



HAL
open science

Regional Shocks, Migration and Homeownership

Florian Oswald

► **To cite this version:**

| Florian Oswald. Regional Shocks, Migration and Homeownership. 2015. hal-03459804

HAL Id: hal-03459804

<https://sciencespo.hal.science/hal-03459804>

Preprint submitted on 1 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Regional Shocks, Migration and Homeownership*

Florian Oswald[†]
University College London

June 22, 2015

To obtain the most recent version of the paper, please go to
<http://floswald.github.io/>

Abstract

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). Using the model I estimate the value of the migration option and the welfare impact of policies that may restrict mobility. The option to move is equivalent to 4.4% of lifetime consumption. I also find that, were the mortgage interest-rate deduction to be eliminated, the aggregate migration rate would increase only marginally by 0.1%. Following a general equilibrium correction, house prices are reduced by 5%, which results in a 1% increase in home ownership. In a new steady state the elimination of the deduction is equivalent to an increase of 2.4% of lifecycle consumption.

*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 and seminar participants at UCL, IFS and Royal Holloway 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. The author would like to acknowledge that the work presented here made use of the IRIDIS High Performance Computing Facility provided by the Science & Engineering South (SES) Centre for Innovation. The Centre is formed by the universities of Oxford, Southampton, Bristol and University College London.

[†]email: florian.oswald@gmail.com.

1 Introduction

Homeownership and the likelihood of moving across regions are negatively correlated. This negative correlation has important implications when considering the insurance mechanisms available to consumers in the face of regional shocks to housing and labor productivity in a world of incomplete markets. The possibility of moving to another region in the event of a shock is a way to self-insure against tail risk.

The aims of this paper are first to investigate the value of this self insurance mechanism, second to analyse the reaction of regional mobility to regional shocks and third to establish how important government policies such as the mortgage interest tax deduction, which encourages ownership, interact with regional migration. To address those issues, I propose a lifecycle model of consumption and savings, housing tenure and location choice, with aggregate uncertainty which affects prices in regions differentially. Migration serves as a partial insurance device because regions are exposed to aggregate shocks to a different degree, consistent with the observed data. The lifecycle considerations have important implications for the analysis of welfare impacts of policy changes, since age groups are differentially affected.

Regional shocks to labor demand and house prices may have profound and longlasting impacts, as we are still witnessing in the aftermath of the 2007 fall in house prices and ensuing Great Recession. In this episode we have seen wide variation in the magnitude of local shocks despite underlying high national correlations. For example, the peak-to-trough decline in the Federal Housing Financing Agency (FHFA) house price index from 2007 to 2011 ranged from -2.3% in Pittsburgh, to -61% in Las Vegas.¹ The decline of the automobile industry in Detroit provides a vivid example of the effects of a permanent shock to labor income in a region: from the 1950 to the 2010 censuses, the population of Detroit declined by 61%.

Shocks to regional labor demand affect renters and owners in a similar way, i.e. labor income of both falls. The effects of a fall in house prices is more nuanced, since renters stand to benefit from cheaper rent, if this is related to prices, while owners may loose wealth invested in the house. Furthermore, if house prices reflect the value of local amenities in some kind of way, the location becomes less attractive for renters as well as for owners. For those reasons, we should expect to see different reactions in the mobility of renters and owners in response to different shocks. Labor market shocks should make both types more likely to leave the region, while price shocks could have differential impacts, depending on the relative importance of cheap rents, low value housing asset, lower amenity values, and moving costs for renters and owners.

¹This refers to FHFA expanded house price index data for the 50 largest MSAs in the USA. Available at http://www.fhfa.gov/DataTools/Downloads/Documents/HPI/HPI_EXP_metro.txt

In 2013, 63% of occupied housing units in the US were owned, while 37% were rented.² At the same time, roughly 1.3% of the population migrate across US census division boundaries per year. Conditioning on ownership we find 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. 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 I develop, there are several mechanisms which affect the home ownership choice of individuals. A downpayment requirement means 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, which also influences the buying decision next to age, the probability of moving, and beliefs about future shocks.

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 an additional 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, as well as wealth accumulation over the lifecycle and by region. After fitting the model to the data, I first illustrate how adding ownership and assets to a dynamic migration model affects behaviour. I then show how the model reacts to large regional shocks, before I move on to compute a measure of the value of the migration option. Finally, the model is used to

²see American Community Survey 2013, table DP04.

perform counterfactual policy simulations.

I find that owners and renters face very different incentives in the model. The probability of moving conditional on being a renter is consistently higher than the one for owners, and it varies considerably with income and assets. The model can be used to quantify the moving cost in terms of dollars, and I find that the estimates for renters and owners bracket the moving cost obtained in [Kennan and Walker \(2011\)](#).

Migration is a low probability event in both data and model. Nevertheless, the welfare implications are large. I conduct an experiment whereby moving away from a certain region is prohibited. I compute the expected lifetime utility of agents in the region under both scenarios, and I find that removing the option to move away from the region reduces expected utility by 5.1%. This implies that residents of this region would demand a 4.4% increase in per period consumption over their lifecycle in order to be indifferent to the baseline environment. Conditioning on age when computing expected utility reveals that the cost of not being able to move is largest for younger individuals, and it gradually fades out as agents grow older.

Government policies might help to increase labor market efficiency, housing market inefficiency or both. The main inefficiencies any policy might target are incomplete asset markets, liquidity constraints, moving and transaction costs, income taxes and imperfect rental markets. One large pre-existing intervention in this context is the mortgage interest deduction. It is interesting to consider how abolishing this tax deduction might interact with borrowing, housing tenure choice and migration.

I find that abolishing the mortgage interest deduction would have a only a negligible impact on the aggregate migration rate. At first sight one might be tempted to think that removing the deduction would result in more renters, which would mechanically translate into an increase in migration, since renters move more. The actual effect is more nuanced. First, because the mortgage tax deduction is a large scale policy, there is likely to be a general equilibrium effect. While my model is a partial equilibrium model, I can approximate the general equilibrium effect by assuming that the policy changes both taxes as well as prices. I use recent results from a stationary GE model by [Sommer and Sullivan \(2013\)](#) who find that house prices fall, while rents keep constant, after the policy is changed. Applying this price and rent correction to my model results in an increase in the ownership rate of roughly 1%, because more households are able to buy at lower prices. The net effect on migration is very small (0.1% of the baseline rate). The reduction in mobility is smaller than what we would expect from a pure change in composition towards more owners, because migration behaviour changes as well. Lower prices and a higher level of disposable income for the poor after redistribution of tax proceeds in all regions change the incentives to exploit regional differences with respect to the baseline. In terms of welfare, households prefer removing the deduction and would agree to giving up 2.4% of period consumption before being indifferent to the status quo.

Literature. My paper builds on [Kenman 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. [Gemici \(2007\)](#) focuses 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 a recent paper that extends [Gemici \(2007\)](#) to include housing choices and focuses on the response of owners to individual labor market shocks.

The main differences to this paper 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 \(2010\)](#). I follow [Kenman and Walker \(2011\)](#) in assuming that individuals must visit a location in order to discover the exact value of their new wage, over and above a predictable part, while [Winkler \(2010\)](#) assumes that job offers arrive in the current location from a random alternative location. My assumption implies firstly that individuals consider all potential locations in each period, and decide to move based on their expectations about how they will fare in each. Secondly it allows for reasons other than job offers to trigger a move, which is a feature of the data, as I will show below. Finally, it is interesting to note that when simulating the effects of abolishing the mortgage interest deduction, I correct for a GE effect on house prices, which is likely to occur as a result of the policy change, and I enforce revenue neutrality by redistributing saved tax receipts, which is not done in his paper. I find that that welfare implications strongly depend on those features.

By considering regional shocks, this paper is also 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. Related to this, [Notowidigdo \(2011\)](#) analyses the incidence of local labor demand shocks on low-skilled workers in a static spatial equilibrium model and finds that they are more likely to stay in a declining city than high-skilled workers to take advantage of cheaper housing.³ 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.

Also related is a recent literature that considers the effects of the 2007 housing bust on labor market mobility. In terms of empirical contributions, [Ferreira et al. \(2010\)](#), [Schulhofer-Wohl \(2011\)](#) and [Demmyanyk et al. \(2013\)](#) look 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 the difference arising from different datasets and definitions of long-distance moves. More theoretical papers like

³See [Moretti \(2011\)](#) for a comprehensive overview of this literature going back to [Roback \(1982\)](#) and [Rosen \(1979\)](#), and [Diamond \(2012\)](#) and [Piyapromdee \(2013\)](#) for recent applications.

Head and Lloyd-Ellis (2012), Nenov (2012), Song et al. (2014) and Karahan and Rhee (2011) use 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 life-cycle framework.⁴

Finally, the paper relates to the literature on tax treatment of housing and ownership. The federal tax code in the United States allows households to deduct mortgage interest payments from Federal taxes. Glaeser (2011) and Glaeser and Shapiro (2002) respectively discuss the benefits and distortions generated by this policy, and Poterba and Sinai (2008) provide an estimate of the financial benefit to owners from it. Sommer and Sullivan (2013) is a recent contribution that analyses the policy in a GE framework. I evaluate the effect of removing the deduction on both homeownership and mobility.

This paper merges housing demand over the lifecycle, as for example Li et al. (2014), with dynamic migration decisions under aggregate and regional shocks. I find that housing is an important feature of the individual migration decision. The likelihood of moving depends not only on whether or not the individual owns their house and on their asset holdings, but also on the price and income levels in all potential destination regions. The value of the migration option to actual migrants is very large.

2 Empirical Background

The amount of regional migration in the US is still high by international standards. According to Molloy et al. (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 1.5% of the entire population moves between two out of four census regions, and about 1.3% move between states within any one region. At a more local level, they find that 5% of the population move 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). An overview of migration rates across different regional delineations and over different time spells is shown in table 1.

It is somewhat unfortunate that none of the datasets employed by Molloy et al. (2011) are very well suited for the purpose of analysing migration and ownership. None of them tracks movers, so it is impossible to know the circumstances of an individual at the moment they decided to move, which is

⁴In 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 Taskin (2011), Battu et al. (2008) or Halket and Vasudev (2014).

Period	Geography	% of US population
Annual	county	5
	state	2
	region	1.5
5-year	county	18.6
	state	8.9
	region	4.8

Table 1: Migration rates at different levels of geographic aggregation and over different time spells. Taken from [Molloy et al. \(2011\)](#), computed from ACS, March CPS and IRS data.

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

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-defined number of regions would be far too large to be computationally feasible. Hence the choice of census divisions.⁷

Descriptive Statistics on Cross Division Migration. I combine four panels of SIPP data (1996, 2001, 2004 and 2008) into a database with 102,529 individuals that I can follow over time and space. From this dataset, a couple of interesting facts emerge. Table 2 presents some summary annual moving rates for both state and Census Division level migration. The overall unconditional migration rate is

⁵It 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.

⁶See the Census bureau’s website at https://www.census.gov/geo/reference/gtc/gtc_census_divreg.html and figure 14 for a map.

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.

1.95% and 1.32% 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 non-reference persons, and potentially because of sample attrition. It is quite clear from table 2 that there is a marked distinction in the likelihood of moving across state as well as division boundaries between renters and owners, with 2.6% (1.85%) of renters versus 0.93% (0.7%) of owners moving across state (division) boundaries on average per year. In total I observe of 2684 cross Division moves made by 2329 unique individuals, implying multiple moves for some movers, see table 3.⁸

Moving on to migration by age, we can see in figure 1 firstly that renters are more likely to move at all ages, and secondly that there is a strongly declining age effect – younger individuals move more. Both of those are highly salient features of the data, and they are the main dimension along which my model’s performance is going to be evaluated. Finally, a summary regarding homeownership rates and median price to income ratios by Census division is presented as an average over the years 1997–2011 in table 4.

Determinants of Migration. The March Supplement to the Current Population Survey (CPS) contains several questions relevant for the study of migration. Here I analyse answers to the 2013 edition of the CPS to the question “What was the main reason for moving”. The results are displayed in table 5. 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 “housing” as their main motivation is roughly 24% of the total population of movers. The table also disaggregates the response to the question by the distance between origin and destination state (not Division), 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.

In table 6 I present estimates from a statistical analysis of the determinants of cross division moves from SIPP data. 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, 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

⁸Just for comparison, the estimation sample in Kennan and Walker (2011) is drawn from the geo-coded version of NLSY79 and contains 124 interstate moves.

	Cross State	Cross Division
Overall	1.95 %	1.32 %
Renter	2.60 %	1.85 %
Owner	0.93 %	0.70 %

Table 2: 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.

	1	2	3	4	5
Renter	1202	98	12	2	1
Owner	936	73	5	0	0

Table 3: Distribution of the number of moves per mover by homeownership status. Households are categorized into “Renter” or “Owner” based on their homeownership status at the *first* move. SIPP data.

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 \$1m 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 does not play a significant role in this specification.

