, & \text{if } j \notin \{\ell^0, \ell^1\}, \tilde{x}' = (j, \ell^0, s_v, x_v^0, s_\xi, x_\xi^0), \\ & (1, 1) \leq (s_v, s_\xi) \leq (n_v, n_\xi), a' = a + 1, \\ 0, & \text{otherwise.} \end{cases}$$

This covers several cases. First, if no migration occurs this period, then the state remains the same except for the age component. If there is a move to a previous location, the current and previous locations are interchanged, and if there is a move to a new location, the current location becomes the previous location and the new location match components are drawn at random. In all cases, age is incremented by one period.

4.6. Data

Our primary data source is the National Longitudinal Survey of Youth 1979 Cohort (NLSY79); we also use data from the 1990 Census. The NLSY79 conducted annual interviews from 1979 through 1994, and subsequently changed to a biennial schedule. The location of each respondent is recorded at the date of each interview and we measure migration by the change in location from one interview to the next. We use information from 1979 to 1994 so as to avoid the complications arising from the change in the frequency of interviews.

To obtain a relatively homogeneous sample, we consider only white non-Hispanic high-school graduates with no post-secondary education, using only

each location. This leads to the following modification of the logit function describing choice probabilities:

$$\rho(x, j) = \frac{\xi_j}{\sum_{k=1} \xi_k}; \quad \xi_k = (1 + \psi n_k) \exp(v_k(\ell, \omega)).$$

Here n_j is the population in location j , and ψ can be interpreted as the number of additional draws per person.

the years after schooling is completed.¹³ Appendix A describes our selection procedures. The NLSY oversamples people whose parents were poor, and one might expect that the income process for such people is atypical and that the effect of income on migration decisions might also be atypical. Thus we use only the “cross-section” subsample, with the poverty subsample excluded. The sample includes only people who completed high school by age 20 and who never enrolled in college. We exclude those who ever served in the military and also those who report being out of the labor force for more than 1 year after age 20. We follow each respondent from age 20 to the 1994 interview or the first year in which some relevant information is missing or inconsistent.

Our analysis sample contains 432 people, with continuous histories from age 20 comprising 4,274 person-years. There are 124 interstate moves (2.9 percent per annum).

In each round of the NLSY79, respondents report income for the most recent calendar year. Wages are measured as total wage and salary income, plus farm and business income, adjusted for cost of living differences across states (using the American Chamber of Commerce Research Association (ACCRA) Cost of Living Index). We exclude observations with positive hours or weeks worked and zero income.

We use information from the Public Use Micro Sample (PUMS) from the 1990 Census to estimate state mean effects (μ_j), since the NLSY does not have enough observations for this purpose. From the PUMS we select white high-school men aged 19–20 (so as to avoid selection effects due to migration).¹⁴ We estimate state mean wage effects using a median regression with age and state dummies.¹⁵ We condition on these estimated state means in the maximum likelihood procedure that jointly estimates the remaining parameters of the wage process, and the utility and cost parameters governing migration decisions.¹⁶

¹³Attrition in panel data is an obvious problem for migration studies, and one reason for using NLSY data is that it minimizes this problem. Reagan and Olsen (2000, p. 339) reported that “Attrition rates in the NLSY79 are relatively low... The primary reason for attrition are death and refusal to continue participating in the project, not the inability to locate respondents at home or abroad.”

¹⁴The parameters governing migration decisions and the parameters of the wage process are estimated jointly to account for selection effects due to migration (although in practice these effects are empirically negligible). The state mean effects are specified as age-invariant and are estimated using wages observed at the beginning of the work life to minimize the potential effects of selection. We include observations for 19-year-olds from the PUMS to increase the precision of the estimated state means.

¹⁵We measure wages as annual earnings and exclude individuals with retirement income, social security income, or public assistance; we also exclude observations if earnings are zero despite positive hours or weeks worked.

¹⁶See Kennan and Walker (2011) for a detailed description of the estimation procedure.

5. ESTIMATION

In this section we discuss the specification and computation of the likelihood function.

5.1. *Discrete Approximation of the Distribution of Location Match Effects*

We approximate the decision problem by using discrete distributions to represent the distributions of the location match components, and computing continuation values at the support points of these distributions. We first describe this approximation and then describe the specification of the other components of wages.

For given support points, the best discrete approximation \hat{F} for any distribution F assigns probabilities so as to equate \hat{F} with the average value of F over each interval where \hat{F} is constant. If the support points are variable, they are chosen so that \hat{F} assigns equal probability to each point.¹⁷ Thus if the distribution of the location match component v were known, the wage prospects associated with a move to state k could be represented by an n -point distribution with equally weighted support points $\hat{\mu}_k + \hat{v}(q_r)$, $1 \leq r \leq n$, where $\hat{v}(q_r)$ is the q_r quantile of the distribution of v , with

$$q_r = \frac{2r - 1}{2n}$$

for $1 \leq r \leq n$. The distribution of v is, in fact, not known, but we assume that it is symmetric around zero. Thus, for example, with $n = 3$, the distribution of $\mu_j + v_{ij}$ in each state is approximated by a distribution that puts mass $\frac{1}{3}$ on μ_j (the median of the distribution of $\mu_j + v_{ij}$), with mass $\frac{1}{3}$ on $\mu_j \pm \tau_v$, where τ_v is a parameter to be estimated. The location match component of preferences is handled in a similar way.

5.2. *Fixed Effects and Transient Wage Components*

Even though our sample is quite homogeneous, measured earnings in the NLSY are highly variable, both across people and over time. Moreover, the variability of earnings over time is itself quite variable across individuals. Our aim is to specify a wage components model that is flexible enough to fit these data, so that we can draw reasonable inferences about the relationship between measured earnings and the realized values of the location match component. For the fixed effect η , we use a (uniform) discrete distribution that is symmetric around zero, with seven points of support, so that there are three parameters

¹⁷See Kennan (2006).

to be estimated. For the transient component ε , we need a continuous distribution that is flexible enough to account for the observed variability of earnings. We assume that ε is drawn from a normal distribution with zero mean for each person, but we allow the variance to vary across people. Specifically, person i initially draws $\sigma_\varepsilon(i)$ from some distribution, and subsequently draws ε_{it} from a normal distribution with mean zero and standard deviation $\sigma_\varepsilon(i)$, with ε_{it} drawn independently in each period. The distribution from which σ_ε is drawn is specified as a (uniform) discrete distribution with four support points, where these support points are parameters to be estimated.

5.3. The Likelihood Function

The likelihood of the observed history for each individual is a mixture over heterogeneous types. Let $L_i(\theta_\tau)$ be the likelihood for individual i , where θ_τ is the parameter vector, for someone of type τ and let π_τ be the probability of type τ . The sample log likelihood is

$$\Lambda(\theta) = \sum_{i=1}^N \log \left(\sum_{\tau=1}^K \pi_\tau L_i(\theta_\tau) \right).$$

For each period of an individual history, two pieces of information contribute to the likelihood: the observed income and the location choice. Each piece involves a mixture over the possible realizations of the various unobserved components. In each location there is a draw from the distribution of location match wage components, which is modeled as a uniform distribution over the finite set $Y = \{v(1), v(2), \dots, v(n_v)\}$. We index this set by ω_v , with $\omega_v(j)$ representing the match component in location j , where $1 \leq \omega_v(j) \leq n_v$. Similarly, in each location there is a draw from the location match preference distribution, which is modeled as a uniform distribution over the finite set $\Xi = \{\xi(1), \xi(2), \dots, \xi(n_\xi)\}$, indexed by ω_ξ . Each individual also draws from the distribution of fixed effects, which is modeled as a uniform distribution over the finite set $H = \{\eta(1), \eta(2), \dots, \eta(n_\eta)\}$, and we use ω_η to represent the outcome of this. Additionally, each individual draws a transient variance from a uniform distribution over the set $\varsigma = \{\sigma_\varepsilon(1), \sigma_\varepsilon(2), \dots, \sigma_\varepsilon(n_\varepsilon)\}$, with the outcome indexed by ω_ε .

The unobserved components of wages and preferences for individual i are then represented by a vector ω^i with $N_i + 3$ elements, $\omega^i = \{\omega_\xi^i, \omega_\eta^i, \omega_\varepsilon^i, \omega_v^i(1), \omega_v^i(2), \dots, \omega_v^i(n_\eta)\}$, where N_i is the number of locations visited by this individual. The set of possible realizations of ω^i is denoted by $\Omega(N_i)$; there are $n_\xi n_\eta n_\varepsilon (n_v)^{N_i}$ points in this set and our discrete approximation implies that they are equally likely. We index the locations visited by individual i in the order in which they appear, and we use the notation κ_{it}^0 and κ_{it}^1 to represent the position of the current and previous locations in this index. Thus $\kappa_{it} = (\kappa_{it}^0, \kappa_{it}^1)$ is a pair of integers between 1 and N_i . For example, in the case of someone who never

moves, κ_{it}^0 is always 1 and κ_{it}^1 is 0 (by convention), while for someone who has just moved for the first time, $\kappa_{it} = (2, 1)$.

