Georgia State University
ScholarWorks @ Georgia State University

UWRG Working Papers | Usery Workplace Research Group

2-3-2015

Cities and Labor Market Dynamics

Kyle Mangum
Georgia State University, kmangum@gsu.edu

Follow this and additional works at: https://scholarworks.gsu.edu/uwrg_workingpapers

Recommended Citation

Mangum, Kyle, "Cities and Labor Market Dynamics" (2015). *UWRG Working Papers*. 133.
https://scholarworks.gsu.edu/uwrg_workingpapers/133

This Article is brought to you for free and open access by the Usery Workplace Research Group at ScholarWorks @ Georgia State University. It has been accepted for inclusion in UWRG Working Papers by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

W. J. Usery Workplace Research Group Paper Series

Working Paper 2015-2-3
February 2015

Cities and Labor Market Dynamics

Kyle Mangum
Georgia State University

This paper can be downloaded at: http://uwrg.gsu.edu

Andrew Young School
of Policy Studies

Cities and Labor Market Dynamics

Kyle Mangum*

February 2015

Abstract

Labor mobility is widely viewed as the principal adjustment mechanism to local labor market shocks. Empirically, however, migration is infrequent, and population adjustment is slow and varies markedly between locations. This paper proposes and estimates a model of heterogeneous local labor markets–"cities"–to understand the mechanisms inhibiting labor market dynamics. To jointly study multiple forms of adjustment dynamics, the proposed model is an innovative dynamic spatial equilibrium, with endogenous local prices, heterogeneous locations, and agents subject to adjustment costs; the methodological contribution of this paper is the empirical implementation of such a model. The model is simulated under counterfactual scenarios in which impediments to labor adjustment are relaxed. The main result is that except in the extreme case of no adjustment costs, labor adjustment is limited not by the total amount of migration but its direction over markets. Population elasticity would be significantly greater if migration were less idiosyncratically motivated or if local shocks were not capitalized into housing prices. The implication is that effective local labor market reallocation does not necessarily require high aggregate mobility rates.

JEL codes: R23, J61, C63
Keywords: labor mobility, migration, local labor markets, regional labor markets, dynamic spatial equilibrium

1 Introduction

Labor mobility has long been viewed as a principal mechanism of adjustment to regional economic changes. The preeminent study of local labor markets, Blanchard and Katz (1992), concluded that workforce migration was the chief response of "regional evolutions." There is a large and long literature in regional analysis studying the geographic adjustment of population to labor market differentials.[1]

Empirically, however, the response of population to local labor market differentials is small and slow. In the U.S. since 1990, about 4 percent of households per year relocate from one metropolitan area to another, and most of this mobility is churning, not reallocation, with gross migration an order of magnitude larger than net migration. A recent structural microeconometric study of migration, Kennan and Walker (2011), found that people were indeed sensitive to income differentials when deciding whether to relocate, but mobility was costly, infrequent, and highly idiosyncratic.

*Department of Economics, Andrew Young School of Policy Studies, Georgia State University. Correspondence address is kmangum@gsu.edu. I thank Pat Bayer, Peter Arcidiacono, Andrew Sweeting, Chris Timmins, David Albouy, Spencer Banzhaf, Paul Beaudry, Patrick Coate, Julio Garin, Aurel Hizmo, Nick Kuminoff, and Nathan Schiff for helpful discussions, and seminar participants at Duke, the Federal Reserve Board, the Federal Trade Commission, Georgia, Georgia State, Penn State, UCLA, and the UEA 2013 meetings in Ottawa for feedback. Remaining errors are my own.

[1] See Greenwood (1985) and Greenwood (1997) for surveys.

Moreover, local labor markets exhibit substantial differences at a point in time and over time. For example, among the 30 largest metropolitan areas in the U.S. from 1990-2011, the unemployment rate gap averaged 8.6 percentage points and was as large 15.2. Roback (1982) establishes reason why some local market differentials might persist, with wages and land prices offering compensating differentials for differences in amenities. Yet much of the observed dispersion, particularly in unemployment, is not permanent. Local labor markets are apparently subject to local productivity shocks, giving workers incentive to move, yet population adjustment to these shocks is incomplete.

The puzzle, therefore, is that labor is mobile, but not quickly nor completely, and is sensitive to labor market conditions, but labor market arbitrage is incomplete. Adjustment frictions are clearly present, but which are important and how do they affect population adjustments? This paper empirically examines several reasons why population may be slow to adjust to local productivity shocks, quantifying the effects of various frictions on labor mobility.

Migration may be costly for financial and psychic reasons, and the paucity of migration suggests the costs are large. Information search may be burdensome, relocation expensive and disruptive, and the loss of local family and social networks costly to psychic utility.[2] Even if labor market differentials exist, in the presence of moving costs, the requisite difference threshold between markets may be very large for many households.

Moreover, in the presence of moving costs, households have incentive to forecast the future payoffs from a location choice. If productivity differences are expected to be short-lived, then relocation may not be worth the cost incurred.

Finally, the apparent labor market differential may not represent the true utility differential. Market amenities and the cost of living likely factor into a households location decision.[3] Workers may have idiosyncratic matches with particular locations, whether in the labor market, preferences for amenities, or personal connections such as family. I present evidence that a preponderance of migration across labor markets is indeed idiosyncratically motivated.

These mobility frictions (to use the term broadly) may have difference salience in different locations. Even if moving costs are large on average, some locations may be more isolated than others. Certainly, amenities, cost of living, and housing elasticity vary between markets.[4] Likewise, productivity may vary not just in levels, but in the process from which shocks are drawn. Thus, markets are not only subject to different shocks, but the implications of a given shock may be different if population responds heterogeneously.

To jointly study these types of mobility frictions, this paper places a dynamic structural model of migration in a spatial equilibrium specification. The dynamic model recognizes the presence of moving costs and the resultant forward-looking behavior, and further, a dynamic discrete choice specification permits rich heterogeneity in the location choice set. The spatial equilibrium allows for compensating differentials and congestion in the labor and housing markets. The emphasis of the model is population

[2] Molho (1995) distinguishes between distance deterrence and "cumulative inertia," or the formation of local personal attachments, in describing disutility to migration.

[3] Roback (1982) and the class of models that followed showed how local labor and housing market prices capitalize amenities as well as productivity.

[4] See, for example, Saks (2008), Saiz (2010), Albouy and Ehrlich (2012).

dynamics, so the point of the spatial equilibrium is not merely to allow for cross sectional differences in locations, but to examine how heterogeneity in locations affects responses to local productivity shocks. In this way, I leverage spatial variation (for instance, in housing markets) to understand how the mechanisms manifest in a dynamic setting.

Embedding a structural dynamic discrete choice model into an equilibrium setting presents a host of challenges in empirical implementation. As in the literature on dynamic games, the state space of such a model grows exponentially in the number of locations, so any reasonably large number of locations quickly becomes impossible to solve. I propose a novel approximation technique that preserves the fundamental elements of the dynamic equilibrium, yet makes the problem tractable enough to be solved repeatedly within an estimation routine. The methodological contribution of this paper is its empirical approach for an equilibrium dynamic discrete choice model with a large choice set.

Using quarterly data on the 30 largest labor markets in the U.S. plus an aggregated outside option, I fit the model to match local population changes and place-to-place migration flows. The locations are individually heterogeneous in the persistence and variance of two dimensions of labor demand shocks, their housing supply, and in the bundle of amenities they offer.

After recovering structural parameters on utility and moving cost, I simulate the model under counterfactual scenarios to examine the population responses to local labor demand shocks when frictions are relaxed. Specifically, I report three simulations: (1) when moving costs are absent, returning to a static model, (2) when idiosyncratic utility is diminished, elevating the importance of local productivity shocks, and (3) when housing demand is national, not local, unlinking house price dynamics from local productivity shocks. The simulation exercise is not merely to test whether these frictions are present–they all matter to some extent–but to examine how the economy would behave in their absence.

Each simulation offers its own insight. First, unsurprisingly, moving costs are present, and static equilibrium models fall far short of the data. A model with no location-adjustment costs responds immediately to labor demand shocks. The population response is greater in markets with elastic housing supply, as in Moretti (2011), since in inelastic locations, smaller population changes more quickly congest the market and resolve utility differentials.

Second, idiosyncratic forces loom large. In an economy in which workers put more weight on common factors vis-a-vis idiosyncratic factors, adjustment to location specific demand shocks is stronger. This simulation preserves underlying market heterogeneity, meaning the population responses change proportionally across locations.

Finally, non-traded goods capitalize local productivity shocks, inhibiting population response. An economy in which demand for housing is determined nationally shows substantially more response to local labor demand shocks, since a major congestion mechanism is removed. The effects are concentrated among markets most constrained by housing and those with longer-lasting shocks, which demonstrates how this effect, present in a static model, is magnified in a dynamic model.