Proportion of Cross-Division movers by age

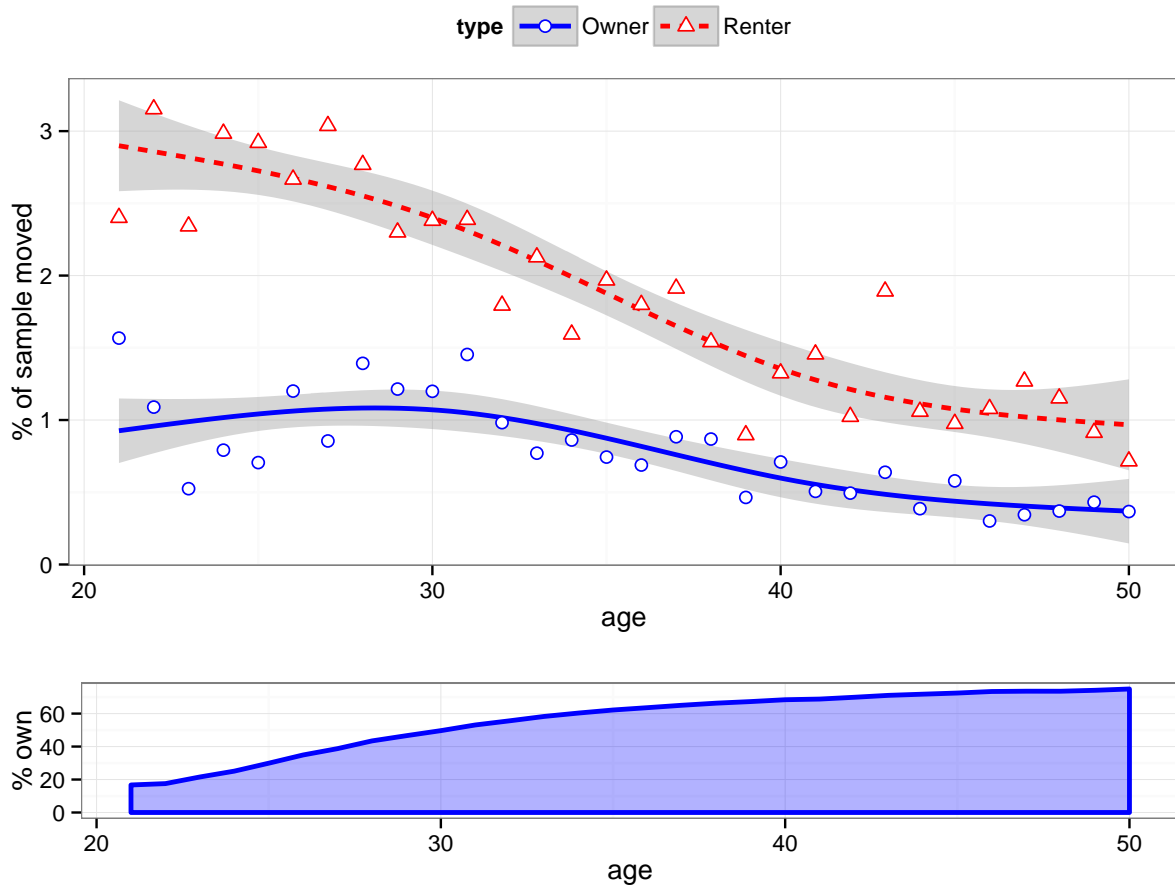


Figure 1: SIPP sample proportion moving across Census Division boundaries by age (upper panel) and proportion of owners by age (lower panel).

Division	Abbreviation	Ownership Rate	$\frac{p}{y}$
South Atlantic	StA	0.63	2.59
West North Central	WNC	0.69	2.08
East North Central	ENC	0.66	2.30
New England	NwE	0.60	2.99
Middle Atlantic	MdA	0.57	2.66
Pacific	Pcf	0.51	3.74
West South Central	WSC	0.60	1.95
East South Central	ESC	0.65	1.85
Mountain	Mnt	0.61	2.83

Table 4: Census Division housing characteristics. Shows average ownership rates over 1997–2011 and median price to income ratios for the same period. The (unobserved) house price for renters is computed assuming an implied user cost of owning of 5%, i.e. $p_{rent} = \frac{rent}{0.05}$.

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 %

Table 5: CPS 2013 data on main motivation of moving, conditional on a cross Division move. 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.

	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	28793.7099	
Dispersion	1.0261	
Num. obs.	294840	

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6: 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, 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%.

3 Model

In the model I view households as a single unit, and I'll use the terms *household* and *individual* interchangeably. Individuals are assumed to live in census Division (or *region*) $d \in D$ for a total of T years of age. At each age t , individual i has to decide whether to move to a different region, 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 an unobservable location preference shock.

Individuals are subject to uncertainty at both the aggregate and individual level. At the aggregate level, regional house price p and average labor productivity levels q fluctuate. This allows some scope for regional migration as an insurance mechanism. The regional fluctuations are driven by a common set of low dimensional stochastic factors denoted P and Q . This reflects the fact the regional shocks to both house prices and average labor income are highly correlated. It also allows for fluctuations in the underlying aggregate factors P and Q to have differential impacts across regions, while maintaining a degree of computational tractability. Every individual in region d faces an identical level of house price p_{dt} and mean labor productivity q_{dt} at time t . In addition to that, q_{dt} enters the individual wage equation as a level shifter. At the individual level uncertainty enters the model through an idiosyncratic component of income risk, a Markovian process that models changes in household size over the lifecycle, and a location-specific preference shock, which is assumed identically and independently distributed across agents, regions and time.

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.

3.1 Individual Labor Income

The logarithm of labor income of individual i at age t , residing in region d , is defined as in equation (1).

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

Here q_{dt} stands for the region specific price of human capital, $f(t)$ is a deterministic age effect 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.⁹

When moving from region d to region d' 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 d' , $q_{d't+1}$, the deterministic age $t + 1$ effect, $f(t + 1)$, 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 non-movers. 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 d' , as mentioned above. The shape of G_{move} determines the probability with which a mover with current value z can expect to draw a new value z' after moving to d' . The data will be informative about whether there is mean reversion or high persistence in G_{move} for movers.

3.2 National factors P and Q

I assume the national state variables Q and P evolve according to a stationary vector autoregression of order one (VAR(1)). At date t , all individuals observe the price vector \mathbf{F}_t containing both factors P_t and Q_t . The VAR(1) process is defined in equation (2), where 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.

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

⁹ Underlying this is an assumption about non-equalizing 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 for example [Bernard et al. \(2013\)](#)).

3.3 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 linear function that depends on both aggregate states Q, P and where the slope coefficients are region dependent, as shown in expression (3). Similarly to the aggregate case, \mathbf{a}_d is a matrix of coefficients specific to region d .

$$\begin{bmatrix} q_{dt} \\ p_{dt} \end{bmatrix} = \mathbf{a}_d \mathbf{F}_t \quad (3)$$

I provide some illustration regarding the fit this model provides to the data in section 5.1, where I describe the estimation of this part of the model in greater detail. 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 .

3.4 Home Ownership Choice

Ownership choice is discrete, $h_t \in \{0, 1\}$, and there is no quantity choice of housing. While renting, i.e. whenever $h_t = 0$, 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 [Atanasio et al. \(2012\)](#), I denote total financial (i.e. non-housing) wealth at age t as “assets” a_t . This includes liquid savings and mortgage debt. There is a terminal condition for net wealth to be non-negative by the last period of life, i.e. $a_T + p_{dT}h_{T-1} \geq 0$, 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, i.e. $a_{t+1} \geq -(1 - \chi)p_{dt}h_t, \forall t$. 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 t , such that they can satisfy the

final period constraint. The setup is formally defined in subsections 3.8.1 and 3.8.2 which describe the budget constraints.

3.5 Moving

Moving Costs. Moving across locations is assumed to be costly 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 psychological costs incurred from moving states which are common between renters and owners. I specify the moving cost function as linear in parameters α :

$$\Delta(d, x) = \alpha_{0,\tau} + \alpha_1 t_{it} + \alpha_2 t_{it}^2 + \alpha_3 h_{it-1} + \alpha_4 s_{it} \quad (4)$$

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_{it} .

The unobserved moving cost type $\tau \in \{0, 1\}$, where $\tau = 1$ indexes the high 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.

Restrictions. I rule out the possibility of owning a home in region d while residing in region d' . 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 d' . In my sample I observe less than 1% of movers for which this appears to be true, most likely as a result of high management fees or a binding liquidity constraint that forces them to sell the house to be able to afford the downpayment in the new region.¹⁰

3.6 Preferences

Period utility u depends on which region k the household chooses, 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

¹⁰SIPP 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.

case.

$$u(c_{it}, h_{it}, d', x_{it}) = \frac{c_{it}^{1-\gamma}}{1-\gamma} + \xi(s_{it}) \times h_{it} - \mathbf{1}[d \neq d'] \Delta(d', x_{it}) + \varepsilon_{ikt} \quad (5)$$

The period t payoff consists of utility from consumption c_{it} , from owning a house, valued differently at different household sizes s_{it} , and the idiosyncratic preference shock for the chosen region d' . Notice the moving cost $\Delta(d', x_{it})$ is only incurred if in fact a move takes place. Household size s at age t 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 3.8.

3.7 Timing

Timing within the period is assumed to proceed in two sub-periods: 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:

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

There is in fact no timing issue in point 3., as the optimization problem is simultaneous. However, the quasi-chronological order is helpful for emphasizing the interdependence of one choice on another.

3.8 Recursive Formulation

It is now possible to formulate the problem recursively. The state vector of individual i at date t is given by

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

where the variables stand for, in order, assets, individual income shock, household size, aggregate price vector, housing status, current region index, moving cost type and age. Following Rust (1987), I assume 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 $\varepsilon \sim \text{EV Type 1}$.

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. 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 d' , net of idiosyncratic location shock ε .

Equation (6) is the top level problem of the consumer which requires to choose one of D potential locations. Notice that the value from each discrete choice is additively seperable in the conditional value v and the choice specific idiosyncratic shock ε . The conditional value function (7) represents the choices that have to be made conditional on being in a given location d while moving to location d' . This formulation conveniently nests all discrete location choice configurations (staying in d and moving from d to d' , $\forall d' \neq d$). The optimization problem is subject to several constraints which are outlined below.

Equation (8) is a result of the distributional assumption on ε , which admits a closed form expression of the expected value function, whereby $\bar{\gamma} \approx 0.577$ is Euler's constant.

$$V_t(x_{it}) = \max_{k \in D} \{v_t(x_{it}, d') + \varepsilon_{ikt}\} \quad (6)$$

$$v_t(x_{it}, d') = \max_{c_{it} > 0, h_{it} \in \{0,1\}} u(c_{it}, h_{it}, d_{it}, d') + \beta \mathbb{E}_{z,s,\mathbf{F}} [\bar{v}_{t+1}(x_{it+1}) | z_{it}, s_{it}, \mathbf{F}_t] \quad (7)$$

$$x_{it+1} = (a_{it+1}, z_{it+1}, s_{it+1}, \mathbf{F}_{t+1}, h_{it+1}, d', t + 1)$$

$$\begin{aligned} \bar{v}_t(x_{it}) &= E_\varepsilon V_t(x_{it}) \\ &= \bar{\gamma} + \ln \left(\sum_{d'=1}^D \exp(v_t(x_{it}, d')) \right) \end{aligned} \quad (8)$$

Another convient by-product of the Type 1 EV assumption is that there is a closed form expression for the conditional choice probability of making a move from d to k when the state is x , denoted as $\mathcal{M}(x, d, d')$.

$$\begin{aligned} \mathcal{M}(x, d, d') &= \Pr [\text{move to } d' | x, d] \\ &= \frac{\exp(v(x, d'))}{\sum_{d'=1}^D \exp(v(x, d'))} \\ &= \frac{\exp(v(x, d'))}{\exp(\bar{v}(x)) / \exp(\bar{\gamma})} \\ &= \exp(\bar{\gamma} + v(x, d') - \bar{v}(x)) \end{aligned} \quad (9)$$

The final period models a terminal value that depends on net wealth and a term that captures future

utility from the house after period T , as shown in equation (10).

$$V_T(a, h_{T-1}, d) = \frac{(a + h_{T-1}p_{dT})^{1-\gamma}}{1-\gamma} + \omega h_{T-1} \quad (10)$$

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.