The likelihood is obtained by first conditioning on the realizations of ω^i and then integrating over these realizations. Let $\psi_{it}(\omega_i, \theta)$ be the likelihood of the observed income for person i in period t . Given ω^i , the transient income component in period t is given by

$$\varepsilon_{it}(\omega^i) = w_{it} - \mu_{\ell^0(i,t)} - G(X_i, a_{it}, \theta_\tau) - v(\omega_v^i(\kappa_{it}^0)) - \eta(\omega_\eta^i).$$

Thus

$$\begin{aligned} \psi_{it}(\omega_i, \theta_\tau) \\ = \frac{1}{\sigma_\varepsilon(\omega_\varepsilon^i)} \phi\left(\frac{w_{it} - \mu_{\ell^0(i,t)} - G(X_i, a_{it}, \theta_\tau) - v(\omega_v^i(\kappa_{it}^0)) - \eta(\omega_\eta^i)}{\sigma_\varepsilon(\omega_\varepsilon^i)}\right), \end{aligned}$$

where ϕ is the standard normal density function.

Let $\lambda_{it}(\omega^i, \theta_\tau)$ be the likelihood of the destination chosen by person i in period t . Recall that $\rho(x, j)$ is the probability of choosing location j when the state is x . Then

$$\begin{aligned} \lambda_{it}(\omega^i, \theta_\tau) = \rho_{h(i)}(\ell(i, t), \omega_v^i(\kappa_{it}^0), \omega_v^i(\kappa_{it}^1), \omega_\xi^i(\kappa_{it}^0), \\ \omega_\xi^i(\kappa_{it}^1), a_{it}, \ell^0(i, t+1), \theta_\tau). \end{aligned}$$

Here the probability that i chooses the next observed location, $\ell^0(i, t+1)$, depends on the current and previous locations, the values of the location match components at those locations, the individual's home location $h(i)$, and the individual's current age. The parameter vector θ_τ includes the unknown coefficients in the flow payoff function and the support points in the sets Y , H , Ξ , and ς .

Finally, the likelihood of an individual history for a person of type τ is

$$\begin{aligned} L_i(\theta_\tau) = \frac{1}{n_\eta n_\varepsilon n_\xi (n_v)^{N_i}} \\ \times \sum_{\omega^i \in \Omega(N_i)} \left(\prod_{t=1}^{T_i} \psi_{it}(\omega^i, \theta_\tau) \lambda_{it}(\omega^i, \theta_\tau) \right). \end{aligned}$$

5.4. Identification

The relationship between income and migration decisions in our model can be identified using the variation in mean wages across locations or by using the variation in the location match component of wages. We assume that the

wage components ($\eta_i, v_{ij}, \varepsilon_{ijt}$) and the location match component of preferences ξ_{ij} are all independently and identically distributed across individuals and states, and that ε_{ijt} is i.i.d. over time. Alternatively, we can allow for unobserved amenities, represented by a component of ξ that is common to all individuals and that may be correlated with the state mean wages (μ_j). In this case identification relies on variation in v . We also assume that the unobserved heterogeneity in moving costs is i.i.d. across individuals, and that it is independent of the wage components and of the preference components.

Our basic empirical results use variation in both μ and v to identify the effect of income differences. In the context of an equilibrium model of wage determination, this can be justified by assuming constant returns to labor in each location, so that wage differences across locations are determined entirely by productivity differences and are thus independent of differences in amenity values. Clearly, this is a strong assumption. Accordingly (in Section 6.6 below), we also present estimates that control for regional differences in unobserved amenity values.

5.4.1. *Nonparametric Identification of the Choice Probabilities*

If the match component of wages could be observed directly, it would be relatively straightforward to identify the effect of wages on migration decisions. But the observed wages include individual fixed effects and transient effects, so that the match component is observed only with error. In addition, there is selection bias in the match component of the observed wages, since an unfavorable draw from the v distribution is more likely to be discarded (because it increases the probability of migration). The maximum likelihood procedure deals with the measurement error problem by integrating over the distributions of η and ε , and it deals with the selection problem by maximizing the joint likelihood of the wage components and the migration decisions. But this of course rests on specific parametric assumptions, and even then it does not give a transparent description of how identification is achieved. Thus it is useful to consider identification in a broader context.¹⁸


The basic identification argument can be well illustrated in a simplified situation in which there are just two observations for each person. Define the wage residual for individual i in period t in location $j(t)$ as

$$y_{it} = w_{it} - \mu_j - G(X_i, a, t) = \eta_i + v_{ij(t)} + \varepsilon_{it}.$$

Recall that the wage components (η, v, ε) are assumed to be independent, with zero means. The probability of moving (in the first period) depends on the

¹⁸Identification of dynamic discrete choice models was analyzed by Magnac and Thesmar (2002), and by Abbring and Heckman (2007); identification of static equilibrium discrete choice models was analyzed by Berry and Haile (2010).

location match component: denote this probability by $\rho(v)$. The process that generates the wage and migration data can then be represented as



The question is whether it is possible to recover the function $\rho(v)$ from these data.

For movers, we have two observations on the fixed effect η , contaminated by errors drawn from distinct distributions. Let \tilde{v}^m denote the censored random variable derived from v by discarding the realizations of v for those who choose to stay and let \tilde{v}^s denote the corresponding censored random variable derived by discarding the realizations of v for those who choose to move. The observed wages for movers are $y_1 = \eta + \tilde{v}^m + \varepsilon_1$ and $y_2 = \eta + v' + \varepsilon_2$, where v' is a random draw from the v distribution (which is independent of \tilde{v}^m). Then, under the regularity condition that the characteristic function of the random vector (y_1, y_2) is nonvanishing, Lemma 1 of Kotlarski (1967) implies that the (observed) distribution of (y_1, y_2) for movers identifies the distributions of η , $\tilde{v}^m + \varepsilon_1$, and $v + \varepsilon_2$.

For stayers, we have two observations on $\eta + \tilde{v}^s$, with measurement errors ε_1 and ε_2 . Thus Kotlarski's lemma implies that the distribution of (y_1, y_2) for stayers identifies the distributions of $\eta + \tilde{v}^s$, ε_1 , and ε_2 . This means that the distributions of η , v , ε_1 , ε_2 , \tilde{v}^m , and \tilde{v}^s are all identified (either directly or by deconvolution).

The choice probabilities $\rho(v)$ are then identified by Bayes theorem,

$$f_{\tilde{v}^m}(v) = \frac{\rho(v)f_v(v)}{\text{Prob}(j(2) \neq j(1))},$$

where $f_{\tilde{v}^m}$ and f_v are the conditional and unconditional density functions, and $j(2) \neq j(1)$ indicates a move.

This argument shows that the effect of income on migration decisions is generically identified under our assumptions. In particular, under the null hypothesis that migration has nothing to do with income, the choice probability function $\rho(v)$ is a constant, while the expected income maximization model

predicts that $\rho(v)$ is an increasing function. Since the shape of this function is identified in the data, the effect of income is identified nonparametrically. This is true even if there is unobserved heterogeneity with respect to moving costs and the location match component of preferences. Moreover, the identification argument does not rely on a particular distribution of the payoff shocks. What is identified is the average relationship between income and migration decisions, after integrating over the distributions of preference shocks and moving costs.

5.4.2. *Identification of the Income Coefficient*

We have shown that the relationship between wages and migration choice probabilities is nonparametrically identified, given panel data on wage and migration outcomes, under the assumption that the wage is a sum of independent components. Although in general this relationship might be quite complicated, in our parametric model it is encapsulated in the income coefficient α_0 . We now illustrate how this parameter is identified in our model, using an argument along the lines of Hotz and Miller (1993).

Fix home location and age, with no previous location. Assume that there is no unobserved heterogeneity in moving costs and that there is no location match component of preferences. Then the state consists of the current location and the location match component of wages, and the choice probabilities corresponding to the function $\rho(v)$ in the nonparametric argument above are given by

$$\rho(\ell, v_s, j) = \begin{cases} \frac{\exp(-\Delta_{\ell j} + \beta \bar{V}_0(j))}{\exp(\beta \bar{V}_s(\ell)) + \sum_{k \neq \ell} \exp(-\Delta_{\ell k} + \beta \bar{V}_0(k))}, & j \neq \ell, \\ \frac{\exp(\beta \bar{V}_s(\ell))}{\exp(\beta \bar{V}_s(\ell)) + \sum_{k \neq \ell} \exp(-\Delta_{\ell k} + \beta \bar{V}_0(k))}, & j = \ell, \end{cases}$$

where $\Delta_{\ell j}$ is the cost of moving from location ℓ to location j , $\bar{V}_s(j)$ is the expected continuation value in j , given the location match component v_s before knowing the realization of ζ , and $\bar{V}_0(j)$ is the expected continuation value before knowing the realization of v :

$$\bar{V}_0(j) = \frac{1}{n} \sum_{s=1}^n \bar{V}_s(j).$$

The probability of moving from ℓ to j , relative to the probability of staying, is

$$\frac{\rho(\ell, v_s, j)}{\rho(\ell, v_s, \ell)} = \exp(-\Delta_{\ell j} + \beta(\bar{V}_0(j) - \bar{V}_s(\ell))).$$

Thus

$$\frac{1}{n} \sum_{s=1}^n \log\left(\frac{\rho(\ell, v_s, j)}{\rho(\ell, v_s, \ell)}\right) = -\Delta_{\ell j} + \beta(\bar{V}_0(j) - \bar{V}_0(\ell))$$

and

$$\frac{1}{n} \sum_{s=1}^n \log\left(\frac{\rho(\ell, v_s, j) \rho(j, v_s, \ell)}{\rho(\ell, v_s, \ell) \rho(j, v_s, j)}\right) = -\Delta_{\ell j} - \Delta_{j\ell}.$$

This identifies the round-trip moving cost between ℓ and j .