Interestingly, the dynamic counterfactuals can generate more responsiveness to local shocks without changing total mobility rates. Thus, these latter two frictions affect the *labor market targeting* of mobility, not the total amount of migration. This result is important because the relevant issue for labor supply reallocation is net, not gross, migration.[5]

The simulations are, of course, tests of primitive mechanisms, and thus a step removed from real-world policy, but the results have at least three implications for policy. First, to the extent that adjustment costs and idiosyncratic preferences represent information asymmetries, population is moving suboptimally due to a market failure. But if these frictions are simply reflective of heterogeneity in preferences or in returns to local shocks, then policy should focus not on spatial reallocation per se, but on within-market matching and aggregate demand. Second, the capitalization of productivity in housing and other non-tradeables suggests that policies or financial instruments relaxing the link between local demand and local prices will lead to improved population allocation over space. Finally, the result that net migration can be improved without affecting gross migration suggests that concern over declining mobility rates may be misplaced.[6]

In any case, however, the model suggests the improvements to aggregate productivity from enhanced spatial reallocation are not large, despite seemingly big differences in local labor market outcomes. Thus, a policy directed toward population reallocation would likely need to be motivated more by concerns about declining-city residents' welfare than aggregate output.[7]

1.1 Related Literature

Motivated in part by large spatial disparities that flared in the Great Recession, this paper revisits the literature on regional evolutions, including the eponymous article of Blanchard and Katz (1992). Factor allocation over space is at the core of regional economics, and thus there is a long history of literature on labor mobility.[8]

Sjaastad (1962) recognized many decades ago that migration is a costly, forward-looking decision subject to many idiosyncrasies. More recently, this intuition has been implemented into estimable structural models of geographic mobility by Kennan and Walker (2011), Bishop (2008), and Gemici (2011).

The spatial equilibrium model also has a long and rich history, popularized by Roback (1982) and developed more recently by David Albouy and co-authors to study spatial variation in federal tax burdens (Albouy (2009)), refinements of quality of life estimates (Albouy (2008)), population density (Albouy and Stuart (2014)), and housing sector productivity (Albouy and Ehrlich (2012)). Diamond (2013) uses a version of the Roback model to measure changing welfare gaps between skill groups. Moretti (2011) outlines a generalized Roback-style framework for studying how population responsiveness to local shocks varies with housing supply and idiosyncratic location preferences.

With the Roback model as foundation, this paper advances the literature on models of segmented

[5]For more discussion on the stylized facts of gross versus net migration, see Davis et al. (2013) and Schulhofer-Wohl and Kaplan (2012).

[6]For detailed documentation of declining internal mobility rates, see Molloy et al. (2011) and Schulhofer-Wohl and Kaplan (2012).

[7]To the extent that local productivity is endogenous (as in a model with agglomeration externalities), my results on mismatch and aggregate production will be understated. See Moretti (2011) for a discussion of this point.

[8]Examples include Partridge and Rickman (2003); Bound and Holzer (2000); Lall et al. (2009). See also Greenwood (1985) and Greenwood (1997) for surveys.

but interrelated markets.[9] The modeling contribution of this paper is an empirically tractable dynamic equilibrium model of an economy composed of a large number of heterogeneous local markets. Lucas and Prescott (1974) developed the theoretical foundations of dynamic equilibrium models of search over segmented markets, or "island economies." Topel (1986) proposed a dynamic spatial equilibrium model of migration, but did not solve or estimate it. Nieuwerburgh and Weill (2010) developed a spatial equilibrium model which evolved over time as a series of static equilibria. Brinkman et al. (2012) structurally estimate a dynamic equilibrium model of firm location choices with agglomeration externalities in a single metropolitan area. Lee and Wolpin (2006), Kline (2008), and Artuc et al. (2010) estimate dynamic equilibrium models of inter-sectoral mobility.[10]

The models most similar to this paper are dynamic spatial equilibrium models from the macroeconomics literature, including Coen-Pirani (2010), Davis et al. (2013, hereafter DFV), Karahan and Rhee (2011), Nenov (2012), and Yoon (2013). Like here, all of these papers start from the premise that productivity varies by location. My approach is complementary because the model herein emphasizes locational heterogeneity,[11] whereas other papers offer more development in household and sectoral details.[12] In this paper, like the others, productivity levels vary between locations, and markets may receive their own shocks. Additionally in this paper, the distribution of productivity shocks varies across markets, so that, for instance, Los Angeles and Washington DC are subject to different processes.[13] Amenities, housing price levels, and housing price elasticities are unique to each market. Mine is foremost a model of unique cities which then aggregate to a national market, whereas those cited are developed as models of an aggregate economy with varying submarket productivity. Thus, my model is better suited to study how population dynamics vary between locations with different attributes, but offers less insight in other dimensions, such as the effects of different types of search or worker heterogeneity.

To obtain a model with rich location heterogeneity, this paper builds on the insights from the literature on demand estimation for differentiated products, including Berry (1994) and Berry et al. (1995). These techniques were applied to models of location choice by Bayer et al. (2007) in a static model and a dynamic choice environment in Bayer et al. (2011). Diamond (2013) is a nice application of these methods to regional choice, adding heterogeneous preferences as well as endogenous productivity and amenities to an evolving static spatial equilibrium.

My approach to the empirical solution of the dynamic equilibrium was inspired by the literature on the solution and estimation of dynamic games, including Aguirregabiria and Mira (2007), which suggests iteratively using policy improvement steps on players' best response mappings, and Weintraub et al. (2008), which suggests an approximation technique when full solutions are infeasible. The approximation technique has antecedent in the heterogeneous agent model of Krusell and Smith (1998).

[9] Spatial/regional analyses are obvious candidates for this type of model, but other applications could include sectors, occupations, trading partners, or some forms of monopolistic competition and oligopoly.

[10] These estimable dynamic equilibrium models manage a handful of submarkets, whereas this paper contains dozens.

[11] The papers cited above, with less richness in locational heterogeneity, are implemented on comparably large or larger sets of locations (especially DFV). Though no dynamic spatial equilibrium model may be considered "easy" as of yet.

[12] For example, other papers offer more detail with respect to: within-location production by sector (DFV, Yoon (2013), Nenov (2012)), housing tenure and equity (Nenov (2012), Karahan and Rhee (2011)), directed versus undirected search (DFV, Nenov (2012)), match quality (Coen-Pirani (2010)), and worker heterogeneity (Yoon (2013)).

[13] Local shocks are permitted to be correlated in heterogeneous ways with other markets.

Lee and Wolpin (2006) and Yoon (2013) also make empirical applications of dynamic discrete choice rational expectations models, but on smaller choice sets.

The rest of the paper proceeds as follows. Section 2 presents motivating empirical facts which guide the model proposed in section 3. Section 4 reports estimation results, and section 5 results from counterfactual simulations. Section 6 discusses the insights offered for policy. Section 7 concludes. Appendices offer additional exhibits, and detail on the equilibrium specification and numerical routines.

2 Motivation

I begin by discussing some empirical patterns in population dynamics and migration between local labor markets that will motivate and guide the analyses.

2.1 Data Sources

I compiled data on local labor markets in the U.S. for the time period 1990-2011 using publicly available data. Unemployment rates come from Local Area Unemployment Statistics (LAU), reported by the U.S. Bureau of Labor Statistics (BLS). Local market labor force populations come from annual estimates by the Census Bureau, using LAU labor force estimates for within year interpolation. The Quarterly Census of Employment and Wages (QCEW) by the BLS is the source for wages of all workers in all industries, set to 2000 dollars using the BLS's consumer price index.[14]

Migration data came from the Internal Revenue Service (IRS). The IRS reports annual county-to-county migration flows for the period 1990-2010 as calculated by address changes on tax returns. Following IRS documentation, I use the count of returns as proxy for the number of households.

All data were collected at the county level and aggregated to the local market level using constant-geography, year 2000 definitions of metropolitan statistical areas (MSAs). I focus on the 30 largest metro areas in the U.S. in 2000, with the remaining smaller metros and rural areas aggregated to an outside option.[15]

These data are used in the following descriptive exercises and then in estimation of the model that follows.

2.2 Basic Facts About Population Dynamics

Local labor market conditions are much more volatile than population. Panel A of Table 1 reports quarterly standard deviations of within city changes in log population, wages and unemployment rates for the 30 largest metro areas in the U.S. for the period 1990-2011. The figures reported are averaged over the 30 metros. A typical city's population changes far less than its income or unemployment rates.

This is in large part because migration is a relatively uncommon event. Panel B of Table 1 reports the annual migration rates for the 30 largest metros, with migration defined as a change in metro area of residence from one year to the next. A typical large city sends away or receives only four percent of its population each year.

[14] "All items, less housing." Series ID: CUSR0000SA0L2.

[15] The smallest specified city is Indianapolis, IN. The 30 named metros comprise about 46 percent of all U.S. population.