3.8.1 Budget constraint for stayers, i.e. $d = d'$

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 $h_{it-1} = 0$) depends on their housing choice, as shown in equation (11). In case they buy at date t , i.e. $h_{it} = 1$, 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 (12) 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 (13), 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 non-negative otherwise. Constraint number (14) 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. The terminal condition constraint is in expression (15).

$$a_{it+1} = (1 + r(a_{it})) (a_{it} + y_{idt} - c_{it} - (1 - h_{it})\kappa_d p_{dt} - h_{it}p_{dt}) \quad (11)$$

$$\ln y_{idt} = \eta_d \ln(q_{dt}) + f(t) + z_{it} \quad (12)$$

$$a_{it+1} \geq -(1 - \chi)p_{dt}h_{it} \quad (13)$$

$$r(a_{it}) = \begin{cases} r & \text{if } a_{it} \geq 0 \\ r^m & \text{if } a_{it} < 0 \end{cases}, r^m = r + q \quad (14)$$

$$a_{iT} + p_T h_{iT-1} \geq 0 \quad (15)$$

Owners. For individuals entering the period as owners ($h_{it-1} = 1$), 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, i.e. an increase in a , arises naturally from the terminal condition $a_{iT} + p_{dT}h_{iT-1} > 0$, as mentioned above. Therefore the budget of the owner is only affected by the house price in case they decide to sell the house, i.e. if $h_{it} = 0$. 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 (14), labor income equation (12) and terminal condition (15) apply.

$$a_{it+1} = (1 + r(a_{it})) (a_{it} + y_{idt} - c_{it} + (1 - h_{it})(1 - \phi - \kappa_d)p_{dt}) \quad (16)$$

$$a_{it+1} \geq -(1 - \chi)p_{dt} \quad (17)$$

3.8.2 Budget constraint for movers, i.e. $d \neq d'$

Renters. For moving renters the budget constraint is close to identical, with the exception that (11) 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.

$$a_{it+1} = (1 + r(a_{it})) (a_{it} + y_{idt} - c_{it} - (1 - h_{it})\kappa_{d'}p_{d't} - h_{it}p_{d't}) \quad (18)$$

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 (19) relates to proceeds from sale of the house in region d , whereas the square brackets describe expenditures in region d' . Notice also that the borrowing constraint (20) now is a function of the value of the new house in d' . 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, i.e. if $y_{idt} < a_{it} + (1 - \phi)p_{dt} - \chi h_{it}p_{d't}$ then the budget set is empty and moving and buying is infeasible.¹¹

¹¹In 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.

$$a_{it+1} = (1 + r(a_{it})) (a_{it} + y_{idt} - c_{it} + (1 - \phi)p_{dt} - [(1 - h_{it})\kappa_{d'} + h_{it}] p_{d't}) \quad (19)$$

$$a_{it+1} \geq -(1 - \chi)p_{d't}h_{it} \quad (20)$$

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 non-differentiabilities in the asset dimension of the value function. This makes using fast first order condition–based approaches to solve the consumption problem more difficult.¹²

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 final period value is computed at all relevant states. From period $T - 1$ 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). After this solution is obtained at a certain state, the discrete housing choice is computed, after which each conditional value function (7) is known. The distributional assumption on ε implies that the discrete location choice does not have to be computed, instead the main object of interest is the probability of moving function (9).

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 I will explain in greater detail 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.

¹²There has recently been a lot of progress on this front. [Clausen and Strub \(2013\)](#) provide 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 \(2014\)](#) are promising avenues. I found my problem not easily amenable to their solution, and focused on a robust solution (i.e. one not subject to potential local minima to the savings problem).

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. A set of summarizing features from the artificial data should then be close to the same features of the the real data. 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 to construct the house value for each year.¹³

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

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

The estimates from this equation are given in table 7.

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

¹³The 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.

	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

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 7: Estimates for Aggregate VAR process

estimates from the census Bureau. The price series by region comes from the same FHFA dataset as used above.

$$\begin{aligned}
 \begin{bmatrix} q_{dt} \\ p_{dt} \end{bmatrix} &= \mathbf{a}_d \mathbf{F}_t + \eta_{dt} \\
 \eta_{dt} &\sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Omega_d \right)
 \end{aligned} \tag{21}$$

The performance of this model in terms of delivered predictions from the aggregate state can be gauged visually in figures 2 and 3. The model parameters are shown in table 18 in the appendix.

Individual Income Process

This part deals with the empirical implementation of equation (12), which models log labor income at the individual level. I estimate the linear regression

$$\begin{aligned}
 \ln y_{idt} &= \beta_0 + \eta_d \ln q_{dt} + \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \beta_3 \text{age}_{it}^3 + u_{it} \\
 &= \beta_0 + (\gamma_d \times \gamma_q) \ln q_{dt} + \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \beta_3 \text{age}_{it}^3 + u_{it}
 \end{aligned} \tag{22}$$

where the region-specific influence of regional mean productivity q_{dt} on individual income is specified as an interaction between a regional fixed effect γ_d and the average effect γ_q of regional income. The results of this are displayed in table 19 in the appendix. Figure 15, also in the appendix, illustrates predicted age profiles from this model.

VAR fit to regional productivity data (q)

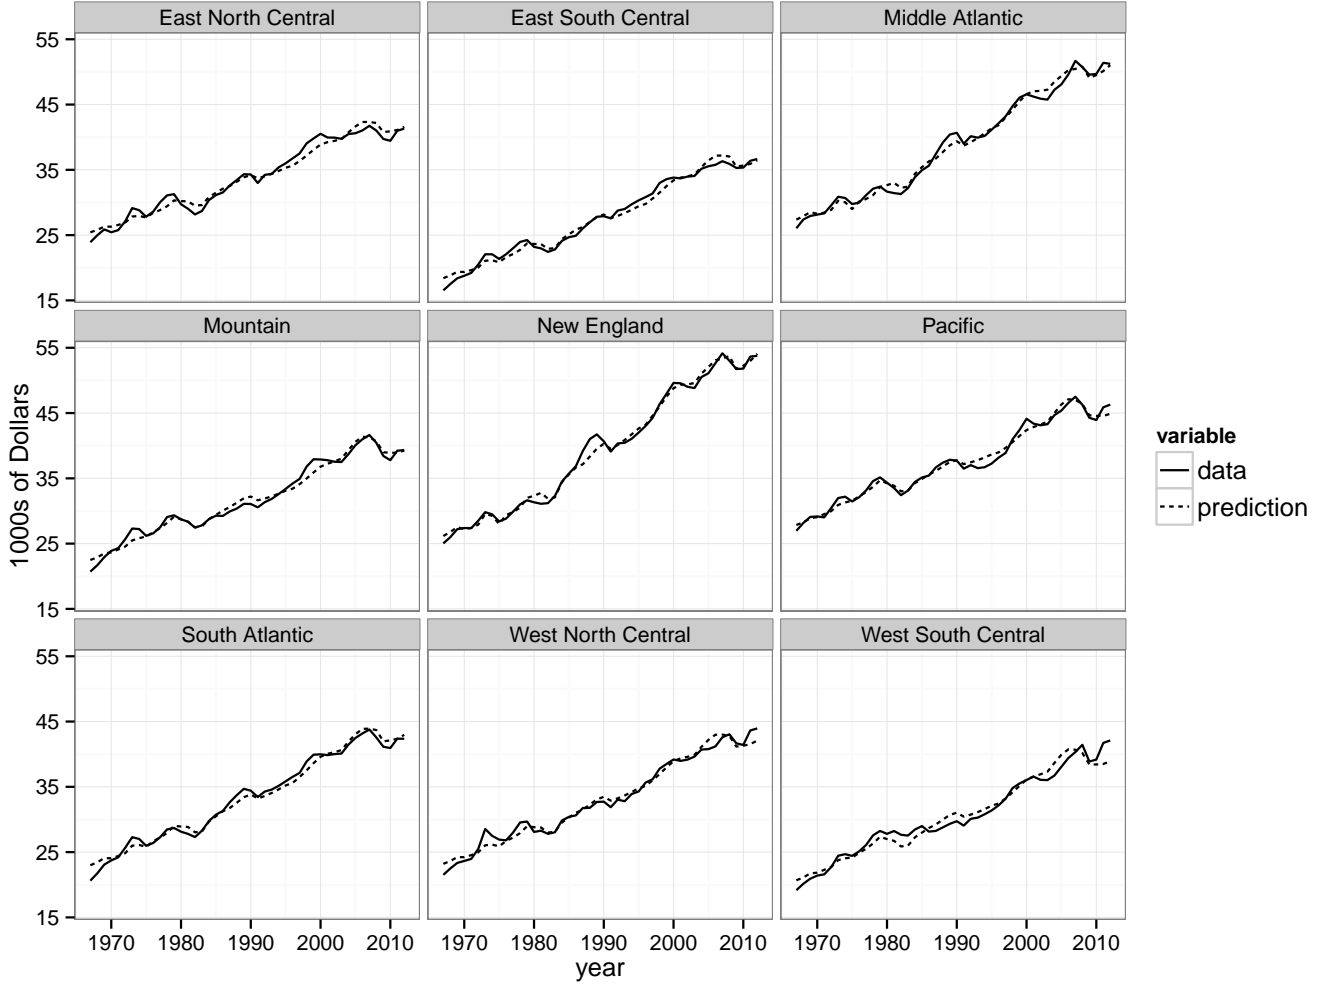


Figure 2: This figure shows the observed and predicted time series for mean income by Census Division. The prediction is obtained from the VAR model in (3), which relates the aggregate series $\{Q_t, P_t\}_{t=1968}^{2012}$ to mean labor productivity $\{q_{dt}\}_{t=1968}^{2012}$ for each region d . Agents use this prediction in the model, i.e. from observing an aggregate value $\mathbf{F}_t = (P_t, Q_t)$ they infer a value for q_{dt} for each region above.

VAR fit to regional price data (p)

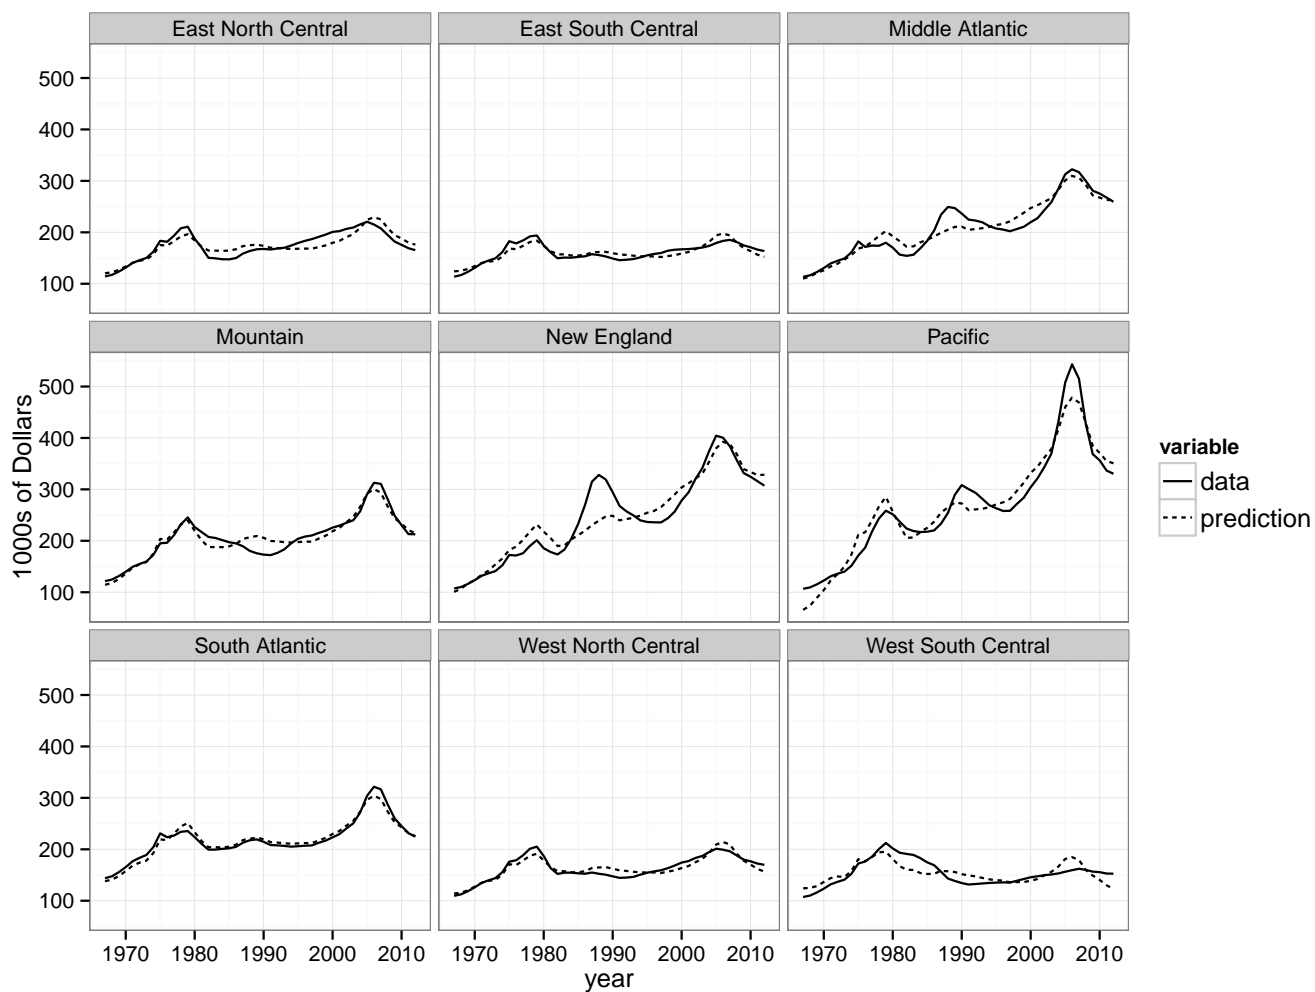


Figure 3: This figure shows the observed and predicted time series for house prices by Census Division. Please refer to the previous figure 2, which uses an identical procedure.

Copula estimates for G_{move}

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.¹⁴ This means I assume that the conditional probability of drawing z' in new region d' is the same regardless the origin location. It would be straightforward to relax this assumption, but data limitations forced me to impose this restriction.

To estimate the parameters of the copula, I view z_{idt} in equation (1) as the residual from an ordinary least squares regression of log wages on time and region effects, as well as person specific demographic variables. The question is whether individuals with a particularly high residual z_{idt} are likely to have a high residual $z_{id't+1}$ after their move to region d' . In other words, we want to investigate the joint distribution of $(z_{idt}, z_{id't+1})$. I describe the full procedure in the appendix.¹⁵

Values for preset parameters

I take several parameters for the model from the literature, as shown in table 8. The estimates for the components of the idiosyncratic income shock process for non-movers, i.e. 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 (2013), 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 thirty 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 $r^m = 0.055$.