The one-way moving costs are identified under weak assumptions on the moving cost function; for example, symmetry is obviously sufficient. In the model, the round-trip moving cost between two nonadjacent locations (for someone aged a with no previous location) is given by

$$\Delta_{\ell j} + \Delta_{j\ell} = 2(\gamma_0 + \gamma_4 a + \gamma_1 D(j, \ell)) - \gamma_5(n_j + n_\ell).$$

Since distance and population vary independently, one can choose three distinct location pairs, such that the three moving cost equations are linearly independent; these equations identify γ_1 , γ_5 , and $\gamma_0 + \gamma_4 a$. Then by choosing two different ages, γ_0 and γ_4 are identified, and by comparing adjacent and nonadjacent pairs, γ_2 is identified.

If the continuation value in all states is increased by the same amount, then the choice probabilities are unaffected, so one of the values can be normalized to zero.¹⁹ We assume $\bar{V}_0(J) = 0$. Then

$$\frac{1}{n} \sum_{s=1}^n \log\left(\frac{\rho(\ell, v_s, J)}{\rho(\ell, v_s, \ell)}\right) = -\Delta_{\ell j} - \beta\bar{V}_0(\ell).$$

This identifies $\bar{V}_0(\ell)$, since we assume that β is known, and $\Delta_{\ell j}$ has already been identified. Once \bar{V}_0 is identified, $\bar{V}_s(\ell)$ is identified by the equation

$$\log\left(\frac{\rho(\ell, v_s, j)}{\rho(\ell, v_s, \ell)}\right) = -\Delta_{\ell j} + \beta(\bar{V}_0(j) - \bar{V}_s(\ell)).$$

¹⁹It might seem that the choice probabilities are also invariant to a rescaling of the continuation values, but we have already normalized the scale by assuming additive payoff shocks drawn from the extreme value distribution.

Given that the expected continuation values in all states are identified, the flow payoffs are identified by

$$\bar{V}_s(\ell) = \bar{\gamma} + \alpha_0 v_s + A_\ell + \log\left(\exp(\beta \bar{V}_s(\ell)) + \sum_{k \neq \ell} \exp(\Delta_{\ell k} + \beta \bar{V}_0(k))\right),$$

where A_ℓ represents amenity values and other fixed characteristics of location ℓ (both observed and unobserved), and where v_s represents the location match component of wages. The income coefficient α_0 is identified by differencing this equation with respect to s (thereby eliminating A_ℓ), and A_ℓ is then identified as the only remaining unknown in the equation.

5.5. Computation

Since the parameters are embedded in the value function, computation of the gradient and hessian of the log-likelihood function is not a simple matter (although, in principle, these derivatives can be computed using the same iterative procedure that computes the value function itself). We maximize the likelihood using a version of Newton's algorithm with numerical derivatives. We also use the downhill simplex method of Nelder and Mead, mainly to check for local maxima. This method does not use derivatives, but it is slow.²⁰

6. EMPIRICAL RESULTS

Our basic results are shown in Table II. We set $\beta = 0.95$, $T = 40$, and $M = 2$; we show below that our main results are not very sensitive to changes in the discount factor or the horizon length.²¹ The table gives estimated coefficients and standard errors for four versions of the model that highlight both the effect of income on migration decisions and the relevance of the location match component of preferences. Unobserved heterogeneity in moving costs is introduced by allowing for two types, one of which is a pure stayer type (representing peo-

²⁰Given reasonable starting values (for example, 50% type probabilities and a fixed cost for the mover type that roughly matches the average migration rate, with variances near 1 for the wage components and all other parameters set to zero), the maximal likelihood is reached, using Nelder–Mead iterations followed by a switch to Newton's method, in less than a day on a cluster of parallel CPUs, with one CPU per home location; each likelihood evaluation requires about 9 seconds. We found the Newton procedure to be well behaved in the sense that it almost always reached the same answer, no matter what starting values were used: we have estimated hundreds of different versions of the model and found very few local maxima; even in these cases, the likelihood and the parameter values were very close to the "true" maximum.

²¹The validity of the estimates is checked in Appendix B: the estimated coefficients were used to simulate 100 replicas of each person in the data and the maximum likelihood procedure was applied to the simulated data. The null hypothesis that the data were generated by the true data generating process (DGP) is accepted by a likelihood ratio test.

TABLE II
 INTERSTATE MIGRATION, YOUNG WHITE MEN^a

	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$
<i>Utility and cost</i>								
Disutility of moving (γ_0)	4.790	0.565	4.514	0.523	4.864	0.601	4.851	0.604
Distance (γ_1) (1000 miles)	0.265	0.182	0.280	0.178	0.311	0.187	0.270	0.184
Adjacent location (γ_2)	0.808	0.214	0.787	0.211	0.773	0.220	0.804	0.216
Home premium (α^H)	0.331	0.041	0.267	0.031	0.332	0.047	0.337	0.045
Previous location (γ_3)	2.757	0.356	2.544	0.300	3.082	0.449	2.818	0.416
Age (γ_4)	0.055	0.020	0.062	0.019	0.060	0.020	0.054	0.020
Population (γ_5) (millions)	0.653	0.179	0.653	0.178	0.635	0.177	0.652	0.179
Stayer probability	0.510	0.078	0.520	0.079	0.495	0.087	0.508	0.082
Cooling (α_1) (1000 degree-days)	0.055	0.019	0.036	0.019	0.048	0.018	0.056	0.019
Income (α_0)	0.312	0.100	–	–	–	–	0.297	0.116
Location match preference (τ_{ξ})	–	–	–	–	0.168	0.049	0.070	0.099
<i>Wages</i>								
Wage intercept	–5.165	0.244	–5.175	0.246	–5.175	0.246	–5.168	0.244
Time trend	–0.035	0.008	–0.033	0.008	–0.033	0.008	–0.035	0.008
Age effect (linear)	7.865	0.354	7.876	0.356	7.877	0.356	7.870	0.355
Age effect (quadratic)	–2.364	0.129	–2.381	0.130	–2.381	0.130	–2.367	0.129
Ability (AFQT)	0.014	0.065	0.015	0.066	0.014	0.066	0.014	0.065
Interaction (Age, AFQT)	0.147	0.040	0.152	0.040	0.152	0.040	0.147	0.040
Transient s.d. 1	0.217	0.007	0.218	0.007	0.218	0.007	0.217	0.007
Transient s.d. 2	0.375	0.015	0.375	0.015	0.375	0.015	0.375	0.015
Transient s.d. 3	0.546	0.017	0.547	0.017	0.547	0.017	0.546	0.017
Transient s.d. 4	1.306	0.028	1.307	0.028	1.307	0.028	1.306	0.028
Fixed effect 1	0.113	0.035	0.112	0.035	0.112	0.035	0.113	0.035
Fixed effect 2	0.298	0.035	0.296	0.035	0.296	0.035	0.298	0.035
Fixed effect 3	0.936	0.017	0.934	0.017	0.934	0.017	0.936	0.017
Wage match (τ_v)	0.384	0.017	0.387	0.018	0.387	0.018	0.384	0.018
<i>Log likelihood</i>								
Exclude income: $\chi^2(1)$	–4,214.880		–4,221.426		–4,218.800		–4,214.834	
				13.09		7.93		
Exclude match preference: $\chi^2(1)$		0.09		5.25				

^aThere are 4,274 (person-year) observations, 432 individuals, and 124 moves.

ple with prohibitive moving costs); little is gained by introducing additional types or by replacing the stayer type with a type with a high moving cost.

We find that distance, home and previous locations, and population size all have highly significant effects on migration. Age and local climate (represented by the annual number of cooling degree-days) are also significant.²² Our main

²²The “cooling” variable is the population-weighted annual average number of cooling degree days (in thousands) for 1931–2000, taken from Historical Climatology Series 5-2 (Cooling Degree Days); see [United States National Climatic Data Center \(2002\)](#). For example, the cooling degree-day variable for Florida is 3.356, meaning that the difference between 65° and the mean daily temperature in Florida, summed over the days when the mean was above 65°, averaged 3,356 degree-days per year (over the years 1931–2000). We explored various alternative specifi-

TABLE III
WAGE PARAMETER ESTIMATES (IN 2010 DOLLARS)

	AFQT Percentile						
	25		50		75		
<i>Average wages</i>							
Age 20 in 1979	25,827		27,522		29,216		
Age 20 in 1989	18,472		20,166		21,861		
Age 30 in 1989	40,360		42,850		45,340		
	Low		Middle		High		
Location match	-8,366		0		8,366		
Fixed effect support	-20,411	-6,498	-2,454	0	2,454	6,498	20,411
State means	Low (WV)	Rank 5 (OK)	Median (MO)	Rank 45 (RI)	High (MD)		
	12,698	14,530	16,978	19,276	22,229		

finding is that, controlling for these effects, migration decisions are significantly affected by expected income changes. This holds regardless of whether the location match component of preferences is included in the specification. Since the estimated effect of this component is negligible and it enlarges the state space by a factor of about 100, we treat the specification that excludes this component as the base model in the subsequent discussion.

6.1. Wages

The estimated parameters of the wage process are summarized in Table III, showing the magnitudes of the various components in 2010 dollars. As was mentioned above, there is a great deal of unexplained variation in wages across people and over time for the same person; moreover, there are big differences in the variability of earnings over time from one individual to the next.²³

The wage components that are relevant for migration decisions in the model are also quite variable, suggesting that migration incentives are strong. For example, the 90–10 differential across state means is about \$4,700 a year, and the value of replacing a bad location match draw with a good draw is almost \$17,000 a year.

cations of the climate amenity variables. Including heating degree-days had little effect on the results (see Table X below). The number of states that are adjacent to an ocean is 23. We considered this as an additional amenity variable, and also estimated models including annual rainfall and the annual number of sunny days, but found that these variables had virtually no effect.

