

This file was downloaded from BI Brage,
the institutional repository (open access) at BI Norwegian Business School
<http://brage.bibsys.no/bi>

Regional reallocation and housing markets in a model of frictional migration

Plamen T. Nenov
BI Norwegian Business School

This is the accepted, refereed and final manuscript to the article published in

Review of Economic Dynamics, 18(2015)4:863-880

Publisher's version available at <http://dx.doi.org/10.1016/j.red.2015.08.002>

Copyright policy of *Elsevier*, the publisher of this journal:

The author retains the right to post the accepted author manuscript on open web sites operated by author or author's institution for scholarly purposes, with an embargo period of 36 months after first view online.

<http://www.elsevier.com/journal-authors/sharing-your-article#>

This manuscript version is made available under the CC-BY-NC-ND 4.0 license

<http://creativecommons.org/licenses/by-nc-nd/4.0/>

Regional Reallocation and Housing Markets in a Model of Frictional Migration

Plamen T. Nenov*

Abstract

Migration frictions are important for understanding key features of gross migration and housing markets. This paper studies a multi-region equilibrium model with frictional migration. Idiosyncratic preference shocks, a mobility cost, and imperfectly directed migration lead to slow worker reallocation in response to changes in local conditions. This leads to a dependence of local house prices on the history of labor market shocks. The model accounts for the comovements of unemployment and rental and house prices with gross migration observed in a panel of U.S. cities. Structural estimation reveals a high mobility cost for unemployed workers and a low probability of directed migration. Both of these imply that regional reallocation has a limited importance for the aggregate labor market and that the effects of housing markets on reallocation are small.

JEL Codes: E24, J61, J64, R23

Keywords: reallocation, housing markets, gross migration, indirect inference

*Norwegian Business School (BI). Nydalsveien 37, 0484 Oslo, Norway e-mail: Plamen.Nenov@bi.no; This paper subsumes much of the analysis contained in previous papers titled "Regional Mismatch and Labor Reallocation in an Equilibrium Model of Migration" and "Labor Market and Regional Reallocation Effects of Housing Busts". I want to thank the managing editor, Matthias Doepke, the associate editor and three anonymous referees for suggestions that helped greatly improve the paper. I also want to thank Daron Acemoglu, George-Marios Angeletos, Ricardo Caballero, Jonathan Halket, Espen Moen, and seminar participants at MIT, UC Louvain, University of Amsterdam, University of Cologne, Aarhus University, Norges Bank, Stockholm University, University of Copenhagen, ECB, BI Norwegian Business School, the Federal Reserve Board, Paris School of Economics, INSEAD, Oxford, University of Exeter, University of Bristol, the Bank of Finland, the Cologne Workshop in Macroeconomics, the 2012 Urban Economics Association Meeting, University of Oslo (ESOP), Mannheim Workshop in Quantitative Macroeconomics, the 2013 SED Meetings, and the 2014 Louvain Workshop on "Housing, Mobility and Labor Market Outcomes" for valuable comments.

1 Introduction

How important is regional labor reallocation - the net flow of workers across regions - for the labor market? Do housing markets affect the labor market through their impact on regional reallocation? The recent U.S. recession and its aftermath have caused renewed interest in these questions.

When a region falls in a deep recession the net worker flow out of that region should serve as adjustment that dampens the labor market effects of the shock (Blanchard and Katz, 1992). However, as population flows out, local housing costs decline, which compensates the workers that remain for the adverse labor market conditions. This equilibrium compensating effect impacts regional labor reallocation, and through that channel, both the local and aggregate labor markets.

This paper studies the importance of regional reallocation for the labor market and of the compensating effect of housing markets on reallocation. Both of these depend on the structure of individual migration decisions. If moving decisions respond strongly (weakly) to a deterioration in local labor market conditions, and moves are directed towards better performing labor markets (untargeted - with many moves to under-performing labor markets), then reallocation is large (small) and populations adjust quickly (sluggishly). In that case the importance of regional reallocation for unemployment is significant (limited), and housing markets exert a large (small) compensating effect.

If individual moves respond weakly and are untargeted, then moving outcomes arise *as if* driven by a frictional migration process. However, this paper argues that this is the empirically relevant case, since such a frictional process is important for understanding key features of gross migration flows and housing markets.

I show this in a spatial equilibrium model that includes an interaction between local labor market conditions, housing markets and migration flows. The model economy I consider consists of a continuum of islands (locations) in the spirit of Lucas and Prescott (1974). Local labor markets are characterized by search and matching frictions, which give rise to

unemployment within islands, while island-specific labor market shocks drive local business cycles. Each region is endowed with a fixed supply of durable housing that workers value for its housing services and rent in a competitive rental market. A downward sloping demand for housing services by workers leads to differences in equilibrium rental prices across locations with different populations.

Workers migrate out in response to local labor and housing market conditions. However, they also move for idiosyncratic reasons due to preference shocks for their current location. Their migration decision is a combination of directed and undirected (random) migration, that is, workers either migrate to regions offering the most favorable labor and housing market conditions or alternatively, to any region of the economy.

This individual migration process allows the model to generate the comovements of unemployment and house prices (or rental prices) with gross migration flows observed in a panel of U.S. cities. Specifically, the unemployment rate in a city correlates positively with migration out of that city and negatively with migration into it, controlling for housing prices. More importantly, housing prices correlate positively with out-migration and negatively with in-migration, controlling for unemployment.

The individual migration process in the model leads to sluggish labor reallocation in response to local labor market differences. This slow reallocation creates a rich equilibrium distribution of regional populations and a dependence of regional house prices on the history of labor market shocks. For example, a region whose labor market is depressed for a longer time has a lower population, and hence, house prices, compared to a region which has experienced a negative shock more recently. This “history dependence” drives the positive comovement between out-migration and house prices, controlling for unemployment, since regions with lower house prices, other things equal, are more attractive to potential emigrant workers. Combined with partially directed migration, the history dependence also drives the negative comovement between in-migration and house prices, controlling for unemployment.

I use the magnitudes of the observed comovements in the data to estimate the individual

migration process using an indirect inference procedure. The migration parameters of interest are the probability of directed migration, a mobility cost for unemployed workers (combining moving costs and a preference for staying in a region), and the dispersion in idiosyncratic regional preferences.

The model can match well the moments used in the estimation and also performs well against a large set of non-targeted moments. It fits particularly well the variability in rental prices observed in the data and also features persistence in house price growth rates. Both of these are hard to generate in models with frictionless mobility, which tend to predict either a counterfactually high dispersion in rental prices (Davis and Ortalo-Magne, 2011) or no persistence in house price growth (Glaeser, Gyourko, Morales, and Nathanson, 2014).¹ The reason why this model produces a lower rental price variability and persistence in price growth is the slow reallocation due to the mobility cost, idiosyncratic regional preferences, and the partially directed migration. The slow reallocation compresses the equilibrium distribution of populations, which leads to a decreased variability in equilibrium rental prices. Also, the smooth out- and in-migration flows lead to persistence in house price growth. Additionally, I directly show that the history dependence in prices and populations in the model is consistent with the observed data.

The estimation reveals a low probability of directed migration and a high mobility cost for unemployed workers, which contribute to a low regional reallocation rate. In particular, the model predicts that around 50% of net flows across U.S. cities are driven by local labor market disparities. The low probability of directed migration is particularly important for the low reallocation rate. Specifically, at the estimated parameters, a small change in the probability of directed migration has a large impact on reallocation.

Next, I consider a set of counterfactual experiments to establish the role of housing for reallocation and the importance of reallocation for unemployment. By comparing the cases with and without a compensating effect of the housing market, I show that, at the

¹In recent work, Head, Lloyd-Ellis, and Sun (2014) study a model with housing markets characterized by search frictions, which features persistence in house price growth rates.

estimated parameters, housing has a limited effect on labor reallocation. The reason for this small effect is that with slow reallocation the equilibrium distribution of regional populations is compressed, so housing market differences play a small role in migration decisions. In contrast, whenever there is more reallocation, particularly, due to a higher probability of directed migration, the housing market also exerts a substantial effect. In that case, the equilibrium distribution of populations across regions is more dispersed, so rental prices differences matter more.

Turning to the importance of reallocation for unemployment, I show that changes in the individual migration process, which lead to a large response of reallocation, have small effects on aggregate unemployment. This suggests a limited role of regional mismatch for the labor market.

Related Literature. The model I consider is related to the literature on regional and sectoral reallocation, initiated by the seminal work of Lucas and Prescott (1974).² Coen-Pirani (2010) and Lkhagvasuren (2012) are two important recent contributions related to my paper. Coen-Pirani (2010) documents a positive correlation between out- and in-migration in the cross-section of states and considers a Lucas-Prescott island model with perfectly directed migration, which can account for it. Lkhagvasuren (2012) shows that gross migration flows across states are substantial compared to net flows despite the existence of unemployment differences. He studies a model with local worker-location productivity shocks, and random migration to explain this pattern.