¹⁴A 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).

¹⁵The procedure relies crucially on the assumption that individuals have to move to the new region before they can discover z_{t+1} . I am investigating ways to account for a potential selection effect on z_t by moving estimation of this part into the structural model and jointly estimate behavioural and wage related parameters. The model provides a set of exclusion restrictions that would allow to do this in theory. Identification of a potential selection effect may be difficult, however, because the sample of movers is small.

		Value	Source
CRRA 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. (2013)
30-year mortgage rate	r^m	0.055	Sommer et al. (2013)

Table 8: Preset parameter values

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), and the utility derived from housing for both household sizes, (ξ_1, ξ_2) . We'll denote the parameter vector $\theta = \{\alpha_0, \alpha_1, \alpha_2, \alpha_3, \alpha_4, \omega, \pi_\tau, \xi_1, \xi_2\}$.

Given θ , the model generates a set of model moments $\hat{m}(\theta)$, where $\hat{m}(\cdot)$ is of dimension K . After obtaining the same set of moments m from the data, the SMM procedure seeks to minimize the criterion function

$$L(\theta) = \frac{1}{2} [m - \hat{m}(\theta)]^T W [m - \hat{m}(\theta)],$$

which delivers point estimate $\hat{\theta} = \arg \min_{\theta} L(\theta)$. The weighting matrix W is formed of the inverse of the diagonal elements of the variance-covariance matrix of data moments Ω , i.e. I set $W = \left(\text{diag}(\hat{\Omega})\right)^{-1}$.

I obtain point estimates for θ by following a modification of the pseudo-likelihood estimator as introduced in Chernozhukov and Hong (2003). The modification defines N parallel instances of Chernozhukov and Hong (2003)–MCMC chains, which are differently “tempered” (i.e. they have different shock variances and rejection criteria), and are able to communicate with each other. This helps to explore large areas of the parameter space and avoids getting stuck in local minima. The procedure is formally defined in Baragatti et al. (2013).¹⁶ The quasi-posterior mean and confidence intervals are computed from the chain with lowest temperature analogously to Chernozhukov and Hong (2003), after accounting for the fact that I do not use the optimal weighting matrix.

¹⁶I've co-authored a software package that implements the procedure at <https://github.com/floswald/MOpt.jl>.

Estimation Sample

My estimation sample is formed mainly out of averages over SIPP data moments covering the period 1997–2012. 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.¹⁷ 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. I describe the procedure in detail in the appendix.

Additional to SIPP data, the moment vector contains three statistics from Kennan and Walker (2011), which relate to the lifetime frequency of moves (“moved never”, “moved once” and “moved more two times or more”). The reason for using external data is that I cannot compute such a statistic from SIPP, where the maximum panel length is four years.

The moment vector m is shown in table 9. It contains conditional means and covariances, which are largely self-explanatory. I introduce here two auxiliary models included in m which relate to the age profiles of both migration and ownership. Both models 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 $\text{move}_{it} = \mathbf{1}[d_{it} \neq d'_{it}]$:

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

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

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.¹⁸

¹⁷<http://research.stlouisfed.org/fred2/series/CPIAUCSL>

¹⁸The model is not non-parametrically identified; Both variation in prices and further restrictions such as functional form are needed, because price variation is at the regional (and not household-) level.

The structural parameters in θ are related to the moment vector $m(\theta)$ in a highly non-linear 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 $\alpha_{0,\tau=0}, \alpha_3, \alpha_4$ 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 $Cov(\text{move}, h)$ and $Cov(\text{move}, s)$, 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 $E[\text{moved twice+}]$ help to identify all moving cost parameters.

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 of owning with household size, and the age-profile parameters from the auxiliary model of ownership in (23).

5.3 Parameter Estimates and Moments (Preliminary)

The model fits the data moments fairly well overall. The fit is displayed in table 9. The upper panel shows moments related to mobility, the lower panel shows moments related to homeownership. Regarding mobility, the fit is very good. The only statistic slightly out of line is the frequency distribution of moves per mover. There is no mechanism in the model that could generate the observed pattern in the data. One possibility would be to augment the set of moving cost types with a third type that has an even lower moving cost. The estimates for the auxiliary model defined in (24) representing the age profile in ownership also provide a good fit to the data.

Moving on to moments related to ownership, we see that the unconditional mean of ownership is identical to the data moment. Condition by region provides a more varied picture, with some regions overestimated and others underestimated. The reason for this is that there is heterogeneity in ownership rates by region which is not easily accounted for by the fundamentals of regional house price and mean income alone.¹⁹ Remember that by taking prices and incomes as given, the model is restricted to only

¹⁹There is large degree of house price heterogeneity at the local level with is not in the model but which contributes to

few levers that affect the homeownership rate. The main parameters in this respect are the utility premia ξ_1, ξ_2 and the weight in the final period utility ω . The model at the moment overpredicts ownership in later periods of life. This is visible from the intercept of the auxiliary model (23), which relates the ownership rate to an age profile. The reason for this is that in a model where age and wealth are the main dimensions of variation across households, as soon as a certain wealth threshold is crossed, all agents become owners. In other words, the model cannot account for wealthy households who prefer not to own. One way to improve in this dimension would be to introduce different types of housing preferences.

Given that the CRRA coefficient γ is taken as fixed in the current implementation of the model, the moments relating to wealth resulting from the model can be viewed as some form of model validation. The model moments in table 10 are not included in the SMM objective function, that is, they are not targeted by the estimation algorithm. Despite this, they are very close the data counterparts.

The estimated parameters and standard errors are shown in table 11.

6 Properties of the model

In this section I use model to illustrate several mechanisms which are implied by the inclusion of homeownership status and assets into a dynamic migration model à la Kennan and Walker (2011).

6.1 Probability of Migrating

In figures 4, 5 and 6 I plot the conditional mean of the moving probability function (9), conditioning on wealth, income (y), income shock (z), and by housing status (h). The plots show throughout a marked difference by housing status h , which is fully expected given the data. The average probability of moving is decreasing in the level of wealth of a household. Figure 4 illustrates that increasing wealth makes households more likely to become owners, and at the same time decreases the probability of moving.

6.2 The role of moving costs

Moving costs in the model are measured in terms of utility. I convert them to dollar values by finding the amount of compensating assets a' which would make an individual at state x indifferent between

the average ownership rate at the regional level. Local building regulations, rent control or certain topographical features all influence the actual house price that the local level; The price index used in the model incurs some unavoidable aggregation error in this respect, and the same holds for my estimate of the average rent to price ratio.

Moments related to mobility		
Moment	Data	Model
$E[\text{move}]$	0.013	0.013
$E[\text{move} T]$	0.005	0.002
$E[\text{move} s = 0]$	0.014	0.017
$E[\text{move} s = 1]$	0.01	0.01
$E[\text{move} h_{t-1} = 0]$	0.019	0.026
$E[\text{move} h_{t-1} = 1]$	0.007	0.005
$Cov(\text{move}, h)$	-0.003	-0.005
$Cov(\text{move}, s)$	-0.001	-0.001
$E[\text{moved never}]$	0.83	0.88
$E[\text{moved once}]$	0.07	0.08
$E[\text{moved twice+}]$	0.09	0.04
Auxiliary model (24): $\text{move}_{it} = \beta_{0,m} + \beta_{1,m}t_{it} + \beta_{2,m}t_{it}^2 + u_{it}$		
$\beta_{0,m}$	0.06	0.05
$\beta_{1,m}$	-0.002	-0.003
$\beta_{2,m}$	0.00001	0.00006
Moments related to homeownership		
$E[h_{t-1}]$	0.61	0.63
$E[h_{t-1} \text{ENC}]$	0.66	0.7
$E[h_{t-1} \text{ESC}]$	0.65	0.64
$E[h_{t-1} \text{MdA}]$	0.57	0.64
$E[h_{t-1} \text{Mnt}]$	0.61	0.63
$E[h_{t-1} \text{NwE}]$	0.6	0.55
$E[h_{t-1} \text{Pcf}]$	0.51	0.5
$E[h_{t-1} \text{StA}]$	0.63	0.62
$E[h_{t-1} \text{WNC}]$	0.69	0.69
$E[h_{t-1} \text{WSC}]$	0.6	0.67
$E[h_{t-1} s = 0]$	0.53	0.58
$E[h_{t-1} s = 1]$	0.66	0.66
$E[h_{t-1} = 1, h_t = 0 T]$	0.01	0.02
$Cov(h_{t-1}, s)$	0.03	0.02
Auxiliary model (23): $h_{it-1} = \beta_{0,h} + \beta_{1,h}t_{it} + \beta_{2,h}t_{it}^2 + u_{it}$		
$\beta_{0,h}$	-1.146	0.051
$\beta_{1,h}$	0.08	0.023
$\beta_{2,h}$	-0.0008	0.0004

Table 9: Empirical targets and corresponding model moments.

Non-targetted moments		
Moment	Data	Model
$E[\text{wealth} t \in [20, 30]]$	51.19	43.949
$E[\text{wealth} t \in (30, 40]]$	130.989	102.34
$E[\text{wealth} t \in (40, 50]]$	209.317	214.823
$E[\text{wealth} \text{ENC}]$	139.125	118.858
$E[\text{wealth} \text{ESC}]$	109.666	104.056
$E[\text{wealth} \text{MdA}]$	165.388	150.903
$E[\text{wealth} \text{Mnt}]$	128.192	138.346
$E[\text{wealth} \text{NwE}]$	203.125	166.753
$E[\text{wealth} \text{Pcf}]$	183.162	182.781
$E[\text{wealth} \text{StA}]$	142.203	139.88
$E[\text{wealth} \text{WNC}]$	142.603	111.702
$E[\text{wealth} \text{WSC}]$	100.025	96.389
$E[\text{wealth} h_{t-1} = 0]$	33.172	30.817
$E[\text{wealth} h_{t-1} = 1]$	219.356	195.733

Table 10: Non-targeted model and data moments. This set of moments does not enter the SMM objective function and can thus be seen as a form of external validation of the model.

		Estimate	Std. error
Utility Function			
Owner premium size 1	ξ_1	0.012	0.00053
Owner premium size 2	ξ_2	0.052	0.00234
Moving Cost Function			
Intercept	α_0	2.77	0.124
Age	α_1	0.017	0.00077
Age ²	α_2	0.001	$9.16e-5$
Owner	α_3	0.26	0.0116
Household Size	α_4	0.36	0.016
Proportion of high type	π_τ	0.68	0.03
Final Period			
Continuation Value	ω_2	5.1	0.22

Table 11: Parameter estimates. Standard errors are still work in progress at this point.

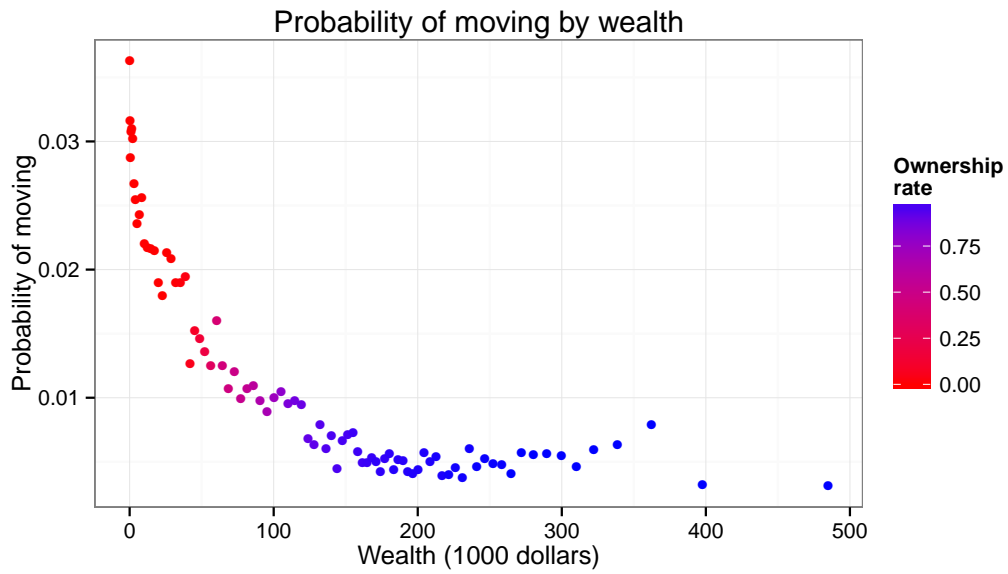


Figure 4: Simulated probability of moving by wealth. The figure shows that as wealth increases, individuals are both more likely to be owners (color shading), as well as less likely to move.