Table 1: Population Volatility Compared to Migration Flows and Labor Market Conditions

Panel A: Average Quarterly Metro-level Standard Deviations

Series	ln(pop)	ln(wage)	Unempl Rate
Linear detrend	0.0109	0.0364	0.0172
HP filter	0.0040	0.0223	0.0080

Panel B: Gross and Net Migration Flows

	(1) Gross outflow (%)	(2) Gross Inflow (%)	(3) |Net flow| (%)
mean	0.0407	0.0422	0.0068
min	0.0251	0.0166	<0.0001
median	0.0390	0.0412	0.0055
max	0.0633	0.0762	0.0172

	(4) $out - net \,\vert\, net < 0$	(5) $in - net \,\vert\, net > 0$	(6) $\Delta pop \vert \Delta pop > 0$	(7) $\Delta pop \vert \Delta pop < 0$
mean	0.0325	0.0434	0.0097	0.0066
min	0.0137	0.0214	<0.0001	<0.0001
median	0.0279	0.0426	0.0068	0.0057
max	0.0713	0.0707	0.0633	0.0518

Not only does a vast majority of workers stay within market, the migration that does occur appears largely idiosyncratic, as the remainder of Panel B shows. In- and out-migration rates are highly correlated spatially and temporally, and gross migration is substantially larger than net migration: the average absolute value of net flow is just 0.6 percent (column 3). The magnitude of the "two-way traffic" necessary to generate net flows is striking. As columns (4) and (5) show, the majority share of out- and in-rates is in excess of net flow.

A restriction of the migration data is that a household needs to be observed in consecutive years, and this may understate population adjustment to the extent it occurs through new household formation. However, columns (6) and (7) show that annual population growth rates are on par with net migration, except at the very tails of the distribution.

2.3 Population Flows and Labor Market Conditions

Thus, local labor market populations are sticky. However, when migration does occur, it is associated with local market conditions. Table 2 reports estimates from a gravity model of migration, a regression of MSA-to-MSA gross population flows as a function of each location's size, distance between locations, and origin and destination attributes:

$$\underbrace{ln(m_{jkt})}_{flows} = \underbrace{ln(pop_{jt}) + ln(pop_{kt}) + Dist_{kj}}_{\text{“}gravity\text{''}} + \underbrace{\upsilon X_{jt}}_{orig.\,attrib.} + \underbrace{\delta X_{kt}}_{dest.\,attrib.} + \varepsilon_{jkt}$$

The signs of the coefficients reported in Table 2 indicate that population flows are repelled by high unemployment rates prices and attracted by higher wages. Specifications (1)-(4) experiment with different forms of a distance function; this functions as a robustness check, and is informative for designing the structural model that follows. Column (1) starts with a dummy for whether the flow is a move ($j \neq k$) and the log miles between city centers, then column (2) adds a measure of workforce dissimilarity–the absolute difference between the MSAs' college attainment shares. Column

(3) uses an index of industrial share dissimilarity, and (4) uses both workforce and industrial share dissimilarities. The findings about attraction and repulsion by market conditions is unchanged, and each form of distance measure reveals that population flows are limited by physical distance as well as differences in labor market composition. Column (5) uses movers only, and the results are very similar. Column (6) adds measures of labor market change in addition to differences in levels, and each are separately significant and of intuitive sign. Finally, columns (7) and (8) address the simultaneity problem by instrumenting for market conditions with lags and Bartik (1991)-style predicted values. The magnitudes change slightly, but the qualitative conclusions remain.

Next, I conduct an exercise following Blanchard and Katz (1992) to trace out the typical response of population following a change to local labor demand. I detail how I define "labor demand" in the next section, but in short, I use one variable each for local incomes and local employment matching efficiency; the latter governs how likely it is to find a job, while the former governs the income offered by a job. I estimate the autoregression

$$p_t = \delta_1 p_{t-1} + \delta_2 w_t + \delta_3 \phi_t \tag{1}$$

where w is the income and ϕ the employment efficiency. The parameter δ_1 captures persistence in population, and δ_2, δ_3 capture elasticity with respect to labor demand innovations. (This exercise uses population stocks, in contrast to the gravity model's population flows.) I used pooled, unweighted quarterly data from the 30 metros, allowing for location fixed effects to capture potential city-specific trends.[16] In Figure 1, I illustrate the impulse response of population following a permanent, one percent improvement to each dimension of labor demand. This illustrates the relative sluggishness with which population responds: after 50 periods (12.5 years), population has not quite adjusted by one percent (and recall that there is an innovation to *both* dimensions of labor demand to generate this). Labor is not completely insensitive, but it is slow to adjust.

Figure 1 is an average across markets and does not distinguish whether certain markets are more responsive than others. Figure 2 splits the sample in two ways, estimating (1) separately by subgroup. In the left panel, the sample of cities is split by an estimate of their housing price elasticity. As argued by Saks (2008) and Saiz (2010), heterogeneity in housing supply elasticity may permit some markets to adjust more in terms of quantities than others. In less elastically supplied markets, housing prices adjust more substantially, choking off population adjustment more quickly. Though Saks (2008) and Saiz (2010) analyze long differences, the effect is borne out in the higher frequency data I use to estimate (1). Figure 2 shows a substantial difference between the rate of population adjustment in elastically supplied markets versus the inelastically supplied. Populations in elastically supplied markets are both initially more responsive and more persistent over time, reaching 1 percent adjustment in under 6 years. Population in inelastically supplied markets are still quite far from 1 percent adjustment even after 12 years.

A second form of heterogeneity is in the persistence of shocks. Some cities may exhibit more cyclical fluctuation than others. If workers are aware of the horizon of labor demand shocks and they face

[16] The requires dynamic panel data methods to be consistently estimated (Arellano and Bond (1991)). Specifically, I use the xtabond procedure in Stata.

Table 2: Gravity Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Move (y/n)	-2.497***	-2.427***	-2.464***	-2.414***		-2.419***	-2.419***	-2.419***
	(0.0344)	(0.0345)	(0.0352)	(0.0352)		(0.0360)	(0.0359)	(0.0359)
Log Miles	-0.814***	-0.812***	-0.811***	-0.810***	-0.827***	-0.810***	-0.810***	-0.810***
	(0.00427)	(0.00425)	(0.00433)	(0.00430)	(0.00388)	(0.00440)	(0.00439)	(0.00439)
College Distance		-1.396***		-1.368***	-1.973***	-1.366***	-1.366***	-1.366***
		(0.0914)		(0.0927)	(0.0860)	(0.0948)	(0.0945)	(0.0946)
Industrial Distance			-0.0717***	-0.0303*	-0.0299*	-0.0326*	-0.0326*	-0.0326*
			(0.0166)	(0.0168)	(0.0159)	(0.0172)	(0.0171)	(0.0171)
Orig Unempl Rate	0.276***	0.276***	0.276***	0.276***	0.271***	0.149***	0.141***	0.134***
	(0.0269)	(0.0267)	(0.0269)	(0.0267)	(0.0243)	(0.0335)	(0.0428)	(0.0389)
Dest Unempl Rate	-0.411***	-0.411***	-0.411***	-0.411***	-0.415***	-0.320***	-0.328***	-0.352***
	(0.0269)	(0.0267)	(0.0269)	(0.0267)	(0.0243)	(0.0335)	(0.0428)	(0.0389)
Orig Wage	-0.230*	-0.230*	-0.230*	-0.230*	-0.208*	-0.0251	0.0223	0.0550
	(0.118)	(0.118)	(0.118)	(0.118)	(0.107)	(0.134)	(0.137)	(0.156)
Dest Wage	0.846***	0.846***	0.846***	0.846***	0.868***	0.666***	0.677***	1.053***
	(0.118)	(0.118)	(0.118)	(0.118)	(0.107)	(0.134)	(0.137)	(0.156)
Orig House Price	0.306***	0.306***	0.306***	0.306***	0.310***	0.198***	0.187***	0.193***
	(0.0386)	(0.0384)	(0.0386)	(0.0384)	(0.0349)	(0.0427)	(0.0460)	(0.0477)
Dest House Price	-0.192***	-0.192***	-0.192***	-0.192***	-0.188***	-0.131***	-0.136***	-0.242***
	(0.0386)	(0.0384)	(0.0386)	(0.0384)	(0.0349)	(0.0427)	(0.0460)	(0.0477)
Orig Wage Growth						-0.454**	-0.483**	-0.506**
						(0.216)	(0.220)	(0.220)
Dest Wage Growth						0.301	0.286	0.109
						(0.216)	(0.220)	(0.220)
Orig Empl Growth						-1.965***	-2.007***	-2.044***
						(0.329)	(0.354)	(0.343)
Dest Empl Growth						1.775***	1.738***	1.580***
						(0.329)	(0.354)	(0.343)
Constant	0.323	0.299	0.328	0.302	-1.928**	-0.104	-0.141	-0.432
	(1.033)	(1.026)	(1.032)	(1.026)	(0.951)	(1.098)	(1.096)	(1.099)
IV?							Bartik	Lags
Observations	18,000	18,000	18,000	18,000	17,400	17,100	17,100	17,100
R-squared	0.947	0.948	0.948	0.948	0.899	0.949	0.949	0.949

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: The dependent variable is the log of monthly city-to-city gross migration flows among the 30 largest metro areas in the U.S., 1991-2009. All specifications contain origin and destination dummies and controls for market size.