²³As indicated in Table II, the individual characteristics affecting wages include age, AFQT score, and an interaction between the two. The interaction effect is included to allow for the possibility that the relationship between AFQT scores and wages is stronger for older workers, either because ability and experience are complementary or because employers gradually learn about ability, as argued by Altonji and Pierret (2001).

TABLE IV
MOVING COST EXAMPLES

	γ_0	α_0	Age	Distance	Adjacent	Population	Previous	Cost
θ	4.790	0.312	0.055	0.265	0.808	0.653	2.757	
Young mover			20	1	0	1	0	\$384,743
Average mover			24.355	0.664	0.427	0.728	0.371	\$312,146

6.2. Moving Costs and Payoff Shocks

Since utility is linear in income, the estimated moving cost can be converted to a dollar equivalent. Some examples are given in Table IV. For the average mover, the cost is about \$312,000 (in 2010 dollars) if the payoff shocks are ignored. One might wonder why anyone would ever move in the face of such a cost and, in particular, whether a move motivated by expected income gains could ever pay for itself. According to the estimates in Table III, a move away from a bad location match would increase income by \$8,366, on average, and a move from the bottom to the top of the distribution of state means would increase income by \$9,531. A move that makes both of these changes would mean a permanent wage increase of \$17,897, or \$311,939 in present value (assuming a remaining work life of 40 years, with $\beta = 0.95$). The home premium is equivalent to a wage increase of \$23,106 and the cost of moving to a previous location is relatively low. Thus in some cases the expected income gains would be more than enough to pay for the estimated moving cost. Of course in most cases this would not be true, but then most people never move.

More importantly, the estimates in Table IV do not refer to the costs of moves that are actually made, but rather to the costs of hypothetical moves to arbitrary locations. In the model, people choose to move only when the payoff shocks are favorable, and the net cost of the move is, therefore, much less than the amounts in Table IV. Consider, for example, a case in which someone is forced to move, but allowed to choose the best alternative location. The expected value of the maximum of $J - 1$ draws from the extreme value distribution is $\bar{\gamma} + \log(J - 1)$ (where $\bar{\gamma}$ is the Euler constant), so if the location with the most favorable payoff shock is chosen, the expected net cost of the move is reduced by $\log(J - 1)/\alpha_0$. Using the estimated income coefficient, this is a reduction of \$271,959. Moreover, this calculation refers to a move made in an arbitrary period; in the model, the individual can move later if the current payoff shocks are unfavorable, so the net cost is further reduced. Of course people actually move only if there is, in fact, a net gain from moving; the point of the argument is just that this can quite easily happen, despite the large moving cost estimates in Table IV. In Section 6.3 below, we analyze the average costs of moves that are actually made, allowing for the effects of the payoff shocks.

Another way to interpret the moving cost is to consider the effect of a \$10,000 migration subsidy, payable for every move, with no obligation to stay

in the new location for more than one period. This can be analyzed by simulating the model with a reduction in γ_0 such that γ_0/α_0 falls by \$10,000, with the other parameters held fixed. We estimate that such a subsidy would lead to a substantial increase in the interstate migration rate: from 2.9% to about 4.9%.

6.2.1. Moving Costs and Payoff Shocks: An Example

To understand the relationship between moving costs and prospective income gains, it is helpful to consider an example in which these are the only influences on migration decisions. Suppose that income in each location is either high or low, the difference being Δy , and suppose that the realization of income in each location is known. Then the odds of moving are given by

$$(1) \quad \frac{1 - \lambda_L}{\lambda_L} = \exp(-\gamma_0)[J_L - 1 + J_H e^{\beta \Delta V}],$$

$$(2) \quad \frac{1 - \lambda_H}{\lambda_H} = \exp(-\gamma_0)[J_H - 1 + J_L e^{-\beta \Delta V}],$$

where λ_L is the probability of staying in one of J_L low-income locations (and similarly for λ_H and J_H), and where ΔV is the difference in expected continuation values between the low-income and high-income locations. This difference is determined by the equation

$$(3) \quad e^{\Delta V} = \frac{e^{\alpha_0 \Delta y} (J_L + (J_H - 1 + e^{\gamma_0}) e^{\beta \Delta V})}{J_L - 1 + e^{\gamma_0} + J_H e^{\beta \Delta V}}.$$

For example, if $\beta = 0$, then $\Delta V = \alpha_0 \Delta y$, while if moving costs are prohibitive ($e^{-\gamma_0} = 0$), then $\Delta V = \frac{\alpha_0 \Delta y}{1 - \beta}$.

These equations uniquely identify α_0 and γ_0 (these parameters are in fact overidentified, because there is also information in the probabilities of moving to the same income level).²⁴ If $\gamma_0 < \beta \Delta V$, then the odds of moving from a low-income location are greater than J_H to 1, and this is contrary to what is seen in the data (for any plausible value of J_H). By making γ_0 a little larger than $\beta \Delta V$ and letting both of these be large in relation to the payoff shocks, the probability of moving from the low-income location can be made small. But then the probability of moving from the high-income location is almost zero, which is not true in the data. In other words, if the probability of moving from a high-income location is not negligible, then the payoff shocks cannot be negligible, since a payoff shock is the only reason for making such a move.

²⁴It is assumed that λ_L , λ_H , J_L , J_H , Δy , and β are given. Dividing (1) by (2) and rearranging terms yields a quadratic equation in $e^{\beta \Delta V}$ that has one positive root and one negative root. Since $e^{\beta \Delta V}$ must be positive, this gives a unique solution for ΔV . Equation (1) then gives a unique solution for γ_0 , and inserting these solutions into equation (3) gives a unique solution for $\alpha_0 \Delta y$.

The net cost of moving from a low-income location to a high-income location is $\gamma_0 - \beta\Delta V$, while the net cost of the reverse move is $\gamma_0 + \beta\Delta V$. The difference is $2\beta\Delta V$, and equations (1) and (2) show that $\beta\Delta V$ determines the relative odds of moving from low-income and high-income locations. Thus $\beta\Delta V$ is identified by the difference between λ_L and λ_H ; this difference is small in the data, so $\beta\Delta V$ must be small. The magnitude of γ_0 is then determined by the level of λ_L and λ_H , and since these are close to 1 in the data, the implication is that γ_0 is large and that it is much larger than $\beta\Delta V$. Since $\beta\Delta V$ is roughly the present value of the difference in income levels, the upshot is that the moving cost must be large in relation to income.

For example, suppose $J_L = J_H = 25$, with $\beta = 0.95$. In our data, the migration probability for someone in the bottom quartile of the distribution of state mean wages is 5.5% (53 moves in 964 person-years); for someone in the top quartile, it is 2.1% (16 moves in 754 person-years). If $1 - \lambda_L = 53/964$ and $1 - \lambda_H = 16/754$, then $\gamma_0 = 7.34$, and $\Delta V = 1.02$, and the implied moving cost is $\gamma_0/\alpha_0 = 85.3\Delta y$. Taking Δy to be the difference in the mean wages for states in the top and bottom quartiles gives $\gamma_0/\alpha_0 = \$304,670$ (in 2008 dollars). On the other hand, if $\lambda_L = 0.7$, the implied moving cost is only $14.4\Delta y$, or $\$51,449$. We conclude that the moving cost estimate is large mainly because the empirical relationship between income levels and migration probabilities is relatively weak.

6.3. Average Costs of Actual Moves

Our estimates of the deterministic components of moving costs are large because moves are rare in the data. But moves do occur and, in many cases, there is no observable reason for a move, so the observed choice must be attributed to unobserved payoff shocks, including random variations in moving costs. Given this heterogeneity in moving costs, both across individuals and over time for the same individual, the question arises as to how large the actual moving costs are, conditional on a move being made.²⁵ Because the payoff shocks are drawn from the type I extreme value distribution, this question has a relatively simple answer.

The cost of a move may be defined as the difference in the flow payoff for the current period due to the move. Since a move to location j exchanges ζ_{ℓ^0} for ζ_j , the average cost of a move from ℓ^0 to j , given state x , is

$$\bar{\Delta}(x, j) = \Delta(x, j) - E(\zeta_j - \zeta_{\ell^0} \mid d_j = 1),$$

where d_j is an indicator variable for the choice of location j . Thus, for example, if a move from ℓ^0 to j is caused by a large payoff shock in location j , the cost

²⁵See Sweeting (2007) for a similar analysis of switching costs in the context of an empirical analysis of format switching by radio stations.

of the move may be much less than the amount given by the deterministic cost $\Delta(x, j)$.