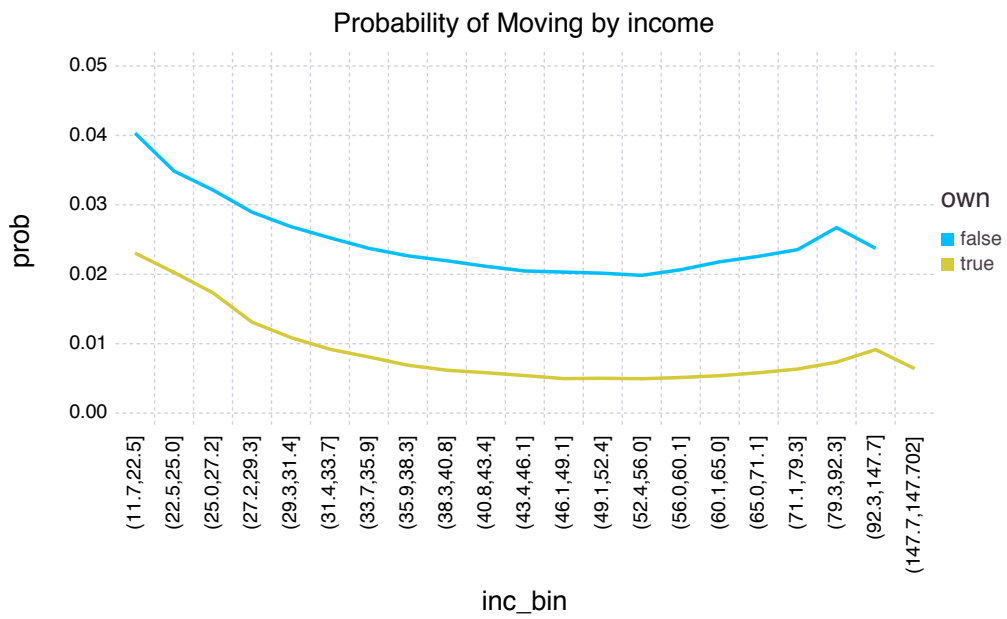


Figure 5: Average probability of moving by income quantiles.

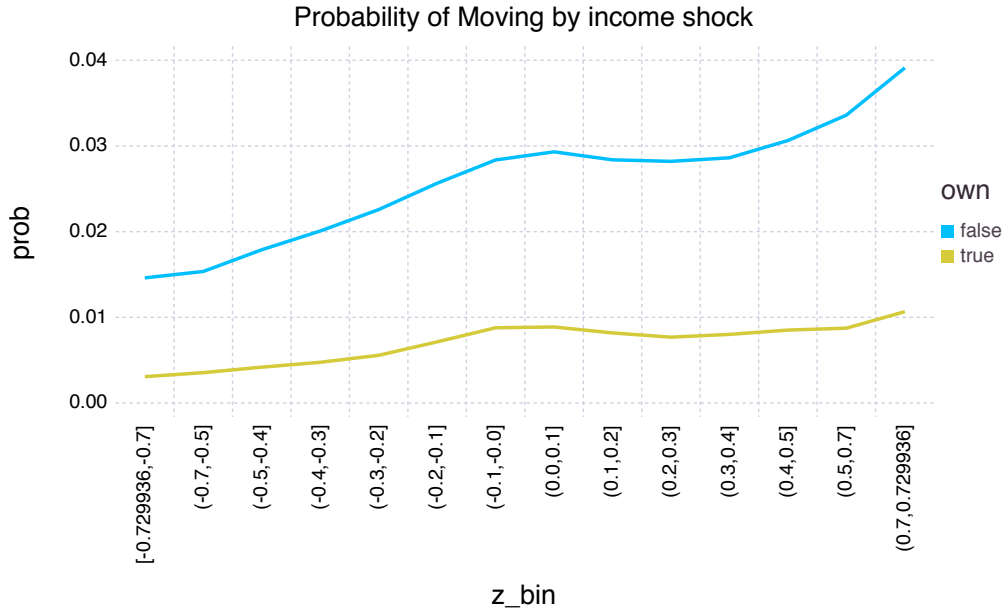


Figure 6: Probability of moving by income shock z

the value of moving with costs (v) (but extra assets) and without (\tilde{v}). I consider moves from region $d = 2$ to $k = 1$. In other words, moving costs are measured by a' which solves

$$v(2, 1, x + a') - \tilde{v}(2, 1, x) = 0.$$

The measure is taken at a state x where the individual has zero asset and where the aggregate factors P and Y are at the midpoint of their respective grids.

Similarly to [Kennan and Walker \(2011\)](#), the moving costs are large. As they explain, this is because “[the estimated moving costs] do not refer to the costs of moves that are actually made, but rather to the costs of hypothetical moves to arbitrary locations.”

The results are displayed in [table 12](#). To calculate moving costs from my nonlinear model, I use a money metric measure. The moving costs for renters at age 20 ($t = 1$) are \$159,000 and \$575,000 for owners. The estimates in period one bracket the average cost of \$312,000 reported in [Kennan and Walker \(2011\)](#).

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), which affects the value of migration. Third, owners

Renter	Owner
159.486	574.97

Table 12: Comparing the moving cost in terms of 1000s of 2012 dollars for owners and renters at $a = 0, z = z_1, Q = Q_2, P = P_2$ in the first period of life (age 20). [Kennan and Walker \(2011\)](#) report an average moving cost of \$384,000 for a young and \$312,000 for an average mover.

	Baseline	$\alpha_3 = 0$	$\phi = 0$	$\alpha_3 = \phi = 0$
Ownership rate	0.602	0.613	0.626	0.64
% Δ Migration rate	0	4.202	4.303	10.354
% Δ Migration Own	0	30.115	21.47	67.664

Table 13: Decomposing owner’s moving costs. 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 ($\alpha_3 = \phi = 0$).

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. Finally, ownership is correlated with larger household size ($s = 1$), which itself carries a higher moving cost (α_4).

In table 13 I decompose the reduced mobility of owners arising from higher utility costs α_3 and the fixed cost of selling ϕ . Ownership increases as we successively remove frictions along the first row of the table, as expected. Perhaps more surprisingly, the second row shows that the average migration rate increases by roughly the same amount if we remove the owner’s moving cost α_3 or the fixed cost of selling ϕ . This does not mean that financial frictions are of equal importance for mobility than non-financial ones. As shown in the final two rows of the table, this result is driven by a change in composition, i.e. the aggregate migration rate is higher because we have more owners, who also move more. The final row shows that the increase in migration resulting from zero non-financial costs is greater than the increase from no financial transaction costs.

6.3 The effects on migration of a regional shock

In this section I illustrate how the model reacts to large economic shocks at the regional level. The presentation focuses on changes in migratory in and outflows to and from a region, when the regional price or income level is reduced unexpectedly at a certain point in time. The implementation of the shock is symmetric for both the case of house price and mean income, and I will use the case of the regional productivity shock to explain the details.

The regional shock consists of reducing the level of the observed mean income series by 10% in the West North Central division only (West North Central (WNC) has region index $d = 8$). The experiment

is conducted under the assumption that agents did not foresee the shock and therefore are surprised in the year 2000 to find that the mapping from aggregate factors (Q, P) into (q_8, p_8) as defined in equation (3) does not apply anymore. In terms of their beliefs about aggregate and regional prices, they observe $\iota_{2000} \times q_{8,2000}, \iota_{2000} = 0.9$ instead of $q_{8,2000} = \mathbf{a}_8 \mathbf{F}_{2000}$, as would be predicted by the model in (3).²⁰ The shock is permanent thereafter, i.e. $\iota_{2000+\kappa} = 0.9, \kappa = 0, 1, \dots$.²¹ Notice that none of the time series properties of the regional price series other than the level are changed. In particular, the growth rate and implied variance remain the same. The other regional series (house price in this instance) is held constant in each experiment, as well as both series in all other regions. Illustrations of the shocked price and income series are shown in figures 7 and 8, respectively.

Shock to average labor productivity. Reducing mean labor productivity q in a given region translates directly into lower disposable income and therefore less utility. The effect of an unexpected reduction in the observed mean income level for WNC by 10% on migration outflows are shown in figure 9. The solid lines trace out the migration flows in the baseline scenario, while the dashed lines correspond to the shock scenario. In the right panel, showing renter emigration, we see clearly that the income reduction leads to a sustained increase in emigrants. The left panel shows owner outflows, and we also see a slight increase in emigration owners. In terms of the above discussion about moving costs and preference shocks, both results come from the fact that smaller payoff shocks are sufficient to trigger a move in the shock scenario than was the case in the baseline.

Looking at the opposite flow of migration, figure 10 shows the corresponding pictures for immigration from the rest of the economy into WNC. Starting with owners moving to WNC in the left panel, lower mean income makes the region is less attractive as a destination relative to other regions, thus we see a slight decrease in arrivals. The effect is much more pronounced for renters in the right hand panel. It is interesting to note the shape of the solid line tracing out the baseline inflows: As prices in the rest of the economy start to increase towards the peak in 2008, WNC, with its relatively low level of house prices becomes more attractive to renters. The increase in renter immigration towards 2005 is large, however, it needs to be looked at together with the corresponding outflows at the time. Over the period 1998–2012, WNC experiences average annual population growth of 1.08% in the baseline model.

Shock to house prices. The pictures in figures 11 and 12 show the changes to outflows and inflows as a result of a 30% reduction in the level of house prices in WNC in the year 2000. Starting with

²⁰The surprise assumption is necessary to observe a causal effect of the price change, since otherwise the response of agents would be partially anticipating the shock. They immediately adjust to the new mapping and behave optimally to the new sequence of prices in WNC, i.e. $\{\hat{p}_{8,t}\}_{t=2000}^{2012}$.

²¹ A separate experiment with a shock that reverts back over 3 years to trend delivered similar results, just limited to a shorter timeframe.

outflows in figure 11, the left panel shows that fewer owners are leaving the region than before. Owners experience a loss in net wealth, but this results in only small changes to migration behaviour. From the perspective of owners, the losses are sunk costs, and to the extent that lower current wealth makes it harder to afford a downpayment in other regions, the benefits of moving are reduced. For renters, we also see a clear reduction in outflows. Given that rent in the region has fallen to lower levels than before, fewer renters decide to leave. It is interesting at this point to remember the result in Notowidigdo (2011), where low skilled individuals are less likely to move away from a depressed region because they want to take advantage of cheap housing. The same mechanism operates here.

Turning finally to the effects of a reduction in house prices on inflows to WNC, we see in both panels that lower prices attract slightly more owners, and significantly more renters, particularly as aggregate prices rise and therefore the cost of living in other regions increases disproportionately for renters. Both owners and renters move to the region to enjoy cheaper housing.

6.4 The value of Migration

In this section I investigate the value of the migration option. The experiment assumes that the situation of West North Central changes in the sense that immigration is allowed, but emigrating is prohibited. First we will discuss changes in migration flows, and the ex-ante welfare effects for residents of WNC under both scenarios. Finally, we will analyze the experience of individuals who would have moved in the baseline environment, but are now prevented from doing so.

Migration flows and are displayed in table 14. The top panel shows changes in immigration flows to the region for both scenarios, and the respective percentage change. Comparing the first two columns as percentage changes in column three, we see that inflows into the region fall dramatically in a world where moving away from WNC is not an option. This shows that agents outside of the region perceive WNC as a much inferior option to the baseline. In the lower panel of the table I show baseline emigration flows for comparison. The bottom row of the table gives an estimate of average lifetime utility across both regimes, and we see that removing the option to leave the region carries a very large penalty. Across both regimes, average expected lifetime utility falls by 5.1%, which implies that individuals would demand a 4.4% increase in per period consumption in a world without the moving option before they would be indifferent to the baseline.

Figure 13 shows the average levels of utility conditional on age inside of WNC and in all other regions. We can see in the right panel, that reducing the value of WNC for potential immigrants affects utility only marginally. This comes from the fact that movers can easily avoid the region by moving somewhere else. For individuals inside the region, however, the changes are substantial and they vary by age. We can see that it is particularly younger individuals who suffer from the removed option to move. This makes intuitive sense, since the forgone expected gains from moving are larger if we consider a longer

	Baseline	No Moving	% Δ
	Immigration per period (percent of local population)		
Total	2.72%	1.76%	-35.3%
Owners	0.52%	0.37%	-28.9%
Renters	2.2%	1.4%	-36.4%
	Emigration per period (percent of local population)		
Total	1.34%	0	-
Owners	0.35%	0	-
Renters	0.99%	0	-
E[Lifetime U]	1.293	1.227	-5.1%
Welfare (<i>c</i>)		4.4%	

Table 14: Removing the migration option from residents of West North Central. The top panel shows immigration flows to West North Central as a percentage of the resident population. The same holds for the second panel shows outflows in the baseline. The final two rows show the average expected lifetime utility across scenarios, and the required consumption compensation in order to make agents indifferent.

time horizon.

To summarize this subsection, restricting the ability to move implies large losses in welfare. This effect stems from the inability of residents to respond to changes in their economic environment over time, as well as their inability to accommodate location preference shocks.

7 Abolishing the Mortgage Interest Deduction

In this section I investigate the impact of abolishing the mortgage interest deduction on migration and ownership rates. In the United States, homeowners are allowed to deduct the interest paid on the mortgage for their primary residence from taxable income. This reduces the user cost of owning because it exempts one of the largest component of an owner's housing cost – mortgage interest – from taxation. This is different for renters, whose housing expenditures are not tax deductible.²²

The mortgage interest deduction has been widely criticized on grounds of equity and distortionary effects in the housing market, see for example Glaeser (2011) for an overview.²³ The common wisdom is that it distorts the housing market by subsidizing owners, thereby leading to a higher rate of ownership as would arise without the subsidy. As discussed at length above, owners are less mobile than renters, so the question investigated in this subsection is whether abolishing the mortgage interest deduction would increase mobility, and if so, by how much. The answer to this question provides some guidance on whether housing policy could serve as an effective means to enhance the efficiency of the labor market by fostering greater mobility, or not.