In contrast to these papers, I conduct the analysis at the city rather than the state level. To the extent that cities rather than states are the relevant local labor markets, examining city-level evidence on migration is important for gaining a more precise understanding of the role of regional reallocation for the labor market. Additionally, the model prominently features the interaction between local labor and housing markets and migration flows, and more specifically, the “history dependence” effects of the labor market on house prices arising

²Recent papers in this literature include Alvarez and Shimer (2011), Coen-Pirani (2010), Lkhagvasuren (2012), Carrillo-Tudela and Visschers (2011), and Kaplan and Schulhofer-Wohl (2012), among others.

from slow reallocation, which is important for understanding key comovements in the data. Finally, a model that nests both directed and random inter-regional migration permits the investigation of the importance of partially directed migration for reallocation, an issue that these previous studies have not addressed.³

The paper is closely related to studies of the interaction between the housing market, labor market, and worker mobility.⁴ Head and Lloyd-Ellis (2012) investigate the impact of home ownership on mobility and unemployment and build a theory that accounts for the reduced mobility of homeowners versus renters observed in micro-data. Van Nieuwerburgh and Weill (2010) document a secular increase in dispersion in regional house prices and wages and propose a model that explains and quantitatively matches this increase. My paper contributes to this literature by studying the link between regional business cycles, housing markets, and migration flows in a model with frictional mobility.⁵

The estimation of the equilibrium model relates this study to the literature on structural models of migration (Kennan and Walker (2011), Bayer and Juessen (2012)). The region preference shocks that agents experience in my model are similar to the shocks considered in Kennan and Walker (2011). Unlike that paper, I examine a simpler migration decision but embed it in a multi-region equilibrium framework with housing markets. The use of an indirect inference procedure to estimate structural parameters from comovements of gross migration flows is similar to the approach in Bayer and Juessen (2012) but with an additional focus on the interaction between local labor and housing markets.

Finally, considering the importance of regional labor reallocation for regional mismatch and aggregate unemployment relates the paper to recent studies that focus on the implications of labor reallocation distortions for aggregate unemployment, particularly in the recent

³The issue of directed versus random migration is related to issues of directed versus random labor market search (Godoy and Moen, 2012).

⁴A non-exclusive list includes Oswald (1996), Eeckhout (2004), Van Nieuwerburgh and Weill (2010), Winkler (2010), Head and Lloyd-Ellis (2012), Rupert and Wasmer (2012), and Davis, Fisher, and Veracierto (2014), among others.

⁵In recent work, Davis, Fisher, and Veracierto (2014) also study a spatial model, in which migration is frictional. However, that paper focuses on explaining a different set of facts regarding the relation between gross and net flows, as well as the dynamic responses of population and gross flows to TFP shocks.

recession (Sterk (2010), Sahin, Song, Topa, and Violante (2014), Karahan and Rhee (2013), and Ravn and Sterk (2013) among others).

The rest of the paper is organized as follows. Section 2 presents a set of motivating empirical facts. In Section 3 presents the general model. Section 4 explains how a simplified version of the model can account for the comovement of gross flows and labor and housing market conditions in the data. Section 5 describes the estimation procedure, while section 6 contains the estimation results, discusses the model mechanisms and performs a number of counterfactual simulations. Section 7 provides brief concluding remarks.

2 Empirical Facts

I first examine how gross migration flows co-move with local labor and housing market conditions in a panel of U.S. cities for the period 1992-2010. Cities are defined as Core Based Statistical Areas (CBSAs), which is the standard definition for a metropolitan area in the U.S. There are a total of 156 such metropolitan areas in the panel.⁶ I first estimate the following regression

$$y_{i,t} = \alpha_i + \zeta_t + \mathbf{x}'_{i,t}\beta + \epsilon_{i,t} \quad (1)$$

where $y_{i,t}$ is the (log of the) out- or in-migration rate for city i and year t , α_i and ζ_t are city and year fixed effects, and $\mathbf{x}_{i,t}$ is a vector of the city-level time-varying characteristics. I consider two specifications with \mathbf{x} containing the (log of) city unemployment rate and (log) house price index in the first and the (log) unemployment rate and (log) rental price in the second. Including city and year fixed effects is particularly important since it allows me to examine comovements that are only due to local shocks by controlling for permanent city-specific differences and common shocks. Additionally, I control for metropolitan area (log)

⁶The Online Appendix contains information on how I construct all the relevant variables, as well as details on the data sources.

income per capita.⁷

[Table 1]

Table 1 shows the results from these regressions. First of all, higher unemployment in a city correlates with both increased migration out of the city and decreased migration into the city. Thus, local labor market conditions affect reallocation, and reallocation serves as an adjustment mechanism in response to local labor market shocks (Blanchard and Katz, 1992). Note that both the out- and in-migration margins co-move with changes in unemployment. This shows that labor reallocation is driven not only by an increase in outflows from a local labor market but also that migration is directed, so better performing labor markets experience higher inflows.

Secondly, higher house prices or higher rental prices correlate with increased out-migration and decreased in-migration, controlling for unemployment.⁸ Thus, housing market conditions drive regional reallocation. This result may not appear surprising at first. After all, it shows that housing markets affect reallocation much like the motivating example in the Introduction. Cities with high house prices are “too expensive”, so expensive housing induces more households to move out and fewer to move in.

What makes this result unexpected is trying to interpret it through the lens of a spatial equilibrium model with frictionless mobility (Rosen (1979) and Roback (1982)). In that framework equilibrium house prices (or rental prices) are never “too high” and always reflect differences in local characteristics (either the state of the local labor market or local amenities). Therefore, with frictionless mobility, a city has a high house price because more people move into it and fewer people move out of it, given city characteristics. However, such instantaneous regional reallocation, should imply the reverse comovements relative to those

⁷Since the focus of this paper is on understanding the extent of reallocation in response to local business cycle fluctuations, I will not be exploring the role of income differences for migration flows beyond the effect of unemployment differences. In the Online Appendix I present estimation results without controlling for income.

⁸Saks and Wozniak (2011) document similar observations for house prices for state-to-state and MSA-to-MSA migration flows. Jackman and Savouri (1992) obtain similar results for the U.K.

observed in the data. In contrast, as I show in this paper, an individual migration process that leads to slow regional reallocation in response to differences in local characteristics can help with explaining these comovements.

What migration process can lead to slow reallocation? One possibility is that city-level inflows and outflows do not respond to differences in local labor markets. However, this runs contrary to the observation of the effects of local unemployment on in- and out-migration. Nevertheless, this response may be weak, due to some frictions in individual migration decisions or due to individual heterogeneity.

To show that this could indeed be the case, I examine the flows between city pairs. I show how flows between an origin city and a destination city depend on the unemployment level in the destination city, specifically, whether the destination city has an unemployment rate above or below the median for a particular year. I estimate the following regression

$$y_{i,j,t} = \gamma D_{j,t} + \alpha_{i,j} + \delta_{i,t} + \zeta_t + \mathbf{x}'_{i,t} \beta + \epsilon_{i,t} \quad (2)$$

where $y_{i,j,t}$ is the (log of the) flow rate between the origin city i and destination city j in year t , and $D_{j,t}$ is an indicator for city j having an unemployment rate below the median in year t . Additionally, I control for a number of fixed effects (origin-destination pair ($\alpha_{i,j}$), origin \times year ($\delta_{i,t}$), and year (ζ_t)) and destination city characteristics such as income per capita (included in the vector $\mathbf{x}_{i,t}$).

[Table 2]

Table 2 shows the results from these regressions. Having an unemployment rate below the median increases the flows to that city by around 7 to 9 percent. Thus, although directed towards cities with better performing labor markets, migration flows respond fairly weakly to such differences. Therefore, imperfectly directed migration in response to local labor market differences must be one important source of slow reallocation.

To summarize, this section shows the following facts about migration flows. First, migration flows, both out of a city and into a city respond to local labor market conditions, and migration is (partially) directed. Second, housing markets have a direct effect on migration flows. Third, the extent of directed migration into a city in response to that city's labor market conditions is small.

3 A Model with Frictional Mobility

This section presents a spatial equilibrium model with housing markets and frictional mobility. In the next section I provide some intuition for the implications of slow reallocation in a simplified version of the model before proceeding to the estimation of the full model in Section 5.

I consider a discrete time infinite horizon economy, with $t = 0, 1, 2, \dots$. The economy consists of a unit measure of islands or regions. The economy is populated by a measure L of infinitely lived workers that reside across islands. Workers are risk neutral and derive utility from consumption, as well as from housing services. They discount future utility flows with a discount factor of $\beta < 1$. Workers have separable flow utility for consumption and housing. Similarly to Van Nieuwerburgh and Weill (2010), the flow utility function is quasi-linear and given by $\tilde{u}(c, h) = c + v(h)$, where $v(\cdot)$ is strictly increasing, twice continuously differentiable, concave, and satisfies Inada conditions. Workers can supply 1 unit of labor.

The initial measure of workers in a region j is given by l_{-1}^j , with $\int l_{-1}^j dj = L$. The end-of-period or post-migration measure of workers in a region j at time t is given by $l_t^j \in [0, \bar{L}]$ for $\bar{L} \geq L$.⁹

⁹The upper bound \bar{L} is a technical restriction necessary for showing analytical results. It can be thought of as a physical limit on the space available within a region.

3.1 Labor markets, job creation, and destruction

In each island, there is a representative firm that can post job vacancies at a per-period cost of k and recruit workers. Each job in island j has productivity a_t^j at time t . Local productivity shocks in the model serve as reduced-form labor market shocks that drive local business cycles and lead to unemployment differences across islands.¹⁰ Specifically, local productivity a_t^j has bounded support over $\mathcal{A} = [\underline{a}, \bar{a}]$, for $\underline{a} < \bar{a}$, and follows a stochastic process with persistence parameter $\rho \geq \frac{1}{2}$. In particular,

$$a_t^j = \begin{cases} a_{t-1}^j & , \text{ with. pr. } \rho \\ \max \{ \underline{a}, \min \{ \bar{a}, A + \eta_t^j \} \} & , \text{ otherwise} \end{cases} \quad (3)$$

where η_t^j is distributed i.i.d. with distribution function F_η and with $E[\eta] = 0$.