In logit models, the expected gain from the optimal choice, relative to an arbitrary alternative that is not chosen, is a simple function of the probability of choosing the alternative (see [Anas and Feng \(1988\)](#) and [Kennan \(2008\)](#)). In the present context, this result means that the average increase in the gross continuation value for someone who chooses to move from ℓ^0 to j is given by

$$E(\tilde{v}(x, j) - \tilde{v}(x, \ell^0) \mid d_j = 1) = -\frac{\log(\rho(x, \ell^0))}{1 - \rho(x, \ell^0)},$$

where $\tilde{v}(x, j)$ is the continuation value when the state is x and location j is chosen, which includes the current flow payoff and the discounted expected continuation value in location j :

$$\begin{aligned} \tilde{v}(x, j) &= v(x, j) + \zeta_j \\ &= u(x, j) + \beta \sum_{x'} p(x' \mid x, j) \bar{V}(x') + \zeta_j. \end{aligned}$$

The deterministic part of the moving cost is

$$\begin{aligned} \Delta(x, j) &= u(x, \ell^0) - u(x, j) \\ &= v(x, \ell^0) - v(x, j) + \beta \sum_{x'} (p(x' \mid x, j) - p(x' \mid x, \ell^0)) \bar{V}(x') \\ &= \tilde{v}(x, \ell^0) - \zeta_{\ell^0} - \tilde{v}(x, j) + \zeta_j \\ &\quad + \beta \sum_{x'} (p(x' \mid x, j) - p(x' \mid x, \ell^0)) \bar{V}(x'). \end{aligned}$$

This implies that the average moving cost, net of the difference in payoff shocks, is

$$\begin{aligned} \bar{\Delta}(x, j) &\equiv \Delta(x, j) - E(\zeta_j - \zeta_{\ell^0} \mid d_j = 1) \\ &= \frac{\log(\rho(x, \ell^0))}{1 - \rho(x, \ell^0)} + \beta \sum_{x'} (p(x' \mid x, j) - p(x' \mid x, \ell^0)) \bar{v}(x'). \end{aligned}$$

Since some of the components of the state vector x are unobserved, we compute expected moving costs using the conditional distribution over the unobservables, given the observed wage and migration history. Recall that the likelihood of an individual history, for a person of type τ , is

$$L_i(\theta_\tau) = \frac{1}{n_\eta n_\varepsilon n_\xi (n_v)^{N_i}} \sum_{\omega^i \in \Omega(N_i)} \left(\prod_{t=1}^{T_i} \psi_{it}(\omega^i, \theta_\tau) \lambda_{it}(\omega^i, \theta_\tau) \right).$$

Thus the conditional probability of ω^i is

$$\begin{aligned}
 Q(\omega^i) &= \frac{\prod_{t=1}^{T_i} \psi_{it}(\omega^i, \theta_\tau) \lambda_{it}(\omega^i, \theta_\tau)}{\sum_{\omega \in \Omega(N_i)} \left(\prod_{t=1}^{T_i} \psi_{it}(\omega, \theta_\tau) \lambda_{it}(\omega, \theta_\tau) \right)} \\
 &= \frac{\prod_{t=1}^{T_i} \psi_{it}(\omega^i, \theta_\tau) \lambda_{it}(\omega^i, \theta_\tau)}{n_\eta n_\varepsilon n_\xi (n_v)^{N_i} L_i(\theta_\tau)}.
 \end{aligned}$$

The unobserved part of the state variable consists of the location match components of wages and preferences. Since the distributions of these components have finite support, there is a finite set $\Omega(N_i)$ of possible realizations corresponding to the observed history for individual i ; this set is indexed by ω^i . Let $x(\omega^i)$ be the state implied by ω^i (including the location match components in the current location, and in the previous location, if any). Then if individual i moves to location j in period t , the moving cost is estimated as

$$\hat{\Delta}_{it} = \sum_{\omega^i \in \Omega(N_i)} Q(\omega^i) \bar{\Delta}(x(\omega^i), j).$$

The estimated average moving costs are given in Table V. There is considerable variation in these costs, but for a typical move the cost is negative. The interpretation of this is that the typical move is not motivated by the prospect of a higher future utility flow in the destination location, but rather by unobserved

TABLE V
AVERAGE MOVING COSTS^a

Previous Location	Move Origin and Destination			
	From Home	To Home	Other	Total
None	-\$147,619 [56]	\$138,095 [1]	-\$39,677 [2]	-\$139,118 [59]
Home	-	\$18,686 [40]	-\$124,360 [10]	-\$9,924 [50]
Other	-\$150,110 [8]	\$113,447 [2]	-\$67,443 [5]	-\$87,413 [15]
Total	-\$147,930 [64]	\$25,871 [43]	-\$97,656 [17]	-\$80,768 [124]

^aThe number of moves in each category is given in brackets.

factors yielding a higher current payoff in the destination location, compared with the current location. That is, the most important part of the estimated moving cost is $\zeta_{\ell^0} - \zeta_j$, the difference in the payoff shocks. In the case of moves to the home location, on the other hand, the estimated cost is positive; most of these moves are return moves, but where the home location is not the previous location, the cost is large, reflecting a large gain in expected future payoffs due to the move.

6.4. Goodness of Fit

To keep the state space manageable, our model severely restricts the set of variables that are allowed to affect migration decisions. Examples of omitted observable variables include duration in the current location and the number of moves made previously. In addition, there are, of course, unobserved characteristics that might make some people more likely to move than others. Thus it is important to check how well the model fits the data. In particular, since the model pays little attention to individual histories, one might expect that it would have trouble fitting panel data.

One simple test of goodness of fit can be made by comparing the number of moves per person in the data with the number predicted by the model. As a benchmark, we consider a binomial distribution with a migration probability of 2.9% (the number of moves per person-year in the data). Table VI shows the predictions from this model: about 75% of the people never move; of those who do move, about 14% move more than once.²⁶ The NLSY data are quite different: about 84% never move and about 56% of movers move more than once. An obvious interpretation of this is mover–stayer heterogeneity: some people are more likely to move than others, and these people account for more than their share of the observed moves. We simulated the corresponding statistics for the model by starting 100 replicas of the NLSY individuals in the observed initial locations and using the model (with the estimated parameters

TABLE VI
GOODNESS OF FIT

Moves	Binomial		NLSY		Model	
None	325.1	75.3%	361	83.6%	36,257	83.9%
One	91.5	21.2%	31	7.2%	2,466	5.7%
More	15.4	3.6%	40	9.3%	4,478	10.4%
Movers with more than one move	14.4%		56.3%		64.5%	
Total observations	432		432		43,201	

²⁶Since we have an unbalanced panel, the binomial probabilities are weighted by the distribution of years per person.

TABLE VII
RETURN MIGRATION STATISTICS

Movers	NLSY	Model
<i>Proportion who</i>		
Return home	34.7%	37.5%
Return elsewhere	4.8%	6.2%
Move on	60.5%	61.9%
<i>Proportion who ever</i>		
Leave home	14.4%	13.7%
Move from not-home	40.0%	42.5%
Return from not-home	25.7%	32.3%

shown in Table II) to generate a history for each replica, covering the number of periods observed for this individual. The results show that the model does a good job of accounting for the heterogeneous migration probabilities in the data. The proportion of people who never move in the simulated data matches the proportion in the NLSY data almost exactly, and although the proportion of movers who move more than once is a bit high in the simulated data, the estimated model comes much closer to this statistic than the binomial model does.

6.4.1. Return Migration

Table VII summarizes the extent to which the model can reproduce the return migration patterns in the data (the statistics in the “Model” column refer to the simulated data set used in Table VI). The model attaches a premium to the home location and this helps explain why people return home. For example, in a model with no home premium, one would expect that the proportion of movers going to any particular location would be roughly 1/50, and this obviously does not match the observed return rate of 35%. The home premium also reduces the chance of initially leaving home, although this effect is offset by the substantial discount on the cost of returning to a previous location (including the home location): leaving home is less costly if a return move is relatively cheap.

The simulated return migration rates match the data reasonably well. The main discrepancy is that the model overpredicts the proportion who ever return home from an initial location that is not their home location. That is, the model has trouble explaining why people seem so attached to an initial location that is not their “home.” One potential explanation for this is that our assignment of home locations (the state of residence at age 14) is too crude; in some cases, the location at age 20 may be more like a home location than the location at age 14. More generally, people are no doubt more likely to put down roots the longer they stay in a location, and our model does not capture this kind of duration dependence.

TABLE VIII
ANNUAL MIGRATION RATES BY AGE AND CURRENT LOCATION^a

Age	All			Not at Home			At Home		
	<i>N</i>	Moves	Migration Rate	<i>N</i>	Moves	Migration Rate	<i>N</i>	Moves	Migration Rate
20–25	2,359	84	3.6%	244	40	16.4%	2,115	44	2.1%
26–34	1,915	40	2.1%	228	20	8.8%	1,687	20	1.2%
All	4,274	124	2.9%	472	60	13.4%	3,802	64	1.7%

^aAt home means living now in the state of residence at age 14.

6.5. *Why Are Younger People More Likely to Move?*

It is well known that the propensity to migrate falls with age (at least after age 25 or so). Table VIII replicates this finding for our sample of high-school men. A standard human capital explanation for this age effect is that migration is an investment: if a higher income stream is available elsewhere, then the sooner a move is made, the sooner the gain is realized. Moreover, since the work life is finite, a move that is worthwhile for a young worker might not be worthwhile for an older worker, since there is less time for the higher income stream to offset the moving cost (Sjaastad (1962)). In other words, migrants are more likely to be young for the same reason that students are more likely to be young.

Our model encompasses this simple human capital explanation of the age effect on migration.²⁷ There are two effects here. First, consider two locations paying different wages and suppose that workers are randomly assigned to these locations at birth. Then, even if the horizon is infinite, the model predicts that the probability of moving from the low-wage to the high-wage location is higher than the probability of a move in the other direction, so that eventually there will be more workers in the high-wage location. This implies that the (unconditional) migration rate is higher when workers are young.²⁸ Second, the human capital explanation says that migration rates decline with age because the horizon gets closer as workers get older. This is surely an important reason for the difference in migration rates between young adult workers and those

²⁷Investments in location-specific human capital might also help explain why older workers are less likely to move. Marriage might be included under this heading, for example, as in Gemici (2008). It is worth noting that if we take marital status as given, it has essentially no effect on migration in our sample in simple logit models of the move–stay decision that include age as an explanatory variable.

