

The effect of homeownership on the option value of regional migration

FLORIAN OSWALD

Department of Economics, SciencesPo Paris

This paper estimates a lifecycle model of consumption, housing choice, and migration in the presence of aggregate and regional shocks, using the Survey of Income and Program Participation (SIPP). The model delivers structural estimates of moving costs by ownership status, age, and family size that complement the previous literature. Using the model, I first show that migration elasticities vary substantially between renters and owners, and I estimate the consumption value of having the option to migrate across regions when there are regional shocks. This value is 19% of lifetime consumption on average, and it varies substantially with household type.

KEYWORDS. Migration, housing, lifecycle consumption, regional shocks.

JEL CLASSIFICATION. J6, R2, R23.

1. INTRODUCTION

Regional migration rates in the USA are relatively low despite the presence of large regional shocks. However, it would be a mistake to conclude from this observation that the option to migrate across regions has a small value to consumers. The goal of this paper is to provide a measure of having the option to migrate in the face of regional income and house price uncertainty, and I show that the value is large. The paper provides a structural interpretation of the insurance value of migration against regional shocks, as proposed first in [Yagan \(2013\)](#). It shows that considering homeowners and renters separately is of first-order importance for this issue, since both have vastly different elasticities of migration with respect to regional shocks. This insight is relevant for labor market and housing policy alike.

Migration probabilities are heterogeneous in the population. Which type of household is likely to move in a regional downturn? In this paper, which is among the first to consider homeownership and migration in an empirical lifecycle model, I provide structural estimates of crucial objects related to this question, for example, moving costs by ownership status, age, and other observables. Modeling homeownership realistically requires modeling asset accumulation and mortgages, and it requires a proper treatment

Florian Oswald: florian.oswald@gmail.com

I would like to thank Costas Meghir, Lars Nesheim, and Mariacristina De Nardi for their continued support and advice throughout this project. I also thank Uta Schönberg, Jeremy Lise, Marco Bassetto, Suphanit Piyapromdee, Richard Blundell, Jose-Victor Rios Rull, Morten Ravn, Fabien Postel-Vinay, Emeric Henry as well as three anonymous referees for helpful comments and suggestions. I would like to acknowledge financial support from cemmap under ESRC grant ES/I034021/1 and ESRC transformative grant ES/M000486/1.

of expectations about house prices, both of which are computationally demanding to integrate in a dynamic model of location choice.

Homeownership and geographical mobility of households are tightly connected: Renters are more mobile than owners. What complicates the analysis, however, is that renters may choose to be renters precisely *because* they are more mobile, in the sense that they might assess their own likelihood of moving to be relatively high. What is more, often the econometrician cannot observe the relevant state variables which would be informative about those considerations; hence, there is unobserved heterogeneity at play. The model introduced below allows to resolve the joint determination of housing tenure status, consumption, savings, and mobility decisions, such that it can be used to structurally estimate deep parameters and to investigate counterfactual policies.

The main counterfactual will be used to shed light on the option value of regional migration under regional price and income risk. How much would households want to pay for a hypothetical migration insurance policy? In other words, what is the value of the option to move? In order to address this, the experiment shuts down migration in the economy, and it reports the compensating consumption stream which would make individuals indifferent between this regime, and the status quo, that is, a world with migration. The results of this exercise differ greatly by type of household considered and their respective current locations.

In 2013, 63% of occupied housing units in the US were owned, while 37% were rented.¹ At the same time, roughly 1.3% of the population migrate across US census division boundaries per year. Conditional on ownership, this implies that 1.9% of renters and 0.67% owners move. A natural question is then to ask why do we observe owners moving less? All else equal, owners face higher moving costs, both in terms of financial as well as time and effort costs. Financial costs occur because of transaction costs in the housing market upon sale of the house (e.g., agency fees or transaction taxes), while costs of effort arise from owners having to spend time finding a suitable buyer, meet with agents and lawyers, etc. A comparable renter is subject to those costs only to a lesser degree. Buying a house means to make a highly local financial investment, which is subject to shocks as discussed above, is relatively illiquid, and in addition may have a location specific flow of utility. Consumers may have preferences for locations. Finally, as already mentioned, there is selection into homeownership based on unobservable moving costs: Individuals with a particular distaste for moving will be more likely to select homeownership, because they anticipate that they are unlikely to ever move in the future. All of these factors interact to shape the joint decision of housing tenure, location choice, and mortgage borrowing. What is more, they all interact to influence the decision to move in response to a shock.

In the model that I develop, there are several mechanisms which affect the homeownership choice of individuals. A downpayment requirement implies that only individuals with sufficient cash on hand are able to buy a house at the current price. The model assumes a preference for owner-occupied accommodation, a local amenity, and a partially unobserved cost of moving, which influence the buying decision in addition to age, the probability of moving, and beliefs about future shocks.

¹See the American Community Survey 2013, Table DP04.

In terms of the decision to migrate to another region, the model predicts that the likelihood of migration is increasing in the difference of discounted expected lifetime utilities between any two regions. Those relative utilities, in turn, depend among other things on the average regional income level and the level of regional house prices, both of which vary over time. Allowing regional characteristics to vary is a significant contribution to the literature on dynamic migration models such as, for example, [Kennan and Walker \(2011\)](#), since it provides a fundamental reason for agents to move in response to a change in their economic environment, rather than as a result of idiosyncratic preference shocks alone. Including time-varying location characteristics, however, increases computational demands substantially. To keep those demands tractable, the model employs a factor structure which allows aggregate shocks to affect regions differently.

I estimate the model using a simulated method of moments estimator. I find that the model fits the data very well along the main dimensions of interest, which are mobility and ownership patterns over the lifecycle, ownership rates by region, migration flows across regions, as well as wealth accumulation over the lifecycle and by region. After fitting the model to the data, I first use the model to compute migration elasticities to regional shocks by tenure status and current location. Then I investigate why owners move less than renters in greater detail. The main result of the paper shows that migration is a low probability event in both data and model, but associated with a large option value for consumers. Shutting down regional migration in an environment with realistic income and price shocks would require a 19% increase in per period consumption to make the average consumer indifferent to the status quo. This number varies greatly by household type (age, housing tenure, persistent income level) as well as location.

Literature. My paper builds on [Kennan and Walker \(2011\)](#), who are the first to develop a model of migration with multiple location choices over the lifecycle. Their main finding is that expected income is an important determinant of migration decisions, and their framework requires large moving costs to match observed migration decisions. The model features location-specific match effects in wages and amenity which are uncertain ex ante, so the consumer has to move to a location in order to discover their values. The distributions of those match effects in each location are stationary. After having learned the value of the current location, the only reason for a move is a favorable realization of an i.i.d. preference shock which might occur in some future period. There is no change in economic fundamentals which might encourage a move, like a shock to wages, for example. Relaxing this feature as well as adding housing and savings decisions are my main contributions to their paper. I am able to let regions experience differential income and price shocks over time, thereby providing an additional reason to move over and above idiosyncratic shocks.

[Gemici \(2007\)](#) focused on migration decisions of couples with two working spouses and finds that, for this subgroup, family ties can significantly hinder migration decisions and wage growth. [Winkler \(2010\)](#) is similar to [Gemici \(2007\)](#) but with housing choices. The main differences to [Winkler \(2010\)](#) are the way I model regional price and income dynamics and the assumption about how job search takes place. Regarding regional dynamics, I am able to allow for shocks which are correlated across regions and with an

aggregate component that is persistent, while they are assumed to be independent in [Winkler \(2010\)](#). The i.i.d. assumption for regional shocks is clearly rejected in the data, as will become clear in the next section. Also, [Winkler \(2010\)](#) assumed that job offers arrive in the current location from a random alternative location. My assumption implies first that individuals consider all potential locations in each period, and decide to move based on their expectations about how they will fare in each. Second, it allows for reasons other than job offers to trigger a move, which is also a feature of the data, as I will show below. [Ransom \(2018\)](#) is another related paper using the [Kennan and Walker \(2011\)](#) setup which allows for shocks to wages and local unemployment rates at the CBSA level, but without considering housing. Finally, [Bishop \(2008\)](#) computed a dynamic migration model using the conditional choice probability setup as proposed by [Arcidiacono and Miller \(2011\)](#) in order to recover willingness to pay for environmental amenities.

By considering regional shocks, the present paper is related to the seminal contribution of [Blanchard and Katz \(1992\)](#). In light of state-specific shocks to labor demand, the authors find that after an adverse shock, the relocation of workers is one of the main mechanisms to restore unemployment and participation rates back to trend in an affected region. [Lkhagvasuren \(2012\)](#) is a more recent paper on the topic, proposing a frictional version of the [Lucas and Prescott \(1974\)](#) island model. Relative to those papers, here we show how the underlying decision maker reacts to regional shocks—in particular, how owners and renters react differently and what this implies for their valuation of the migration option. Related to this, [Notowidigdo \(2011\)](#) analyzed the incidence of local labor demand shocks on low-skilled workers in a static spatial equilibrium model and found that they are more likely to stay in a declining city than high-skilled workers to take advantage of cheaper housing.² The same mechanism operates in my model. Furthermore, the dynamic nature of my model allows me to evaluate the response of migration to shocks over time. The present paper can be seen as a complement to the exercise proposed in [Yagan \(2013\)](#) or (2019), where the question is how much insurance against local labor market shocks is offered by migration. The author finds migration insures against 7% of an average local labor demand shock. I implement a fully structural analysis of the same question, with the added benefit that I can measure a value of the migration option in terms of consumption. In this sense, the present paper offers a more direct answer to the question of *how much consumption would I forgo today in order to be insured in an adverse future state*, which describes an insurance contract fairly well. [Bartik \(2018\)](#) is a recent paper which extends [Yagan \(2013\)](#) to consider the influence of the China trade shock as well as the fracking boom, abstracting from a detailed model of housing.

Another related literature considers the effects of the 2007 housing bust on labor market mobility. In terms of empirical contributions, [Ferreira, Gyourko, and Tracy \(2010\)](#), [Schulhofer-Wohl \(2011\)](#), and [Demyanyk, Hryshko, Luengo-Prado, and Sorensen \(2017\)](#) looked at whether negative equity in the home reduces the mobility of owners and report mixed findings. The first paper finds an effect, whereas the next two do not, with

²See [Moretti \(2011\)](#) for a comprehensive overview of this literature going back to [Roback \(1982\)](#) and [Rosen \(1979\)](#), and [Diamond \(2016\)](#) and [Piyapromdee \(2019\)](#) for recent applications.

TABLE 1. Percent of US population migrating across different geographic boundaries over different time spells. Taken from Molloy, Smith, and Wozniak (2011), computed from ACS, March CPS, and IRS data.

	Annually	Over 5 Years
County	5%	18.6%
State	2%	8.9%
Division	1.5%	4.8%

the difference arising from different datasets and definitions of long-distance moves. More theoretical papers like Head and Lloyd-Ellis (2012), Nenov (2012), Şahin, Song, Topa, and Violante (2014), and Karahan and Rhee (2019) used search models of labor and housing markets to look at geographical mismatch in order to understand how a fall in house prices affects unemployment and migration rates. The last paper, in particular, formalizes the negative equity lock-in notion in a model with two locations and finds only a moderate effect of lock-in on the increase in unemployment. The present paper differs from this group of contributions by assuming multiple locations and by adopting a lifecycle framework.³

In the remainder of this paper, I will first present a set of facts from aggregate and microdata about regional migration in the US in Section 2 before introducing a structural model which can speak to those fact in Section 3. I will then discuss solution and estimation of the model in Sections 4 and 5 in order to finally present the results regarding the option value of migration in Section 6.

2. FACTS

According to Molloy, Smith, and Wozniak (2011), who use three publicly available datasets (American Community Survey (ACS), the Annual Social and Economic Supplement to the CPS (March CPS), and Internal Revenue Service (IRS) data), each year roughly 5% of the population moves between counties each year, which amounts to roughly one-third of the annual flows into and out of employment according to the measure in Fallick and Fleischman (2004). The cross-state figure is 2%, and the cross-census division rate is estimated at 1.5% of the population, per year (see Table 1).

It is somewhat unfortunate that none of the datasets employed by Molloy, Smith, and Wozniak (2011) are well suited for the purpose of analyzing migration and ownership. None of them track movers, so it is impossible to know the circumstances of an individual at the moment they decided to move, which is ultimately of interest in this

³In general, the relationship between homeownership and labor market mobility or unemployment has been discussed in many other places, and an incomplete list might include Oswald (1996), Blanchflower and Oswald (2013), Coulson and Fisher (2002), Güler and Taskın (2018), Battu, Ma, and Phimister (2008), or Halket and Vasudev (2014).

paper.⁴ I therefore use the Survey of Income and Program Participation (SIPP) in this paper, a longitudinal and nationally representative dataset.⁵

Before presenting statistics from SIPP data, I will explain the geographic concept I will be using in this paper, which is a US census division. Census divisions are nine relatively large regions which separate the United States into groups of states “for the presentation of census data.”⁶ To a first approximation, those regions represent areas with a common housing and labor market. In the model, a move within any region is not considered as migration and, therefore, does not contribute to the overall migration rate. This implies that there is a proportion of moves across markets that do happen in the data, but which are not picked up by my geographic definition of a market.