There is free entry into job creation, and as in the standard search and matching framework, jobs have stochastic lives. In particular, at the end of each period, after production takes place, with probability $s \in (0, 1)$ a job becomes unproductive and is destroyed. The firm is owned by all agents in the economy and discounts payoffs at the same discount factor of β .

The labor market of each region is characterized by a search and matching friction as in the standard Diamond-Mortensen-Pissarides framework (Pissarides, 2000). In particular, in a region j at time t , after migration takes place, a measure u_t^j of unemployed workers and measure v_t^j of vacancies try to match. Matching is described by a reduced-form constant returns to scale matching function $m^j(u_t^j, v_t^j)$ with standard properties, which gives the total number of local matches per period. I assume that matching functions are identical across regions, i.e. $m^j(u_t^j, v_t^j) = m(u_t^j, v_t^j) = u_t^j m\left(1, \frac{v_t^j}{u_t^j}\right)$. Letting $z_t^j \equiv \frac{v_t^j}{u_t^j}$ be the regional labor

¹⁰I do not attempt to identify separate drivers of local business cycles in this paper (for example, TFP shocks). Rather, in the estimation I (implicitly) consider all shocks that drive local business cycle fluctuations by looking directly at the labor market outcome of these shocks (i.e. fluctuations in local unemployment rates). Identifying separate drivers of local business cycles is beyond the scope of the paper, since the focus is on labor reallocation in response to any shocks that cause local fluctuations.

market tightness and defining $\mu(z) \equiv m(1, z)$, it follows that $m(u_t^j, v_t^j) = u_t^j \mu(z_t^j)$. This translates into a job finding probability for a worker in a given period of $\mu(z_t^j)$ and a job filling probability for a vacancy of $\frac{\mu(z_t^j)}{z_t^j}$. Additionally, I assume that $\lim_{z \rightarrow 0} \frac{\mu(z_t^j)}{z_t^j} = 1$ and $\lim_{z \rightarrow \infty} \frac{\mu(z_t^j)}{z_t^j} = 0$. Workers that are matched with a job obtain wage w , while those that remain unmatched in a given period receive a period payoff of e .¹¹

I allow for wages to be determined either by Nash bargaining or to be rigid as in Hall (2005). The particular wage determination rule does not affect the qualitative results in the next section. However, as I discuss in Section 5, the calibrated model features wage rigidity.

3.2 Housing markets

Every region j has a fixed housing stock, H , which is perfectly durable (no depreciation). The assumption of a durable housing stock is similar to the assumption in Glaeser and Gyourko (2005). However, while these authors explore the amplification effects of durable housing on labor market shocks in a static context, I consider its dynamic implications. Additionally, the fixed housing stock creates congestion effects, which combined with the concave utility from housing makes housing markets affect migration decisions.

The housing stock is owned by a sector of real estate firms that trade it in a competitive housing market and also rent it to workers residing in the region in a competitive rental market. Real estate firms are equally owned by all workers in the economy and also discount future payoffs with a factor β .¹²

The regional rental price is given by r_t^j and the price of a unit of housing is p_t^j . There is a no arbitrage relationship for the price of housing:

$$p_t^j = r_t^j + \beta E_t [p_{t+1}^j], \quad (4)$$

¹¹Notice that free entry and a linear production technology imply that the labor market tightness z will be a function of local productivity a only.

¹²Allowing for a local bias in the ownership of the local housing stock reduces the effects of the housing market on regional reallocation since in that case workers are (partially) compensated for their housing costs in equilibrium. Therefore, the current set-up without a local bias in ownership gives an upper bound for the effect of the housing market on reallocation.

where $E_t[\cdot]$ denotes expectation with respect to information available at time t . Together with a transversality condition on p_t^j , $\lim_{T \rightarrow \infty} \beta^T E_t [p_{t+T}^j] = 0 \forall t, j$, equation (4) implies that

$$p_t^j = E_t \left[\sum_{s=0}^{\infty} \beta^s r_{t+s}^j \right] \quad (5)$$

3.3 Migration

Even though workers are assumed to be identical there is heterogeneity over their employment state, i.e. in the beginning of each period some workers are unemployed, while others are still employed at a particular job. Therefore, I make migration assumptions for both employed and unemployed workers.

I assume that unemployed workers have an idiosyncratic region preference ϵ for the region they currently reside in. At the beginning of each period an unemployed worker draws a new ϵ from a continuous distribution F_ϵ with $E[\epsilon] = \chi \geq 0$ and support over $[-B + \chi, B + \chi]$ for some $B > 0$.¹³ After observing his match quality, the worker decides whether to move to a different region. Moving is instantaneous and entails a fixed cost of c .¹⁴ Unemployed workers that move search for a job in their new region of residence.

In contrast, employed workers move in response to an exogenous idiosyncratic shock, which arrives with per period probability ϕ . Once an employed worker is hit by a moving shock he separates from his job and immediately moves to another region as an unemployed worker. This exogenous moving assumption is reasonable given the moving behavior of employed workers observed in the data.

Worker migration for both the (previously) employed and the unemployed is a combination of directed and random migration. In particular, with probability $\lambda > 0$ a worker

¹³I also assume that employed workers do not draw a stochastic region preference ϵ but directly obtain a per-period payoff equal to $E[\epsilon] = \chi$.

¹⁴Such idiosyncratic region preferences are similar to those in Kennan and Walker (2011). I model the regional preference shocks as i.i.d. rather than including some persistence, since a persistent regional preference shock essentially acts as a higher cost of moving and the two cannot be separately identified with the aggregated data that I use. See Bayer and Juessen (2012) for a discussion of the dynamic selection effect due to persistent regional preference shocks and for identification using individual income data.

migrates to regions that offer the best labor and housing market conditions (directed migration), while with probability $1 - \lambda$ the worker is equally likely to migrate to any region in the economy (undirected migration). Setting $\lambda = 1$ leads to perfect directed migration by unemployed workers.

3.4 Timing

The timing within a period is as follows. 1. Agents observe the realization of regional productivity, a ; 2. Workers make migration decisions; 3. Housing and rental markets opens; 4. Firms make vacancy posting decisions; 5. Matching takes place between unemployed workers and vacant jobs; 6. Production takes place and wages for the period are paid; 7. Finally, there is job destruction.

4 Equilibrium

I will be focusing on stationary symmetric recursive equilibria, in which each region j is fully characterized by a vector of state variables $X_t^j = (a_t^j, l_{t-1}^j, u_{t-1}^j)$ and there is an invariant distribution over regional characteristics X denoted by ω^* . The relevant state vector contains the current period productivity a_t^j , as well as the beginning-of-period labor force (l_{t-1}^j), and the measure of unemployed workers (u_{t-1}^j).

A stationary symmetric recursive equilibrium will then be defined by laws of motion for the endogenous state variables, $l'(X_t^j)$, $u'(X_t^j)$, and for the distribution over X , and by functions, $z(a_t^j)$, $w(a_t^j)$, $p(X_t^j)$, $r(X_t^j)$, giving local market tightness, wages, house prices, and rental prices as a function of the payoff-relevant state vector, such that (i) worker migration decisions are optimal given the laws of motions for a , and the endogenous state variables, (ii) the law of motion for the endogenous state variables are consistent with worker migration decisions and with population constancy in the economy, (iii) the invariant distribution ω^* is a fixed point of the law of motion for the distribution over X , and (iv) there is rental and

housing market clearing.

In the rest of this section, I discuss the equilibrium and regional dynamics for a simplified version of the model. The purpose of this section is to show how the mobility assumptions in the model allow it to produce the comovements discussed in Section 2. In particular, I assume that jobs last for only one period, i.e. that $s = 1$, so that all workers in a given region are unemployed at the beginning of a period. This leads to only one relevant endogenous state variable, namely the beginning-of-period labor force l_{t-1} . Secondly, I assume that there can be only two productivity realizations, $a \in \mathcal{A} = \{\underline{a}, \bar{a}\}$, so productivity follows a two-state Markov chain with persistence $\rho \geq \frac{1}{2}$. Given the stationary distribution for this Markov chain, at any time t , one half of regions have $a_t^j = \bar{a}$ and the other half have $a_t^j = \underline{a}$. I refer to the former as (relatively) “booming” regions and to the latter as (relatively) “depressed” regions. Lastly, I define the unemployment rate as the fraction of unemployed workers after the matching stage, i.e. $U_t^j = (1 - \mu(z_t^j))$.¹⁵

4.1 Worker migration decisions

Let $V(X)$ be a worker’s end-of-period (post-migration) value, given the regional state X . Then

$$\bar{V} = \lambda \max_{\tilde{x}} \{V(\tilde{x})\} + (1 - \lambda) \int V(a, l) d\omega^* \quad (6)$$

is the migration value of the worker. The first term captures the value from directed migration - with probability λ the worker moves to regions giving him the highest expected utility. The second term is the random migration component - with probability $1 - \lambda$ the worker is equally likely to migrate to any region of the economy. We have that

$$V(X) = \max_h \{v(h) - r(X)h\} + e + \mu(z(a))(w(a) - e) + \beta E_X \left[W(X') \right], \quad (7)$$

¹⁵I assume that wages are determined by Nash bargaining. In particular, letting workers’ bargaining power be $\iota \in [0, 1)$, the wage rate is $w_t^j = e + \iota(a_t^j - e)$.

where

$$W(X) = \max_{\bar{\epsilon}} \left\{ F_{\epsilon}(\bar{\epsilon}) (\bar{V} - c) + (1 - F_{\epsilon}(\bar{\epsilon})) V(X) + \int_{\bar{\epsilon}} \epsilon dF_{\epsilon} \right\} \quad (8)$$

is the beginning-of-period (pre-migration) value function for the worker. These value functions follow directly from the value functions of workers in the general model, which can be found in the Online Appendix.