²⁸One way to see this is to consider the extreme case in which there are no payoff shocks. In this case, all workers born in the low-wage location will move to the high-wage location at the first opportunity (if the wage difference exceeds the moving cost), and the migration rate will be zero from then on.

within sight of retirement. But the workers in our sample are all in their 20s or early 30s, and the prospect of retirement seems unimportant for such workers.

We find that the simple human capital model does not fully explain the relationship between age and migration in the data. Our model includes age as a state variable to capture the effects just discussed. The model also allows for the possibility that age has a direct effect on the cost of migration; this can be regarded as a catchall for whatever is missing from the simple human capital explanation. The results in Table II show that this direct effect is significant.

6.6. *Decomposing the Effects of Income on Migration Decisions*

Migration is motivated by two distinct wage components in our model: differences in mean wages (μ_j) across locations and individual draws from the location match distribution (v_{ij}). The relevance of these components can be considered separately, first by suppressing the dispersion in v , so that wages affect migration decisions only because of differences in mean wages across locations, and alternatively by specifying the wage distribution at the national level, so that migration is motivated only by the prospect of getting a better draw from the same wage distribution (given our assumption that location match effects are permanent).

Consider an economy in which everyone has the same preferences over locations and also the same productivity in each location. In a steady state equilibrium, everyone is indifferent between locations: there are wage differences, but these just equalize the amenity differences. People move for other reasons, but there are just as many people coming into each location as there are going out. There should be no correlation between wages and mobility in the steady state. Nevertheless, if moving costs are high, at any given time one would expect to see flows of workers toward locations with higher wages as part of a dynamic equilibrium driven by local labor demand shocks. As was mentioned above (in footnote 6), there is some evidence that local labor market shocks have long-lasting effects. So in a specification that uses only mean wages in each location (with no location match effects), we should find a relationship between mean wages and migration decisions. This is, in fact, what we find in Table IX (in the “State Means” column). But we also find that the exclusion of location match wage effects is strongly rejected by a likelihood ratio test.

Even if differences in mean wages merely equalize the amenity differences between locations, the model predicts a relationship between wage realizations and migration decisions, because of location match effects: if the location match component is bad, the worker has an incentive to leave. This motivates the “National Wages” column of Table IX, where it is assumed that mean wages are the same in all locations (as they would be if measured wage differences merely reflect unmeasured amenities). We find that workers who have unusually low wages in their current location are indeed more likely to move.

Finally, the “Regional Amenities” column shows that the results are robust to the inclusion of regional amenity differences. Section 5.4.2 shows that the

TABLE IX
ALTERNATIVE INCOME SPECIFICATIONS

	Base Model		State Means		National Wages		Regional Amenities ^a	
	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$
Disutility of moving	4.790	0.565	4.567	0.533	4.755	0.568	4.761	0.611
Distance	0.265	0.182	0.253	0.183	0.269	0.183	0.353	0.236
Adjacent location	0.808	0.214	0.810	0.213	0.805	0.213	0.777	0.234
Home premium	0.331	0.041	0.273	0.032	0.328	0.040	0.373	0.051
Previous location	2.757	0.356	2.558	0.299	2.728	0.345	2.730	0.423
Age	0.055	0.020	0.061	0.019	0.055	0.020	0.052	0.020
Population	0.653	0.179	0.661	0.181	0.649	0.179	0.682	0.234
Stayer probability	0.510	0.078	0.517	0.079	0.513	0.078	0.495	0.093
Cooling	0.055	0.019	0.040	0.019	0.055	0.019	0.031	0.044
Income	0.312	0.100	0.318	0.181	0.315	0.099	0.290	0.130
Wage intercept	-5.165	0.244	-5.440	0.238	-4.024	0.270	-5.133	0.249
Time trend	-0.035	0.008	-0.051	0.005	-0.011	0.009	-0.033	0.008
Age effect (linear)	7.865	0.354	8.112	0.366	7.445	0.381	7.842	0.360
Age effect (quadratic)	-2.364	0.129	-2.321	0.134	-2.386	0.128	-2.364	0.131
Ability (AFQT)	0.014	0.065	0.062	0.060	0.024	0.064	0.011	0.066
Interaction (age, AFQT)	0.147	0.040	0.158	0.041	0.140	0.039	0.144	0.041
Transient s.d. 1	0.217	0.007	0.231	0.007	0.217	0.007	0.217	0.007
Transient s.d. 2	0.375	0.015	0.385	0.016	0.372	0.015	0.375	0.015
Transient s.d. 3	0.546	0.017	0.559	0.018	0.544	0.017	0.546	0.017
Transient s.d. 4	1.306	0.028	1.331	0.028	1.305	0.027	1.306	0.028
Fixed effect 1	0.113	0.035	0.253	0.013	0.357	0.039	0.113	0.036
Fixed effect 2	0.298	0.035	0.547	0.011	0.167	0.041	0.297	0.036
Fixed effect 3	0.936	0.017	1.028	0.014	0.905	0.023	0.933	0.017
Wage match	0.384	0.017	-	-	0.363	0.023	0.384	0.019
Log likelihood	-4,214.88		-4,267.88		-4,215.83		-4,203.58	

^aThe "Regional Amenities" model includes 12 regional dummy variables (coefficients not shown).

model is identified even if each location has an unobserved amenity value that is common to all individuals. In practice, we do not have enough data to estimate the complete model with a full set of fixed effects for all 50 locations. As a compromise, we divide the states into 13 regions and present estimates for a model with fixed amenity values for each region.²⁹ This has little effect on

²⁹The regions are as follows: (i) Northeast (NE, ME, VT, NH, MA, RI, CT); (ii) Atlantic (DE, MD, NJ, NY, PA); (iii) Southeast (SE, VA, NC, SC, GA, FL); (iv) North Central (NC, MN, MI, WI, SD, ND); (v) Midwest (OH, IN, IL, IA, KS, NB, MO); (vi) South (LA, MS, AL, AR); (vii) South Central (OK, TX); (viii) Appalachia (TN, KY, WV); (ix) Southwest (AZ, NM, NV); (x) Mountain (ID, MT, WY, UT, CO); (xi) West (CA, HI); (xii) Alaska and (xiii) Northwest (OR, WA).

the estimated income coefficient; moreover, a likelihood ratio test accepts the hypothesis that there are no regional amenity differences.³⁰

6.7. *Sensitivity Analysis*

Our empirical results are inevitably based on some more or less arbitrary model specification choices. Table X explores the robustness of the results with respect to some of these choices. The general conclusion is that the parameter estimates are robust. In particular, the income coefficient estimate remains positive and significant in all of our alternative specifications.

The results presented so far are based on wages that are adjusted for cost of living differences across locations. If these cost of living differences merely compensate for amenity differences, then unadjusted wages should be used to measure the incentive to migrate. This specification yields a slightly lower estimate of the income coefficient without much effect on the other coefficients and the likelihood is lower (mainly because there is more unexplained variation in the unadjusted wages). Thus, in practice, the theoretical ambiguity as to whether wages should be adjusted for cost of living differences does not change the qualitative empirical results: either way, income significantly affects migration decisions.

The other specifications in Table X are concerned with sensitivity of the estimates to the discount factor (β), the horizon length (T), heterogeneity in moving costs and the inclusion of a second climate variable (heating degree-days).³¹ Again, the effect of income is quite stable across these alternative specifications.

7. MIGRATION AND WAGES

7.1. *Spatial Labor Supply Elasticities*

We use the estimated model to analyze labor supply responses to changes in mean wages for selected states. We are interested in the magnitudes of the migration flows in response to local wage changes and in the timing of these responses. Since our model assumes that the wage components relevant to migration decisions are permanent, it cannot be used to predict responses to wage innovations in an environment in which wages are generated by a stochastic

³⁰Alaska is the only region that has a significant (positive) coefficient; this is perhaps not surprising given that the model specifies the utility flow as a linear function of average temperature and Alaska is an outlier in this respect.

³¹Table X is a sample of many alternative specifications that were tried. As was mentioned earlier, size (as measured by population) may affect migration either as a scaling factor on the payoff shocks or as a variable affecting the cost of migration. We experimented with these alternatives and also expanded the moving cost specification to allow quadratic effects of distance, location size, and climate variables; none of these experiments changed the results much.