Figure 1: Population Impulse Response

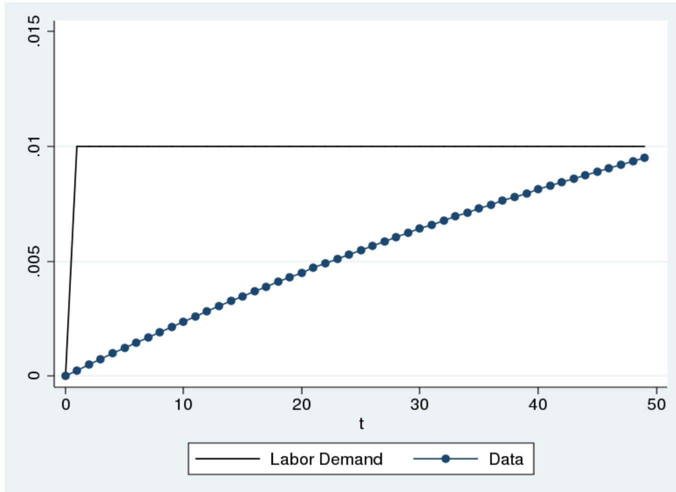

Figure 2: Heterogeneity in Population Impulse Responses

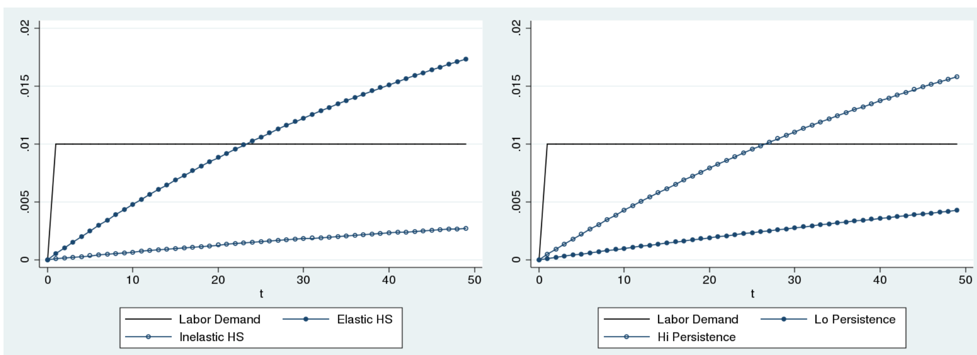

adjustment costs in migrating, they have less incentive to move for short-lived premia.[17] In the right panel, the sample of cities is split by a separate estimate of the persistence of its labor market shocks. (To clarify, each type of city is given the same permanent labor demand shock in the impulse response calculation, but the point estimates come from an underlying data generating processes in which the perceptions of persistence vary.) Again, a striking heterogeneity is evident. Cities with more persistent shocks are nearly twice as elastic. Note that the initial, first period response is larger, meaning that population is more elastic to persistent-looking shocks even in the short run.

2.4 Implications for Modeling and Testing

The forgoing establishes that population adjustment is slow, yet not nonexistent or unimportant. There are potentially many mechanisms in play, and their relative importance may vary by market. The model that follows incorporates several classes of potential frictions to the labor market. First, there will be explicit frictions in the form of mobility costs. Second–and resulting from the first–are dynamic considerations: if agents face mobility costs, they have incentive to forecast the future, and therefore

[17]This was argued by Topel (1986) in combination with life-cycle considerations.

sensitivity to labor demand differentials may vary with the expected duration. Third, there will be non-labor market considerations, such as the cost of non-tradeables such as housing, the availability of amenities, and the idiosyncratic preferences that drive gross migration in excess of net. All of these will be present in a model which preserves market heterogeneity, recognizing that the relative importance of the various mechanisms may depend on the markets' unique features. After fitting the model to data, the counterfactual exercises offer thought experiments that remove the frictions in the model and measure how population adjusts in their absence.

3 Model

I now lay out the model I use to study the impacts of spatial frictions. Several subsections contain additional discussion of how the portion of the model is empirically implemented.

3.1 Setting

The economy consists of a fixed number J of distinct markets, or "cities."[18] Time is discrete and the horizon infinite. A measure of atomistic, identical, infinitely-lived workers is divided amongst the markets, with each worker living and supplying labor to one market per period. Workers pay a moving cost if they relocate, and are imperfectly substitutable across markets.

At the start of the period, all local productivity shocks are realized. Then the worker decides on her location. Every location is available to a worker each period, but the size of the moving costs may depend on the origin-destination pair. After locating, workers are matched to jobs, produce their output, and consume goods. Housing prices in each location adjust to clear the market.

Within a market, workers compete for jobs and housing; that is, all else equal, a worker would prefer the market have lower population.[19]

3.2 Local Labor Market

Within a location, workers inelastically supply one unit of their labor, matching to jobs via a frictional mechanism. I use a one-period matching function:

$$n = m(a, p, v) \tag{2}$$

where the number of jobs n depends via $m(\cdot)$ on the labor supply in the market (p), the number of vacancies posted (v), and match efficiency parameters (a).

All workers are equally productive and produce y if they find match to a job. Workers earn a share δ of their productivity, so that

[18] I use the terms "cities," "locations," "markets," and "metros" interchangeably to refer to metropolitan areas. The empirical counterpart is the metropolitan statistical area (MSA).

[19] The model ignores agglomerative externalities because of its emphasis on short-run labor supply dynamics. Agglomeration would likely be more relevant if considering longer horizons, as in Diamond (2013) or Yoon (2013), or firm location decisions, as in Brinkman et al. (2012).

$$w = \delta y \tag{3}$$

This form means I assume that wages are exogenous, but employment levels and rates vary endogenously, as typical in models of search.

Thus, each market has two forms of "productivity." I incorporate the two dimensions of labor demand to be able to match the multiple forms of heterogeneity exhibited in local labor markets. Shimer (2005) shows the difficulty in using one form of productivity to match both wage and employment data at the national level; the difficulty is equally prevalent at the local labor market level, yet spatial equilibrium models generally ignore employment rates.[20] Furthermore, some markets show strong trends in employment growth, while others in wage growth. The model allows that workers may be choosing locations for different reasons, trading off various forms of productivity and forming different expectations about each type.

I model each local labor market as containing a large number of firms which can freely enter or exit. A "firm" is a single job posting. The firms' profit function is

$$\pi = \frac{n}{v}(y - w) - cn \tag{4}$$

where y is the productivity of a matched worker, w is her wage, and cn is the cost of a vacancy. Note that the vacancy posting cost depends on the level of employment (n) at a rate governed by a parameter c. I use this form for two reasons. First, it seems reasonable that recruiting costs would rise as employment grows, especially in a local market with a fixed (for the period) supply of labor; this is a notion labor market tightness. Second, a constant vacancy cost would cause the employment rate to be independent of population under the zero-profit model of firms elaborated below.[21] That is, without some form of scale-dependent vacancy cost or deceasing returns in the matching function, the model would completely lack any form of labor market congestion, an assumption I am unwilling to make.[22] Therefore congestion in the labor market occurs through the job finding rate.

Beaudry et al. (2012) argues that a search and bargaining framework better describes local labor market conditions than a traditional Walrasian demand/supply framework. Additionally, I choose to treat wages as exogenous but allow unemployment rates to vary endogenously because of the stability of cross sectional variation in wages, and the comparative lack of stability in relative unemployment rates. From theory, Roback (1982) provides ample motivation why wages, by compensating for amenities, would vary cross-sectionally in a consistent way. Table 3 reports R^2 from regressions of pooled quarterly wage and unemployment rates on metro dummies. The table shows that most variation in wages across locations is effectively permanent, whereas a city's ranking by unemployment rate is far more likely to turnover. Therefore, in the model, when workers migrate to follow better labor demand, they will be affecting the location's unemployment risk margin relative to other markets, without affecting its

[20] Exceptions include Kline and Moretti (2013), Yoon (2013).

[21] This would occur because vacancies would increase in proportion to labor supply as the job-filling probability increased.

[22] Parameters are permitted to vary by market, so the congestion may occur at different rates. Thus, the model has no trouble rationalizing markets of different size. For example, New York can be larger than Boston because it has higher values of a or lower values of c, in addition to potential job productivity differences (y).

Table 3: Variance Explained by Permanent Cross Sectional Differences

Data Series	Labor market variable	Variance Explained by Location Fixed Effects (R^2)
Raw	Wage	0.75
	Unemployment Rate	0.16
Relative to Natl	Wage	0.95
	Unemployment Rate	0.49

relative income rank.