The interpretation of these value functions is straightforward. For V , the first three terms capture the flow utility, net of housing costs. The second is the expected value in the next period, which takes into account the migration option of the worker. In particular, for a given region preference ϵ , a worker compares the value of staying in the region to the value of moving. Given the structure of the problem, migration will follow a cutoff rule for ϵ , which I denote by $\bar{\epsilon}$, and which is given by:

$$\bar{\epsilon}(X) = \bar{V} - V(X) - c \quad (9)$$

Then for $\epsilon < \bar{\epsilon}(X)$ the worker migrates and for $\epsilon > \bar{\epsilon}(X)$ the worker stays, which means that the fraction of workers migrating from a region with state X is:

$$q(X) = Pr(\epsilon \leq \bar{\epsilon}(X)) = F_{\epsilon}(\bar{\epsilon}(X)), \quad (10)$$

which is also the *ex ante* probability of worker migration prior to the realization of ϵ .

4.2 Laws of motion for the endogenous state variable

Given the worker migration decisions above, it follows that the end-of-period measure of workers in a given region j , l_t^j , is:

$$l_t^j = l'(X_t^j) \equiv (1 - q(X_t^j)) l_{t-1}^j + \Psi(X_t^j), \quad (11)$$

where $\Psi(X_t^j)$ is a function that gives the measure of workers migrating into region j at time t . In particular,

$$\begin{aligned}\Psi(X_t^j) &\geq (1 - \lambda)M, & X_t^j \in \arg \max_x \{V(x)\} \\ \Psi(X_t^j) &= (1 - \lambda)M, & o.w.\end{aligned}$$

where

$$M = \int q(a, l) l d\omega^* \tag{12}$$

is the aggregate measure of workers migrating in a given period. Note that due to partially directed migration, some regions will experience higher inflows than others. The set $\mathcal{X} \equiv \arg \max_x \{V(x)\}$ of regions that offer the highest expected utility, $\max_{\tilde{x}} \{V(\tilde{x})\}$, and the exact form of $\Psi(X)$ for $X \in \mathcal{X}$ are determined in equilibrium.

4.3 Regional Dynamics

The following Proposition characterizes regional dynamics in a stationary equilibrium of this economy:

Proposition 1. *There exist two population levels \underline{l}^* and \bar{l}^* such that:*

1. *regional populations lie in the set $[\underline{l}^*, \bar{l}^*]$ and the stationary distribution ω^* is discrete;*
2. *a depressed region's population declines towards \underline{l}^* , and the out-migration rate decreases as the population declines towards \underline{l}^* ;*
3. *transitioning from depressed to booming, a region's population first increases to a value $\tilde{l}(\bar{a}) \leq \bar{l}^*$, and after that increases towards \bar{l}^* , provided $\tilde{l}(\bar{a}) < \bar{l}^*$. Also, the out-migration rate increases as the population increases towards \bar{l}^* .*

Proof. See Online Appendix B. □

Figure 1 summarizes these implications (for the case, in which $\tilde{l}(\bar{a}) < \bar{l}^*$) and compares them against the outcome of a model with frictionless mobility and perfectly directed migration in the spirit of a classical spatial equilibrium model (Rosen (1979) and Roback (1982)).

In the first case (Figure 1a), idiosyncratic region preferences and partially directed migration lead to slow population adjustments, creating a rich stationary distribution of populations even with only two labor market states. This population distribution, combined with durable housing and a concave utility from housing services, induces a dependence of housing prices on the history of labor market shocks. In contrast, in the second case (Figure 1b), the equilibrium distribution of populations has a two-point distribution, $\{\underline{l}^{**}, \bar{l}^{**}\}$, for depressed and booming regions, respectively. Depressed regions that experience a labor market improvement also experience a jump in population from \underline{l}^{**} to \bar{l}^{**} , and vice versa. As a result, variation in housing prices fully reflects variation in current labor market conditions.

[Figure 1]

The dependence of housing prices on the history of labor market shocks allows the model to produce the comovements between prices and gross migration documented in Section 2.

Proposition 2. *Consider a cross-sectional sample of J regions from the model economy. Let $out^j = out(a^j, l^j)$ and $in^j = in(a^j, l^j)$ be the out-migration and in-migration rates for region $j \in \{1, 2, \dots, J\}$. Also, let $U^j = U(a^j, l^j)$ be the unemployment rate and $\tilde{p}^j = \tilde{p}(a^j, l^j)$ be the beginning-of-period house price prior to worker migration. Then:*

1. *for a given \tilde{p}^j , out^j is increasing in U^j and in^j is decreasing in U^j ;*
2. *for a given U^j , out^j is increasing in \tilde{p}^j and in^j is decreasing in \tilde{p}^j .*

Proof. See Online Appendix B. □

The following features of the model drive these results. First, the heterogeneity in equilibrium housing prices arising from the regional histories of labor market shocks drives the positive comovement between out-migration and house prices, holding current labor market conditions fixed. Frictional out-migration also leads to a comovement between out-migration and unemployment, which reflects current labor market conditions only. On the other hand,

partially directed migration implies that regions with booming labor markets and lower populations have larger population inflows, which leads to a negative comovement of both house prices and unemployment with in-migration.

5 Estimation

I now estimate the general model introduced in Section 3. I use indirect inference (Smith (1993), Gourieroux, Monfort, and Renault (1993)) to estimate the model parameters. The Online Appendix contains additional details on the indirect inference estimation procedure. Below I discuss the features of the data that I use to estimate the parameters of the model, as well as specific functional form and parametric assumptions.

5.1 Functional form assumptions and parameters

I jointly estimate a subset of the model parameters. The rest of the parameters I calibrate to values in data sets different from the data I will use for estimation or to values that are standard in the literature. Additionally, I make a number of functional form assumptions.

First, I set the supply of housing in each region to one ($H = 1$).¹⁶ I assume that the flow utility from housing services $v(h)$ takes the form

$$v(h) = \begin{cases} \kappa \frac{h^{1-\nu}}{1-\nu} & , \nu \neq 1 \\ \kappa \log(h) & , \nu = 1 \end{cases} \quad (13)$$

where κ is a scaling parameter that affects the expenditure share on housing and the importance of housing in agents' utility, and ν is the inverse of the elasticity of demand for housing. I use standard estimates in the literature to pre-set the housing utility parameters. Specifically, I set $\nu = 2$ using the estimate for the inverse elasticity of housing demand from

¹⁶I also normalize the total population of workers in the economy to $L = 1$. Since I do not match rental price levels but only the dispersion in rental prices, these normalizations are natural.

Hanusheck and Quigley (1980). Also, I use the estimate for the housing expenditure share of 24% from Davis and Ortalo-Magne (2011) and set $\kappa = 0.24$. This value gives a housing expenditure share of approximately 24% in a version of the model without unemployment differences between regions.

Turning to the matching technology, I use a Cobb-Douglas matching function, $m(u, v) = Mu^{1-\alpha}v^\alpha$, which implies that $\mu(z) = Mz^\alpha$ for $z \equiv \frac{v}{u}$. I calibrate α and M from JOLTS using monthly data from December, 2000 to December, 2007.¹⁷ The estimates I obtain for the matching function are $\alpha = 0.605$ and $M \approx 1$. The value of α obtained lies in the middle of the set of estimates reported by Petrongolo and Pissarides (2001).

I assume that the labor productivity shock η is normally distributed with variance σ_η^2 and normalize the average labor productivity to $A = 1$. Similarly, I assume that regional preferences are drawn from a truncated normal distribution with zero mean and variance σ_ϵ^2 , with support given by $[-B + \chi, B + \chi]$, where I set $B = 5\sigma_\epsilon$.

The assumption of a normally distributed preference shock implies that one cannot separately identify the moving cost c and the average preference for staying in a region χ . Since relying only on a functional form assumption to identify a parameter is unsatisfactory, in the estimation below I estimate the sum $c + \chi$, which I denote by \tilde{c} . For brevity, below I refer to \tilde{c} as a mobility cost, even though it contains both a pecuniary and non-pecuniary component.

A time period in the model corresponds to a month, and a region in the model corresponds to a metropolitan area. I set the discount factor β to 0.995, which gives an annual discount rate of around 6%. For the estimation, I set the flow benefit from unemployment, e , to 0.65, which lies between the values proposed by Shimer (2005) and Hagedorn and Manovskii (2008). The job destruction probability s is set to 0.032. This is slightly lower than the

¹⁷JOLTS contains information on total hires per month, which when divided by the total stock of unemployed gives the job finding probability $\mu(z)$. The value of z is similarly obtained as the total vacancies divided by the stock of unemployed. I estimate the aggregate matching function for a time period with low regional dispersion to ensure that estimate would be close to a matching function estimate at a lower level of aggregation (Barnichon and Figura, 2011).

number used in Shimer (2005) and Hall (2005). However, note that there is additional job destruction with probability ϕ induced by the moving shocks of employed workers. At the estimated parameters, the overall job destruction probability is comparable to the number in these two studies. I summarize the predetermined parameters in Table 3a.

[Table 3]

Given the set of predetermined parameters, there are 7 model parameters that remain to be estimated. I denote these by the vector θ . In particular,

$$\theta = \{\rho, \sigma_\eta, k, \lambda, \tilde{c}, \sigma_\epsilon, \phi\} \in \Theta \quad (14)$$

where the parameter space Θ is a compact subset of a Cartesian product of positive real intervals. The parameters to be estimated are summarized in Table 3b.