TABLE X
ALTERNATIVE SPECIFICATIONS

	Base Model		No COLA		$\beta = .9$		$T = 20$		1 Cost Type		Heating	
	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$
Disutility of moving	4.790	0.565	4.702	0.556	4.694	0.578	4.487	0.613	5.287	0.560	4.764	0.556
Distance	0.265	0.182	0.282	0.182	0.297	0.198	0.266	0.186	0.262	0.182	0.274	0.193
Adjacent location	0.808	0.214	0.797	0.214	0.840	0.232	0.811	0.224	0.781	0.214	0.794	0.220
Home premium	0.331	0.041	0.324	0.040	0.465	0.052	0.382	0.041	0.185	0.022	0.323	0.039
Previous location	2.757	0.356	2.711	0.366	2.809	0.347	2.806	0.332	3.373	0.322	2.770	0.363
Age	0.055	0.020	0.056	0.019	0.060	0.020	0.068	0.022	0.074	0.020	0.058	0.019
Population	0.653	0.179	0.639	0.181	0.696	0.186	0.650	0.179	0.645	0.163	0.683	0.190
Stayer probability	0.510	0.078	0.512	0.079	0.486	0.080	0.494	0.079	0	–	0.505	0.079
Cooling	0.055	0.019	0.057	0.019	0.069	0.027	0.056	0.022	0.020	0.014	0.099	0.031
Heating	–	–	–	–	–	–	–	–	–	–	0.020	0.012
Income	0.312	0.100	0.263	0.095	0.451	0.139	0.359	0.110	0.148	0.069	0.304	0.098
Wage intercept	–5.165	0.244	–5.125	0.270	–5.172	0.244	–5.168	0.244	–5.199	0.243	–5.166	0.244
Time trend	–0.035	0.008	–0.030	0.010	–0.035	0.008	–0.035	0.008	–0.035	0.008	–0.035	0.008
Age effect (linear)	7.865	0.354	7.826	0.385	7.875	0.354	7.869	0.355	7.906	0.353	7.867	0.355
Age effect (quadratic)	–2.364	0.129	–2.379	0.128	–2.366	0.129	–2.366	0.129	–2.375	0.129	–2.365	0.129
Ability (AFQT)	0.014	0.065	0.053	0.067	0.016	0.065	0.017	0.065	0.012	0.065	0.014	0.065
Interaction (age, AFQT)	0.147	0.040	0.133	0.040	0.144	0.040	0.144	0.040	0.153	0.040	0.147	0.040
Transient s.d. 1	0.217	0.007	0.220	0.007	0.217	0.007	0.217	0.007	0.218	0.007	0.217	0.007
Transient s.d. 2	0.375	0.015	0.380	0.016	0.375	0.015	0.375	0.015	0.375	0.015	0.375	0.015
Transient s.d. 3	0.546	0.017	0.553	0.016	0.546	0.017	0.546	0.017	0.546	0.017	0.546	0.017
Transient s.d. 4	1.306	0.028	1.322	0.029	1.306	0.028	1.306	0.028	1.306	0.028	1.306	0.028
Fixed effect 1	0.113	0.035	0.138	0.034	0.112	0.035	0.112	0.035	0.111	0.035	0.113	0.035
Fixed effect 2	0.298	0.035	0.311	0.034	0.298	0.035	0.298	0.035	0.298	0.035	0.298	0.035
Fixed effect 3	0.936	0.017	0.966	0.020	0.937	0.017	0.936	0.017	0.938	0.017	0.936	0.017
Wage match	0.384	0.017	0.400	0.019	0.384	0.017	0.384	0.017	0.382	0.017	0.384	0.017
<i>Log likelihood</i>	–4,214.880		–4,282.473		–4,214.038		–4,214.034		–4,231.700		–4,214.082	

process. Instead, it is used to answer comparative dynamics questions: we use the estimated parameters to predict responses in a different environment. First we do a baseline calculation, starting people in given locations and allowing them to make migration decisions in response to the wage distributions estimated from the Census data. Then we do counterfactual simulations, starting people in the same locations, but facing different wage distributions.

We take a set of people who are distributed over states as in the 1990 Census data for white male high-school graduates aged 20–34. We assume that each person is initially in the home state at age 20 and we allow the population distribution to evolve over 15 years by iterating the estimated transition probability matrix. We consider responses to wage increases and decreases that represent a 10% change in the mean wage of an average 30-year-old for selected states. First, we compute baseline transition probabilities using the wages that generated the parameter estimates. Then we increase or decrease the mean wage in a single state and compare the migration decisions induced by these wage changes with the baseline. Supply elasticities are measured relative to the supply of labor in the baseline calculation. For example, the elasticity of the response to a wage increase in California after 5 years is computed as $\frac{\Delta L}{\Delta w} \frac{w}{L}$, where L is the number of people in California after 5 years in the baseline calculation, and ΔL is the difference between this and the number of people in California after 5 years in the counterfactual calculation.

Figure 1 shows the results for three large states that are near the middle of the one-period utility flow distribution. The supply elasticities are above 0.5. Adjustment is gradual, but is largely completed in 10 years. Our conclusion from this exercise is that despite the low migration rate in the data, the supply of labor responds quite strongly to spatial wage differences.

7.2. Migration and Wage Growth

Our model is primarily designed to quantify the extent to which migration is motivated by expected income gains. Interstate migration is a relatively rare event and our results indicate that many of the moves that do occur are motivated by something other than income gains. This raises the question of whether the income gains due to migration are large enough to be interesting.

One way to answer this question is to compare the wages of the mover and the stayer types as time goes by, using simulated data. Table XI shows results for a simulation that starts 1,000 people at home in each of the 50 states at age 20 and measures accumulated income and utility gains at age 34 (the oldest age in our NLSY sample).³² Migration increases the total utility flow by a modest but nontrivial amount. Most of the gain comes from improved location

³²The results are weighted by the state distribution of white male high-school graduates aged 19–20 from the 1990 Census.




FIGURE 1.—Responses to wage changes.

matches; even though there is considerable dispersion in mean wages across states, the estimated dispersion in the location match component of wages is much larger and, therefore, a much more important source of income gains due to migration. The dollar value of the nonpecuniary gains due to (climate) amenities is also larger than the gains from moving toward high-wage states.

The importance of the home location can be seen by simulating migration decisions with the home premium parameter set to zero. The results are shown in Table XII. With no attachment to a home location, the annual migration rate increases to 7.3% and the mover type moves about once every 6 years. By age 34, the accumulated gains due to migration exceed 20% of the base utility level. Given that people are willing to forgo gains of this magnitude to stay in their home location, it follows that the costs of forced displacements (due to natural disasters such as hurricane Katrina, for example) are very high.

TABLE XI
MIGRATION GAINS^a

	Migration Rate	Mean μ	Match v	Amenity $\frac{\alpha_1 Y_1}{\alpha_0}$	Total
Mover type	5.23%	17,222	966	4,677	22,865
Stayer type	0	17,192	-58	4,374	21,508
Gain		31	1024	303	1,357
Percentage gain		0.1%	4.8%	1.4%	6.3%
Standard deviation		1,460	6,835	3,018	7,572

^aMigration gains are measured in 2010 dollars.

TABLE XII
MIGRATION GAINS WITH NO HOME LOCATION

	Migration Rate	Mean	Match	Amenity	Total
Mover type	13.58%	17,223	2,876	5,847	25,946
Stayer type	0	17,192	-21	4,350	21,521
Gain		31	2,897	1,497	4,424
Percentage gain		0.1%	13.5%	7.0%	20.6%

8. CONCLUSION

We have developed a tractable econometric model of optimal migration in response to income differentials across locations. The model improves on previous work in two respects: it covers optimal sequences of location decisions (rather than a single once-for-all choice) and it allows for many alternative location choices. Migration decisions are made so as to maximize the expected present value of lifetime income, but these decisions are modified by the influence of unobserved location-specific payoff shocks. Because the number of locations is too large to allow the complete dynamic programming problem to be modeled, we adopt an approximation that truncates the amount of information available to the decision-maker. The practical effect of this is that the decisions of a relatively small set of people who have made an unusually large number of moves are modeled less accurately than they would be in the (computationally infeasible) complete model.

Our empirical results show a significant effect of expected income differences on interstate migration for white male high-school graduates in the NLSY. Simulations of hypothetical local wage changes show that the elasticity of the relationship between wages and migration is roughly 0.5. Our results can be interpreted in terms of optimal search for the best geographic match. In particular, we find that the relationship between income and migration is partly driven by a negative effect of income in the current location on the probability of out-migration: workers who get a good draw in their current location tend to stay, while those who get a bad draw tend to leave.

The main limitations of our model are those imposed by the discrete dynamic programming structure: given the large number of alternative location choices, the number of dynamic programming states must be severely restricted for computational reasons. Goodness of fit tests indicate that the model nevertheless fits the data reasonably well. From an economic point of view, the most important limitation of the model is that it imposes restrictions on the wage process, implying that individual fixed effects and movements along the age-earnings profile do not affect migration decisions. A less restrictive specification of the wage process would be highly desirable.

APPENDIX A: THE SAMPLE

In this appendix, we describe the selection rules use to construct the analysis sample of 432 respondents with 4,274 person-years. As noted in the text, we applied strict sample inclusion criteria to obtain a relatively homogeneous sample. In Table XIII we report the selection rules and the number of respondents deleted by each rule. The NLSY79 contains three subsamples, a nation-

TABLE XIII
SAMPLE SELECTION

	Respondents	Person-Years
White non-Hispanic males (cross-section sample)	2,439	39,024
Restrictions applied to respondents		
Ever in military	-246	
High school dropouts and college graduates	-1,290	
Attended college	-130	
Older than age 20 at start of sample period	-134	
Missing AFQT score	-41	
Attend or graduate from high school at age 20	-87	
Not in labor force for more than 1 year after age 19	-44	
Location at age 20 not reported	-20	
Income information inconsistent	-1	
Died before age 30	-4	
Residence at age 14 not reported	-2	
In jail in 1993	-1	
Subtotal	-2,000	439
		6,585
Restrictions applied to periods		
Delete periods after first gap in history	-1	-1,104
Delete periods before age 20	-6	-1,207
Analysis sample	432	4,274
<i>Years per person</i>		
1	14	14
2	16	32
3	19	57
4	14	56
5	14	70
6	14	84
7	13	91
8	9	72
9	34	306
10	61	610
11	53	583
12	44	528
13	45	585
14	44	616
15	38	570
Total	432	4,274

ally representative cross-section sample, a supplemental sample of minorities and economically disadvantaged youth, and a sample of individuals in the military in 1979. We start with the 2,439 white non-Hispanic males in the cross-section sample. We exclude respondents who ever served in the military and we include only those with exactly a high-school education.