Implementation

In practice I use a Cobb-Douglas matching function, $n = m(a,p,v) = ap^{\alpha}v^{1-\alpha}$ for (2). The matching function allows the model to generate unemployment, which can vary substantially between cities and over time in heterogeneous ways.

I leverage an assumption on the local labor market equilibrium to circumvent data limitations and empirically implement the model. In (4), entry by firms–that is, the posting of a new vacancy–causes the probability of a match to fall and the cost of posting a vacancy to rise, lowering expected profits. Firms will enter when profits are positive and exit when negative, ensuring zero expected profits[23] in the local labor market equilibrium. The equilibrium number of vacancies (at $\pi = 0$) is

$$v = \frac{1}{c}\frac{1-\delta}{\delta}w \tag{5}$$

In equilibrium, the vacancies posted depends negatively on the posting cost parameter and positively on the surplus captured by the firm. Combining (5) and (2), the equilibrium number of employed workers in the local market is

$$n = \phi w^{1-\alpha}p^{\alpha} \tag{6}$$

where $\phi = \frac{a}{c}\frac{1-\delta}{\delta}$. Data on population, employment, employment rates, and wages are readily available at the local labor market level; information on vacancies v and "true" labor productivity y is not. Without more information, α, c, a and δ cannot be separately identified. My approach is to take w directly to be one labor demand state, and then use data on n and p to derive the implied values of ϕ for each labor market at each time period, taking ϕ to be an exogenous parameter. I calibrate α to be 0.72, following Shimer (2005).[24] The parameter ϕ subsumes all remaining features determining how well a market allocates available workers to jobs,[25] and therefore I refer to it as the "employment efficiency."

The levels and the transition processes for ϕ and w are allowed to vary and are estimated separately for each market, as described in 4 below. For expositional convenience, I refer to these jointly as "labor

[23] "Expected" refers to the fact that the cost of posting the vacancy is paid with probability one, but the probability of finding a match is $\frac{n}{v}$.

[24] Shimer (2005) calibrated from national data. If assuming match elasticity is the same in all markets, the submarket elasticity aggregates to the same parameter for national elasticity. (Derivation available upon request.) Heterogeneity in local employment efficiency is reflected in other parameters: a, c, δ

[25] Note that ϕ includes the bargaining power of workers, δ, which affects firms' surplus and therefore willingness to post vacancies.

demand."

3.3 Local Housing Market

A worker is constrained to reside in the market in which she supplies her labor. The worker consumes one unit of constant-quality housing in any market. Each location is endowed with a housing supply function which represents the inherent costs of making housing in the location. The costs vary according to the availability of suitable land and the local regulation of its use, which can vary substantially between cities (Saiz (2010)). Absentee landlords supply housing according to this function and rent it to residents. The rental rate is endogenously determined as demand moves up or down the supply function.

$$r_{jt} = \rho_{0j} + \rho_{1j} I_{jt} \tag{7}$$

The index for housing demand is the total income in the market in the period, $I_{jt} = n_{jt} w_{jt}$. Thus, demand will vary with wages, employment efficiency, and the endogenous population. The parameters ρ_{0j} and ρ_{j1} govern the level of housing rental price and its elasticity with respect to demand. These parameters may vary across markets.

3.4 Workers' Utility

Workers receive flow utility from consuming the numeraire good (normalized to a price of 1 in all locations) and from an amenity specific to their location. Workers use all of their income to pay for housing or consumption; there are no savings. All workers pay for housing,[26] but only matched workers draw an income. At the time of location decision, the worker does not know whether she will match to a job but accounts for it probabilistically. The expected utility function is

$$u(er, w, r) = er \cdot u(w, r) + (1 - er) \cdot u(0, r) \tag{8}$$

where er is the employment rate. That is, the effective utility is the employment probability-weighted expectation of consumption after paying for housing, and the (dis)utility of paying for housing only. This is akin to full insurance within a location.

The worker's utility within a location j at time t is

$$U_{ijt} = u(er_{jt}, w_{jt}, r_{jt}, \gamma) + \mu_j + \varepsilon_{ijt} \tag{9}$$

where γ are parameters, μ_j is the flow value of amenities within location j and ε_{int} is a temporary idiosyncratic preference shock for the location. The amenity parameter captures the extent to which local prices capitalize something desirable, such as climate, cultural amenities, local public goods, and

[26] Workers do not receive utility from housing services; rather, housing is the "ticket price" for access to the labor market and amenities. In the discrete choice framework, only relative differences between cities are identified, and hence utility from housing is superfluous to the choice of market, except to the extent that housing quality varies between locations. Average differences in housing quality will be picked up by the location's amenity term.

so on.[27] Empirically, the amenity terms help to match population sizes conditional on local prices, as in Berry (1994) and Berry et al. (1995).

Preference shocks ε are introduced to match the empirical fact that gross migration is much larger than net migration. In essence, the magnitude of ε relative to the other features of utility determines the extent to which population mobility is directed towards markets with unambiguously better features (vertical preferences) versus purely idiosyncratic preferences for a location (horizontal preferences). Like Kennan and Walker (2011) but unlike DFV, all migration is the "undirected" type. As is standard in multinomial choice settings, I assume ε is distributed Type I extreme value.

Implementation

In practice, I use the constant absolute risk aversion function,

$$u(c) = \gamma_1[1 - exp(-\gamma_2 c)]$$

This is a fairly standard utility function, with, as is explained below, the useful property that it is bounded above. The parameter γ_1 governs the scale of utility from common utility sources relative to idiosyncratic location preferences in ε. The parameter γ_2 governs the worker's risk aversion.

3.5 The Dynamic Decision Problem

Each period, the worker decides on her location. All markets are available, but relocating from her previous market incurs a moving cost. The worker's dynamic discrete choice problem is to choose location k to solve:

$$\begin{aligned} V_k(X_t) &= max_j[u(er_{jt}, w_{jt}, r_{jt}) + \mu_j + \varepsilon_{ijt} + mc_{kj} + \beta EV_j(X_{t+1}|X_t)] \\ &= max_j[v_{jk}(X_t) + \varepsilon_{ijt}] \end{aligned} \tag{10}$$

where X collectively refers to all the state variables in the problem, and $E(\cdot|X)$ denotes the expectation formed conditional on current states.[28] v_{jk} represents the choice-specific value function separable from the idiosyncratic preference draws. With Type I extreme idiosyncratic draws, the choice probability conditional on current location is

$$Pr(locate\ in\ j|origin\ k) = \sigma_{kj} = \frac{v_{kj}(X)}{\sum_i v_{ki}(X)} \tag{11}$$

In words, (10) shows the value of being in a location is derived from the utility perceived commonly by all workers, $u_{jt}(\cdot)$ and μ_k, the worker's idiosyncratic temporal preference for the location, ε_{kit}, the moving cost between locations j and k, and the continuation value of being faced with a similar decision problem in k next period. Note that because of the moving cost term, the value will depend on one's

[27] Amenities are taken to be exogenous. Diamond (2013) proposes a model with endogenous production of amenities.
[28] This includes the E_{max} of future idiosyncratic draws. See also Rust (1987),Kennan and Walker (2011).

current location (hence the subscript for k on the left side of the expression).[29]

A typical discrete choice problem includes attributes of the choice and the chooser. The attribute of the chooser in this context is the worker's location at time of choice. Because relocation may occur, the chooser's attribute evolves endogenously.

It is the presence of moving costs that make the problem dynamic. The worker's location is itself a state variable, and there are adjustment costs to change one's state. These adjustment costs fundamentally change the problem from Roback (1982) or its more recent applications. The dynamics are not merely from a series of static equilibria on evolving locations,[30] but from the forward-looking decision making of the agents in the model. The stickiness these impose means that utility may no longer be equalized between markets.

I have made several simplifications to the model in the interest of tractability; namely, job matches last for only one period, and workers may only rent and not own their residences. Consider, however, that as as long as the decision problem is dynamic, these simplifications are not as restrictive as they seem at first. Job matches may be destroyed and recreated each period, but a worker choosing a location takes into account expected future unemployment risk, much like a worker forecasting a job's duration before returning to search. On the housing tenure choice, the worker choosing a location takes into account the expectation of a flow of future rental rates; this is a calculation of net present value, which a contracted home price should reflect.[31]

3.6 Moving Costs

The moving cost function can be specified flexibly. In practice, I parameterize it as function of several notions of distance.

$$mc_{jk} = mc_0 + mc_{oo} + mc_d Log\_mi + mc_c Coll\_dist + mc_i Ind\_dist \tag{12}$$

This is a parsimonious specification that nevertheless accounts for a wide degree of variation in place-to-place mobility rates. When $j = k$, $mc_{jk} = 0$. Each form of distance is strictly positive and symmetric; hence, $mc_{jk} = mc_{kj}$.