The aggregation of states into this particular grouping is but one of many possibilities, and I adopt this particular partition based on computational constraints. In many respects, the ideal concept of a region is what economists would refer to as a local labor market, and metropolitan statistical areas (MSA) or commuting zones (CZ) come close to this. Unfortunately, for the purpose of the model in this paper, the so-called defined number of regions would be far too large to be computationally feasible. Hence the choice of census divisions.⁷ I will demonstrate below what the choice of divisions implies for the captured state-level variation. In the Online Appendix, Figure H.1 presents a map, and Table H.1 lists division abbreviations and the member states.⁸

2.1 *The main reasons to move are: Work, housing, and family*

The March Supplement to the Current Population Survey (CPS) contains several questions relevant for the study of migration. Here, I analyze answers to the 2013 edition of the CPS to the question “What was the main reason for moving” where respondents are offered 19 options to choose from. The results are displayed in Table 2. It is striking to note that even though we are conditioning on moves across division boundaries (and thus think of long-range moves), the percentage of people citing category “housing” as their main motivation is roughly 24% of the total population of movers. The table also

⁴It is possible to construct a panel dataset from the CPS, but only with postal address as unit identifier. If an individual moves out, this can be inferred from the data, however, the destination of the move cannot—in particular, it is unknown whether they relocated within the city, or somewhere else.

⁵The PSID is a natural competitor to the SIPP for this kind of study, with the PSID’s main advantage being the fact it is a long panel. I found that cell sizes got extremely small, however, after conditioning on the most important covariates in the PSID. Even unconditionally there are only 1560 unique cross-division moves in the PSID 1994–2011, and four cells in the region-by-region transition matrix have no observations for this entire period. I have 2512 unique cross-division moves in SIPP 1996–2012 and the corresponding transition matrix is dense.

⁶See the Census bureau’s website at https://www.census.gov/geo/reference/gtc/gtc_census_divreg.html.

⁷The model presented below contains 25.4 million different points in the state space at which to solve a savings problem. Increasing the number of regions to 51 (to represent US states) increases this to 815 million points in the state space. Given that estimation requires evaluation of the model solution many times over, the former state space can be handled with code that is highly optimized for speed, while the latter cannot.

⁸The online appendix is available in the Online Supplemental Material at <http://qeconomics.org> (see Oswald (2019)) and <https://floswald.github.io/pdf/homeownership-appendix.pdf>

TABLE 2. CPS 2013 data on main motivation of moving, conditional on a cross-division move. The purpose of this table is to show that the distribution of responses is stable conditional on quartiles of *distance moved*. This selects a sample of 20–50-year olds and aggregates the response to the question “What was the main reason for moving” (variable NXTRES) as follows. Work = {new job/transfer, look for job, closer to work, retired}, Housing = {estab. own household, want to own, better house, better neighborhood, cheaper housing, foreclosure, other housing}, family = {change marstat, other fam reason}, other = {attend/leave college, climate change, health, natural disaster, other}. The distance of a move is computed as the distance between geographic center of the *state* of origin (not division) and the center of the destination state. The rows of the table categorize the distance measure into its quartiles.

Distance Moved (KM)	Main Reason			
	Work	Housing	Family	Other
< 718	47.9%	23.2%	22.7%	6.1%
(718, 1348]	55.3%	25.7%	16.7%	2.3%
(1348, 2305]	51.6%	24.1%	22.5%	1.8%
(2305, 8087]	65.5%	22.7%	11.1%	0.7%
Total	55%	23.9%	18.3%	2.7%

disaggregates the response to the question by the distance between origin and destination State, and we can see that the proportion of respondents does vary with distance moved, but not to an extent that would suggest that housing becomes irrelevant as a motivation with increasing distance. Summing up in the bottom row of the table, we see that 55% say work was the main reason, 24% refer to housing, and the remaining 21% is split between family and other reasons. The model to be presented below addresses each of these categories: Individuals can move out of work-related concerns (regional and individual level income fluctuations), because of housing considerations (regional house price fluctuation), for family reasons (stochastic age-dependent arrival of children) as well reasons classified as “other,” which are accounted for by an idiosyncratic preference shock.

2.2 Homeownership and college education are important predictors for migration

Putting somewhat more structure onto this, I next present estimates from a statistical analysis of the determinants of cross-division moves from household-level SIPP data. I combine four panels of SIPP data (1996, 2001, 2004, and 2008) into a database with 102,529 household heads that I can follow over time and space. This will be the central estimation sample in main analysis below. Table 3 shows the results.⁹ I regress a binary indicator for whether or not a cross-division move took place in a given year on a set of explanatory variables, which relate to the household in question in a probit regression. The table shows marginal effects computed at the sample mean of each variable,

⁹It is worth emphasizing that at this point I am abstracting away from the severe endogeneity issues which the structural model below will account for.

TABLE 3. Determinants of cross-census division moves in SIPP data. Household income and wealth are measured in 100,000 USD. This regresses a binary indicator for whether a cross-division move takes place at age t on a set of variables relevant at that date. The first column shows marginal effects and the second column shows the marginal effects relative to the unconditional baseline mobility rate of 0.0132. The interpretation of this column is, for example, that the effect of being a homeowner is equivalent to reducing the baseline probability of migration by 51%.

	Marginal Effects	ME/Baseline
Intercept	-0.0250 (0.0020)	
Age	-0.0008 (0.0001)	-0.06
Age squared	0.0000 (0.0000)	0.0
Children in HH	-0.0008 (0.0003)	-0.06
Homeowner	-0.0067 (0.0004)	-0.51
Household income	0.0006 (0.0003)	0.05
Total wealth	0.0000 (0.0001)	0.0
College	0.0063 (0.0004)	0.48
Price/Income	0.0000 (0.0000)	0.0
Deviance	28,793.7099	
Dispersion	1.0261	
Num. obs.	294,840	

Note: Standard errors in parentheses.

as well as the ratio of marginal effects to the baseline unconditional probability of moving (1.32%). The results indicate that there is a pronounced age effect, with each additional year of age implying a reduction that is equal to 6% of the baseline probability. The same effect is found for whether or not children are present in the household. The effect of being a homeowner is very large and equivalent to a reduction in the propensity to move of 51% of the baseline probability. Increasing household income by \$100,000 is equivalent to a 5% baseline increase. Finally, having a college degree has an effect of equal magnitude than being a homeowner, but in the opposite direction: a college degree amounts to an increase of the baseline of 49%. According to this model, the effect of being a homeowner on the baseline moving probability is equal to an age increase of 8.3 years, thus taking a 30-year old to age 38; also, a household which owns the house would have to experience an increase in household income of \$1 m in order to make up for the implied loss in the probability of moving across divisions from being an owner. The house price to income ratio and total household wealth are not statistically significant in this specification.

TABLE 4. Annual moving rate in percent of the population. Households are categorized into “Renter” or “Owner” based on their homeownership status at the beginning of the period in which they move. SIPP data subset to noncollege degree holders.

	Cross-State	Cross-Division
Overall	1.51	0.99
Renter	2.07	1.49
Owner	0.82	0.64

Sample selection: Noncollege degree. Even though the estimates in Table 3 only measure statistical associations, they highlight an important feature of the data: moving and having a college degree are strongly correlated. While this paper specifically aims to investigate the other strong correlation in that table, that is, between ownership status and mobility, a full treatment which endogenizes education choices is too ambitious. A pragmatic solution to this problem is to condition the data on a certain education group and disregard education choices, as is done in the previous literature (e.g., Bishop (2008), Ransom (2018), Bartik (2018), Kennan and Walker (2011) all impose this restriction). In what follows, therefore, all SIPP data will refer to household heads without a college degree, which selects 62% of the original sample, resulting in 65,482 unique household heads.

2.3 Renters move at twice the rate of owners at all ages

In order to give a sense of the magnitude of migration rates by ownership status in this selected sample, Table 4 presents summary annual moving rates for both the state and census division level migration. The overall unconditional migration rate is 1.51% and 0.99% of households per year for cross-state and cross-division, respectively. The cross-state figure differs from the 2% in Table 1 because I set up the SIPP data in terms of household heads, thereby missing some moves of nonreference persons, and because I condition on noncollege. It is quite clear from Table 4 that there is a marked distinction in the likelihood of moving across state as well as division boundaries between renters and owners, with 2.07% (1.49%) of renters versus 0.82% (0.64%) of owners moving across state (division) boundaries on average per year. In total, I observe 1259 cross-division moves made by 1069 unique individuals in my noncollege sample, implying multiple moves for some movers.¹⁰

Reconsidering homeownership and migration by age gives rise to Figure 1. It is clear that renters are more likely to move at all ages, with a strongly declining age effect— younger individuals move more. At the same time, homeownership is increasing with

¹⁰By way of comparison, the estimation sample in Kennan and Walker (2011) is drawn from the geocoded version of NLSY79 and contains 124 interstate moves. The disadvantage of SIPP is that I can track an individual for at most four years.

Proportion of Cross-Division movers by age




FIGURE 1. SIPP sample proportion moving across the census division boundaries by age (upper panel) and proportion of owners by age (lower panel). Conditions on individuals without a college degree.

age. These are highly salient features of the data, and they are among the key dimensions along which this model's performance is going to be evaluated.

2.4 Regional income and house price risk are not IID

The time series of regional disposable income and regional house prices are each strongly correlated across divisions. Additionally, they exhibit high degrees of autocorrelation, that is, shocks to regional incomes and prices are persistent. To illustrate the degree of cross-correlation of both prices and incomes, consider Figure 2. The top panels show the detrended version of each time series, by region, while the bottom panels show the pairwise correlation of those detrended time series across regions. The figure highlights that deviations from trend are highly correlated between divisions, for both

average regional incomes (q) and regional prices (p).¹¹ Regarding persistence of those time series, the average autocorrelation coefficients are 0.91 for p and 0.92 for q , respectively (for details, see Online Appendix Table B.3) over the considered time period. Modeling regional risk as an IID process seems like an unjustifiably strong assumption given those high degrees of cross-correlation and persistence. Therefore, the model introduced below will take both correlation and persistence in regional prices seriously and will propose a method to solve and estimate the resulting high-dimensional problem. Online Appendix B.1 contains detailed descriptions of the raw data.

3. MODEL

In the model, I view households as a single unit, and I will use the terms *household* and *individual* interchangeably. Individuals are assumed to live in census division (or *region*) $d \in D$ in any given period at date t , and we let $j \in \{1, \dots, J\}$ index age. At each age j , individual i has to decide whether to move to a different region k , whether to own or rent, and how much of his labor income to save. Individuals derive utility from consumption c , from owning a house h and from local amenity A_d .

Every individual in region d faces an identical level of house price p_{dt} and mean labor productivity q_{dt} at time t , where q_{dt} enters the individual wage equation as a level shifter. At the individual level, uncertainty enters the model through a Markovian idiosyncratic component of income risk z_{ij} , a Markovian process that models changes in household size over the lifecycle s_{ij} , and a location-specific preference shock ε_{idt} , which is assumed identically and independently distributed across agents, regions and time. In short, region d is characterized by a tuple (q_{dt}, p_{dt}, A_d) , households can move to a different region subject to a moving cost, and they hold expectations about the evolution of regional prices (q_{dt+1}, p_{dt+1}) , $\forall d$ in such a way that is compatible with the evidence from Figure 2 (i.e., correlated shocks across region and high persistence) and is at the same time computationally feasible, as detailed below.¹²

The job search process is modeled as in Kennan and Walker (2011). Individuals do not know the exact wage they will earn in the new location. The new wage is composed of a deterministic, and thus predictable, part and a component that is random. Over and above an expectation about some prevailing average level of wages the mover can expect in any given region at time t , it is impossible to be certain about the exact match quality of the new job ex ante. The new job can be viewed as an experience good where quality is revealed only after an initial period. This setup gives rise to income risk associated with moving. I do not attempt to explain return migrations, which Kennan and Walker

¹¹Data for q comes from the BEA series "Personal Income by State," p is the FHFA house price index by the census division. Both sets of series are a direct input to the structural model to be introduced below. Data are available via <https://github.com/floswald/EconData>.

¹²Let it suffice for now to state that taking into account 9 different house price and labor income processes would not be feasible and, therefore, the solution will seek to reduce the number of relevant dimensions of these series, similar to what a principal component analysis would try to do.