We assume that permanent labor force attachment begins at age 20; thus we exclude respondents who were born in 1957 and who were, therefore, not interviewed until they were already more than 20 years old. We drop those who are in school or report graduating from high school at age 20. Since we use the AFQT (conducted in 1980) to help explain wages, we drop individuals with missing AFQT scores. Respondents who report being out of the labor force for more than 1 year after age 19, due to disability, tending house, or "other," are dropped on the grounds that they are not typical of this population. We use residence at age 14 as the home location, so we drop people for whom this variable is missing; we also drop people whose location at age 20 is unknown. We dropped one person who never reported income after age 19. We also dropped four people who died in their 30s, again on the grounds that they are atypical. Finally, we dropped one individual who was incarcerated in 1993 (after reporting remarkably high incomes in earlier years). Application of these criteria produced a sample of 439 individuals and 6,585 person years.

We apply two period-level restrictions. The first is that the histories must be continuous: we follow individuals from age 20 to their first noninterview or the 1994 interview. Since a missed interview means that location is unknown, we discard all data for each respondent after the first missed interview. Finally, we delete observations before age 20 from the analysis sample. Seven respondents have information only during their teenage years.

Our final sample contains 4,274 periods for 432 men. There are 124 interstate moves, with an annual migration rate of 2.9 percent. More than a one-third of the moves (43) were returns to the home location. There are 361 people who never moved, 31 who moved once, 33 who moved twice, and 7 who moved three times or more. The median age is 25, reflecting the continuous-history restriction.

APPENDIX B: VALIDATION OF ML ESTIMATES

The parameter estimates from Table II were used to generate 100 replicas of each NLSY observation, starting from the actual value in the NLSY data and allowing the model to choose the sequence of locations. Table XIV gives maximum likelihood estimates using the simulated data. The last column reports the t -value that tests the difference between the estimates and the individual parameters of the data generating process; the last row reports a likelihood ratio test of the hypothesis that the data were generated by the process that did, in fact, generate them (assuming that the simulation program works). The estimated coefficients are close to the true values and the χ^2 test accepts the

TABLE XIV
ESTIMATES FROM SIMULATED MIGRATION HISTORIES

	Base Model		100 Reps		<i>t</i>
	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	$\hat{\theta}$	$\hat{\sigma}_{\theta}$	
Disutility of moving	4.790	0.565	4.732	0.058	-1.00
Distance	0.265	0.182	0.272	0.015	0.42
Adjacent location	0.808	0.214	0.808	0.017	0.04
Home premium	0.331	0.041	0.333	0.004	0.42
Previous location	2.757	0.356	2.709	0.031	-1.52
Age	0.055	0.020	0.056	0.002	0.25
Population	0.653	0.179	0.648	0.017	-0.30
Stayer probability	0.510	0.078	0.517	0.008	1.00
Cooling	0.055	0.019	0.055	0.002	-0.07
Income	0.312	0.100	0.307	0.008	-0.56
Wage intercept	-5.165	0.244	-5.158	0.033	0.21
Time trend	-0.035	0.008	-0.035	0.001	0.40
Age effect (linear)	7.865	0.354	7.853	0.050	-0.23
Age effect (quadratic)	-2.364	0.129	-2.359	0.018	0.28
Ability (AFQT)	0.014	0.065	0.017	0.010	0.35
Interaction (age, AFQT)	0.147	0.040	0.142	0.007	-0.67
Transient s.d. 1	0.217	0.007	0.217	0.001	0.07
Transient s.d. 2	0.375	0.015	0.376	0.002	0.37
Transient s.d. 3	0.546	0.017	0.545	0.002	-0.84
Transient s.d. 4	1.306	0.028	1.306	0.004	0.14
Fixed effect 1	0.113	0.035	0.114	0.003	0.40
Fixed effect 2	0.298	0.035	0.300	0.003	0.68
Fixed effect 3	0.936	0.017	0.938	0.002	1.37
Wage match	0.384	0.017	0.383	0.002	-0.17
Log likelihood, $\chi^2(24)$	-4,214.88		-473,830.3		9.95

truth. We take this as evidence that our estimation and simulation programs work.

REFERENCES

- ABBING, J. H., AND J. J. HECKMAN (2007): "Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation," in *Handbook of Econometrics*, Vol. 6B, ed. by J. J. Heckman and E. Leamer. New York: Elsevier. [225]
- ALTONJI, J. G., AND C. R. PIERRET (2001): "Employer Learning and Statistical Discrimination," *Quarterly Review of Economics*, 116, 313-350. [231]
- ANAS, A., AND C. M. FENG (1988): "Invariance of Expected Utilities in Logit Models," *Economics Letters*, 27, 41-45. [235]
- BANKS, J. S., AND R. K. SUNDARAM (1994): "Switching Costs and the Gittins Index," *Econometrica*, 62, 687-694. [215]
- BARRO, R. J., AND X. SALA-I MARTIN (1991): "Convergence Across States and Regions," *Brookings Papers on Economic Activity*, 1, 107-158. [214]

- BERRY, S. T., AND P. A. HAILE (2010): "Nonparametric Identification of Multinomial Choice Demand Models With Heterogeneous Consumers," Unpublished Manuscript, Yale University. [225]
- BLANCHARD, O. J., AND L. F. KATZ (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 1, 1–37. [212,214]
- CRAWFORD, G. S., AND M. SHUM (2005): "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73, 1137–1173. [217]
- DAHL, G. B. (2002): "Mobility and the Return to Education: Testing a Roy Model With Multiple Markets," *Econometrica*, 70, 2367–2420. [212]
- ERDEM, T., AND M. P. KEANE (1996): "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15, 1–20. [217]
- GALLIN, J. H. (2004): "Net Migration and State Labor Market Dynamics," *Journal of Labor Economics*, 22, 1–21. [212]
- GEMICI, A. (2008): "Family Migration and Labor Market Outcomes," Unpublished Manuscript, New York University. [212,213,239]
- GREENWOOD, M. J. (1997): "Internal Migration in Developed Countries," in *Handbook of Population and Family Economics*, Vol. 1B, ed. by M. R. Rosenzweig and O. Stark. North Holland. [211]
- HOLT, F. (1996): "Family Migration Decisions: A Dynamic Analysis," Unpublished Manuscript, University of Virginia. [212]
- HOTZ, V., AND R. MILLER (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 60, 497–529. [227]
- KEANE, M. P., AND K. I. WOLPIN (1994): "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," *Review of Economics and Statistics*, 76, 648–672. [217]
- KEANE, M. P., AND K. I. WOLPIN (1997): "The Career Decisions of Young Men," *Journal of Political Economy*, 105, 473–522. [212]
- KENNAN, J. (2006): "A Note on Discrete Approximations of Continuous Distributions," Unpublished Manuscript, University of Wisconsin–Madison. [222]
- (2008): "Average Switching Costs in Dynamic Logit Models," Unpublished Manuscript, University of Wisconsin–Madison. [235]
- KENNAN, J., AND J. R. WALKER (2010): "Wages, Welfare Benefits and Migration," *Journal of Econometrics*, 156, 229–238. [213]
- (2011): "Supplement to 'The Effect of Expected Income on Individual Migration Decisions'," *Econometrica Supplementary Material*, 79, http://www.econometricsociety.org/ecta/Supmat/4657_data_and_programs.zip. [221]
- KOTLARSKI, I. I. (1967): "On Characterizing the Gamma and Normal Distribution," *Pacific Journal of Mathematics*, 20, 69–76. [226]
- LUCAS, R. E. B. (1997): "Internal Migration in Developing Countries," in *Handbook of Population and Family Economics*, Vol. 1B, ed. by M. R. Rosenzweig and O. Stark. North Holland. [211]
- MAGNAC, T., AND D. THESMAR (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801–816. [225]
- McFADDEN, D. D. (1974), "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka. Academic Press. [216]
- NEAL, D. (1999): "The Complexity of Job Mobility of Young Men," *Journal of Labor Economics*, 17, 237–261. [212]
- PESSINO, C. (1991): "Sequential Migration Theory and Evidence From Peru," *Journal of Development Economics*, 36, 55–87. [218]
- REAGAN, P. B., AND R. OLSEN (2000): "You Can Go Home Again: Evidence From Longitudinal Data," *Demography*, 37, 339–350. [221]

- RUST, J. P. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033. [216]
- (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, Vol. 4, ed. by R. Engle and D. McFadden. Elsevier. [212,215]
- SCHULTZ, T. P. (1982): "Lifetime Migration Within Educational Strata in Venezuela: Estimates of a Logistic Model," *Economic Development and Cultural Change*, 30, 559–593. [219]
- SJAASTAD, L. A. (1962): "The Costs and Returns of Human Migration," *Journal of Political Economy*, 70, 80–89. [239]
- SWEETING, A. (2007): "Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commerical Radio Industry," Working Paper 13522, NBER. [234]
- TOPEL, R. H. (1986): "Local Labor Markets," *Journal of Political Economy*, 94, S111–S143. [212, 214]
- TUNALI, I. (2000): "Rationality of Migration," *International Economic Review*, 41, 893–920. [212]
- UNITED STATES NATIONAL CLIMATIC DATA CENTER (2002): "Historical Climatology Series 5-1 and 5-2," Report, Asheville, NC. Available at <http://lwf.ncdc.noaa.gov/oa/climate/normal/usnormals.html>. [230]

Dept. of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, U.S.A. and NBER; jkennan@ssc.wisc.edu
and

Dept. of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, U.S.A. and NBER; walker@ssc.wisc.edu.

Manuscript received May, 2003; final revision received December, 2009.