Lastly, as already mentioned in Section 3, I assume that wages in the calibrated model are rigid in the sense of Hall (2005). As pointed out by Shimer (2005), the canonical search model with Nash bargaining leads to a large response of wages to changes in labor productivity, unless the flow payoff from unemployment is close to productivity (Hagedorn and Manovskii, 2008). The large sensitivity of the bargained wage implies that changes in labor productivity are mostly absorbed by changes in the wage, resulting in small effects on the job finding probability and from there on unemployment. An analogous problem arises in the environment with regional labor markets that I consider. Specifically, a model with Nash bargaining has troubles producing the observed metropolitan area unemployment dispersion without assuming an implausibly high variance for labor market shocks. Furthermore, a large response of regional wages to local business cycles appears counterfactual (Blanchard and Katz (1992), Mangum (2012)).

The modification of the standard search model that Hall (2005) proposes is to include a rigid wage arising, for example, from a social norm, which does not vary with the aggregate business cycle, thus breaking the strong link between productivity and the wage. Further-

more, the wage lies in the bargaining set of a worker-job pair for every value of productivity over the cycle, and hence, does not violate individual rationality. This is the approach I adopt, as well.¹⁸ Similarly to Hall (2005), the regional wage rate is set at the wage rate from a standard search and matching model with symmetric Nash bargaining and no regional productivity dispersion but otherwise parametrized as my model.

5.2 Data and Moments

I use the same data-set for the estimation as the one used in Section 2 for the motivating facts, since several of the arguments of the binding function used in the indirect inference procedure are regression coefficients from the estimation in Section 2.

I use 7 moments in an exactly-identified estimation procedure. These moments reflect important features of the data that the model can account for. First, I use the two regression coefficient from the regression of out-migration on unemployment and rental prices given in Table 1 of Section 2. I denote these by β_u^{out} and β_r^{out} , respectively. Additionally, I include the coefficient on the indicator variable for local unemployment below the median in the regression with city-to-city flows (denoted by γ). Specifically, I use the last estimate in Table 2.

Next, I include the annual coefficient of autocorrelation for (the log of) local unemployment rate (denoted by ρ_u), after controlling for metropolitan area and year fixed effects, as well as for (log) per capita income. I also include the average residual standard deviation of (log) unemployment (σ_t^u) obtained after controlling for metropolitan area and year fixed effects, and (log) per capita income. Finally, I include the average aggregate unemployment rate (u_t^{agg}) for the 156 metropolitan areas in my sample, as well as the average

¹⁸See the Online Appendix for the exact conditions. There is a large subsequent literature dealing with rigid wages in search models (Gertler and Trigari (2009), Shimer (2010)).

population-weighted migration rate (q_t), defined as

$$q_t \equiv \sum_i \frac{l_{i,t}}{l_t} out_{i,t} \quad (15)$$

where $out_{i,t}$ is the out-migration rate from metropolitan area i and year t , $l_{i,t}$ is the population in metropolitan area i and year t , and l_t is the aggregate population for the metropolitan areas in my sample in year t . Therefore, q_t measures the fraction of the population that moves between metropolitan areas in a year. All of the aggregate series are time-averaged. Table 4 in the next section summarizes the moments used in the estimation, their observed values, as well as their values at the identified model parameters.

5.3 Identification

As long as the derivative matrix of the moment function in the indirect inference procedure has full column rank, the parameter vector θ will be locally identified by the estimation procedure. Additionally, we want the moment function to be sensitive to changes in θ , which is ensured by using features of the data that are informative for the underlying model parameters.

Proposition 2 in Section 4.3, and the discussion after it, touched on these issues. In particular, the coefficients in the migration regressions, as well as the migration rate are informative for the parameters in the model that affect the migration of employed and unemployed workers ($\lambda, \tilde{c}, \sigma_\epsilon, \phi$). For example, a higher coefficient on the indicator variable for local unemployment below the median in the regression with city-to-city flows corresponds to a higher probability of directed migration λ . The out-migration regression coefficients and the migration rate are informative about \tilde{c} and σ_ϵ , as well as the migration probability of employed workers, ϕ . For example, a higher value of \tilde{c} for a given value of ϕ and σ_ϵ leads to lower out-migration in response to an adverse labor market shock in a region, decreasing both the out-migration regression coefficients and the migration rate. Additionally, apart

from the direct effects on the migration rates, the parameters that affect migration have indirect effects through the whole equilibrium distribution of populations and rental prices. These general equilibrium effects also influence the model-generated regression coefficients and aggregate averages. Finally, the autocorrelation of regional unemployment rates and the dispersion of (log) unemployment are informative about the productivity process (ρ and σ_η), while the aggregate unemployment rate is informative about the vacancy posting cost k .

The estimation procedure relies on a simulation of the stationary equilibrium of the model economy and on constructing moments from model-generated data similarly to the way they are constructed using the observed data-set. The Online Appendix contains information on the numerical algorithm I use. The algorithm is standard apart from two features. First of all, rather than simulating a nested fixed point model, I adjust all endogenous objects simultaneously, which greatly speeds up the computation of equilibrium. Second, I use a tatonnement procedure to adjust the law of motion for the labor force, since, in essence, it constitutes a continuum of housing market clearing conditions.

6 Results and Counterfactual Experiments

6.1 Estimation results

[Table 4]

Table 4 presents the observed and simulated moments from the estimation procedure and the parameter estimates. The model does well in terms of matching the targeted moments observed in the data. The estimates for the mobility process λ , c , σ_ϵ , and ϕ in Table 4 are of particular interest.¹⁹ First of all, the estimation reveals a very low probability of directed migration with λ only around 0.07. Therefore, only 7% of all migrants in the model move towards the “best” regions of the economy (i.e. the regions with most favorable labor and

¹⁹The estimated local productivity process has properties that are similar to related estimates in the literature. For example, Lkhagvasuren (2012) calibrates a state-level autoregressive productivity process with a (weekly) persistence of 0.988 and a standard deviation of 0.0047.

housing market conditions), and the vast majority instead move in a non-targeted way. The main driver of this very low estimate is the low value of the regression coefficient γ (the coefficient on the indicator variable for local unemployment below the median in a regression of city-to-city flows).

The low estimate for the probability of directed migration implies that very few inter-city moves are related to labor or housing market differences between cities. There are several interpretations for this finding. First, a substantial part of inter-city moves may be the result of unmodeled heterogeneity of preferences for other regions (for example, due to a preference for return migration (Kennan and Walker (2011) and Kaplan (2012))). Second, agents may not be perfectly informed about local characteristics when making their migration decisions. Lastly, there may be belief heterogeneity about the future evolution of local labor and housing markets.

The estimated mobility cost \tilde{c} (the sum of a moving cost, c , and average preference for staying in a region, χ) is substantial. It equals around 2 years of wage income in the model. This is similar in magnitude to the estimates obtained in Kennan and Walker (2011), though using different sources of variation. In particular, while in my model differences in the present discounted value of income across locations are due to different job-finding probabilities only, in their framework it is due to differences in wages. Nevertheless, there is a similar reason for the identified high value of the mobility cost, \tilde{c} , in both cases, namely the weak relation between mobility and the local unemployment rate in my data set and between mobility and income levels in Kennan and Walker (2011).

The dispersion in regional preferences σ_η is low compared to \tilde{c} . This translates into a lower sensitivity of out-migration in response to changes in labor or housing market conditions. Finally, the migration probability for employed workers, ϕ , and the associated job separation probability are low compared to the baseline calibrated separation probability s . Therefore, separations induced by moves of employed workers have a small effect on the vacancy posting decisions of firms.

6.2 Model validation

I compare the performance of the model against a set of untargeted moments. These comparisons are presented in Table 5. Table 5a compares the regression coefficients from the in- and out-migration regressions from Section 2 (for the period 2001-2010) against the regression coefficients obtained from a model generated data-set. The model is generally able to replicate well the regression coefficients. The coefficient on house prices in an out-migration regression (the β_p^{out} coefficient) is a notable exception, with the model generating a substantially larger coefficient estimate compared to that obtained from the observed data. This large estimate is linked to the small dispersion in house prices in model-generated data compared to the observed dispersion, as evident in Table 5b.

[Table 5]

Table 5b presents observed and simulated standard deviations and autocorrelations for (the log of) out-migration and in-migration as well as the standard deviations for (log of) rental and house prices. As the table shows, the model is very close in generating a dispersion in rental prices similar to that observed in the data given standard parameter values for the utility from housing. The main reason for this is the slow reallocation that arises from the individual migration process given the estimated mobility parameters. In contrast, a model with frictionless mobility tends to produce counterfactually high variability of rental prices - an observation which has been termed a “puzzle” in urban economics (Davis and Ortalo-Magne, 2011).

The model produces a substantially lower house price variability compared to the data. The reason for this difference is that in the model house price variability is driven only by rental price variability. However, this is insufficient to generate substantial house price fluctuations. This provides indirect support for the notion that local house price variability both in the time series but also cross-sectionally is driven by variability in discount factors beside rental price variability (Campbell, Davis, Gallin, and Martin, 2009). Since the model

does not include a mechanism that would create variability in discount factors, it also cannot generate large price variability.

Comparing the standard deviations and autocorrelations for out- and in-migration, these are generally close apart from the autocorrelation of in-migration. The autocorrelation of out-migration is very close, while the standard deviation of out-migration is higher in the data than in the model. One reason for this is that in the data out-migration responds to house price differences independently of rental price differences, whereas in the model there is no independent source of variation in house prices beyond variation in rental prices. This is related to the discussion about the “non-fundamental” volatility of house prices above. The model also produces slightly more variable and substantially less persistent in-migration compared to the data. These are both due to the non-trivial fraction of directed migration in the model, which leads to occasional jumps in the population. Jumps in population in the model occur in response to rental price differences arising between locations with similar labor market conditions.