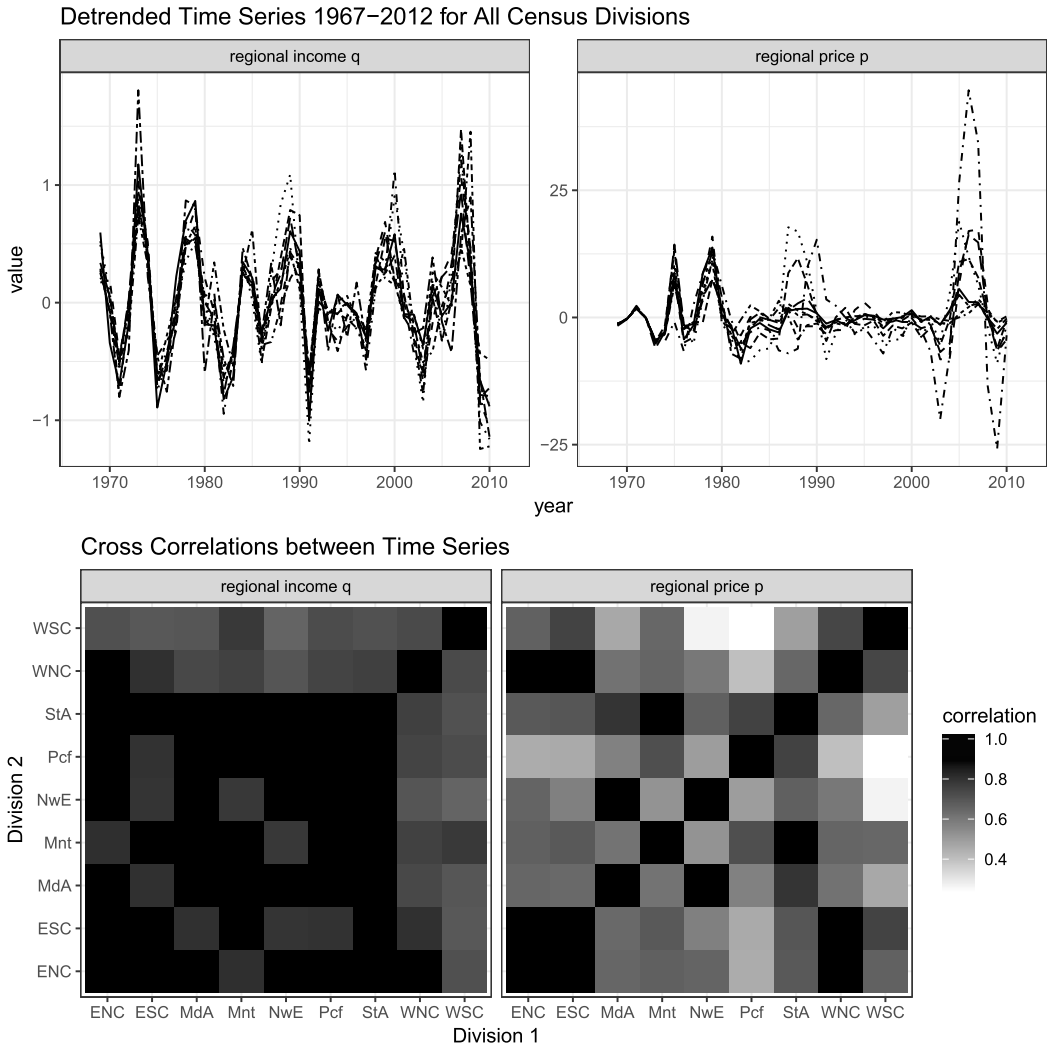


FIGURE 2. The time series for regional incomes $\{q_{dt}\}_{t=1967}^{2012}$ and house prices $\{p_{dt}\}_{t=1967}^{2012}$ are strongly correlated across divisions d . The top panels of this figure show the detrended series for all divisions for both q and p . The bottom panel further emphasizes that the cross-correlations across states between regional trend deviations are substantial. For instance, in the bottom left panel, the top left tile indicates that the correlation between the time series of q for East North Central (ENC) and West South Central (WSC) is around 0.75 (raw numbers in Online Appendix Tables B.1 and B.2). Detrended with fourth-order moving average.

(2011) achieve with a region-person specific match effect and by including this match effect from the last location in the state space.¹³

¹³Adding this feature would increase the computational burden of the model to make it infeasible, even with the limited memory assumption employed in Kennan and Walker (2011). I do not expect return migration to be of first order for the questions addressed here.

The model describes the partial equilibrium response of workers to regional wage and price shocks, as well as idiosyncratic income and family size shocks. The fairly detailed description of the consumer's decision problem rules out a full equilibrium analysis where house prices and wages clear local markets for computational feasibility reasons; hence, (q_{dt}, p_{dt}) are exogenous to the model.

3.1 Individual labor income

The logarithm of labor income of individual i depends on age j , time t , and current region d and is defined as in equation (1):

$$\begin{aligned} \ln y_{ijdt} &= \eta_d \ln q_{dt} + f(j) + z_{it}, \\ z_{it} &= \rho z_{it-1} + e_{it-1}, \\ e &\sim N(0, \sigma^2). \end{aligned} \tag{1}$$

Here, q_{dt} stands for the region specific price of human capital, $f(j)$ is a deterministic age effect modeled as a nonlinear function, and z_{it} is an individual specific persistent idiosyncratic shock. The coefficient η_d allows for differential transmission of regional shocks into individual income by region d . The log price of human capital q_{dt} is allowed to differ by region to reflect different industry compositions by region, which are taken as given.¹⁴

When moving from region d to region k at date t , I assume that the timing is such that current period income is earned in the origin location d . The individual's next period income is then composed of the corresponding mean income at that date in the new region k , q_{kt+1} , the deterministic age $j + 1$ effect, $f(j + 1)$, and a new draw for z_{it+1} conditional on their current shock z_{it} . For a mover, this individual-specific idiosyncratic component is drawn from a different conditional distribution than for nonmovers. Let us denote the different conditional distributions of z_{it+1} given z_{it} for stayers and movers by G_{stay} and G_{move} , respectively. This setup allows for some uncertainty related to the quality of the match with a job in the new region k , as mentioned above. In the model, I use G_{stay} and G_{move} as transition matrices from state z today to state z' tomorrow for stayers and movers, respectively.

3.2 Dimensionality reduction: National factors P and Q

As stated above, allowing $(q_{dt+1}, p_{dt+1}), \forall d$ to vary in an unrestricted fashion would make computation of this model infeasible. To solve this problem, I assume that agents

¹⁴Underlying this is an assumption about nonequalizing factor prices across regions. It is plausible to think that within a single country, wages should tend to converge to a common level, particularly in the presence of large migratory flows from one region to the next. In assuming no relative factor price equalization across US regions, I rely on a host of evidence showing that relative wages vary considerably across regions over a long time horizon (see, e.g., Bernard, Redding, and Schott (2013)).

use a 2-dimensional factor model to infer regional prices.¹⁵ To this end, I define aggregate state variables Q and P , which evolve according to a stationary vector autoregression of order one. At date t , all individuals observe the price vector \mathbf{F}_t containing both P_t and Q_t . The process is formally defined in equation (2), where \mathbf{A} is a matrix of coefficients and Σ is the variance-covariance matrix of the bivariate normal innovation ν . Agents in the model have rational expectations concerning this process:

$$\begin{aligned}\mathbf{F}_t &= \mathbf{A}\mathbf{F}_{t-1} + \nu_{t-1}, \\ \nu_t &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma\right), \\ \mathbf{F}_t &= \begin{bmatrix} Q_t \\ P_t \end{bmatrix}.\end{aligned}\tag{2}$$

3.2.1 Mapping aggregate factors to regional prices I assume that there is a deterministic mapping from the aggregate state \mathbf{F}_t into the price and income level of any region d which is known by all agents in the model. This means that once the aggregate state is known, agents know the price p_{dt} and income level q_{dt} in each region with certainty. The mapping is defined in terms of a function that depends on both aggregate states Q, P and where the coefficients are region dependent, as shown in expression (3). Similar to the aggregate case in (2), \mathbf{a}_d is a 2×2 matrix of coefficients specific to region d .

$$\begin{bmatrix} q_{dt} \\ p_{dt} \end{bmatrix} = \mathbf{a}_d \mathbf{F}_t.\tag{3}$$

Notice that the great virtue of this formulation is that the relevant price and income related state variables in each region are subsumed in \mathbf{F}_t , given the assumption that \mathbf{a}_d is known for all d . To be completely clear, equation (3) shuts down any uncertainty at the regional level once \mathbf{F}_t and \mathbf{a}_d are known.¹⁶ Shocks materialize in region d as a transformation of aggregate shocks to Q and P . The implications of this will be discussed in greater detail in Section 5.1 when I describe estimation of this part of the model and where I also provide some illustration regarding the fit of this model to the data.

3.3 Home ownership choice

Ownership choice is discrete, $h_j \in \{0, 1\}$, and there is no quantity choice of housing. While renting, that is, whenever $h_j = 0$, individuals must pay rent which amounts to a constant fraction κ_d of the current region- d house price p_d . Similar to the setup in Attanasio, Bottazzi, Low, Nesheim, and Wakefield (2012), I denote total financial (i.e.,

¹⁵The method of Krusell and Smith (1998) is conceptually similar to what I am doing. Instead of mean and variance of a distribution, consumers here track the value of two aggregate state variables.

¹⁶One could say that the formulation is missing a shock, for example, $\mathbf{a}_d \mathbf{F}_t + \epsilon_{dt}$, $\epsilon_{dt} \sim N(0, \sigma^2)$. Adding such a shock would increase the state space by a factor equal to the number of integration nodes to be used for the approximation of the resulting integral, which is a big cost. I do not expect any major difference in my results. The fit of this approximation is very good, as will be shown further below.

nonhousing) wealth at age j as assets a_j , which include liquid savings and mortgage debt. There is a terminal condition for net wealth to be nonnegative by the last period of life, that is, $a_J + p_{dt}h_{J-1} \geq 0, \forall t$, which translates into an implicit borrowing limit for owners. Additional to that, in order to buy, a proportion χp_{dt} of the house value needs to be paid up front as a downpayment, while the remainder $(1 - \chi)p_{dt}$ is financed by a standard fixed rate mortgage with exogenous interest rate r^m . The mortgage interest rate is assumed at a constant markup $\hat{r} > 0$ above the risk-free interest rate r , such that $r^m = r + \hat{r}$. The markup captures default risk incurred by a mortgage lender.

The equity constraint must be satisfied in each period, that is, $a_{ij+1} \geq -(1 - \chi)p_{dt}h_j, \forall t$. This means that only owners are allowed to borrow, with their house as a collateral. Selling the house incurs proportional transaction cost ϕ , such that given current house price p_t , upon sale the owner receives $(1 - \phi)p_t$.

This setup implies that owners will choose a savings path contingent on the current price, their income and debt level, the mortgage interest rate, and their current age j , such that they can satisfy the final period constraint. Sections 3.7.1 and 3.7.2 below describe the budget constraints in greater detail.

3.4 Moving

Moving costs. Moving is costly both in monetary terms (see the budget constraints below in 3.7) and in terms of utility. Denote $\Delta(d, x)$ the utility costs of moving from d at a current value of the state vector x (defined below). Moving costs differ between renters and owners. Moving for an owner requires to sell the house, which in turn requires some effort and time costs. This is in addition to any other utility costs incurred from moving regions which are common between renters and owners. I specify the moving cost function as linear in parameters α :

$$\Delta(d, x) = \alpha_{0,\tau} + \alpha_1 j + \alpha_2 j^2 + \alpha_3 h_{ij-1} + \alpha_4 s_{ij}. \tag{4}$$

In expression (4), $\alpha_{0,\tau}$ is an intercept that varies by unobserved moving cost type τ , α_1 and α_2 are age effects, α_3 measures the additional moving cost for owners, and α_4 measures moving cost differential arising from family size s_{ij} .

The unobserved moving cost type $\tau \in \{0, 1\}$, where $\tau = 1$ indexes the high-cost type, is a parsimonious way to account for the fact that in the data, some individuals never move. This is of particular relevance when thinking about owners, who may self-select into ownership because they know they are unlikely to ever move. In the model, this selection mechanism, together with any other factor that implies a high unobservable location preference, is collapsed into a type of person that has prohibitively high moving costs ($\alpha_{0,\tau=1}$ is large), and thus is unlikely to move. Providing some real-world context for this setup, Koşar, Ransom, and Van der Klaauw (2019) used consumer expectations data to find that for close to 50% of the population, nonpecuniary moving costs approach infinity.

Restrictions. I rule out the possibility of owning a home in region d while residing in region k . This would apply, for example, for households who keep their home in d , rent it out on the rental market, and purchase housing services either in rental or owner-occupied sector in the new region k . In my sample, I observe less than 1% of movers for which this is the case. Most likely, this is a result of high management fees or a binding liquidity constraint that forces households to sell the house to be able to afford the downpayment in the new region.¹⁷

3.5 Preferences

Period utility u depends on the choice of region k , and whether this is different from the current region d . A move takes place in the former case, and the household stays in d in the latter case:

$$u(c, h, k; x_{it}) = \eta \frac{c^{1-\gamma}}{1-\gamma} + \xi(s_{ij}) \times h - \mathbf{1}[d \neq k] \Delta(d', x_{it}) + A_k + \varepsilon_{ikt}. \quad (5)$$

Notice that (c, h, k) are *current* period choices of consumption, housing status, and location that affect utility. Those choices interact with the value of the state vector x_{it} , hence they depend on household sizes s_{ij} , and an additively separable idiosyncratic preference shock for the chosen region k , ε_{ikt} . Parameter η measures the scale at which consumption enters utility, while ξ measures the importance of ownership at various household sizes s . Household size s at age j is a binary random variable, $s \in \{0, 1\}$, relating to whether or not children are present in the household. It evolves from one period to the next in an age-dependent way as described in Section 3.7. Moving costs $\Delta(k, x_{it})$ are only incurred if in fact a move takes place. Finally, amenities in region k are given by the fixed effect A_k .

3.6 Timing and state vector

The state vector of individual i at date t when they are of age j is given by

$$x_{it} = (a_{ij}, z_{ij}, s_{ij}, \mathbf{F}_t, h_{ij-1}, d, \tau, j),$$

where the variables stand for, in order, assets, individual income shock, household size, aggregate price vector, housing status coming into the current period, current region index, moving cost type, and age.¹⁸

Timing within the period is assumed to proceed in two subperiods: in the first part, stochastic states are realized and observed by the agent, and labor income is earned; in the second part, the agent makes optimal decisions regarding consumption, housing, and location. The chronological order within a period is thus as follows:

¹⁷SIPP allows me to verify whether individuals possess any real estate other than their current home at any point in time. Fewer than 1% of movers provide an affirmative answer to this.