The parameter mc_0 is the intercept, the cost paid for a move of any distance. The parameters mc_d, mc_c, mc_i govern the cost with respect to the several distance metrics presented in Table 2:

- mc_d, the physical distance, the log of miles between city centers

- mc_c, the absolute value of the difference between markets in the share of workers who are college educated

- mc_i, an index of the industrial dissimilarity between the two markets.

[29] This means that locations "farther away" in distance have a lower option value, ceteris paribus, because future mobility is more costly.

[30] See Nieuwerburgh and Weill (2010), Diamond (2013).

[31] The abstraction that remains is that the workers do not save or accumulate assets, of which one's home would be a large share. For a treatment of home price risk in a dynamic model, see Hizmo (2010). In principle, asset accumulation including home tenure choice could be added to the model at the cost of a nontrivial expansion of the state space.

The latter two measures are intended as proxies for worker heterogeneity.[32] Workers are ex ante identical, but not equally substitutable between markets; that is, worker heterogeneity enters through the locations themselves. The distance measures ensure that, for instance, workers in Detroit are more exchangeable with Cleveland than San Francisco. As evidenced in Table 2, each of the distance metrics is associated with lower mobility between the pairs of markets.[33]

Finally, mc_{oo} is an intercept shifter if either the k or j location is the consolidated outside option, which is of indeterminate location.

3.7 Equilibrium

Intuitively, the equilibrium definition is that each worker is solving (10) conditional on the location decisions of other workers in the economy. The worker knows the equilibrium rules within in each market–the local labor market matching mechanism and the housing supply function–as well as the transition processes governing the evolution of the labor demand states.

Using capital letters to denote the vector of variables over the cross section of locations, $Y = \{y_1, y_2, ..., y_J\}$, I define a rational expectations equilibrium.

Equilibrium. Given a vector of labor demand states, $\{W\}$ and $\{\Phi\}$, the national population of $\mathcal{P}$, the transition processes for each and the initial allocation of $\mathcal{P}$ over markets J, the rational expectations equilibrium is an allocation of $\mathcal{P}$ over local markets J with associated employment rate vector ER and housing rental price vector R such that:

1. Local labor markets match according to (2), (3), and (4) with $\pi = 0$
2. Local rental prices are determined by (7)
3. Each worker is optimally solving the dynamic discrete choice problem given by (10)

The first two conditions are the definitions of the local market equilibrium. Workers choose locations knowing how the local markets evolve, match, and clear, and, for any given vector of states and expected population in each location, the utility and continuation values are known exactly. What is left to satisfy equilibrium is for each worker's location decision to be optimal *conditional on the location decisions of other workers in the economy.*

The equilibrium concept is a rational expectations equilibrium because workers know the transition process of the state variables and form expectations about the strategies (i.e. the location choice probabilities) of the other workers in the economy. In equilibrium, the workers' beliefs about location choice probabilities are correct. With atomistic workers, equilibrium is achieved at a fixed point in the population distribution over locations

[32] The form of worker heterogeneity present in the model is through the substitutability between locations. In principle, worker types could be added via a suitable definition of local labor markets. For instance, each location could have a local labor market for the college educated and another for non-college educated. This would come at an expansion of the state space ($M \times J$ for M submarkets), but is left as a topic for future research.

[33] To the extent that differences in industries or occupations are regional, the physical distance term soaks up much of the market heterogeneity as well.

$$P\Sigma(W, \Phi, P^*) = P^* \tag{13}$$

where $\Sigma(\cdot)$ represents the aggregated worker's location decisions (with matrix elements given by (11)) as a function of the states and believed choice probabilities P'.

The equilibrium has path dependence because, with specific market-to-market adjustment costs, the current distribution of population over locations will determine migration decisions and hence the next period's distribution. The path dependence is rationalized in one period, however, as the current vector of population is a sufficient statistic for the history of economy's evolution over markets.[34]

There is no proof that an equilibrium of this nature is unique in general. The interactions between workers are congestive, so the problem does not have the obvious multiplicity problem of, for instance, a model with agglomeration. Bayer and Timmins (2005) have developed proofs of uniqueness in similar congestive locational sorting problems in static contexts, but their proofs rely on matrix non-singularity conditions which are not necessarily satisfied in a nonlinear dynamic problem like this. However, an appendix offers a proof of existence and uniqueness for an approximated version of the dynamic equilibrium which I now describe.

Implementation of Dynamic Equilibrium

The problem as specified is impossible to solve analytically or numerically. Local employment and housing rental rates are determined endogenously as a function of the local population, and because moving costs introduce stickiness to future decision processes, the market's population size is an endogenously determined state variable. Because the labor demands of each location determines its current and future attractiveness among a menu of options, and these labor demand states evolve in heterogeneous ways, *every state variable from every location* enters as a state variable to the worker's decision problem in (10). That is, with K states, there are KJ state variables for every location. Even worse, these value functions must all converge jointly, as agents can change locations. For any number of locations finer than very large regions, the full numerical solution becomes utterly infeasible. And yet, I contend, this is unnecessarily complicated, at least for the major questions at hand. To have any hope of empirical implementation, I use an approximated version of the dynamic equilibrium.

The major components of the valuation of a market are its labor demand and current and expected future population. The approximation summarizes information about other markets into aggregates, yielding two variables, "own" and "others" for each type of state. Not only does this greatly reduce the number of states from KJ to $2K$, it has an appealing bounded rationality intuition appealing for this context. The full solution treats all agents as if they interacted strategically with agents of other locations. While the equilibrium model necessitates an accounting for competition from other workers,

[34] This is in contrast to a model with memory over past locations, as employed in a limited way by Kennan and Walker (2011), in which a worker who had visited location k may more easily migrate back to it. This is a form of worker heterogeneity which could in principle be accommodated with a suitable definition of markets, though at significant cost to the state space.

there is little reason in the model to track from where the population inflows are coming, nor is there reason to believe that the actual workers who generated the data behaved according to such detailed decision rules. Hence, the full solution may be artificial in addition to empirically difficult. This technique was inspired by other state-space reduction approaches, including the "Oblivious Equilibrium" concept of Weintraub et al. (2008) and the aggregated-moments approach of Krusell and Smith (1998).

The remaining approximated problem can be solved using standard methods. In practice, the problem is still relatively large ($J = 31$ value functions at $2K = 6$ states each), and I found parametric policy iteration (PPI) to be most expedient.[35] PPI uses a standard policy iteration step, with the value functions approximated as linear-in-parameters polynomial functions of the state space. With X standing in for the listing of the state space, this is

$$v_j(X) \approx \hat{v}_j(X) = \sum_{q}^{Q} \lambda_{j,q} g_q(x_j) \tag{14}$$

There are Q polynomial functions of the states, and approximating parameters λ are found to best fit the implied recursive problem.[36] The problem converges when all cities' approximating functions $v(X)$ have jointly converged.

The approximation algorithm was fast enough to nest inside an estimation routine, shortened significantly by the generation of starting values from estimation of a simpler version of the model.[37] In general with this class of models, other dynamic estimation methods (e.g. Aguirregabiria and Mira (2007), Bajari et al. (2007)) are available and may perform superiorly in some instances. In some applications a model may be calibrated using out-of-sample information. However, with the sheer size of the state space owing to the large number of markets, the state summarization would be necessary to solve the model even once.

An appendix contains additional detail on the procedure for approximating the equilibrium.

3.8 Estimation

Most of the data were described at the outset of section 2. Additionally, I calculated annual house prices using quarterly all-transactions house price indices from the Federal Housing Finance Administration (FHFA) and median home values in the 2000 census,[38] converted to rents by the user cost method (Poterba (1992)).

To match population data, the model needs to allow for population growth, both from natural increase and immigration. I operationalize this by allowing a small flow of new population to arrive without a home location; that is, they choose an initial labor market without paying any moving cost.

[35]For an example of PPI used in dynamic games estimation, see Sweeting (2013).

[36]This is a standard application of projection methods. See also Judd (1998) ch. 11 and Miranda and Fackler (2002), ch. 6.

[37]The equilibrium solution took about two minutes to achieve joint convergence of all policy functions. With 37 parameters at good starting values, the estimation needed to find the equilibrium several thousand and evaluate the objective function several hundred times. Estimation took about a week on a 16x2.9ghz core machine (limited to 12 Matlab workers). Opportunities for parallelization were limited, as the computationally expensive steps were matrix operations; i.e. the approximated problem is not "massively parallel."

[38]Downloaded from NHGIS, www.nhgis.org.

The initially "homeless" group includes both new arrivals from immigration and a retiree-replacement cohort, found using Current Population Survey data.