Finally, the slow reallocation process that arises from the individual migration decisions can generate positive autocorrelation in house price growth rates. The reason for this is the smooth adjustment in house prices when labor market conditions worsen, due to slow migration out of a region or when labor market conditions improve, due to the low probability of directed migration. This creates persistence in house price changes, with an annual autocorrelation of 0.22 compared to 0.71 in the data. In contrast, a spatial equilibrium model with frictionless mobility cannot generate any autocorrelation in house price growth rates at annual frequency (Glaeser, Gyourko, Morales, and Nathanson, 2014).

6.3 Model mechanisms

A key feature of the model is that the slow reallocation process that results from individual migration decisions leads to variation in housing prices beyond the variation in current labor market conditions. As explained in Section 4 for a simplified version of the model, this

additional variation comes from the different histories of labor market shocks, which given the slow reallocation influence regional populations and housing prices.²⁰

To show this property in the estimated model I perform the following exercise. I generate an artificial data set for 1000 regions over 5 years and identify regions that have unemployment rates persistently above (below) the median for the first four years. I then sort these regions into deciles based on their unemployment rates in the fifth year of the simulated data set. Finally, I compute deviations of the (log of) regional house prices, rental prices and populations from their decile-specific mean and compare the average of these quantities between the two groups of regions. Table 6aa presents these comparisons. In the model, regions with low unemployment rates for the previous four years tend to have higher house prices, rental prices and populations, compared to regions with high unemployment in the previous four years, controlling for current labor market conditions.

I perform a similar exercise using the observed data. Specifically, I create an indicator variable for cities with unemployment below (above) the median in the previous four years. I then use these indicator variables in a regression of (log) house prices, rental prices, and population, additionally controlling for (log) unemployment, (log) income per capita, and city and year fixed effects. Therefore, I estimate the following regression

$$y_{i,t} = \gamma^{below} D_{i,t}^{below} + \gamma^{above} D_{i,t}^{above} + \mathbf{x}_{i,t}'\beta + \alpha_i + \zeta_t + \epsilon_{i,t}, \quad (16)$$

where $y_{i,t}$ is the dependent variable of interest (log of house price index, rental price or population) for city i and year t , α_i and ζ_t are city and year fixed effects, $\mathbf{x}_{i,t}$ is a vector of city-level time-varying characteristics, and γ^{below} and γ^{above} are the coefficients on the respective indicator variables. Table 6bb contains the estimation results. As in the model generated

²⁰Note that while in the one-period job model from Section 4 local labor productivity and the unemployment rate are perfectly correlated, this need not be the case in the general model. However, a notable feature of the canonical search and matching model is that unemployment is a fast moving state variable (Pissarides (2000), Shimer (2005)), so unemployment and productivity are nearly perfectly correlated. Local labor markets in my estimated model inherit this property of near perfect correlation between local productivity and unemployment.

data, controlling for the contemporaneous unemployment rate, cities with persistently low (high) unemployment have higher (lower) house prices, rental prices, and populations. The differences in the data are larger than in the model, particularly for house prices. This is related to the limited variation in house prices that the model generates, as discussed in Section 6.2. Overall, these results lend strong support for the main mechanism in the model.²¹

[Table 6]

6.4 Regional reallocation implications

Next, I turn to the implications of the estimated model for regional reallocation. I define the regional reallocation rate as:

$$realloc_t = \frac{1}{2} \sum_i \frac{l_{i,t}}{l_t} |in_{i,t} - out_{i,t}|, \quad (17)$$

where $in_{i,t}$ and $out_{i,t}$ are the in- and out-migration rates for city i in year t , and $\frac{l_{i,t}}{l_t}$ are population weights. This is a standard measure of the net flows across entities (cities, states, sectors) in a given year as a fraction of the total population. Its average value over the period 1992-2010 is 0.31%. Inter-city migration during 1992-2010 averages 3%. Therefore, on average the net reallocation rate is about 10 times lower than the gross migration rate across cities. This is similar in magnitude to the state-level evidence obtained in other studies (e.g. Coen-Pirani (2010), Lkhagvasuren (2012)).

The estimated model predicts a reallocation rate of 0.156%. This is around half of the average reallocation rate between 1988-2009. Therefore, the model predicts that around 50% of the observed net flows in the data is the result of reallocation in response to differences in local labor market conditions. Through the lens of the model, the remaining fraction of the

²¹In the Online Appendix I provide an additional test of the model's mechanisms based on the link between house prices and populations in the model.

observed reallocation is thus due to unmodeled persistent differences in amenities or income levels.²²

The level of reallocation in the model is influenced by the structure of individual migration decisions, namely the large value of the mobility cost, \tilde{c} , and the low probability of directed migration, λ . The mobility cost affects the extent to which reallocation is driven by out-migration, while the probability of directed migration regulates the importance of in-migration for reallocation. Additionally, housing markets exert a compensating effect for labor market differences. Below, I investigate the reallocation effects of these factors.

I first compare the importance of out-migration versus in-migration for reallocation. Figure 2, Panel A plots the regional reallocation rate for different levels of λ (left column) and \tilde{c} (right column). Unsurprisingly, increasing λ (reducing \tilde{c}) increases reallocation. What is more interesting is the shape of these curves. In particular, for low values of λ , a change in λ leads to a strong reallocation effect. In contrast, the reallocation effect from changes in \tilde{c} is limited given large values of \tilde{c} .

The compensation effects of housing markets affect the shape of the reallocation response to changes in λ and to a smaller extent the response to changes in \tilde{c} . To show this, in Figure 2 Panel B, I compare how regional reallocation depends on λ (left column) and \tilde{c} (right column) in two cases: 1) in the benchmark case with $\kappa = 0.24$ (solid line); and 2) in a “no housing market” case with $\kappa = 0.01$ (dashed line), which essentially removes housing from agents’ utility. As is evident from the Figure, while the reallocation response to changes in λ is concave in the benchmark case, it is linear in the “no housing market” case. Therefore, there is a strong complementarity between the reallocation effects of partially directed migration and the housing market. While at the estimated parameters the difference in reallocation is only 0.074% between the benchmark (0.156%) and the “no housing market” (0.23%) cases, for higher values of λ that differences increases substantially.

²²It should not be surprising that local labor market shocks cannot explain all of the observed net flows. For example, regressing city-level net migration rates only on city and year fixed effects explains around 52% of the variation in these rates.

The reason for this interaction arises from general equilibrium effects that operate through the equilibrium distribution of regional populations and which depend strongly on the value of λ . For example, when λ is low and there is little reallocation, the equilibrium distribution of regional populations is more compressed, and as a result there are small differences in regional rental prices. As λ increases, so the extent of reallocation is substantial, the equilibrium distribution of regional populations becomes more dispersed - there are now more regions with very low and very high populations in equilibrium. In that case housing markets compensate to a larger extent for labor market differences, and so also affect reallocation to a greater extent. There is a similar complementarity between \tilde{c} and the housing market, although, as evident from the figure, it is much smaller.

[Figure 2]

6.5 Unemployment implications

The previous section showed that the combination of a high mobility cost, a low probability of directed migration and the compensating effects of housing markets act to slow down reallocation. Slow reallocation in response to labor market differences naturally leads to regional mismatch - the co-existence of relatively tight and slack regional labor markets that unemployed workers search for jobs in.

To get a sense of the magnitude of this mismatch for the labor market and of the importance of regional reallocation for aggregate unemployment, I examine the joint response in aggregate unemployment and reallocation to changes in mobility and housing market parameters. Figure 2 Panel C plots the equilibrium unemployment rate for different values of λ and \tilde{c} for the cases with $\kappa = 0.24$ (benchmark) and $\kappa = 0.01$ (no housing market). Therefore, the figure is analogous to Figure 2 Panel B. As the Figure shows, aggregate unemployment changes very little even when reallocation is increased substantially. For example, even when reallocation goes up to around 2.5% (for the case with $\lambda \approx 1$ and $\kappa = 0.01$), unemployment decreases by only 20 basis points (from 5.5% to 5.3%).

Therefore, given the observed differences in local labor market conditions, regional mismatch does not play an important role for aggregate unemployment. Since (log) unemployment dispersion (after controlling for city and year fixed effects) remained fairly constant even during and after the recent recession, this result is suggestive of a limited role for regional mismatch in that recession. This observation confirms the conclusions from Sahin, Song, Topa, and Violante (2014) and Karahan and Rhee (2013), who show that contemporaneous events in the housing market did not matter much for the high level of unemployment in the aftermath of the recession.

Given the limited role of reallocation for unemployment, any other parameter change that affects reallocation will not affect unemployment significantly. For example, consider an increase in unemployment benefits e . An increase in e affects aggregate unemployment through its effect on reallocation by reducing the mobility of unemployed workers, since as e is increased employment becomes less valuable, and so moving to a region with a tighter labor market is less valuable. However, at the estimated parameters an increase in e from the benchmark value of 0.65 to 0.95 leads to an increase in unemployment of only 1 basis point.²³

7 Concluding comments

Individual migration decisions shape the process of regional labor reallocation in response to local labor market shocks, and ultimately, the whole equilibrium distribution of regional populations. This paper argues that a frictional migration process is important for accounting for the comovements between unemployment, rental and house prices and gross migration flows in a panel of U.S. cities, as well as for a number of other empirical facts. I build and estimate a spatial equilibrium model, in which workers experience idiosyncratic region

²³In the standard search and matching model, a change in unemployment benefits also has an indirect equilibrium effect on unemployment through firms' job creation decisions, as workers have a higher outside option and extract more of the match surplus. In my model this channel is not present, since wages are assumed to be fixed and do not respond to changes in workers' outside options.

preference shocks, and unemployed workers incur a moving cost when migrating. Worker moves are a combination of directed and undirected (random) migration. The estimation reveals that mobility costs are large and the probability of directed migration is very low. These act to limit regional reallocation. The resulting equilibrium distribution of populations also implies a limited effect of housing markets on reallocation.