¹⁸A word of caution regarding the two time indices j and t : For large parts of the exposition, this distinction is irrelevant, that is, saying a_{ij} or a_{it} is equivalent. However, in the estimation I will allow different cohorts C_1, \dots, C_N to experience different sequences of prices $\mathbf{F}_{C_1}, \mathbf{F}_{C_2}, \dots$ and, therefore, separating time and age will become necessary.

1. observe $\mathbf{F}_t, s_{it}, z_{ij}$ and $\varepsilon_{it} = (\varepsilon_{i1t}, \varepsilon_{i2t}, \dots, \varepsilon_{iDt})$, i.i.d. location taste shock
2. earn labor income in current region d , as a function of q_{dt} and z_{ij}
3. given the state, compute optimal behavior in all D regions, that is,
 - (a) choose optimal consumption c_h^* conditional on housing choice $h \in \{0, 1\}$ in all regions k
 - (b) choose optimal housing $h_d^*(c_h^*)$
 - (c) choose optimal location, based on the value of optimal housing in each location

3.7 Recursive formulation

It is now possible to formulate the problem recursively. Following Rust (1987), I have assumed additive separability between utility and idiosyncratic location shock ε as well as independence of the transition of ε conditional on x . Furthermore, I assume that ε is distributed according to the standard Type 1 extreme value distribution.¹⁹

The consumer faces a nested optimization problem in each period. At the lower level, optimal savings and housing decision must be taken conditional on any discrete location choice, and at the upper level the discrete location choice with the maximal value is chosen; see (6). It is useful to define the conditional value function $v(x, k)$, which represents the optimal value after making housing and consumption choices at state x , while moving to location k , net of idiosyncratic location shock ε , in (7). Equation (8) is a result of the distributional assumption on ε , which admits a closed-form expression of the expected value function (also known as the *Emax* function in this model class), whereby $\bar{\gamma} \approx 0.577$ is Euler’s constant:

$$V(x_{it}) = \max_{k \in D} \{v(x_{it}, k)\}, \tag{6}$$

$$v(x_{it}, k) = \max_{c > 0, h \in \{0, 1\}} u(c, h, k; x_{it}) + \varepsilon_{ikt} + \beta \mathbb{E}_{z, s, \mathbf{F}} [\bar{v}(x_{it+1}) | z_{ij}, s_{ij}, \mathbf{F}_t], \tag{7}$$

$$x_{it+1} = (a_{ij+1}, z_{ij+1}, s_{ij+1}, \mathbf{F}_{t+1}, h, k, j + 1),$$

$$\begin{aligned} \bar{v}(x_{it+1}) &= E_\varepsilon V(x_{it+1}) \\ &= \bar{\gamma} + \ln \left(\sum_{k=1}^D \exp(v(x_{it+1}, k)) \right). \end{aligned} \tag{8}$$

The final period models a terminal value that depends on net wealth and a term that captures future utility from the house after age J , as shown in equation (9).

$$V_J(a, h_{J-1}, d) = \frac{(a_J + h_{J-1} p_{dt})^{1-\gamma}}{1-\gamma} + \omega h_{J-1}, \quad \forall t \tag{9}$$

The maximization problem in equation (7) is subject to several constraints, which vary by housing status and location choice. It is convenient to lay them out here case by case.

¹⁹This is also called the standard Gumbel distribution. Notice that the *Standard* part implies that location and scale parameters of the Gumbel distribution are chosen such that $E[\varepsilon] = \bar{\gamma}$, a constant known as the Euler–Mascheroni number, and that the its standard deviation is fixed at $\sqrt{\text{Var}(\varepsilon)} = \frac{\pi}{\sqrt{6}}$.

3.7.1 Budget constraint for stayers, i.e. $d = k$ Starting with the case for stayers, the relevant state variables in the budget constraint refer only to the current region d . In particular, given (p_{dt}, q_{dt}) , renters may choose to become owners, and owners may choose to remain owners or sell the house and rent.

Renters. The period budget constraint for renters (i.e., individuals who enter the period with $h_{ij-1} = 0$) depends on their housing choice, as shown in equation (10). In case they buy at date t , that is, $h_{ij} = 1$, they need to pay the date t house price in region d , p_{dt} , otherwise they need to pay the current local rent, $\kappa_d p_{dt}$. Labor income is defined in equation (11) and depends on the regional mean labor productivity level q_{dt} as introduced in Section 3.1. Buyers can borrow against the value of their house and are required to make a proportional downpayment amounting to a fraction χ of the value at purchase, while renters cannot borrow at all. This is embedded in constraint (12), which states that if a renter chooses to buy, their next period assets must be greater or equal to the fraction of the purchase price that was financed via the mortgage, or nonnegative otherwise. Constraint number (13) defines the interest rate function, which simply states that there is a different interest applicable to savings as opposed to borrowing, both of which are taken as exogenous parameters in the model. \hat{r} stands for the exogenous risk premium of mortgages charged over the risk-free rate. The terminal condition constraint is in expression (14):

$$a_{ij+1} = (1 + r(a_{ij}))(a_{ij} + y_{ijdt} - c_{ij} - (1 - h_{ij})\kappa_d p_{dt} - h_{ij}p_{dt}), \quad (10)$$

$$\ln y_{ijdt} = \eta_d \ln(q_{dt}) + f(j) + z_{ij}, \quad (11)$$

$$a_{ij+1} \geq -(1 - \chi)p_{dt}h_{ij}, \quad (12)$$

$$r(a_{ij}) = \begin{cases} r & \text{if } a_{ij} \geq 0, \\ r^m & \text{if } a_{ij} < 0, \end{cases} \quad r^m = r + \hat{r}, \quad (13)$$

$$a_{iJ} + p_t h_{iJ-1} \geq 0, \quad \forall t. \quad (14)$$

Owners. For individuals entering the period as owners ($h_{ij-1} = 1$), the budget constraint is similar except for two differences which relate to the borrowing constraint and transfers in case they sell the house. Owners are not required to make a scheduled mortgage payment—a gradual reduction of debt, that is, an increase in a , arises naturally from the terminal condition $a_{iJ} + p_t h_{iJ-1} \geq 0, \forall t$, as mentioned above. Therefore, the budget of the owner is only affected by the house price in case they decide to sell the house, that is, if $h_{ij} = 0$. In this case, they obtain the house price net of the proportional selling cost ϕ , plus they have to pay rent in region d . Apart from this, the same interest rate function (13), labor income equation (11), and terminal condition (14) apply:

$$a_{ij+1} = (1 + r(a_{ij}))(a_{ij} + y_{ijdt} - c_{ij} + (1 - h_{ij})(1 - \phi - \kappa_d)p_{dt}), \quad (15)$$

$$a_{ij+1} \geq -(1 - \chi)p_{dt}. \quad (16)$$

3.7.2 Budget constraint for movers, that is, $d \neq k$

Renters. For moving renters, the budget constraint is close to identical, with the exception that (10) needs to be slightly altered to reflect that labor income is obtained in the current period in region d before the move to k is undertaken:

$$a_{ij+1} = (1 + r(a_{ij}))(a_{ij} + y_{ijdt} - c_{ij} - (1 - h_{ij})\kappa_k p_{jt} - h_{ij} p_{kt}). \tag{17}$$

Owners. The budget constraint for moving owners depends on the house price in both current and destination regions d and k since the house in the current region must be sold by assumption. The expression $(1 - \phi)p_{dt}$ in (18) relates to proceeds from sale of the house in region d , whereas the square brackets describe expenditures in region k . Notice also that the borrowing constraint (19) now is a function of the value of the new house in k . It is important to note that this formulation precludes moving with negative equity if labor income is not enough to cover it. This is exacerbated in cases where the mover wants to buy immediately in the new region, since in that case the downpayment needs to be made as well, that is, if $y_{ijdt} < a_{ij} + (1 - \phi)p_{dt} - \chi h_{ij} p_{kt}$ then the budget set is empty and moving and buying is infeasible.²⁰

$$a_{ij+1} = (1 + r(a_{ij}))(a_{ij} + y_{ijdt} - c_{ij} + (1 - \phi)p_{dt} - [(1 - h_{ij})\kappa_k + h_{ij}]p_{kt}), \tag{18}$$

$$a_{ij+1} \geq -(1 - \chi)p_{kt}h_{ij}. \tag{19}$$

4. SOLVING AND SIMULATING THE MODEL

The model described above is a typical application of a mixed discrete-continuous choice problem. In the next section, I will introduce a nested fixed point estimator, which requires repeated evaluation of the model solution at each parameter guess, thus placing a binding time-constraint on time each solution may take.

The consumption/savings problem to be solved at each state, and its combination with multiple discrete choices and borrowing constraints, introduces several nondifferentiabilities in the asset dimension of the value function. This makes using fast first-order condition-based approaches to solve the consumption problem more difficult.²¹

²⁰In my sample, I observe 29 owners who move with negative equity (amounting to 3.4% of moving owners). 78% of those do buy in the new location, the rest rent. I do not observe whether or not an owner defaults on the mortgage. Accounting for this subset of the population would require to (1) assume that they actually defaulted and (2) it would substantially increase the computational burden. For those reasons, the model cannot account for this subset of the mover population at the moment.

²¹There has recently been a lot of progress on this front. Clausen and Strub (2013) provided an envelope theorem for the current case, and the endogenous grid point method developed by Carroll (2006), further extended to accommodate (multiple) discrete choice as in Fella (2014) and Iskhakov, Jørgensen, Rust, and Schjerning (2017) are promising avenues. I did not further pursue conditional choice probability (CCP) methods as in Arcidiacono and Miller (2011) or Bajari, Chan, Krueger, and Miller (2013), for example, because of data limitations. I experimented in particular with the latter paper's approach but soon had to give up because of too many empty cells in the empirical choice probability matrix (e.g., an entry like $\Pr(\text{own}, \text{save} = s, \text{move}|X)$ would be empty for many values of X ; in general, my problem was to recover the first stage decision rules from the data in a satisfactory kind of way).

I solve the model in a backward-recursive way, starting at maximal age 50 and going back until initial age 20. In the final period, the known value is computed at all relevant states. From period $J - 1$ onwards, the algorithm in each period iterates over all state variables and computes a solution to the savings problem at each combination of state and discrete choices variables (including housing and location choices). Notice that this state space spans all values for \mathbf{F}_t observed over the sample period. After this solution is obtained at a certain state, the discrete housing choice is computed, after which each conditional value function (7) is known.

Once the solution is obtained, simulation of the model proceeds by using the model implied decision rules and the observed aggregate prices series \mathbf{F}_t as well as their regional dependants (q_{dt}, p_{dt}) to obtain simulated lifecycle data. As will become clear in the next section, this procedure needs to replicate the time and age structure found in the data, which is achieved by simulating different cohorts, starting life in 1967 and all successive years up until 2012. The model moments are then computed using the empirical age distribution found in the estimation sample as sampling weights.

5. ESTIMATION

In this section, I explain how the model is estimated to fit some features of the data. There is a set of preset model parameters, the values of which I either take from other papers in the literature or I estimate them outside of the structural model and treat them as inputs. The remaining set of parameters are estimated using the simulated method of moments (SMM) approach, whereby given a set of parameters, the model is used to compute decision rules of agents, which in turn are used to simulate artificial data. In what follows, I will first discuss estimation of the exogenous stochastic processes, and then turn to the estimation of the model preference parameters.

5.1 Estimation of exogenous processes

VAR process for aggregates Q_t and P_t The VAR processes at the aggregate and regional level are estimated using a seemingly unrelated regression with two equations, one for each factor Q_t and P_t , $t = 1967, \dots, 2012$. I use real GDP per capita as a measure for Q_t , and the Federal Housing and Finance Association (FHFA) US house price index for P_t . Given that I am interested in the level of house prices (i.e., a measure of house *value*), I compute the average level of house prices found in SIPP data for the year 2012 and then apply the FHFA index backwards to construct the house value for each year.²²

I reproduce equation (2) here for ease of reading:

$$\mathbf{F}_t = \mathbf{A}\mathbf{F}_{t-1} + \nu_{t-1},$$

$$\nu_t \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma\right),$$

²²The GDP series is as provided by the Bureau of Economic Analysis through the FRED database. All non-SIPP data used in this paper are provided in an R package at <https://github.com/floswald/EconData>, documenting all sources and data-cleaning procedures.

TABLE 5. Estimates for aggregate VAR process. $\{P_t, Q_t\}_{t=1967}^{2012}$ are time series for FHFA national house prices, and real GDP per capita, respectively.

	Q_t	P_t
Intercept	0.86 (0.58)	19.13 (7.31)
Q_{t-1}	1.00 (0.02)	0.16 (0.28)
P_{t-1}	0.00 (0.01)	0.89 (0.06)
R^2	0.99	0.94
Adj. R^2	0.99	0.94
Num. obs.	94	94

Note: Standard errors in parenthesis.

$$\mathbf{F}_t = \begin{bmatrix} Q_t \\ P_t \end{bmatrix}.$$

The estimates from this equation are given in Table 5 and are used by agents in the model to predict \mathbf{F}_{t+1} given \mathbf{F}_t .