At this point it is important to emphasize that depending on the elasticity of housing supply in a given market, we expect to see a general equilibrium effect of house prices in response to the policy change. The subsidy implied by the interest deduction can be viewed much like a characteristic of the house, like for instance distance to schools, or access to transportation networks. As such, it is capitalized into house price. Therefore, removing the subsidy would change long-run equilibrium prices, in the same way as moving a certain property to a location with a different set of available amenities would change its value. In order to partially accommodate this effect, I will refer to the recent paper by Sommer and Sullivan (2013), in which the authors compute the stationary general equilibrium in a model with infinitely lived, heterogeneous agents who are either renters, owner-occupiers or owner-landlords. They compute equilibrium prices and rents arising from different assumption about the tax system. They find that house prices decrease on the order of 5% when the mortgage interest deduction is abolished, with rents keeping roughly constant and the homeownership rate increasing from 65% to 71%.²⁴

²² The Joint Committee on Taxation estimates that tax expenditure to finance the mortgage interest deduction, which is the largest housing policy program in the states, was \$71.7 billion in 2014 [Joint Committee of Taxation \(2013\)](#).

²³It is often held that if there were an offsetting taxation of imputed rent which owner-occupiers pay to themselves in place, the negative aspects of the subsidy could be greatly reduced. Implementing such taxes is politically difficult in most circumstances, probably highlighted by the small number of nations which have adopted some form of this tax. For the case of Switzerland, see [Bourassa and Hoesli \(2010\)](#).

²⁴The mechanism is as follows: removing the subsidy lowers the equilibrium house price, since this is no longer capitalized into the value of the house; lower house prices makes housing affordable to more low wealth people who couldn't buy before, thereby increasing the ownership rate. Their model does not feature multiple labor or housing markets, and resting on the stationary equilibrium concept it does not allow for variation in house prices over time as is the case here; nevertheless, using their result to adjust the level of house prices by 5% in all locations when removing the mortgage deduction seems like a worthwhile exercise to at least approximate a potential general equilibrium effect

The experiment is implemented as follows. The model being in after tax terms, it does not feature a formal tax system which could be changed in an experiment. The role of taxes and, in general, the user cost of housing is subsumed in the implicit mortgage repayment path brought about by the final period restriction on assets, as explained in section 3.4. In reality the interest rate deduction reduces taxable income, thereby increasing disposable income of owners. Therefore, to simulate the removal of the deduction, I reduce disposable income of owners by an amount equivalent to their average implied tax savings. To do this I use data from [Poterba and Sinai \(2008\)](#), which provides estimates of the average tax savings from the interest deduction for several income and age brackets. Table 15 shows my version of their table.

age	$y < 40K$	$y \in [40K, 75K]$	$y \in [75K, 125K]$	$y \in [125K, 250K]$	$y > 250K$
< 34	250	720	2220	4400	8650
35-50	260	880	1810	4400	7130

Table 15: Average annual tax savings in dollars implied by the mortgage interest rate deduction for owners at various age and income groups. Data from [Poterba and Sinai \(2008\)](#) adjusted to 2012 dollars.

To provide some guidance, it is useful to start with a simple calculation to gauge the likely effect of the policy change on migration rates under the assumption that the conditional choice probabilities (i.e. whether to move or not) in the model do not change. This is equivalent to saying that there is a pure composition effect, but no effect on behaviour. To simplify the example, let's assume that the baseline ownership rate is $\pi_o = 0.63$, and that the annual migration rates for owners and renters are given by $m_o = 0.7\%$ and $m_r = 2\%$ respectively. This would give rise to a baseline unconditional migration rate of $\pi_o m_o + (1 - \pi_o) m_r = 1.81\%$. Assuming that the conditional choice probabilities which give rise to m_o and m_r do not change due to the policy change, and assume, for example, that removing the subsidy increases the ownership rate by 1% such that $\pi'_o = 0.64$. In that case we would expect to see a new migration rate of $\pi'_o m_o + (1 - \pi'_o) m_r = 1.68\%$, implying a decrease in the migration rate of 0.013 percentage points (or a 35% reduction of the rate).

The choice probabilities do change after the policy is implemented. In fact, it is easy to show that the only case under which the relative moving probabilities would remain constant is the one where the relative lifetime utilities in each region would change by exactly the same amount. This is not very likely to happen. Remember that the value in any location is determined by the utility derived from owning (enters utility directly), by the level of house prices, and by the level of consumption. Given a vector of house prices and incomes in the baseline environment, we can determine the relative differences in utilities and the corresponding probabilities of moving. Removing the mortgage deduction affects the price vector. Furthermore, by redistributing tax receipts we increase disposable income of a set of the

in my model.

population that was not able to buy before – not in their current region, and possibly not in another region. The overall effect on the migration rate will therefore be composed of a composition effect (shifting individuals from renter to owner status or vice versa), and an effect that arises from changes in behaviour, influenced by different price levels, different price to rent ratios, and redistribution of tax receipts from previous owners to the entire population.

The results are listed in table 16, and the columns under the subheading *Remove Deduction* show the results of the experiment with and without the GE correction. Both cases redistribute the taxes saved by removing the deduction on a per capita, per period basis. In other words, for total tax revenue X , and N individuals observed for T periods, each individual receives $X/(TN)$ additional dollars per period. This scheme approximates a reduction in income taxes financed by the removal of the mortgage deduction.

In the column without the GE correction, we see that ownership decreases by roughly 0.8%. This is a result of some marginal buyers being unable to satisfy the downpayment without the deduction. The migration rate increase only marginally despite this, because the change in composition towards more renters is offset to some degree by harder access to ownership in all regions without the benefits of the deduction. This reduces the incentives to move for some households who would have bought in different regions before. In terms of welfare effects, agents are close to indifferent between this policy and the baseline.²⁵

The final column shows the results with the GE correction, which lowers house prices by 5% in all regions, while keeping rents at a constant level. Keeping rents constant is achieved by changing the price to rent ratios κ_d in proportion to the reduction in prices. We find that ownership increases by about 1.3% as a result of lower house prices. The reverse mechanism to the previous paragraph is at work, implying that more people can buy, in this case despite the removal of the subsidy to owners. The change in composition towards more owners combined with access to housing in more regions produces a net effect on migration that is slightly positive at 0.115% of the baseline rate. The welfare implications for this experiment are larger than in the previous case and amount to 2.4% of per period consumption over the lifecycle.

In concluding this section, it is instructive to note the importance of the general equilibrium effect in this case. Under the assumption of constant prices, the effect of removing the deduction is that ownership is reduced and that migration slightly increases. Removing the deduction means fewer low wealth households can buy housing, which reduces welfare, and this offsets any gains made from slightly more migration. Taking the price change and the implications for the price to rent ratio into account, however, shows that ownership would increase. This would affect the migration rate only marginally,

²⁵As a metric of welfare I compute the implicit consumption tax δ which would make agents indifferent in terms of expected lifetime utility between the baseline and policy regimes, as detailed in appendix A.

but have a significant impact on welfare.

8 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 leave a certain region in the model implies an associated reduction in expected lifetime utility of 5.1%, or 4.4% of per period consumption. 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 behaviour 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.

In a policy experiment where I remove the mortgage interest deduction for owners, I find that accounting for a likely general equilibrium effect is crucial for the resulting effects. Assuming that house prices would not change as a result of removing the deduction implies that migration rates would increase slightly, mainly due to a shift in composition towards more renters. Correcting prices with an external estimate of the likely GE effect (a 5% reduction as reported in [Sommer and Sullivan \(2013\)](#)) results in a different conclusion: lower prices after the removal of the deduction allow more people to buy, resulting in an increase of ownership. The net effect on migration is only marginally positive, however, since part of the shift in composition towards more owners is offset by changes in migration behaviour brought about by lower prices in all regions and redistribution of income to the poor. The welfare implications depend greatly on whether or not prices are adjusted, with average welfare gains 1% in the former, and of the order of 2.4% of period consumption in the latter case.

In its current state, the model is silent with regards to location amenities. Dissentangling the effects of price changes on the consumer's budget constraint from its effects on local amenities, and therefore utility, is clearly important to understand migration decisions better. I view the current implementation as a first step in this direction, and more work is necessary to incorporate amenities.

	Baseline	Remove Deduction	
		w/o GE	w GE
Ownership Rate	0.631	0.625	0.639
Migration Rate	0.01282	0.01285	0.01284
% Δ Ownership	0	-0.882%	1.36%
% Δ Migration	0	0.247%	0.115%
Welfare	0	1%	2.4%

Table 16: Abolishing the mortgage interest deduction. Depending on whether or not a GE correction to house prices is applied (house prices are 5% reduced), the migration rate increases slightly. Cheaper housing everywhere and higher disposable income from redistribution enables more previously liquidity constraint households to buy – in their home region, but crucially also in other regions. This leads to a very small increase in the migration rate.

Shocking p_{dt} for $d = \text{West North Central}$

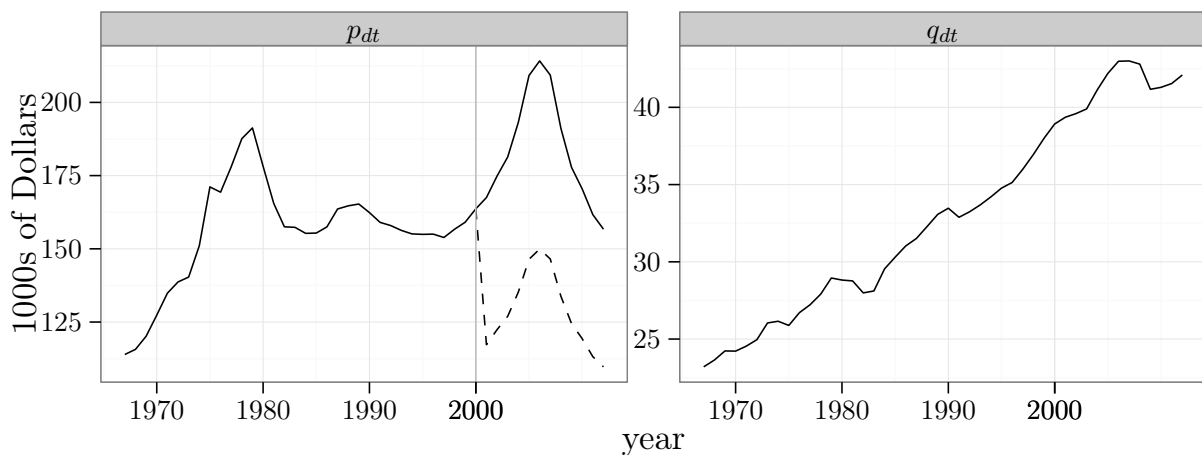


Figure 7: Shocking the house price series for West North Central from 2000 onwards. The shock multiplies the data series with the factor $\iota = 0.7$, i.e. reduces the level by 30%. The dashed line in the right panel is the path of house prices after the shock.

Shocking q_{dt} for $d = \text{West North Central}$




Figure 8: Shocking the level of average labor productivity in region West North Central by 10% in 2000.

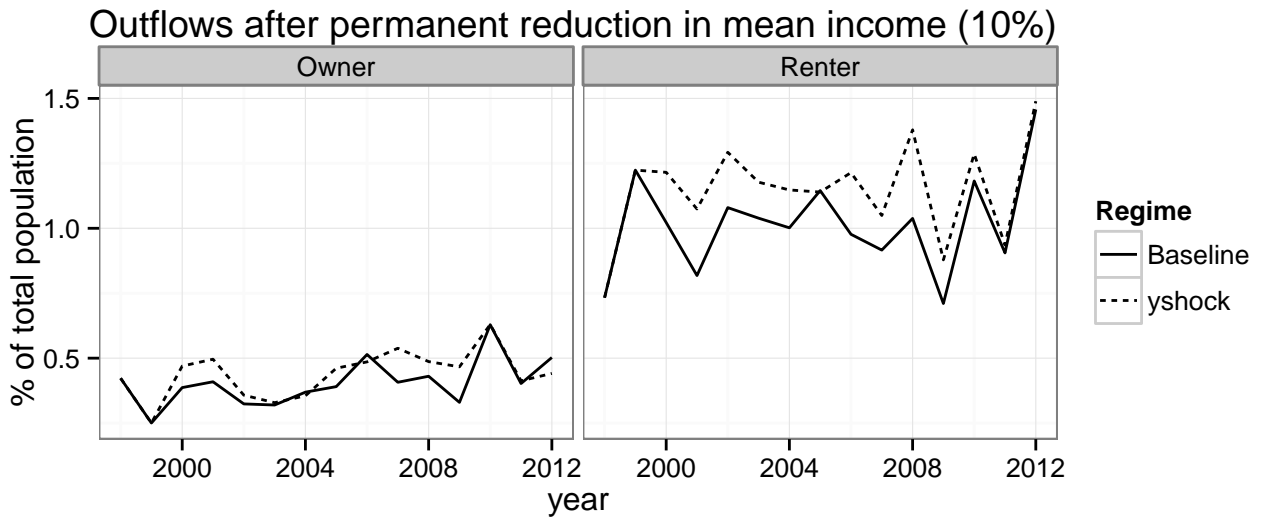


Figure 9: The effect of a permanent 10% reduction in productivity q in 2000 in the WNC region only. The experiment holds the regional price series in WNC and price and incomes in all other regions fixed at their observed values.




Figure 10: Changes in inflow patterns to WNC after the mean productivity level is permanently reduced by 10% in 2000.




Figure 11: The effect of a 30% decrease of the house price level in 2000 in the West North Central Division on migratory outflows from the region.