Estimation proceeds in two stages. In the first stage, I estimate the transition processes for two labor demand states, the wage process and the derived employment efficiency, ϕ, and the housing price elasticity for each city. For each labor demand state, I use an AR(1) specification, which is parsimonious with respect to the state space. Each series is tested for a unit root and trend. While each state in each market has its own transition process, I allow for correlation in the error terms by using the full $2J \times 2J$ covariance matrix when finding expected future states. For the housing price elasticities, I estimate (7) for each city, using lagged demand as an instrument.[39]

With the first stage parameters in hand, I turn to estimation of the primitive parameters in the utility function and moving costs: γ, mc, and the amenity terms, μ. I use Method of Simulated Moments, minimizing a weighted least squares objective function. Using $\theta = \{\gamma_1, \gamma_2, \{mc\}, [\mu_2, \cdots, \mu_J]\}$, the estimator is

$$
\begin{aligned}
e &= \begin{pmatrix} \hat{p}_{jt}(\theta) \\ \vdots \\ \hat{m}_{jkt}(\theta) \\ \vdots \end{pmatrix} - \begin{pmatrix} p_{jt} \\ \vdots \\ m_{jkt} \\ \vdots \end{pmatrix} \\
L &= min_{\Theta} e'We \\
\hat{\theta} &= argmin_{\Theta} L
\end{aligned}
\tag{15}
$$

where Θ is restricted to the real numbers (nonnegative real numbers for CARA utility and move cost parameters), and W is a weight matrix. In practice I use the inverse of the covariance matrix of the population shares p_{jt} and migration flows m_{jkt}. Pesendorfer and Schmidt-Dengler (2008) develop the properties of least squares estimators for dynamic models predicting choice probabilities. The moments to match are the population shares of each city, of which there are $(J-1)T = 30 \cdot 22 \cdot 4 = 2,640$, and the annual migration probabilities, of which there are $J(J-1)\frac{T-8}{4} = 31(30) \cdot 19 = 17,670$.[40] The moments exclude one city as a degrees of freedom correction, since population shares and migration probabilities add to one.

The estimation procedure is a nested fixed point routine, which solves and simulates forward the model for each guess of parameters. Within each iteration of the estimator, I use a Berry (1994) contraction to fit the amenity parameters μ to the average city size over the sample period.

The "amenity" parameters represent unobserved quality of the cities, much like a differentiated products setting. The express utility in the model includes opportunities for wages and employment

[39] Bartik (1991)-style predicted demand IV yields similar results.
[40] The IRS migration data were available for two years less than the population data.

and disutility from cost of housing, similar to Roback (1982) but with unemployment rates, and the dynamic specification includes the valuation of future utility, which captures expectations about future wages, employment, and housing costs, including competition from future population. The amenity is residual to all of these.

Identification of the utility function and moving cost parameters come from the association of within-city population dynamics to local labor demand conditions. Scale parameter γ_1 is governed by how much mobility in the data results from common, market driven utility (including the dynamic valuation) versus idiosyncratic reasons; essentially, this is governed by differences between gross and net migration. Risk aversion parameter γ_2 is governed by the apparent valuation of unemployment risk relative to utility from income less housing costs. The moving cost parameters are identified by (i) the average amount of mobility in the data and the persistence of population, conditional on valuation of the local market, and (ii) the amount of gross population flows between markets as a function of their relative positions in physical and industry space.

Since the migration probabilities are essentially conditional choice probabilities (CCPs), a computationally simpler CCP estimator (Hotz and Miller (1993)) is available to fit the model to annual data. In practice, I use these results as starting values, refining the parameters using the available quarterly moments. This allows me to use the more detailed quarterly frequency, but generates good guesses that significantly shorten time to convergence.

4 Estimation Results

4.1 Parameter Estimates

Table 4 reports estimates of the housing supply function, sorted from least-elastic supply (i.e. most volatile prices) to most-elastic. Unsurprisingly, land constrained coastal markets make up most of the top of the list, along with some Rustbelt markets below their historical average population sizes, and thus possibly on the inelastic portion of their supply curves (Glaeser and Gyourko (2005)). The bottom of the list consists of interior markets with few geographic impediments. Note the magnitude of the differences: a one percent increase in location demand in Dallas generates one-seventh the price response of the same demand shift in Los Angeles. These differences, especially when compounded over many periods, could generate significant heterogeneity in response to local labor demand shocks.

Results for the transition process and associated covariance matrices are presented in Table 9 in the appendix. In brief, I merely point out that cities exhibited heterogeneity not just in the levels of their incomes and employment efficiencies, but also in the persistence and volatility of the shocks.

Table 5 reports the structural parameter estimates from the estimation's second stage. The table includes an abbreviated moving cost matrix depicting the estimated utility cost of a move from location k to j. Since these parameters govern the specification of utility in the model, they can be somewhat difficult to interpret on their own. However, a few conclusions about labor mobility emerge.

First, moving is quite costly on average, evidenced by the size of the moving costs relative to the γ_1 parameter. As Kennan and Walker (2011) emphasize, however, it is important to consider that the utility cost of moving is *on average.* The size of the moving cost dictates that a relatively large

Table 4: Housing Supply Elasticity Estimates

MSA	Elasticity	(se)	Intercept	(se)	Mean implied rent ($/wk)
Los_Angeles_CA	2.34	0.18	-45.72	3.96	421
Detroit_MI	1.79	0.11	-32.55	2.24	203
New_York_NJ_NY	1.54	0.12	-28.88	2.79	445
Miami_Hialea_FL	1.45	0.11	-24.82	2.27	256
Baltimore_MD	1.35	0.09	-22.13	1.84	235
San_Francisc_CA	1.23	0.10	-20.17	2.22	506
St_Louis_MO	1.21	0.08	-19.61	1.54	166
Boston_MA	1.18	0.08	-19.46	1.71	301
Philadelphia_PA	1.15	0.10	-19.20	2.16	217
Cleveland_OH	1.13	0.11	-17.61	2.21	200
Portland_Van_OR	1.09	0.05	-16.36	1.10	274
Minneapolis_MN	1.08	0.07	-17.15	1.44	229
San_Diego_CA	1.08	0.08	-16.08	1.70	406
Chicago_Gary_IL	1.04	0.10	-17.14	2.13	267
Tampa_St_Pe_FL	0.97	0.09	-14.43	1.73	178
San_Jose_CA	0.94	0.09	-13.06	1.84	571
Washington_DC	0.92	0.07	-13.94	1.53	348
Denver_Boulder_CO	0.91	0.03	-13.25	0.72	256
Seattle_Ever_WA	0.82	0.06	-11.39	1.32	331
Kansas_City_MO	0.80	0.04	-11.08	0.87	170
Sacramento_CA	0.78	0.10	-10.04	2.09	306
Riverside_Sa_CA	0.74	0.09	-9.27	1.84	293
Pittsburgh_B_PA	0.70	0.04	-9.17	0.88	147
Phoenix_AZ	0.65	0.06	-8.08	1.28	229
Orlando_FL	0.62	0.08	-7.12	1.50	211
Cincinnati_OH	0.48	0.04	-4.42	0.82	191
Houston_Braz_TX	0.44	0.02	-4.29	0.47	161
Indianapolis_IN	0.43	0.04	-3.45	0.85	182
Atlanta_GA	0.40	0.05	-3.00	0.99	229
Dallas_FW_TX	0.31	0.02	-1.44	0.48	174

idiosyncratic shock is necessary to allow a relocation. This is not necessarily a cost borne directly by a marginal mover, but a quantitative statement of how atypical migration is.[41]

Furthermore, much of this cost comes through simply changing labor markets, even if they are near in the region. Comparing the intercept ($move(y/n)$) parameter to the total moving costs in the city-to-city matrix, one can see that about 70 percent of a move's utility costs comes merely from the changing of labor markets.[42]

Between cities, distance looms important, meaning within-region moves are more common than cross country. However, labor market differences, as measured by industry and education dissimilarity, matter at the margin. For instance, Riverside, California is "farther" from many of the large east coast cities than is Los Angeles, despite L.A. actually being to the west; between L.A. and Riverside, the marginal contribution of distance has decayed, but the Los Angeles labor market composition is more like that of New York and Boston than is Riverside. As another example, Detroit is slightly farther away from other cities on average than is Chicago, despite both being comparably central geographically.

[41] The costs estimated here may differ from Kennan and Walker (2011) for several reasons. First and likely most importantly, this paper's model offers a quarterly opportunity to move, instead of annual; since mobility rates at a quarterly frequency are mechanically lower, the costs will *appear* larger than Kennan and Walker (2011). Also, Kennan and Walker (2011) restrict their sample to high school educated white males, whereas my data encompass the U.S. population. Finally, the equilibrium elements to this model may remove components of moving costs picked up in the partial equilibrium model of Kennan and Walker (2011). Some reluctance to move observed in the data actually may come from competition in the labor or housing markets, a feature modeled explicitly in this paper.

[42] This is even an understatement since the distance in miles between any two markets is strictly greater than zero.