Aggregate to regional price mappings The series for q_{dt} is constructed as per capita personal income by region, with a measure of personal income obtained from the Bureau of Economic Analysis and population counts by state from intercensal estimates from the Census Bureau. The price series by region, p_{dt} , comes from the same FHFA dataset as used above:

$$\begin{bmatrix} q_{dt} \\ p_{dt} \end{bmatrix} = \mathbf{a}_d \mathbf{F}_t + \eta_{dt}, \tag{20}$$

$$\eta_{dt} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Omega_d \right).$$

The performance of this model in terms of delivered predictions from the aggregate state can be gauged visually in Figures 3 and B.3 in the Online Appendix B, which also contains the respective parameters estimates in Table B.4. It is important to understand the purpose of models (2) and (20): I do not want to make statistical inference based on the estimates from those models, which is something they may be ill-suited for, given the nature of the data. I am purely interested in their ability to replicate the observed regional prices, when fed the observed aggregate series for the purpose of approximating the evolution of the prices state space during simulation. In that regard, and by looking at 3 and B.3, I find they perform well.

A different concern that might arise from looking at the models in (20) is that it is unclear a priori how they in fact transform aggregate shocks into regional counterparts,

VAR fit to regional price data (p)

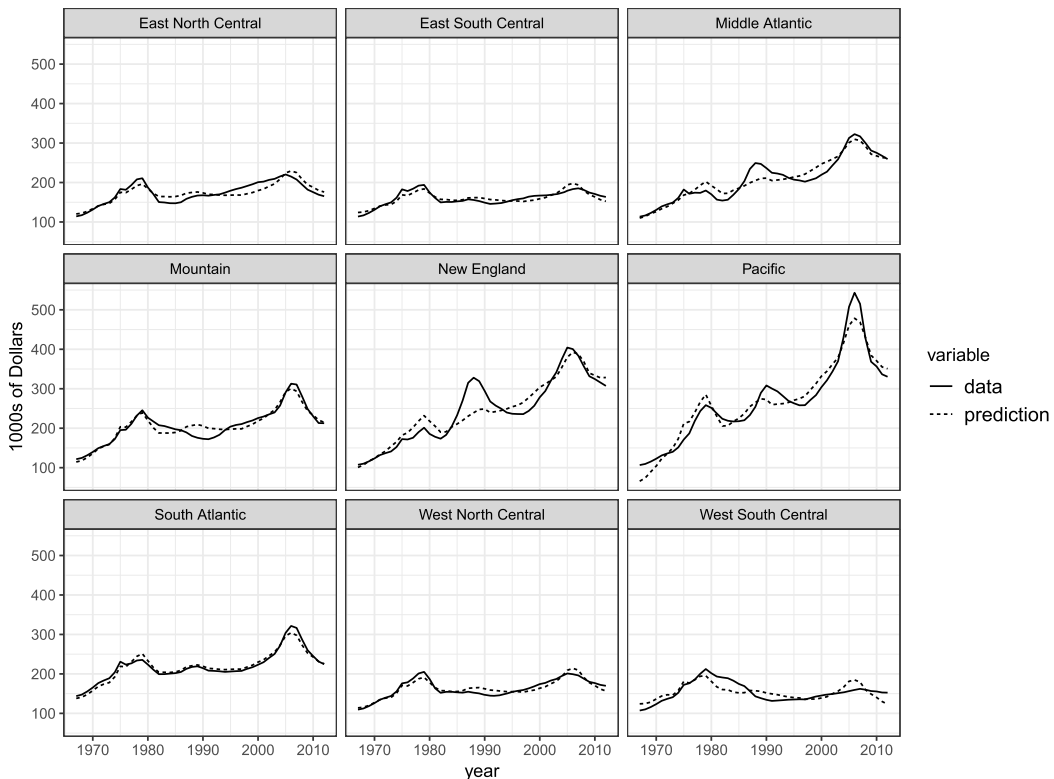


FIGURE 3. This figure shows the observed and predicted time series for mean income by the census division. The prediction is obtained from the VAR model in (3), which relates the aggregate series $\{Q_t, P_t\}_{t=1968}^{2012}$ to mean labor productivity $\{q_{dt}\}_{t=1968}^{2012}$ for each region d . Agents use this prediction in the model, that is, from observing an aggregate value $F_t = (P_t, Q_t)$ they infer a value for q_{dt} for each region above.

as this of course depends on the value of the estimated parameters \mathbf{a}_d . An illustration of this translation of shocks is shown in Online Appendix B.4.

Finally, a reasonable concern is how good an approximation of a more fine-grained geography such as state-level this setup based on divisions is. In order to shed some light on this, I run pooled OLS on my entire prices dataset, where on the left-hand side I have the price index for state s in period t , p_{st} , and as explanatory variable the corresponding census division level index, p_{dt} . In Table 6, we see the R^2 measured from each regression, implying that the division index explains a very large fraction of state-level variation throughout. The full regression output is in Online Appendix B.5.

Individual income process This part deals with the empirical implementation of equation (11), which models log labor income at the individual level. I estimate the linear regression

$$\ln y_{ijt} = \beta_0 + \eta_d \ln q_{dt} + \beta_1 j_{it} + \beta_2 j_{it}^2 + \beta_3 j_{it}^3 + \beta_4 \text{college}_{it} + z_{it}, \tag{21}$$

TABLE 6. R^2 from pooled OLS regression of state level indices p_{st}, q_{st} on corresponding division level indices p_{dt}, q_{dt} .

	$R^2 : p_{st} \sim p_{dt}$	$R^2 : q_{st} \sim q_{dt}$
East North Central	0.68	0.95
East South Central	0.93	0.96
Middle Atlantic	0.93	0.93
Mountain	0.68	0.83
New England	0.89	0.85
Pacific	0.72	0.83
South Atlantic	0.65	0.72
West North Central	0.73	0.96
West South Central	0.91	0.95

where $college_{it} = 1$ if i has a college degree, zero else, and where z_{it} are the regression residuals. The results of this are shown in the Online Appendix in Table C.1 and Figure C.1. The estimated residuals are used together with parameters β to generate an income grid for individuals without a college degree.

Copula estimates for z transitions The conditional distribution of z for movers is specified as the density of a bivariate normal copula G_{move} , which is invariant to date and region.²³ This means I assume that the conditional probability of drawing z' in new region k is the same regardless the origin location.²⁴ A copula is a multivariate distribution function with marginals that are all uniformly distributed on the unit interval. For example, if F is a bidimensional CDF, and if F_i is the CDF of the i th margin, then the bivariate copula is given by

$$C(u_1, u_2) = F(F_1^{-1}(u_1), F_2^{-1}(u_2)),$$

where F_i^{-1} is the quantile function. There are different families of copulae, and I will use a normal copula.

To estimate the parameters of the copula, I take residuals z_{it} from equation (21) and I want to study their joint distribution for movers, that is, $(z_{it}, z_{it+1})|d \neq k$. This object is informative for the question of whether individuals with a particularly high residual z_{it} are likely to have a high residual z_{it+1} after their move to region k , or not. In other words, we want to investigate the joint distribution of stayers $(z_{it}, z_{it+1})|d = k$ and of movers $(z_{it}, z_{it+1})|d \neq k$ separately. I obtain an estimate for the copula parameter ρ_s of 0.58, indicating substantial positive dependence for mover's z .²⁵ I report estimates and

²³A copula is a multivariate probability distribution function which connects univariate margins by taking into account the underlying dependence structure. For example, a finite state Markov transition matrix is a nonparametric approximation to a bivariate copula, and they converge as the number of states goes to infinity; see Bonhomme and Robin (2006).

²⁴It would be straightforward to relax this assumption, but data limitations forced me to impose this restriction.

²⁵ ρ_s is also called Spearman's rho, and it is related to Pearson's correlation coefficient ρ_p via $2 \sin(\frac{\pi}{6} \rho_s) = 2 \sin(\frac{\pi}{6} 0.58) = \rho_p = 0.598$ in this case of a Gaussian copula. In particular, $\rho_s \in [-1, 1]$.

TABLE 7. Preset parameter values.

		Value	Source
CRRRA coefficient	γ	1.43	Attanasio and Weber (1995)
Discount factor	β	0.96	Assumption
AR1 coefficient of z	ρ	0.96	French (2005)
SD of innovation to z	σ	0.118	French (2005)
Transaction cost	ϕ	0.06	Li and Yao (2007)
Downpayment proportion	χ	0.2	Assumption
Risk-free interest rate	r	0.04	Sommer et al. (2018)
30-year mortgage rate	r^m	0.055	Sommer et al. (2018)

describe the full procedure in the Online Appendix C.1. The conditional distribution of z for nonmovers will be parameterized externally as explained next.

Values for preset parameters I take several parameters for the model from the literature, as shown in Table 7. The estimates for the components of the idiosyncratic income shock process for nonmovers, that is, the autocorrelation $\rho = 0.96$ and standard deviation of the innovation $\sigma = 0.118$ are taken from French (2005). I set the financial transaction cost of selling a house, ϕ , to 6% in line with Li and Yao (2007) and conventionally charged brokerage fees. The time discount factor β is set to 0.96 which lies within the range of values commonly assumed in dynamic discrete choice models (e.g., Rust (1987)). The downpayment fraction χ is set to 20%, which is a standard value on fixed rate mortgages and used throughout the literature. The coefficient of relative risk aversion could be estimated, but is in this version of the model fixed to 1.43 as in Attanasio and Weber (1995).

To calibrate the interest rate for savings and for mortgage debt, I follow Sommer and Sullivan (2018), who use the constant maturity Federal Funds rate, adjusted by headline inflation as measured by the year on year change in the CPI. They obtain an average value of 4% for the period of 1977–2008, and I thus set $r = 0.04$. For the markup q of mortgage interest over the risk-free rate, they use the average spread between nominal interest on a 30-year constant maturity Treasury bond and the average nominal interest rate on 30-year mortgages. This spread equals 1.5% over 1977–2008, therefore, $\hat{r} = 0.015$, and $r^m = 0.055$.

5.2 Estimation of preference parameters

The parameter vector to be estimated by SMM contains the parameters of the moving cost function (α), the parameter in the final period value function ω , the population proportion of high moving cost types (π_τ), the scale of consumption η , and the utility derived from housing for both household sizes, (ξ_1, ξ_2). We denote the parameter vector of length K as $\theta = \{\alpha_0, \alpha_1, \alpha_2, \alpha_3, \alpha_4, \omega, \pi_\tau, \eta, \xi_1, \xi_2\}$.

Given θ , the model generates a set of M model moments $\hat{m}(\theta) \in \mathbb{R}^M$. After obtaining the same set of moments m from the data, the SMM procedure seeks to minimize the

criterion function

$$L(\theta) = [m - \hat{m}(\theta)]^T W [m - \hat{m}(\theta)], \tag{22}$$

which delivers point estimate $\hat{\theta} = \arg \min_{\theta} L(\theta)$. Given that this is a tightly parameterized model, I cannot use the theoretically optimal weighting matrix W , because a range of economically important moments vanish in the objective function because they enter at different scales. This is equally true if I use the common strategy of assigning the inverse of the variances of the data moments. To solve this problem, I prespecify a W as the identity matrix, but I modify the diagonal entries for some moments so that the corresponding derivative of the moment function is not negligible.²⁶

The maximization of the objective in (22) is performed with a cyclic coordinate search algorithm, where cycle $n + 1$ is defined as follows:

$$\begin{aligned} \theta_1^{(n+1)} &= \arg \min_{\theta_1} L(\theta_1, \theta_2^{(n)}, \dots, \theta_K^{(n)}), \\ \theta_2^{(n+1)} &= \arg \min_{\theta_2} L(\theta_1^{(n)}, \theta_2, \theta_3^{(n)}, \dots, \theta_K^{(n)}), \\ \theta_3^{(n+1)} &= \arg \min_{\theta_3} L(\theta_1^{(n)}, \theta_2^{(n)}, \theta_3, \theta_4^{(n)}, \dots, \theta_K^{(n)}), \\ &\vdots \\ \theta_K^{(n+1)} &= \arg \min_{\theta_K} L(\theta_1^{(n)}, \theta_2^{(n)}, \dots, \theta_K). \end{aligned}$$

This procedure is repeated until θ has converged. Convergence was not affected by different starting values and occurred in all cases after less than 10 iterations over the above scheme.²⁷

Denoting θ_0 the true parameter vector, by θ^* the optimizer of the above program and Σ the variance-covariance matrix of the asymptotic distribution of moment function errors as in

$$\sqrt{n}(m - \hat{m}(\theta^*)) \rightarrow \mathcal{N}(0, \Sigma),$$

the distribution of the parameter estimates $\hat{\theta}$ is given by the standard sandwich formula

$$\sqrt{n}(\hat{\theta} - \theta_0) \rightarrow \mathcal{N}(0, [dWd']^{-1} dW \Sigma Wd' [dWd']^{-1}),$$

where $d \equiv \frac{\partial m(\theta)}{\partial \theta}$ is the derivative of the moment function, given as a $K \times M$ matrix in this case. The derivative is approximated via finite differences, and Σ is obtained by obtaining 400 draws from the moment function.²⁸

²⁶Notice that this procedure still leads to valid standard errors, since W appears together with the covariance matrix of moments in the sandwich formula (see below). The weights are given by the values 10 for moments `cov_move_h`, `mean_move`, `mean_move_ownFALSE`, `mean_move_ownTRUE` and `lm_h_age2`, 1.5 for all migratory flow moments `flow_move_to_j`, and finally by 2 for `lm_mv_intercept` and `cov_own_kids`. This adjustment is similar to what is done in Lamadon (2014).

²⁷This optimization takes around 16 hours on a 10-instance cluster on AWS of type `t3.xlarge`. The procedure uses the function `optSlices` in Julia package <https://github.com/floswald/SMM.jl>.