The identified low probability of directed migration is particularly important for the observed low reallocation rate. Since existing models of regional reallocation assume either fully directed or fully random mobility, allowing for both directed and random mobility in an estimable model and bringing it to the data provides an important insight into the migration decisions of individuals. However, there are additional important determinants of migration decisions that the model has not considered, which may explain part of the random migration component, and more importantly, account for the remaining fraction of the observed reallocation rate across U.S. states. Therefore, enriching the model by considering persistent differences across regions, such as, for example, differential growth rates in income or amenities would allow the model to account for other important cross-sectional facts about migration. Considering such an extension is an important step for future research.

References

- ALVAREZ, F., AND R. SHIMER (2011): “Search and Rest Unemployment,” *Econometrica*, 79(1), 75–122.
- BARNICHON, R., AND A. FIGURA (2011): “What Drives Matching Efficiency? A Tale of Composition and Dispersion,” Federal Reserve Board Working Paper 2011-10.
- BAYER, C., AND F. JUESSEN (2012): “On the Dynamics of Interstate Migration: Migration Costs and Self-Selection,” *Review of Economic Dynamics*, 15:3, 377–401.
- BLANCHARD, O., AND L. KATZ (1992): “Regional Evolutions,” *Brookings Papers on Economic Activity*, pp. 1–75.
- CAMPBELL, S., M. DAVIS, J. GALLIN, AND R. MARTIN (2009): “What Moves Housing


- Markets: A Variance Decomposition of the Rent-Price Ratio,” *Journal of Urban Economics*, 66:2, 90–102.
- CARRILLO-TUDELA, C., AND L. VISSCHERS (2011): “Unemployment and Endogenous Reallocation over the Business Cycle,” mimeo.
- COEN-PIRANI, D. (2010): “Understanding gross worker flows across U.S. states,” *Journal of Monetary Economics*, 57(7), 769–784.
- DAVIS, M., J. FISHER, AND M. VERACIERTO (2014): “Gross Migration, Housing, and Urban Population Dynamics,” Federal Reserve Bank of Chicago Working Paper WP2013-19.
- DAVIS, M., AND F. ORTALO-MAGNE (2011): “Household Expenditures, Wages, Rents,” *Review of Economic Dynamics*, mimeo.
- ECKHOUT, J. (2004): “Gibrat’s Law for (All) Cities,” *American Economic Review*, 94:5, 1429–1451.
- GERTLER, M., AND A. TRIGARI (2009): “Unemployment Fluctuations with Staggered Nash Wage Bargaining,” *Journal of Political Economy*, 117(1), 38–86.
- GLAESER, E., AND J. GYOURKO (2005): “Urban Decline and Durable Housing,” *Journal of Political Economy*, 113:2, 345–375.
- GLAESER, E., J. GYOURKO, E. MORALES, AND C. NATHANSON (2014): “Housing Dynamics: An urban approach,” *Journal of Urban Economics*, 81, 45–56.
- GODOY, A., AND E. MOEN (2012): “Mixed search,” mimeo.
- GOURIEROUX, C., A. MONFORT, AND E. RENAULT (1993): “Indirect Inference,” *Journal of Applied Econometrics*, 8, S85–S118.
- HAGEDORN, M., AND I. MANOVSKII (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” *American Economic Review*, 98(4), 1692–1706.
- HALL, R. (2005): “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1), 50–65.
- HANUSHECK, E., AND J. QUIGLEY (1980): “What is the Price Elasticity of Housing Demand?,” *Review of Economics and Statistics*, 62:3, 449–454.
- HEAD, A., AND H. LLOYD-ELLIS (2012): “Housing Liquidity, Mobility and the Labour Market,” *Review of Economic Studies*, 79:4, 1559–1589.
- HEAD, A., H. LLOYD-ELLIS, AND H. SUN (2014): “Search, Liquidity and the Dynamics of House Prices and Construction,” *American Economic Review*, 104:4, 1172–1210.
- JACKMAN, R., AND S. SAVOURI (1992): “Regional Migration in Britain: An Analysis of Gross Flows Using NHS Central Register Data,” *Economic Journal*, 102, 1433–1450.

- KAPLAN, G. (2012): “Moving Back Home: Insurance Against Labor Market Risk,” *Journal of Political Economy*, 120:3, 446–512.
- KAPLAN, G., AND S. SCHULHOFER-WOHL (2012): “Understanding the Long-Run Decline in Interstate Migration,” Federal Reserve Bank of Minneapolis Working Paper 697.
- KARAHAN, F., AND S. RHEE (2013): “Geographical Reallocation and Unemployment during the Great Recession: The Role of the Housing Bust,” mimeo.
- KENNAN, J., AND J. WALKER (2011): “The Effect of Expected Income on Individual Migration Decisions,” *Econometrica*, 79:1, 211–251.
- LKHAGVASUREN, D. (2012): “Big Locational Unemployment Differences Despite High Labor Mobility,” *Journal of Monetary Economics*, 59:8, 798–814.
- LUCAS, R., AND E. PRESCOTT (1974): “Equilibrium Search and Unemployment,” *Journal of Economic Theory*, 7, 188–209.
- MANGUM, K. (2012): “A Dynamic Model of Cities and Labor Markets,” mimeo.
- OSWALD, A. (1996): “A Conjecture on the Explanation for High Unemployment in the Industrialized Nations: Part 1,” Warwick Economic Research Paper 475.
- PETRONGOLO, B., AND C. PISSARIDES (2001): “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39(2), 390–431.
- PISSARIDES, C. (2000): *Equilibrium Unemployment Theory*. MIT Press.
- RAVN, M., AND V. STERK (2013): “Job Uncertainty and Deep Recessions,” mimeo.
- ROBACK, J. (1982): “Wages, rent, and the quality of life,” *Journal of Political Economy*, 90(2), 191–229.
- ROSEN, S. (1979): “Wage-based Indexes of Urban Quality of Life,” in *Current Issues in Urban Economics*, ed. by P. Miezkowski, and M. R. Straszheim. Johns Hopkins University Press.
- RUPERT, P., AND E. WASMER (2012): “Housing and the labor market: Time to move and aggregate unemployment,” *Journal of Monetary Economics*, 59:1, 24–36.
- SAHIN, A., J. SONG, G. TOPA, AND G. VIOLANTE (2014): “Mismatch Unemployment,” *American Economic Review*, 104:11, 3529–3564.
- SAKS, R., AND A. WOZNIAK (2011): “Labor Reallocation over the Business Cycle: New Evidence from Internal Migration,” *Journal of Labor Economics*, 29(4), 697–739.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95(1), 25–49.
- (2010): *Labor Markets and Business Cycles*. Princeton University Press.

- SMITH, A. (1993): “Estimating Nonlinear Time-series Models Using Simulated Vector Autoregressions,” *Journal of Applied Econometrics*, 8, S63–S84.
- STERK, V. (2010): “Home Equity, Mobility and Macroeconomic Fluctuations,” mimeo.
- VAN NIEUWERBURGH, S., AND P.-O. WEILL (2010): “Why Has House Price Dispersion Gone Up?,” *Review of Economic Studies*, 77(4), 1567–1606.
- WINKLER, H. (2010): “The Effect of Homeownership on Geographic Mobility and Labor Market Outcomes,” mimeo.

Figure 1: Regional population dynamics

(a) With frictional mobility



(b) With frictionless mobility






Figure 2: Counterfactual experiments: dependence on λ (left column) and \tilde{c} (right column).

Panel A: Regional reallocation and mobility parameters



Panel B: Housing market effects of reallocation



Panel C: Housing market effects on unemployment

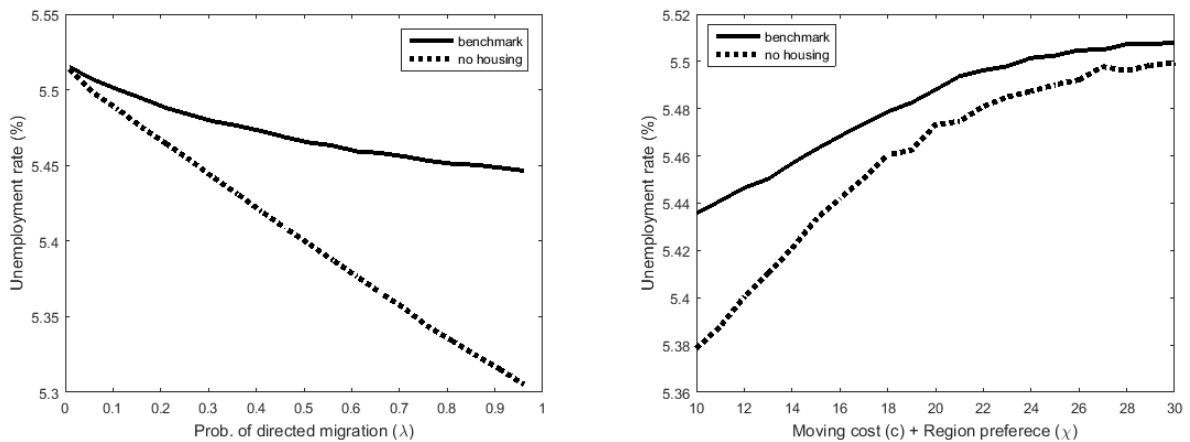


Table 1: Gross migration and metropolitan area characteristics

Dep. Variable:	out-migration rate (log)			in-migration rate (log)		
unemployment rate (log)	0.107*** (0.0263)	0.237*** (0.0678)	0.157** (0.0721)	-0.152*** (0.0357)	-0.248*** (0.0494)	-0.136*** (0.0481)
house price (log)	0.227*** (0.0339)	0.258*** (0.0393)		-0.218*** (0.0535)	-0.306*** (0.0527)	
rental price (log)			0.306*** (0.0985)			-0.421*** (0.154)
Adjusted R^2	0.945	0.944	0.941	0.961	0.964	0.96
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes	Yes	Yes
Period	1992-2010	2001-2010	2001-2010	1992-2010	2001-2010	2001-2010
Observations	2,964	1,560	1,338	2,964	1,560	1,338