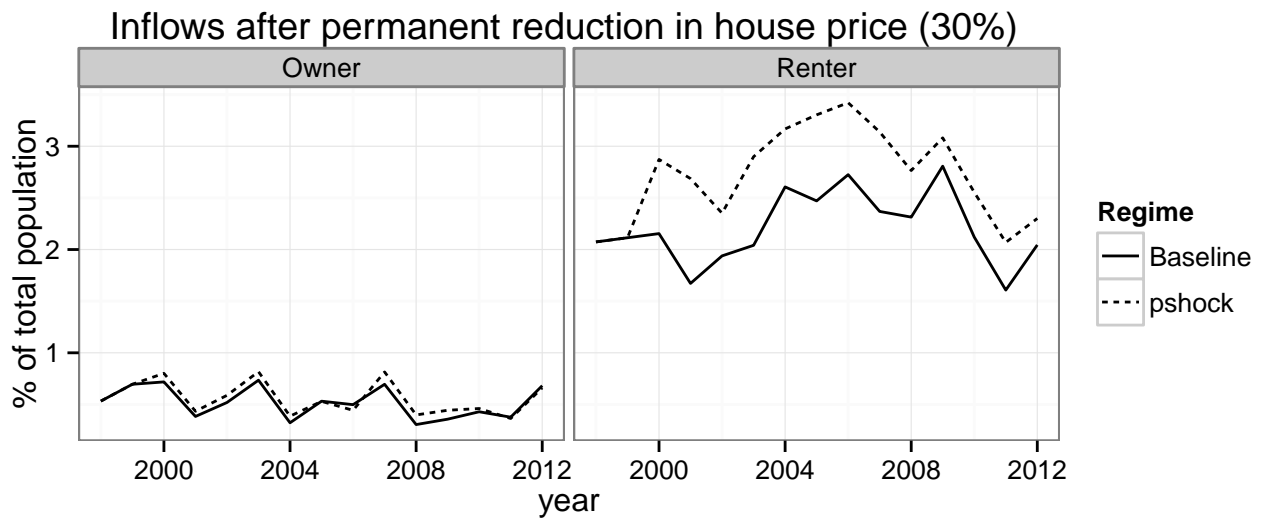


Figure 12: The effect of a 30% decrease of the house price level in 2000 in the West North Central Division on migratory inflows to the region.

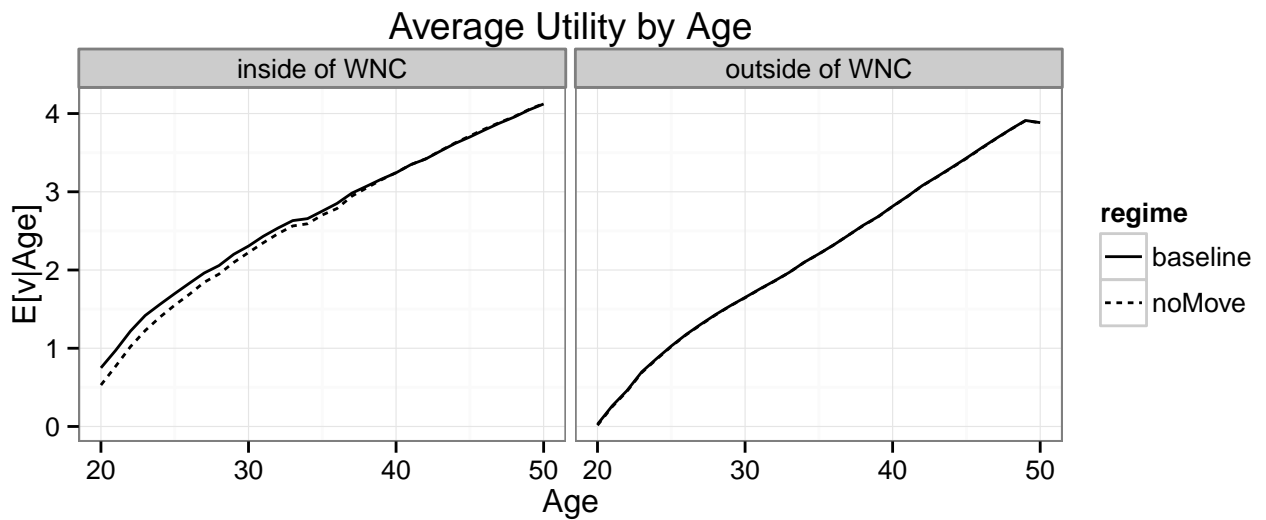


Figure 13: Removing the option to leave WNC. This plot shows average utility levels conditional on age, for the population residing in WNC, and the population residing elsewhere.

A Welfare Measure

Denoting the expected lifetime utility from the baseline and policy regimes under consumption tax δ by U and $\hat{U}(\delta)$ respectively, the equalizing consumption tax δ^* solves

$$\begin{aligned} U - \hat{U}(\delta^*) &= 0 \\ U &= \frac{1}{TN} \sum_{i=1}^N \sum_{t=1}^T \beta^{t-1} u(c_{it}) \\ \hat{U}(\delta) &= \frac{1}{TN} \sum_{i=1}^N \sum_{t=1}^T \beta^{t-1} u(\delta c_{it}) \end{aligned}$$

where N is the number of simulated individuals and $u(c)$ stands, with some abuse of notation, for the stream of utility from equation (5) as implied by optimal behaviour. Accordingly, a value $\delta^* > 1$ implies that agents would be indifferent between any proposed policy change if consumption were scaled up in every period, i.e. they would demand a subsidy. In the opposite case of $\delta^* < 1$ they would be happy to give up a fixed proportion δ^* of period consumption if they were given the opportunity to participate in the policy.

B Initial Conditions and Cohort Setup

The SIPP estimation sample runs from 1998 through 2012. The data moments the model is supposed to replicate are weighted averages over this period, where the weights are the SIPP sampling weights. When reconstructing an artificial sample from the model simulation, care must be taken to replicate the shocks experienced by each cohort in the data leading up to the point where they are observed.

The data is subset to the ages allowed for in the model, i.e. 20–50. I compute data moments, for example the average homeownership rate in region d , or the average total wealth of age group 40–45 in d , as averages over the entire sample period:

$$\begin{aligned} \text{mean_own_data}_d &= \frac{1}{15} \sum_{t=1998}^{2012} \left(\frac{1}{N_{dt}} \sum_{i \in d, t} \omega_{it} \mathbf{1}[h_{it} = 1] \right) \\ \text{mean_wealth_data_40_45}_d &= \frac{1}{15} \sum_{t=1998}^{2012} \left(\frac{1}{N_{dt, j \in [40, 45]}} \sum_{i \in d, t, j \in [40, 45]} \omega_{it} w_{ijt} \right) \end{aligned}$$

where N_{dt} is the number of people in d at date t , and ω_{it} is a person's cross-sectional weight, and $i \in d, t$ stands for i is in d at date t . Similarly, $i \in d, t, j \in [40, 45]$ stands for i is in d at date t and age j in

[40,45].

This means that for the second data moment, for example, 40 year-olds from 1998 contributed as well as 40 year-olds from the 2012 cohort. Needless to say, those cohorts faced a different sequence of house price shocks leading up the point of observation. For individuals “born” before the first data period, i.e. 1998, I construct regional house price and regional income series going back until 1968. Simulating individuals from the 1968 cohort for a full lifetime of $J=30$ years until they reach age 50 brings them into the year 1998, where they form the group of 50 year-olds in that particular year. This sort of staggered simulation is carried out until the final cohort is born in 2012 at age 20. No simulation needs to take place for any individual alive at years after 2012.

C Estimation of G_{move}

In a first step I estimate the marginal distributions of z_{idt} and z_{ikt+1} for all movers. These are the cross-sectional distributions of residuals u_{it} and u_{it+1} from the regression in expression (25), which is estimated for all movers. The move takes place in period t , such that by assumption, u_{it} is the residual wage in origin location d , and u_{it+1} is the residual wage in the new location k .

$$\ln y_{idt} = \beta_0 + \beta_1 \text{college}_{it} + \delta p(\text{age}_{it}) + \beta_2 \text{numkids}_{it} + \beta_3 \text{sex}_{it} + \beta_4 \text{metro}_{it} + \gamma_d + u_{it} \quad (25)$$

Here $p(\text{age})$ is a third order polynomial in age, metro is an indicator for metropolitan status and γ_d is a Division fixed effect. Then I convert the residuals into their respective rank in the empirical distributions of before and after move residuals. Denoting the standardized values by $(\hat{u}_{it}, \hat{u}_{it+1})$, the next step involves fitting the a normal copula via maximum likelihood to this data. The results are shown in table 17, and they indicate a correlation between \hat{u}_{it} and \hat{u}_{it+1} of 0.59. This estimate together with the marginal distributions of u_{it} and u_{it+1} are used in the structural model, where I use the current value of z , evaluated in the marginal distribution of u_{it} for a mover together with the copula estimate \hat{G}_{move} to draw the next value of z' .

D Map

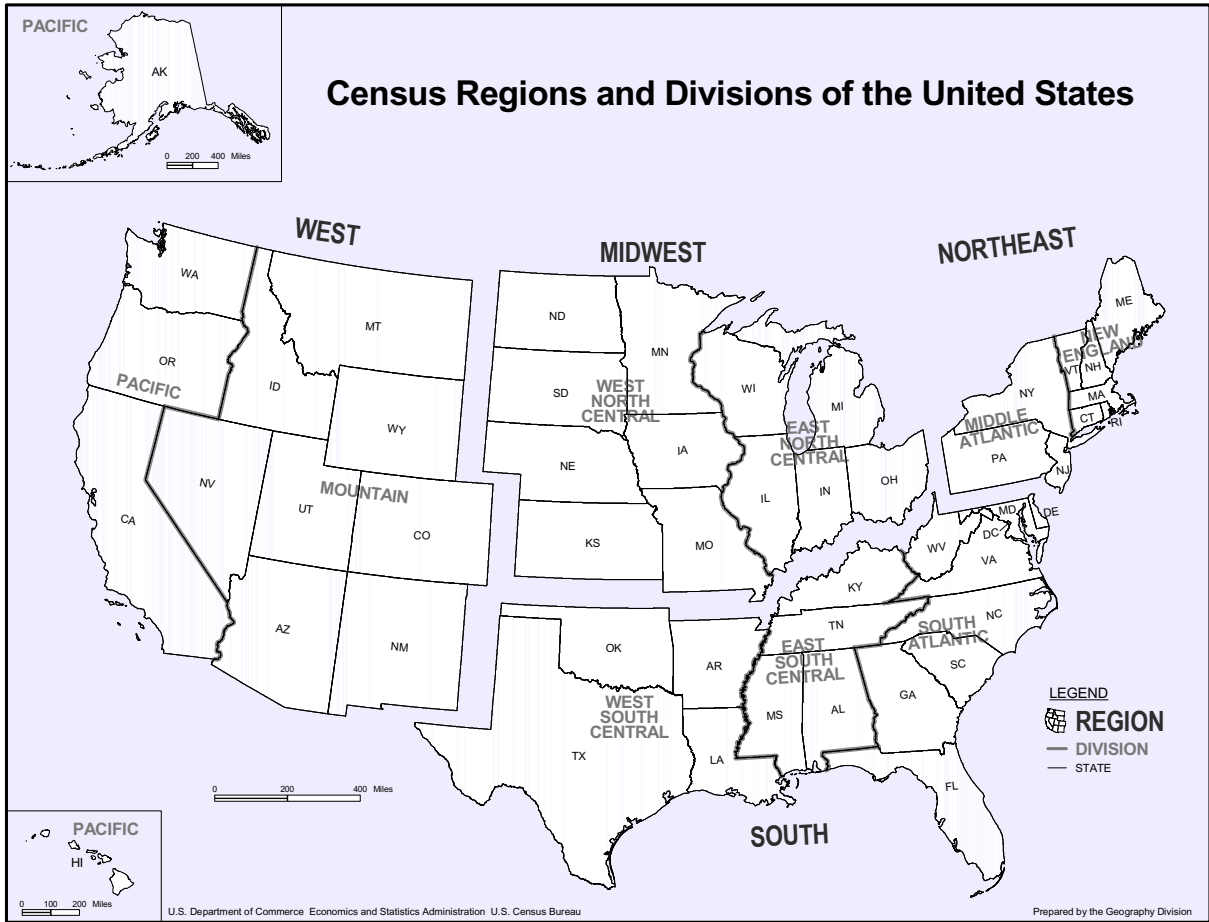


Figure 14: Census Division Map, taken from https://www.census.gov/geo/maps-data/maps/pdfs/reference/us_regdiv.pdf. The Divisions are from left to right Pacific, Mountain, West North Central, West South Central, East North Central, East South Central, New England, Middle Atlantic and South Atlantic.

E Additional tables and figures

	ρ	S.E.
$G_{\text{move}}(z_t, z_{t+1})$	0.59762	0.01795

Table 17: Normal Copula estimates for the rank of wage residuals u_{it} and u_{it+1} for individuals who move in period t .