²⁸See function `get_stdErrors` in the same Julia package <https://github.com/floswald/SMM.jl>.

Estimation sample My estimation sample is formed mainly out of averages over SIPP data moments covering the period 1997–2012, conditional on noncollege as described above. All moments are constructed using SIPP cross-sectional survey weights, and all dollar values have been inflated to base year 2012 using the BLS CPI for all urban consumers.²⁹ Averaging over years was necessary to preserve a reasonable sample size in all conditioning cells. However, it also introduces an initial conditions and cohort effects problem, since, for example, a 30-year old in 1997 faced a different economic environment over their lifecycle than a similar 30-year old in 2012 would have. The challenge is to construct an artificial dataset from simulated data, which has the same time and age structure as the sample taken from the data—in particular, agents in the model should have faced the same sequence of aggregate shocks as their data counterparts from the estimation sample. This requires to simulate individuals starting in different calendar years, taking into account the actual observed time series for regional house prices and incomes.

Identification Identification is achieved by comparing household behaviour under different price regimes. The variation comes from using the observed house price and labor productivity series in estimation, which vary over time and by region. The identifying assumption is that, conditional on all other model features, households must be statistically identical across those differing price regimes. In particular, this requires that household preferences be stable over time and do not vary by region.

The structural parameters in θ are related to the moment vector $m(\theta)$ in a highly nonlinear fashion. In general, all moments in $m(\theta)$ respond to a change in θ . However, it is possible to use graphical analysis to show how some moments relate more strongly to certain parameters than others.

Regarding parameters of the moving cost function, parameters $\alpha_{0,\tau=0}$, α_3 , α_4 represent the intercept for low moving cost types, the coefficient on ownership, and the effect of household size on moving costs, respectively. They are related to, in order, the average moving rate $E[\text{move}]$, the moving rate conditional on owning $E[\text{move}|h_t = 1]$, and the moving rate conditional on household size $E[\text{move}|s_t = 1]$. The age effects α_1 , α_2 are related to the age coefficients of the auxiliary model for moving, defined in expression (24), as well as the the average proportion of movers in the last period of life $E[\text{move}|T]$. The relationship between mobility and ownership, as well as mobility and household size are also captured by the covariances $\text{Cov}(\text{move}, h)$ and $\text{Cov}(\text{move}, s)$, both of which are again related to the moving cost parameters α_3 and α_4 .

The proportion of high moving cost types π_τ is related to the data moments concerning the number of moves per person, and in particular the fraction of individuals who never moved, $E[\text{moved never}]$. The other two moments on the frequency of moves, $E[\text{moved once}]$ and $E[\text{moved twice+}]$ are not part of the moment function, hence provide out of sample tests.

Given that the house price processes in each region are exogenous to the model, the parameters measuring utility from ownership, ξ_1 , ξ_2 are related to a relatively large number of moments: ownership rates by region and by household size, the covariance

²⁹<http://research.stlouisfed.org/fred2/series/CPIAUCSL>.




FIGURE 4. Graphical device to show model fit. These plots show how moments from data (x axis) line up with moments from simulated data (y axis). Ideally, all points would lie on the 45-degree line. A detailed listing and additional plots are available in the Online Appendix D.

of owning with household size, and the age-profile parameters from the auxiliary model of ownership in (23). A crucial parameter in the model is η , which measures the scale of consumption in utility: It informs us how changes in consumption and, therefore, changes in income induced by migration, affect payoffs. η is nonparametrically identified from differences in regional mean wages and moving probabilities, as demonstrated in Kennan and Walker (2011), Section 5.4.2.³⁰

5.3 Parameter estimates and moments

The model fits the data moments fairly well overall. Figures 4 and D.1 in the Online Appendix provide a quick overview of how the model moments line up with their data counterparts.

The moment vector m contains conditional means and covariances, which are largely self-explanatory. I introduce two auxiliary models included in m which relate to the age profiles of both migration and ownership. Both are linear probability models, where the dependent variable is either ownership status at the beginning of the period, h_{it-1} , or whether a move took place, denoted by $\text{move}_{it} = \mathbf{1}[d_{it} \neq d'_{it}]$:

$$h_{it} = \beta_{0,h} + \beta_{1,h}t_{it} + \beta_{2,h}t_{it}^2 + u_{h,it}, \tag{23}$$

$$\text{move}_{it} = \beta_{0,m} + \beta_{1,m}t_{it} + \beta_{2,m}t_{it}^2 + u_{m,it}. \tag{24}$$

Several tables contrasting model and data values as well as a detailed discussion of the fit this provides has been relegated to the Online Appendix D.

³⁰Thanks to a referee and the editor for pointing this out to me.

TABLE 8. Parameter estimates and standard errors.

		Estimate	Std. Error
Utility Function			
Owner premium size 1	ξ_1	-0.009	6.50e-05
Owner premium size 2	ξ_2	0.003	4.92e-05
Scale of c	η	0.217	0.0003
Continuation value	ω	4.364	1.76e-05
Moving cost function			
Intercept	α_0	3.165	9.29e-07
Age	α_1	0.017	0.1731
Age ²	α_2	0.0013	0.0002
Owner	α_3	0.217	2.08e-05
Household size	α_4	0.147	0.0007
Proportion of high type	π_τ	0.697	7.90e-05
Amenities			
New England	A_{NwE}	0.044	0.00946
Middle Atlantic	A_{MdA}	0.112	0.00029
Middle Atlantic	A_{StA}	0.168	2.12e-07
West North Central	A_{WNC}	0.090	6.24e-05
West South Central	A_{WSC}	0.122	7.45e-09
East North Central	A_{ENC}	0.137	0.0014
East South Central	A_{ESC}	0.063	0.0099
Pacific	A_{Pcf}	0.198	0.0002
Mountain	A_{Mnt}	0.124	3.37e-05

The estimated parameter vector and standard errors are shown in Table 8. It is not possible to attach a simple interpretation to parameter values in this nonlinear model, however, it is interesting to identification by looking at the standard errors. For most parameters, the gradient of the moment function is nonnegligible, and hence, we get precisely estimated coefficients at conventional levels of statistical significance. The age coefficient in the cost function, α_1 , is the main nonstatistically significant exception to this.

6. RESULTS

I will now move on to describe the results of this paper. In order to fully appreciate the results, it is useful to first illustrate a set of migration elasticities, before answering the question of why owners move less through the lens of the model. Then I will present the main set of results pertaining to the value of the migration option.

6.1 Elasticities with respect to regional shocks

The model can be used to compute elasticities of population size and migratory flows with respect to regional income shocks. Those elasticities are an important precursor to the main result of the paper, because they illustrate the incentives of agents in the event of such a shock. To measure the elasticity of population or migratory inflows, I simulate the economy and apply an unexpected and permanent shock to q_d in division d in

TABLE 9. Elasticities with respect to an unexpected and permanently positive income shock by region. This table reports elasticities of population (i.e., the stock of individuals present in each period) and migration *inflows* and *outflows* elasticities. For example, the percentage change in renter inflows is defined as $\frac{\#[\text{move to } d \text{ as renter}|\text{shock}] - \#[\text{move to } d \text{ as renter}|\text{no shock}]}{\#[\text{move to } d \text{ as renter}|\text{no shock}]}$ in each period, similarly for owners and for outflows. Elasticities are computed as averages over all years after the shock occurs.

Division	Population	Inflows			Outflows		
		Total	Buyers	Renters	Total	Owners	Renters
Aggregate	0.1	1.2	0.2	1.2	-0.4	-0.0	-0.4
East North Central	0.1	1.2	0.4	1.2	0.0	-0.8	0.1
East South Central	0.1	0.9	0.6	0.8	-0.0	0.7	-0.0
Middle Atlantic	0.1	1.4	-0.1	1.5	-0.4	-1.2	-0.4
Mountain	0.2	1.1	-0.1	1.1	-0.7	0.9	-0.7
New England	0.1	0.9	0.00	0.9	-0.2	1.6	-0.2
Pacific	0.2	1.3	0.8	1.2	-1.4	-0.3	-1.5
South Atlantic	0.1	1.3	0.6	1.3	-0.8	-0.5	-0.9
West North Central	0.1	0.8	0.0	0.9	-0.1	0.1	-0.1
West South Central	0.1	1.9	-0.7	2.1	0.0	-0.8	0.1

the year 2000. The elasticities are computed by comparing population size or migration flows across shocked and baseline scenarios, normalizing the result by the size of the shock.³¹

The results by region are shown in Table 9. First, we observe that the average of population elasticities across regions is a value around 0.1, implying that on average, a 1% permanent increase to regional income will lead to a 0.1% increase of population size of the shocked area. Total inflows into regions increase in the range of approximately 0.8% to 1.9%, the inflow rate of renters increases more than the one of incoming buyers throughout. The next set of columns looks at the complement to those statistics, that is, the elasticity of outflows. In the present case of a positive income shock, outflows decline in general as both renters and owners are less likely to move away. Table E.1 in the Online Appendix shows the corresponding elasticities for the case of a positive regional house price shock.

6.2 Why do owners move less?

There are several reasons for why owners move less than renters. First, they have higher moving costs as implied by a positive estimate for parameter α_3 . Second, owners pay a transaction cost each time they sell the house (proportional cost ϕ), so any (expected)

³¹Notice that given the cohort setup of the simulator, in this and all other experiments that involve some notion of a “shock”, it is necessary to simulate the model as many times as there are cohorts. This is so because each cohort experiences the shock at a different age, and the p_{dt} and q_{dt} are predetermined in the data. Members of the cohort born in 1985 reach the shock year t^* at a different age than those from the 1984 cohort. The shock is implemented by immediately changing the policy functions when the shock arises, and expectations adapt to the new setting. Hence for cohort 1985, the policy functions look different than for the 1984 cohort, and so on.

TABLE 10. Decomposing owners' moving costs. This table shows the percentage increase in three key model moments related to ownership and migration as we successively remove moving costs for owners. Compares baseline statistics to scenarios with no additional moving cost for owners ($\alpha_3 = 0$), no financial transaction costs from selling the house ($\phi = 0$), and neither of the two ($\alpha_3 = \phi = 0$). The top panel is for the entire population, the bottom panel conditions on mover types ($\tau = 0$) only.

All	$\alpha_3 = 0$	$\phi = 0$	$\alpha_3 = \phi = 0$
%Δ Ownership rate	5.6	2.7	8.2
%Δ Migration rate	4.0	0.5	4.9
%Δ Migration Own	140.0	9.3	155.4
Mover types: $\tau = 0$			
%Δ Ownership rate	21.4	5.1	26.3
%Δ Migration rate	4.0	0.5	4.9
%Δ Migration Own	108.8	6.7	118.9

gains from migration need to be traded off against this financial cost. Third, owners have to comply with the downpayment constraint if they wish to buy in the new region, which puts restrictions on the consumption paths of movers. Most owners will indeed return to ownership status in the new region, given the utility benefits, and given that in principle they are above the downpayment constraint. Fourth, ownership is correlated with larger household size ($s = 1$), which itself carries a higher moving cost (α_4). Last but certainly not least, a large proportion of owners is of the stayer type because they self-select into ownership as was discussed in Section 3.4. The ownership rate conditional only on moving cost type is 0.59 for movers ($\tau = 0$) versus 0.64 for stayers ($\tau = 1$).

To investigate those issues in more detail, I now sequentially remove owner-specific moving costs in Table 10 from the model. We should imagine an approximation to the partial derivative of the moment function of the model with respect to the parameters α_3 and ϕ . If we see a large reaction of the model-generated moments after setting a certain cost component to zero, we can conclude that this component is relatively important to explain the data. Starting therefore in the top panel of Table 10, we see changes in three key moments when considering first all types of households. The first row shows the percentage change in the aggregate ownership rate for different configurations of the model parameters. Setting the utility cost of moving for owners to zero in the column labeled $\alpha_3 = 0$ increases the ownership rate by 5.6%, because this makes owning a more attractive option in case of the need to migrate. Similarly, abolishing the financial transaction cost from selling in column $\phi = 0$ leads to an increase of 2.7%. In the final column $\alpha_3 = \phi = 0$, combining both changes, we see an increase of 8.2% in ownership.

The second row shows the same experiment for the overall migration rate. Here, the direct impact of $\alpha_3 = 0$ is larger than removal of the financial transaction cost, and it leads to roughly a 4% increase in overall migration. The third row finally conditions on owners in order to emphasize that most migratory movement comes from renters to start with, so the previous experiment masks a great deal of heterogeneity. Here, we see

a large increase of 140% in the migratory propensity of owners after removing the α_3 cost component. Overall, the first panel shows that both cost components affect both ownership and migration simultaneously, and that the impact of removing α_3 is more important.

The bottom panel of the table repeats this exercise while conditioning only on the *mover* population, that is, types $\tau = 0$. This helps to further clarify the top panel, where the results are from a mixture of stayers and movers. It is interesting to note that among the movers, the impacts on ownership rate are greater than for the population at large, whereas those on migration of owners are smaller throughout. The former is a consequence of *stayers* being overrepresented in the group of owners already, hence we see a smaller increase in ownership in the top panel. The latter effect has to do with this change in composition of renters versus owners: given that in the mover population we register a larger increase in ownership, the denominator in *Migration|Own* becomes larger, hence the value decreases. Finally, the fact that the increase in migration rate is identical across both panels shows that any additional mobility can come only from the subpopulation of mover types.