Notes: Bootstrapped standard errors with clustering on the CBSA level in parenthesis. 'FE' denotes Fixed Effects. The out- and in-migration rates are constructed from the IRS migration data. See Appendix A for detailed description. House price is the log of the FHFA repeat sales price index. Unemployment is the log of the unemployment rate for the metropolitan area taken from the BLS LAUS database. Rental price is the log of the median contract rent in a metropolitan area taken from the American Community Survey. An additional control is the log income per capita (from the Bureau of Economic Analysis). *** denotes significance at 1% and ** denotes significance at 5%.

Table 2: City-to-city flow regression

Dep. Variable:	flow rate (log)		
unemployment below median (indicator)	0.0883*** (0.00275)	0.0698*** (0.00291)	0.0709*** (0.00270)
Origin-destination FE	Yes	Yes	Yes
Additional controls	Year FE	Year FE, Income (log)	Year FE, Income (log), Origin×Year FE
Observations	122,580	122,580	122,580

Notes: Robust standard errors with clustering on origin-destination CBSA pair in parenthesis. 'FE' denotes Fixed Effects. The city-to-city flows are constructed from the IRS migration data. See Appendix A for detailed description. The unemployment indicator takes the value 1 if the unemployment rate in a destination CBSA is below the median unemployment rate for the given year, and the value zero otherwise. Unemployment rates for metropolitan areas are taken from the BLS LAUS database. Additional controls include year fixed effects, origin CBSA×year fixed effects, and the log of income per capita (from the BEA). *** denotes significance at 1%.

Table 3: Model Parameters: pre-set (a) and estimated (b)

(a)		
parameter	description	value
β	discount factor	0.995
e	period unemployment payoff	0.65
s	job destruction probability	0.032
A	average regional productivity	1
H	regional housing supply	1
ν	inverse elasticity of housing demand	2
κ	weight on utility from housing	0.24
$\mu(z)$	matching function	$z^{0.605}$
(b)		
parameter	description	
ρ	persistence of labor productivity	
σ_η	dispersion of labor market shocks	
k	vacancy posting cost	
λ	probability of directed migration	
\tilde{c}	mobility cost ($c + \chi$)	
σ_ϵ	dispersion in regional preferences	
ϕ	probability of migration for an employed worker	

Table 4: Data and model-generated moments (a) and parameter estimates (b)

(a)

Description	Observed	Simulated
β_u^{out}	0.1573	0.1445
β_r^{out}	0.3059	0.3113
γ	0.0709	0.0617
ρ_u	0.7991	0.7805
σ^u	0.1530	0.1675
unempl. rate (u^{agg}) (%)	5.51	5.5
migration rate (q) (%)	3.01	3.23

(b)

ρ	σ_η	k	λ	\tilde{c}	σ_ϵ	ϕ
0.968	0.0036	0.6313	0.0722	26.6542	9.9276	0.0026
(0.0040)	(0.0005)	(0.0877)	(0.0505)	(2.2767)	(2.0308)	(0.0011)

Notes: Calculations are based on metropolitan area level data between 1992-2010 (2001-2010 for β_u^{out} and β_r^{out}). Estimation is performed by Indirect Inference. The simulated moments are obtained from 100 simulated panels of 156 regions over 220 time periods (months) keeping the last 120 time periods (10 years). β_u^{out} and β_r^{out} are regression coefficients from a regression of (log of) out-migration on (log of) unemployment rate and (log of) rental price, controlling for city and year fixed effects, and (log of) per-capita income. γ is the coefficient from a regression of city-to-city migration flow between cities on an indicator of whether unemployment in the destination city is below the median, controlling for origin-destination pair, origin \times year, and year fixed effects, and (log of) per capita income in the destination city. ρ_u is the autoregressive coefficient on (log of) unemployment, controlling for city and year fixed effects and (log of) per-capita income. σ^u is the (time-averaged) residual standard deviation of (log) unemployment, obtained after controlling for city and year fixed effects and (log of) per-capita income. u^{agg} is the (time-averaged) unemployment rate q is the inter-MSA migration rate. Parameter estimates are from Indirect Inference estimation. The estimation procedure chooses a vector of parameter values to minimize the distance between moments observed in the data and moments generated by simulated data from the model. Standard errors derived from numerical derivatives in parenthesis.

Table 5: Untargeted moments: regression coefficients (a) and standard deviation and auto-correlations (b)

(a)

Regression coefficient	Observed	Simulated
<i>Regression: in-migration on unempl.rate and rental price</i>		
β_u^{in}	-0.136	-0.201
β_r^{in}	-0.421	-0.334
<i>Regression: out-migration on unempl.rate and house price</i>		
β_u^{out}	0.237	0.153
β_p^{out}	0.258	0.856
<i>Regression: in-migration on unempl.rate and house price</i>		
β_u^{in}	-0.248	-0.178
β_p^{in}	-0.306	-0.266

(b)

Quantity	Observed	Simulated
σ_{out}	0.0956	0.0226
σ_{in}	0.0975	0.151
σ_{rent}	0.0378	0.0396
σ_{price}	0.0926	0.0151
ρ_{out}	0.5244	0.7211
ρ_{in}	0.7106	0.047

Notes: For Table 5a, β_u^{in} and β_r^{in} are regression coefficients from a regression of (log of) in-migration on (log of) unemployment rate and (log of) rental price, controlling for city and year fixed effects, and (log of) per-capita income. β_u^{out} and β_p^{out} are regression coefficients from a regression of (log of) outmigration on (log of) unemployment rate and (log of) house price index, controlling for city and year fixed effects, and (log of) per-capita income. β_u^{in} and β_r^{in} are regression coefficients from a regression of (log of) outmigration on (log of) unemployment rate and (log of) house price index, controlling for city and year fixed effects, and (log of) per-capita income. All are estimated in the period 2001-2010. The simulated regression coefficients are for data simulated from the estimated model. Specifically, I simulate 100 data-sets with 156 regions over 10 years and run regressions on the simulated data. I then average the coefficients over the 100 estimations. For Table 5b, σ_x is the standard deviation for (the log of) out-migration rate ($x=out$), in-migration rate ($x=in$), rental price ($x=rent$), house price ($x=price$) and ρ_x is coefficient of autocorrelation (annual) for (log of) out-migration rate ($x=out$) and in-migration rate ($x=in$). Quantities for the observed data (left column) are estimate for 1992-2010 (2001-2010 for rental prices), after controlling for city and year fixed effects, and (log of) per-capita income. Simulated quantities and for data simulated from the estimated model. Specifically, I simulate 100 data-sets with 156 regions over 10 years and compute the respective standard deviations and autocorrelations. I then average the coefficients over the 100 estimations.

Table 6: Comparisons between regions (cities) with historically low and high unemployment

(a) Model			
	house price (log)	rent (log)	population (log)
Unemployment below median in 4 previous years	0.0011	0.0032	0.0011
Unemployment above median in 4 previous years	-0.0031	-0.0096	-0.0042
Difference	0.0042	0.0128	0.0054
(b) Data			
	house price (log)	rent (log)	population (log)
Unemployment below median in 4 previous years	0.0452*** (0.00932)	0.0232*** (0.00437)	0.0091 (0.00561)
Unemployment above median in 4 previous years	-0.0338*** (0.00977)	-0.0200** (0.00683)	-0.015** (0.00739)
Difference	0.079***	0.0432***	0.0241***
CBSA FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Sample Period	1990-2013	2001-2010	1990-2010
Observations	3,657	1,350	3,339

Notes for Table 6b: Robust standard errors with clustering on the CBSA level in parenthesis. 'FE' denotes Fixed Effects. See Appendix A for detailed description. "Unemployment below median in 4 previous years" is an indicator variable with a value of 1 if the unemployment rate in the metropolitan area is below the (annual) median for each of the 4 previous years and 0, otherwise. "Unemployment above median in 4 previous years" is an indicator variable with a value of 1 if the unemployment rate in the metropolitan area is above the (annual) median for each of the 4 previous years and 0, otherwise. "Difference" refers to the difference between the two coefficient estimates. House price is the log of the FHFA repeat sales price index. Rent is the log of the median contract rent in a metropolitan area taken from the American Community Survey. Population is the log population for a metropolitan area (from the Bureau of Economic Analysis). Additional controls are the log unemployment rate and log income per capita (from the Bureau of Economic Analysis). *** denotes significance at 1%, ** denotes significance at 5%, and * denotes significance at 10%.