	East North Central		East South Central		Middle Atlantic		Mountain	
	q_{dt}	p_{dt}	q_{dt}	p_{dt}	q_{dt}	p_{dt}	q_{dt}	p_{dt}
$\mathbf{a_{0d}}$	12.30*** (0.72)	61.10*** (10.51)	3.74*** (0.60)	88.19*** (7.15)	8.42*** (0.64)	-34.84** (12.00)	8.38*** (0.67)	5.89 (10.85)
Q_t	0.62*** (0.03)	-0.84 (0.49)	0.70*** (0.03)	-1.53*** (0.34)	1.00*** (0.03)	2.87*** (0.56)	0.56*** (0.03)	-1.23* (0.51)
P_t	0.01 (0.01)	0.70*** (0.10)	0.01 (0.01)	0.61*** (0.07)	-0.01* (0.01)	0.75*** (0.12)	0.03*** (0.01)	1.20*** (0.10)
R^2	0.97	0.74	0.98	0.73	0.99	0.92	0.98	0.89
Adj. R^2	0.97	0.72	0.98	0.72	0.99	0.91	0.97	0.89
Num. obs.	92	92	92	92	92	92	92	92

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

	New England		Pacific		South Atlantic		West N Central		West S Central	
	q_{dt}	p_{dt}	q_{dt}	p_{dt}	q_{dt}	p_{dt}	q_{dt}	p_{dt}	q_{dt}	p_{dt}
$\mathbf{a_{0d}}$	3.77*** (0.64)	-114.58*** (20.60)	13.32*** (0.56)	-214.09*** (17.11)	6.54*** (0.64)	39.23*** (5.32)	7.75*** (0.71)	62.46*** (7.80)	5.46*** (0.93)	106.64*** (12.75)
Q_t	1.18*** (0.03)	4.54*** (0.97)	0.55*** (0.03)	3.08*** (0.81)	0.75*** (0.03)	-1.47*** (0.25)	0.72*** (0.03)	-1.69*** (0.37)	0.63*** (0.04)	-3.73*** (0.60)
P_t	-0.01* (0.01)	1.05*** (0.20)	0.03*** (0.01)	1.91*** (0.16)	0.01 (0.01)	1.14*** (0.05)	0.01 (0.01)	0.78*** (0.07)	0.02* (0.01)	0.85*** (0.12)
R^2	0.99	0.89	0.98	0.95	0.98	0.97	0.98	0.81	0.96	0.53
Adj. R^2	0.99	0.89	0.98	0.95	0.98	0.96	0.98	0.80	0.96	0.51
Num. obs.	92	92	92	92	92	92	92	92	92	92

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 18: Aggregate to Regional price mappings. This table shows the estimated coefficients from equation (3), which relates the aggregate factors (Q_t, P_t) to regional income and house price (q_{dt}, p_{dt}).

	$\log y_{it}$
Intercept	-0.684*** (0.135)
age	0.211*** (0.006)
age ²	-0.004*** (0.000)
age ³	0.000*** (0.000)
East North Central	0.273*** (0.029)
East South Central	0.214*** (0.030)
Middle Atlantic	0.273*** (0.028)
Mountain	0.276*** (0.029)
New England	0.294*** (0.027)
Pacific	0.287*** (0.028)
South Atlantic	0.265*** (0.029)
West North Central	0.280*** (0.029)
West South Central	0.248*** (0.030)
R ²	0.043
Adj. R ²	0.043
Num. obs.	288952
RMSE	0.904

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 19: Regional Mean Income to Individual level income mapping. This is the empirical implementation of equation (1), as explained in section 5.1. The estimated equation is $\log y_{idt} = \beta_0 + \eta_d \log \bar{y}_{dt} + \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \beta_3 \text{age}_{it}^3 + u_{it}$ and the coefficients η are denoted with the Division names.

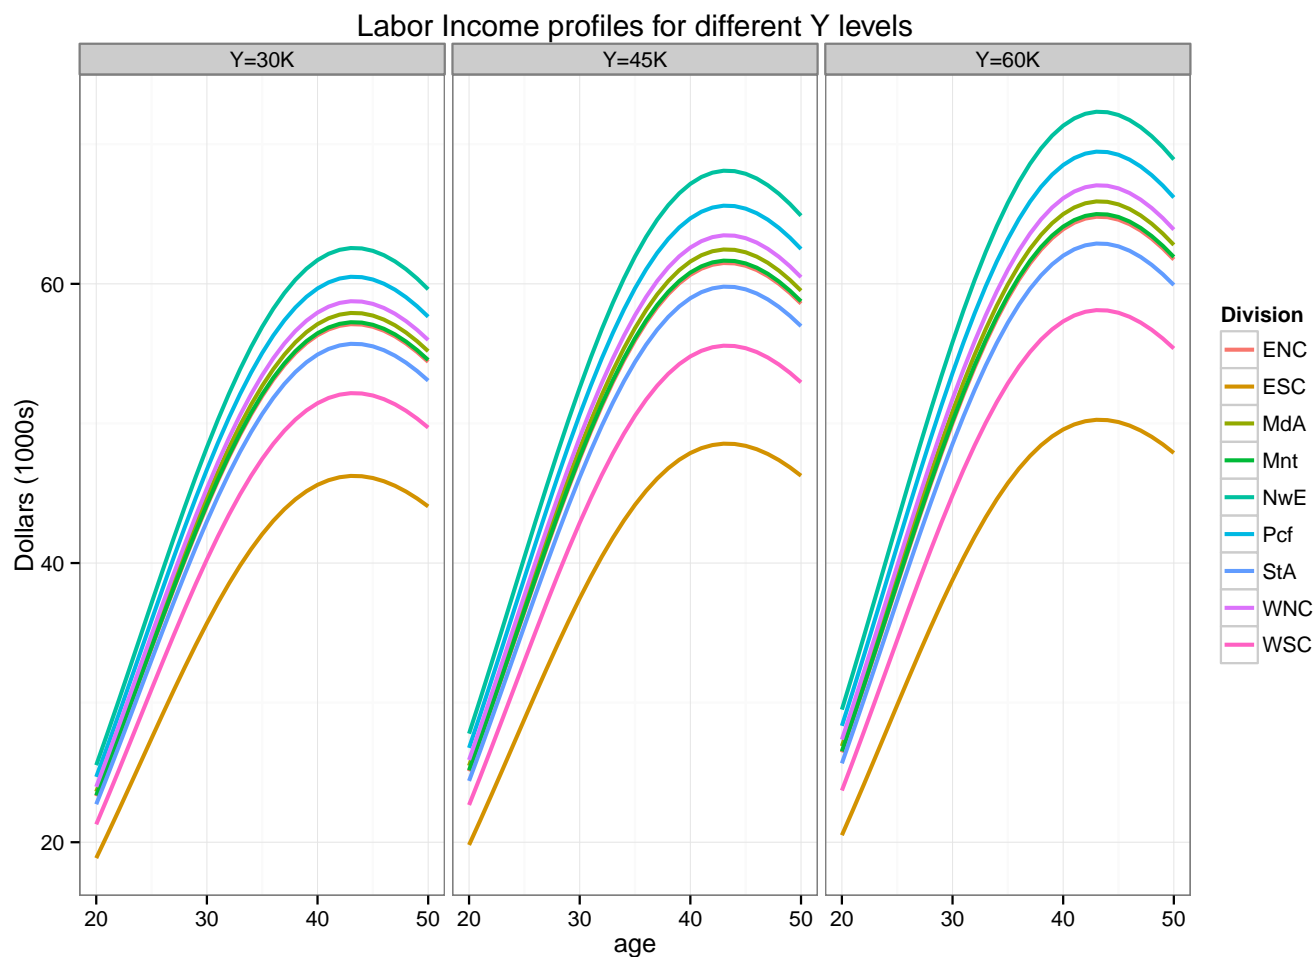


Figure 15: Age profiles as predicted by the empirical implementation of individual labor income, equation (22), for three different levels of regional mean productivity q . Notice that in the model as well as in the data it is never the case that all regions have the same level of average income.

References

- 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.
- 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.
- BARAGATTI, M., A. GRIMAUD, AND D. POMMERET (2013): “Likelihood-free parallel tempering,” *Statistics and Computing*, 23, 535–549.
- BATTU, H., A. MA, AND E. PHIMISTER (2008): “Housing Tenure, Job Mobility and Unemployment in the UK*,” *The Economic Journal*, 118, 311–328.
- 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, 135–63.
- BLANCHARD, O. J. AND L. F. KATZ (1992): “Regional evolutions,” in *Brookings Papers on Economic Activity 2*, ed. by O. J. Blanchard, L. F. Katz, R. E. Hall, and B. Eichengreen, Brookings Institution Press, vol. 1992, 1–75.
- BLANCHFLOWER, D. G. AND A. J. OSWALD (2013): “Does High Home-Ownership Impair the Labor Market?” Tech. rep., National Bureau of Economic Research.
- BONHOMME, S. AND J. ROBIN (2006): “Modeling individual earnings trajectories using copulas: France, 1990–2002,” *Structural Models of Wage and Employment Dynamics*.
- BOURASSA, S. AND M. HOESLI (2010): “Why Do the Swiss Rent?” *The Journal of Real Estate Finance and Economics*, 40, 286–309.
- CARROLL, C. D. (2006): “The method of endogenous gridpoints for solving dynamic stochastic optimization problems,” *Economics letters*, 91, 312–320.
- CHERNOZHUKOV, V. AND H. HONG (2003): “An MCMC approach to classical estimation,” *Journal of Econometrics*, 115, 293–346.
- CLAUSEN, A. AND C. STRUB (2013): “Envelope Theorems for Non-Smooth and Non-Concave Optimization,” <https://andrewclausen.net/research.html>.

- COULSON, N. E. AND L. M. FISHER (2002): “Tenure choice and labour market outcomes,” *Housing Studies*, 17, 35–49.
- DEMYANYK, Y. S., D. HRYSHKO, M. J. LUENGO-PRADO, AND B. E. SORESENSEN (2013): “Moving to a job: The role of home equity, debt, and access to credit,” .
- DIAMOND, R. (2012): “Housing Supply Elasticity and Rent Extraction by State and Local Governments,” *Harvard Job Market Paper*.
- 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.
- FELLA, G. (2014): “A generalized endogenous grid method for non-concave problems,” *Review of Economic Dynamics*.
- FERREIRA, F., J. GYOURKO, AND J. TRACY (2010): “Housing busts and household mobility,” *Journal of Urban Economics*, 68, 34 – 45.
- FRENCH, E. (2005): “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour,” *The Review of Economic Studies*, 72, 395–427.
- GEMICI, A. (2007): “Family migration and labor market outcomes,” Ph.D. thesis, University of Pennsylvania.
- GLAESER, E. L. (2011): “Rethinking the federal bias toward homeownership,” *US Department of Housing and Urban Development*.
- GLAESER, E. L. AND J. M. SHAPIRO (2002): “The benefits of the home mortgage interest deduction,” Tech. rep., National Bureau of Economic Research.
- GÜLER, B. AND A. TASKIN (2011): “Homeownership and Labor Market: The Effect of Market Size,” Tech. rep., mimeo.
- HALKET, J. AND S. VASUDEV (2014): “Saving up or settling down: Home ownership over the life cycle,” *Review of Economic Dynamics*, 17, 345–366.
- HEAD, A. AND H. LLOYD-ELLIS (2012): “Housing liquidity, mobility and the labour market,” *The Review of Economic Studies*, rds004.
- JOINT COMMITTEE OF TAXATION (2013): “Estimates Of Federal Tax Expenditures For Fiscal Years 2012-2017,” <https://www.jct.gov/publications.html>.

- KARAHAN, F. AND S. RHEE (2011): "Housing and the Labor Market: The Role of Migration on Aggregate Unemployment," *University of Pennsylvania, mimeo*.
- KENNAN, J. AND J. WALKER (2011): "The effect of expected income on individual migration decisions," *Econometrica*, 79, 211–251.
- LI, W., C. MEGHIR, AND F. OSWALD (2014): "Consumer Bankruptcy and Mortgage Default," *working paper*.
- 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.
- MOLLOY, R., C. L. SMITH, AND A. WOZNAK (2011): "Internal Migration in the United States," *The Journal of Economic Perspectives*, 25, 173–96.
- MORETTI, E. (2011): "Chapter 14 - Local Labor Markets," Elsevier, vol. 4, Part B of *Handbook of Labor Economics*, 1237 – 1313.
- NENOV, P. (2012): "Regional mismatch and labor reallocation in an equilibrium model of migration," *Norwegian Business School*.
- NOTOWIDIGDO, M. J. (2011): "The incidence of local labor demand shocks," Tech. rep., National Bureau of Economic Research.
- OSWALD, A. J. (1996): "A conjecture on the explanation for high unemployment in the industrialized nations: part 1," .
- PIYAPROMDEE, S. (2013): "'The Impact of Immigration on Wages, Internal Migration and Welfare'," Ph.D. thesis, U Wisconsin.
- POTERBA, J. AND T. SINAI (2008): "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income," *The American Economic Review*, 98, pp. 84–89.
- ROBACK, J. (1982): "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 90, pp. 1257–1278.
- ROSEN, S. (1979): "Wage-based indexes of urban quality of life," *Current issues in urban economics*, 3.
- RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.

- SCHULHOFER-WOHL, S. (2011): “Negative equity does not reduce homeowners’ mobility,” Tech. rep., National Bureau of Economic Research.
- SOMMER, K. AND P. SULLIVAN (2013): “Implications of US Tax Policy for House Prices, Rents and Homeownership,” <http://kamilasommer.net/Taxes.pdf>.
- SONG, J., G. TOPA, G. VIOLANTE, AND A. SAHIN (2014): “Mismatch Unemployment,” *American Economic Review*. *forthcoming*.
- WINKLER, H. (2010): “The effect of homeownership on geographic mobility and labor market outcomes,” *SSRN eLibrary*.