The conclusion from this exercise is that including utility cost α_3 over and above the assumed financial transaction cost ϕ (6%) in the model is important in order to fit the data, both in terms of propensity to move and in terms of changes in the ownership rate.

6.3 Owner regret

Individuals in the model decide whether to buy a house or rent based on a multifaceted tradeoff involving the size of available houses, their preference for owned property, and their expectations about future moves and prices. Buying a house is consequential for later mobility decisions, as we have seen, because owners face higher moving costs. A pertinent question in this context is whether owners actually *regret* having bought their house, if the state of the world changes in an unforeseen way? For example, imagine that expectations about future prices were wrong, in the sense that there is an unexpected shock. Imagine further that that owner finds themselves with a greatly devalued house—how big is the cost of being an owner in this new circumstance?

We want to assess this cost by looking at how much an owner in the shocked scenario would be willing to pay to become a comparable renter. Establishing comparability is important, because in general owners are of higher net worth than renters (by virtue of the downpayment requirement and subsequent capital gains). Therefore, the experiment proceeds by establishing an asset level a^* in the baseline such that, with some abuse of notation, $V^{\text{own}}(a^*) = V^{\text{rent}}(0)$, in other words where an owner's value is equal to a renter's value with zero assets. a^* is in general a negative number, proportional to the greater net worth of owners and the additional utility derived from owning. Then we shock either q or p at a certain age in a given region, and we want to know how much the owner would be willing to pay in order to convert to a renter, measured at the preshock reference level a^* .³²

³²Note that this experiment measures something different from the value of the option *sell* of the owner. Since, if it was optimal to become a renter by selling, they would of course make this choice. Here, we know

TABLE 11. Owners willingness to pay to convert to a comparable renter after an unexpected 10% reduction in price or income arises. In 1000 s of dollars.

Division	10% p -Shock	10% q -Shock
East North Central	16.4	-0.1
East South Central	15.6	-3.6
Middle Atlantic	19.7	-4.8
Mountain	19.3	-4.0
New England	23.4	-3.0
Pacific	25.8	-0.4
South Atlantic	20.5	0.5
West North Central	14.7	2.8
West South Central	14.9	-2.5

The results of this exercise are displayed in Table 11 for a permanent and unexpected reduction in the level of either q or p of 10% in a given division in year 2000. We observe first that in the column corresponding to the q shock, there are relatively small dollar amounts (table is in 1000 of dollars), and some are negative. The negative entries imply that the owner would not be willing to pay anything after income drops by 10%. The intuition here is that in the case of an income shock, the renter is equally affected, hence there is only a small desire to become, or no desire at all to become, a renter with zero assets and a significantly reduced income. This is also the case because owners still have their housing capital to fall back on, which is unaffected from the shock in this scenario. While we saw previously in Table 9 that renters respond stronger to shocks than owners in terms of increased emigration, for obvious reasons, the expected gains in terms of lifetime value do not seem to compensate an owner to want to switch to that particular type of renter in most regions and take advantage of less costly migration.

Quite different to this is the outcome of a price shock. We see in the first column of Table 11 that owners would be willing to pay amounts ranging from about 14,700 and up to 25,800 dollars in order to be converted to a renter with zero assets after the unexpected price reduction. This experiment shares some features of mortgage default (which is not modeled), in that the owner evaluates a “reset” option here: the mortgage debt burden is increased substantially by the price drop, so much so that the owner would be willing to pay substantial amounts to get converted to a renter with zero assets. Thus, to give an answer to the initial question of how much owners regret to having bought when things turn out not as expected, or on the contrary, how much they would value being a “free but asset-poor” renter again in case of a price shock, one could say that this lies in between 15 and 25 thousand dollars.

that before the shock, the owner is as well-off at a^* as the zero asset renter, and we want to know how this relationship changes in the shocked economy, again measured at a^* .

6.4 *The value of migration*

This section presents a measure of how much individuals care about having the option to migrate across regions. The question is motivated by relatively low migratory flows across regions, 1.32% of households per year, as initially stated. Do low flows imply low value? And how does this valuation depend on age, location, ownership status, and current state of the business cycle? This value is related to what in [Yagan \(2013\)](#) is called *migratory insurance*. Here, I do not infer this from the amount of migration after a local shock has occurred, *ex post*, but I consider how much individuals value to have the option to migrate *ex ante*, in case a shock were to occur.

I will attempt to answer this question by first simulating the model in the baseline equilibrium, that is, at observed regional prices and incomes and importantly, *with migration* as an option. Subsequently, I will compare this to a counterfactual equilibrium *without* the migration option, that is, migration is shut down in the entire economy.³³ The welfare measure in terms of compensating consumption is defined in Online Appendix F.

Before delving into the results, it is paramount to clarify the role of the partial equilibrium assumption in this counterfactual.³⁴ Assuming that the model is well specified, the structural parameters are such that given prices and incomes, the resulting decision rules of agents in the model are correct. This mapping from model to data was shown to be satisfactory in Section 5.3. What this model cannot deliver is a prediction of how the *exogenous* series for q and p would change if we were to abolish migration. To address this concern at least to some degree, I will present different scenarios of this counterfactual, a baseline version with prices unchanged, and a set of experiments with changed prices in the Online Appendix. There is no direct empirical guidance on the effects of such a drastic experiment on regional house prices and labor income levels, except maybe that *in general*, wage effects of immigration are *small*.³⁵ Absent such empirical evidence, I try to cover the most relevant cases. Version two thus decreases both regional incomes and prices by 1% relative to their observed trajectories after the shut-down of migration. This could arise as a result of decreased firm productivity from the lack of suitable (migrant) skills, which leads to lower disposable income in a given region, and hence a reduction in house prices. Version three simulates a bust scenario, where incomes decline by 5% and house prices by 10%. Summing up, this experiment is available in three versions:

³³Even though trade is absent from the model, it helps the interpretation to assume that existing trade channels between regions remain intact throughout the experiment, that is, we only expect changes from the absence of individuals changing location.

³⁴Thanks to an anonymous referee for helping me to clarify this point.

³⁵I take this insight from the literature that assesses the impact of immigrants on native wages, exemplified by, for example, [Dustmann, Glitz, and Frattini \(2008\)](#), [Dustmann, Frattini, and Preston \(2013\)](#), [Card \(2012\)](#), which finds negligible negative impacts of immigrants on wages of low-skilled workers, and slightly positive ones elsewhere along the income distribution. Needless to say that the current model is a much simplified version of those studies in terms of skill composition of the labor force and indeed wage determination—it abstracts from immigrants (i.e., different skills groups and wage determination) altogether, hence it is probably too remote in order to directly use their results.

1. Baseline $\{q_{dt}, p_{dt}\}_{t=1997}^{2012}$: Loss of migrants has negligible impact on regional prices.
2. -1% shock to $\{q_{dt}, p_{dt}\}_{t=1997}^{2012}$: Local productivity suffers a small loss.
3. $-5\%/ -10\%$ shock: Large productivity decline and amplified effect on house prices.

Starting with the scenario where prices are unaffected by the experiment, Table 12 provides the main results of the paper. In the first row, the consumption compensation demanded in the entire economy, for different subsets of the population, including young, old, owners at age 30, renters at age 30, individuals whose average z history is lower than the 20th percentile of the distribution of z histories (i.e., *poor* individuals), same for people above the 80th percentile, and finally the entire population. The values in the table stand for the per period increase in consumption that would make individuals indifferent between baseline and migration shutdown, as a percentage of what they had optimally chosen to consume in the policy environment. Hence it is a measure for their willingness to pay to maintain migration. For example, in the first row of Table 12, the group of young people (with age below half of their lifespan) would demand an increase of 30.8% of optimal per period consumption. This amount is lower for old people at 8%, which is intuitive since they forgo fewer periods where migration could have been optimally chosen. We can look at the same measure by ownership status at age 30: owners at that age demand 4.2% more consumption, while same aged renters demand more, that is, 21.7%. Part of this difference comes from the fact that stayer types know that they will not move, hence are more likely to buy, hence suffer less from a removal of the migration option. The next two columns condition the measure on position in the distribution of realized z draws. It is evident that people who have relatively favorable draws of z , value migration more in most regions, which has to do with the shape of the estimated transition matrix for movers, G_{move} as illustrated in Online Appendix C.1. (High z movers can expect another high z draw in the new location.) As an average, overall individuals treated in this experiment (column labeled ATE), the corresponding number is 19.2% of consumption compensation demanded.

In the same table, we continue with this exercise and split the sample by region to understand how heterogeneous those valuations are distributed. Perhaps not surprisingly, there is a lot of variation in how individuals feel about removal of the migration option by region. In general, living in a high-income, high-priced region like the Pacific accentuates the difference between young and old even more (54.7% versus -7.8% compensation demanded). Put simply, this is because high prices are good for owners, who are more likely to stay in the region no matter what, but bad for renters who cannot afford to buy a house. They would prefer to migrate at some point if necessary and, therefore, suffer disproportionately from the removal of the migration option. By way of summary of this table, individuals value having the option of migration to a large degree and in the range of -11% (owners aged 30 in South Atlantic) and up to 60% of per period consumption. Negative entries imply an *unwillingness* to pay for the migration option: this applies to groups who are (young) owners in high-priced regions.

TABLE 12. Consumption compensation demanded after migration shutdown in the baseline scenario (regional prices are unaffected by the shutdown of migration). The numbers in this table represent the required percentage scaling factor Δc by which optimal consumption under the shutdown policy would have to be increased in order for individuals to be indifferent to the baseline. Positive values indicate people disapproving the policy, negative values indicate the opposite. The columns are, in order, *young* (population below half of total lifetime J), *old* (complement to *young*), *own*, 30 (population who owns at age 30), *rent*, 30 (population who rents at age 30), z_q (population below/above the q -quantile of the lifetime idiosyncratic income shock distribution). A value of 30.8% as in the first row of the second column means that young people would demand an increase of 30.8% of the consumption level which they had optimally chosen under the policy, in order for them to be indifferent.

Region	Young	Old	own,30	rent,30	$z_{0.2}$	$z_{0.8}$	ATE
Aggregate	30.8%	8.0%	4.2%	21.7%	11.8%	21.0%	19.2%
East North Central	16.2%	0.4%	6.4%	11.2%	3.7%	7.9%	9.7%
East South Central	47.9%	1.4%	-0.6%	14.9%	25.0%	35.6%	39.7%
Middle Atlantic	34.7%	6.2%	3.7%	16.5%	7.0%	41.4%	29.9%
Mountain	21.5%	9.1%	3.0%	15.3%	13.3%	22.5%	11.7%
New England	60.9%	0.6%	6.4%	28.1%	12.9%	35.0%	40.6%
Pacific	54.7%	-7.8%	0.2%	19.1%	5.6%	-5.1%	-0.4%
South Atlantic	13.5%	4.0%	-11.0%	10.8%	6.9%	11.5%	9.5%
West North Central	28.8%	9.1%	0.3%	23.2%	13.5%	31.9%	24.3%
West South Central	17.2%	3.2%	0.6%	11.2%	15.3%	12.8%	15.1%

Versions two and three of this experiment are described in the Online Appendix E.3, where the general conclusions from this series of results goes through, with the qualification that the bigger the price shock after migration shutdown, the more individuals value the baseline. This sequence of counterfactuals tells us that the estimates from the baseline scenario with observed prices provide a lower bound. If in reality a general equilibrium effect would change regional incomes and prices downwards, the valuation of the migration option would be larger than displayed in Table 12.

6.5 Effects along the lifecycle and by ownership status

We have shown that lifetime utility changes dramatically with the removal of migration. We now go further and investigate where those changes come from, that is, how do the state variables of individuals in the model change? In the following, we focus only on the main counterfactual with constant prices.

Starting out with the lifecycle considerations, the results are presented in Figure 5. The first panel in the top row in some ways repeats the insights from the previous section: younger individuals suffer particularly, experiencing a loss as in lifetime utility of almost 4% in the first period of life. As time goes by, the losses get smaller until they vanish toward the end of the lifecycle. We observe in the next panel that removing migration implies a substantial drop in average income at all ages. This implies that some profitable moves in terms of better wage draw could not be completed as a result of the policy.




FIGURE 5. Impact of *migration shutdown* along the lifecycle. Each line shows the percentage difference between baseline conditional mean of a certain variable of interest and the conditional mean under the policy. The means are conditional on age. The labels stand for, in order: v value function, y is individual labor income, a is asset position, and h is the ownership status.

The next two panels for a and h are best viewed in conjunction, as they are tightly connected: After migration is abolished, the aggregate homeownership rate increases strongly for 30-year olds, and with it the outstanding mortgage balance as measured by negative net assets a . In light of the fixed regional price series $\{q_{dt}, p_{dt}\}_{t=1997}^{2012}$, this result may be somewhat surprising. We observe that after migration is abolished, homeownership rises by about 39% at age 30. Notice that this increase implies that far more than only the previous migrants now choose to buy, as this was relatively small group of people. The increase in ownership comes from the fact that potential future moves of *all individuals* have been ruled out and, therefore, a much larger number of them finds it profitable to take out a mortgage and buy in the current (constant) location. This suggests that absent the option to move to a better region in response to shocks, the best thing to do is to invest in more enjoyable housing in the current region. In particular, this implies that the aggregate mortgage balance increases dramatically, hence the $\% \Delta a$ panel shows negative percentage changes for most ages. The nonmonotonicity in that figure occurs because around age 37 the group of mortgage holders is considerably poorer than in the baseline, where the asset balance starts to turn positive for individuals with high labor incomes of that age.

In Figure 6, I look at the experiment by measuring the effects for individuals who did and did not own their house by the age of 30. It shows that the bulk of the previous lifecycle results must be driven by young renters, as opposed to owners. Young owners

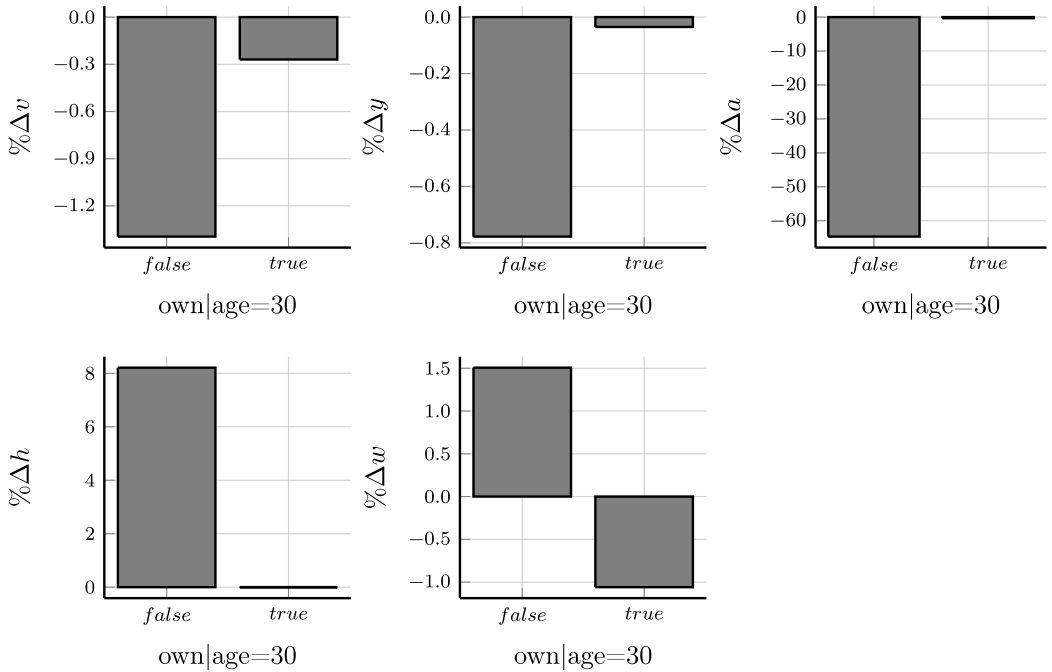


FIGURE 6. Impact of migration shutdown by ownership status at age 30. Bars show the percentage difference between the conditional mean of the respective variable conditional on ownership status at age 30 in the baseline, and the shutdown policy. The labels stand for, in order: v value function, u period payoff, y is individual labor income, a is asset position, h is the ownership status, and w is total wealth.

in Figure 6 have a smaller utility loss, as shown in panel v which displays lifetime utility. The second panel again illustrates that it is young renters who miss out on profitable moves in terms of individual income y . The remaining panels for a , h , and total wealth w all show once more that young renters take on more mortgage debt in order to buy houses in their current (now, permanent) region.

In summary, the results in this section show that moving costs differ greatly by ownership status; that the resulting elasticities of migration with respect to regional shocks differ greatly as well; and finally, that individuals place a large consumption value on having the option to migrate across census division borders, ranging from negative values (some groups prefer not to have the option) and up to about 19% of per period optimal consumption.

7. CONCLUSION

The main result of this paper is to show that despite average migration rates being low, the option value associated with the possibility to leave a location in a world with regional shocks to house prices and labor income is large. Removing the option to migrate in the model implies an associated reduction in per period consumption that ranges from 0% (or even implied increases in consumption) up to 19% depending on the type

of household one considers. Variations in this measure vary widely with age, ownership status, original location, and point in the business cycle. To arrive at this result, I construct a lifecycle model which includes homeownership as a choice variable next to savings and location choices, which I then fit to SIPP data and use to make counterfactual experiments. Considering homeownership is motivated by the fact that well over 60% of the US population are owners, and the observation that owners exhibit vastly different migratory behavior than renters. The model places particular emphasis on a close representation of the observed house price and income series, both of which exhibit strong correlation of regional shocks. These results resonate with the findings in Yagan (2013) in the sense that both papers provide an estimate of *migratory insurance*. Here, I provide a well-defined value in terms of consumption, whereas the unit of measurement is more abstract in Yagan (2013).

The present model delivers structural estimates for differential moving costs between owners and renters, over and above financial transaction costs. It is shown that this moving cost component is first order in explaining differential moving rates between renters and owners. The results are important for policy makers in terms of both housing and labor markets, since they illustrate the implied costs of policies which might reduce mobility, such as implicit or explicit subsidies to homeownership, for example.

REFERENCES

- Arcidiacono, P. and R. Miller (2011), “Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity.” *Econometrica*, 79, 1823–1867. [1456, 1471]
- Attanasio, O. and G. Weber (1995), “Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey.” *Journal of Political Economy*, 103, 1121–1157. [1476]
- Attanasio, O. P., R. Bottazzi, H. W. Low, L. Nesheim, and M. Wakefield (2012), “Modelling the demand for housing over the life cycle.” *Review of Economic Dynamics*, 15, 1–18. [1466]
- Bajari, P., P. Chan, D. Krueger, and D. Miller (2013), “A dynamic model of housing demand: Estimation and policy implications.” *International Economic Review*, 54, 409–442. [1471]
- Bartik, A. W. (2018), “Moving costs and worker adjustment to changes in labor demand: Evidence from longitudinal census data.” Technical Report, University of Illinois at Urbana-Champaign. [1456, 1461]
- Battu, H., A. Ma, and E. Phimister (2008), “Housing tenure, job mobility and unemployment in the UK.” *The Economic Journal*, 118, 311–328. [1457]
- Bernard, A. B., S. J. Redding, and P. K. Schott (2013), “Testing for factor price equality with unobserved differences in factor quality or productivity.” *American Economic Journal: Microeconomics*, 5, 135–163. [1465]

Bishop, K. C. (2008), "A dynamic model of location choice and hedonic valuation." Technical Report, Unpublished, Washington University in St. Louis. [1456, 1461]

Blanchard, O. J. and L. F. Katz (1992), "Regional evolutions." In *Brookings Papers on Economic Activity 2*, Vol. 1992 (O. J. Blanchard, L. F. Katz, R. E. Hall, and B. Eichengreen, eds.), 1–75, Brookings Institution Press. [1456]

Blanchflower, D. G. and A. J. Oswald (2013), "Does high home-ownership impair the labor market?" Technical Report, National Bureau of Economic Research. [1457]

Bonhomme, S. and J. Robin (2006), *Chapter 18: Modeling Individual Earnings Trajectories Using Copulas: France, 1990–2002*. Contributions to Economic Analysis, Vol. 275. Elsevier. [1475]

Card, D. (2012), "Comment: The elusive search for negative wage impacts of immigration." *Journal of the European Economic Association*, 10, 211–215. [1485]

Carroll, C. D. (2006), "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics letters*, 91, 312–320. [1471]

Clausen, A. and C. Strub (2013), "Envelope theorems for non-smooth and non-concave optimization." <https://andrewclausen.net/research.html>. [1471]

Coulson, N. E. and L. M. Fisher (2002), "Tenure choice and labour market outcomes." *Housing Studies*, 17, 35–49. [1457]

Demyanyk, Y., D. Hryshko, M. J. Luengo-Prado, and B. E. Sorensen (2017), "Moving to a job: The role of home equity, debt, and access to credit." *American Economic Journal: Macroeconomics*, 9, 149–181. [1456]

Diamond, R. (2016), "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000." *The American Economic Review*, 106, 479–524. [1456]

Dustmann, C., T. Frattini, and I. P. Preston (2013), "The effect of immigration along the distribution of wages." *The Review of Economic Studies*, 80, 145–173. [1485]

Dustmann, C., A. Glitz, and T. Frattini (2008), "The labour market impact of immigration." *Oxford Review of Economic Policy*, 24, 477–494. [1485]

Fallick, B. and C. A. Fleischman (2004), "Employer-to-employer flows in the US labor market: The complete picture of gross worker flows." Division of Research, Statistics and Monetary Affairs, Federal Reserve Board., 4. [1457]

Fella, G. (2014), "A generalized endogenous grid method for non-concave problems." *Review of Economic Dynamics*. [1471]

Ferreira, F., J. Gyourko, and J. Tracy (2010), "Housing busts and household mobility." *Journal of Urban Economics*, 68, 34–45. [1456]

French, E. (2005), "The effects of health, wealth, and wages on labour supply and retirement behaviour." *The Review of Economic Studies*, 72, 395–427. [1476]

- Gemici, A. (2007), *Family Migration and Labor Market Outcomes*. Ph.D. thesis, University of Pennsylvania. [1455]
- Güler, B. and A. Taskin (2018), “Homeownership and unemployment: The effect of market size.” *Labour Economics*, 54. [1457]
- Halket, J. and S. Vasudev (2014), “Saving up or settling down: Home ownership over the life cycle.” *Review of Economic Dynamics*, 17, 345–366. [1457]
- Head, A. and H. Lloyd-Ellis (2012), “Housing liquidity, mobility, and the labour market.” *Review of Economic Studies*, 79, 1559–1589. [1457]
- Iskhakov, F., T. H. Jørgensen, J. Rust, and B. Schjerning (2017), “The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks.” *Quantitative Economics*, 8, 317–365. [1471]
- Karahan, F. and S. Rhee (2019), “Geographic reallocation and unemployment during the Great Recession: The role of the housing bust.” *Journal of Economic Dynamics and Control*, 100, 47–69. [1457]
- Kennan, J. and J. Walker (2011), “The effect of expected income on individual migration decisions.” *Econometrica*, 79, 211–251. [1455, 1456, 1461, 1463, 1464, 1479]
- Koşar, G., T. Ransom, and W. Van der Klaauw (2019), “Understanding migration aversion using elicited counterfactual choice probabilities.” FRB of New York Staff Report. [1467]
- Krusell, P. and A. A. Smith (1998), “Income and wealth heterogeneity in the macroeconomy.” *The Journal of Political Economy*, 106, 867–896. [1466]
- Lamadon, T. (2014), “Productivity shocks, dynamic contracts and income uncertainty.” Manuscript, University of Chicago. [1477]
- Li, W. and R. Yao (2007), “A life-cycle model of housing and mortgage choices with labor income.” *Journal of Money, Credit, and Banking*, 1375–1409. [1476]
- Lkhagvasuren, D. (2012), “Big locational unemployment differences despite high labor mobility.” *Journal of Monetary Economics*, 59, 798–814. [1456]
- Lucas, R. E. and E. C. Prescott (1974), “Equilibrium search and unemployment.” *Journal of Economic Theory*, 7, 188–209. [1456]
- Molloy, R., C. L. Smith, and A. Wozniak (2011), “Internal migration in the United States.” *The Journal of Economic Perspectives*, 25, 173–196. [1457]
- Moretti, E. (2011), “Chapter 14—local labor markets.” Part B of *Handbook of Labor Economics*, Vol. 4, 1237–1313. [1456]
- Nenov, P. (2012), “Regional mismatch and labor reallocation in an equilibrium model of migration.” Norwegian Business School. [1457]
- Notowidigdo, M. J. (2011), “The incidence of local labor demand shocks.” Technical Report, National Bureau of Economic Research. [1456]

Oswald, A. J. (1996), "A conjecture on the explanation for high unemployment in the industrialized nations: Part 1." [1457]

Oswald, F. (2019), "Supplement to 'The effect of homeownership on the option value of regional migration.'" *Quantitative Economics Supplemental Material*, 10, <https://doi.org/10.3982/QE872>. [1458]

Piyapromdee, S. (2019), "The impact of immigration on wages, internal migration and welfare." Working paper, <https://sites.google.com/site/suphanitpiyapromdee/research>. [1456]

Ransom, T. (2018), "Labor market frictions and moving costs of the employed and unemployed." Technical Report, University of Oklahoma. [1456, 1461]

Roback, J. (1982), "Wages, rents, and the quality of life." *Journal of Political Economy*, 90, 1257–1278. [1456]

Rosen, S. (1979), "Wage-based indexes of urban quality of life." *Current Issues in Urban Economics*, 3. [1456]

Rust, J. (1987), "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." *Econometrica*, 55, 999–1033. [1469, 1476]

Şahin, A., J. Song, G. Topa, and G. L. Violante (2014), "Mismatch unemployment." *American Economic Review*, 104, 3529–3564. [1457]

Schulhofer-Wohl, S. (2011), "Negative equity does not reduce homeowners' mobility." Technical Report, National Bureau of Economic Research. [1456]

Sommer, K. and P. Sullivan (2018), "Implications of US tax policy for house prices, rents, and homeownership." *American Economic Review*, 108, 241–274. [1476]

Winkler, H. (2010), "The effect of homeownership on geographic mobility and labor market outcomes." SSRN eLibrary. [1455, 1456]

Yagan, D. (2013), "Moving to opportunity? Migratory insurance over the great recession." Harvard Job Market Paper. [1453, 1456, 1485, 1490]

Yagan, D. (2019), "Employment hysteresis from the great recession." *The Journal of Political Economy*. [1456]

Co-editor Peter Arcidiacono handled this manuscript.

Manuscript received 18 April, 2017; final version accepted 13 May, 2019; available online 24 May, 2019.