Table 5: Structural Parameter Estimates, With Selected Moving Costs

	Scale parameter	Risk aversion
	γ_1	γ_2
Est	0.257	0.430
Std Error	0.154	0.260

	Moving Costs				
	Intercept		Distance		
	Move (y/n)	To/from O.O.	Physical (log miles)	College Share	Industrial Distance
Est	6.914	0.607	0.467	0.136	0.054
Std Error	1.029	0.896	0.167	4.641	0.187

	Moving Costs								
	Atlanta	Baltimore	Boston	Chicago	Detroit	Los Angeles	New York	Riverside	Washington DC
Atlanta	0.00	10.01	10.23	9.97	10.02	10.56	10.06	10.65	10.02
Baltimore	10.01	0.00	9.75	10.01	9.86	10.65	9.38	10.65	8.71
Boston	10.23	9.75	0.00	10.14	10.02	10.65	9.42	10.75	9.81
Chicago	9.97	10.01	10.14	0.00	9.50	10.41	10.07	10.56	10.11
Detroit	10.02	9.86	10.02	9.50	0.00	10.48	9.93	10.66	9.95
Los Angeles	10.56	10.65	10.65	10.41	10.48	0.00	10.66	9.00	10.73
New York	10.06	9.38	9.42	10.07	9.93	10.66	0.00	10.74	9.50
Riverside	10.65	10.65	10.75	10.56	10.66	9.00	10.74	0.00	10.78
Washington DC	10.02	8.71	9.81	10.11	9.95	10.73	9.50	10.78	0.00

While hypothesis testing is not an objective here, the standard errors indicate the precision with which the data reveal these patterns. The move cost intercept and distance coefficients are very precisely estimated, the utility function parameters slightly less so, and the industry/labor market differences are substantially more noisy.

The amenity parameters μ_j are reported in Table 10 in the appendix. These are identified off of population share, conditional on wages and rents, and thus they are positively correlated with size, negatively with wages and employment rates, and positively with housing price levels. New York City and most California cities rank highly,[43] whereas Cleveland, Indianapolis, and Kansas City rank low.

4.2 Model Fit

The model is simulated at the estimated parameter values to produce a baseline simulation. Figure 3 shows the baseline model's predicted moments compared to the data. By design, the model does well at replicating the distribution of population and the associated unemployment rates (quarterly by city, top panel), and the population growth rates over the sample period 1990-2011(by city, bottom left). The model also does reasonably well at matching the annual out-migration rates for each city (annual by city, bottom right).[44]

Figures 4 and 5 show how well the model produces the reduced-form population responses of section 2. (Note these were not targeted moments.) Figure 4 displays the population impulse response to labor demand shocks generated by the baseline model. The model does very well at replicating the stickiness of population response. Figure 5 decomposes the migratory responses using the gravity model of Table 2. The model matches the basic pattern of response to labor demand shocks of each type. However, the data exhibit an asymmetry to origin and destination labor demand conditions not present in the

[43]One notable exception is San Jose, California, which ranks low, owing to its high income (the highest in the sample) and relatively small size.

[44]Some cities have persistent but spatially idiosyncratic average migration rates which is not captured by the specification of the moving cost function. In principle, one could estimate city-specific moving cost intercepts, although this would come either at a substantial expansion of the parameter set, or from ad hoc designations.

Figure 3: Model Fit

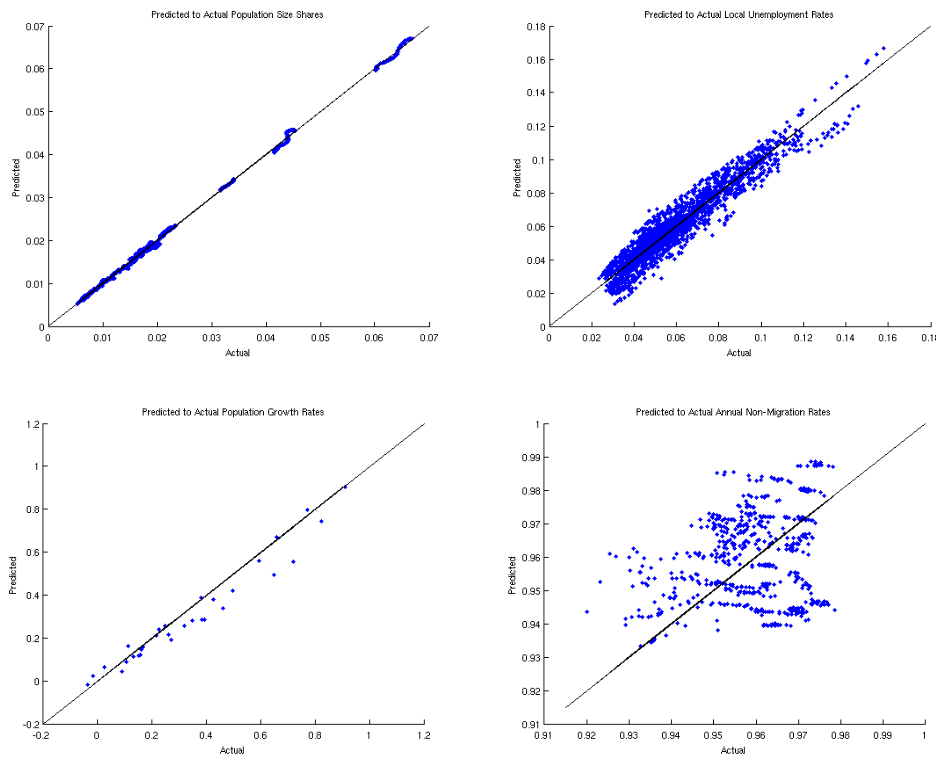

model. Migration is positively associated with both forms of labor demand, but more sensitive to wages than employment efficiency.[45]

The last test of the model is its ability to generate the heterogeneity in population responses between cities shown in section 2. Figure 6 reports the population impulse response for cities of high and low housing supply elasticity (left) and high and low shock persistence (right). The model generates heterogeneity comparable to the data, matching well the differences in shock persistence, and understating somewhat the heterogeneity between markets of varying housing supply elasticity.

5 Counterfactual Simulations

The advantage of a structural equilibrium model is its ability to simulate counterfactual scenarios. I use the model to conduct "thought experiments" in which the various frictions are selectively reduced, detailing how population response to labor demand changes in these alternative economies. In this section, I highlight the results of three specific counterfactual experiments, comparing them to the baseline simulation of section 4.2 and to each other.

One counterfactual reduces moving costs to zero. This returns the model to the static free mobility scenario of Roback-style models, in which the typical equilibrium condition of equal utility across space applies. I refer to this as the "Static" scenario. Because this scenario reduces the problem to a

[45]Recall that wages affect utility in two ways: directly in consumption and indirectly by influencing the employment level (see equation 6). The regression represented by Figure 5 is the reduced form of all effects.

Figure 4: Population Impulse Responses, Data and Baseline Model

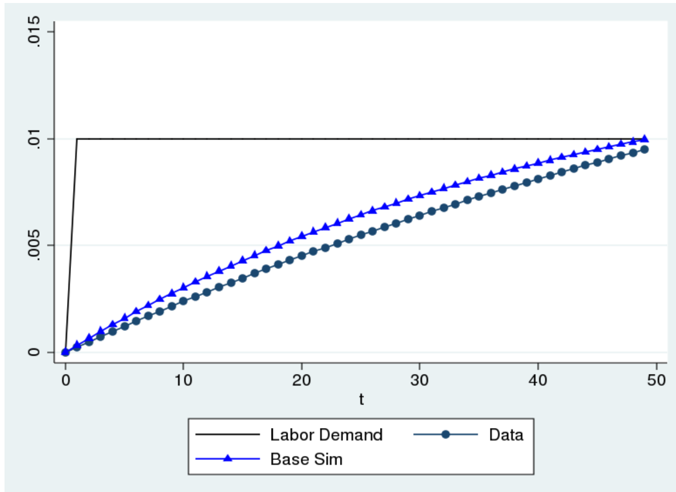

Figure 5: Gravity Model Parameter Estimates, Data and Baseline Model

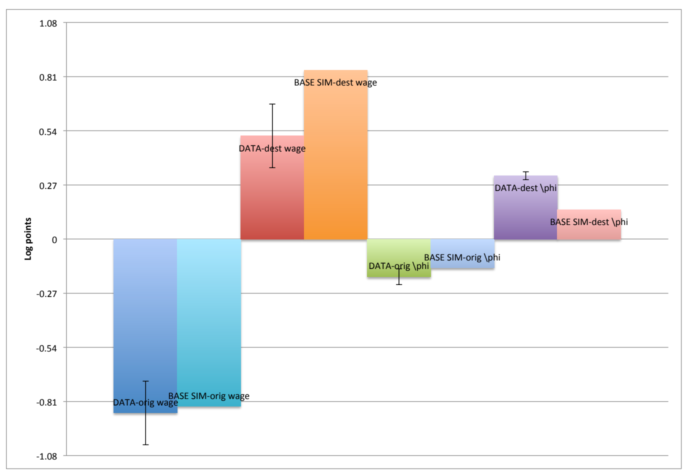

Figure 6: Heterogeneity in Population Impulse Responses, Data and Baseline Model

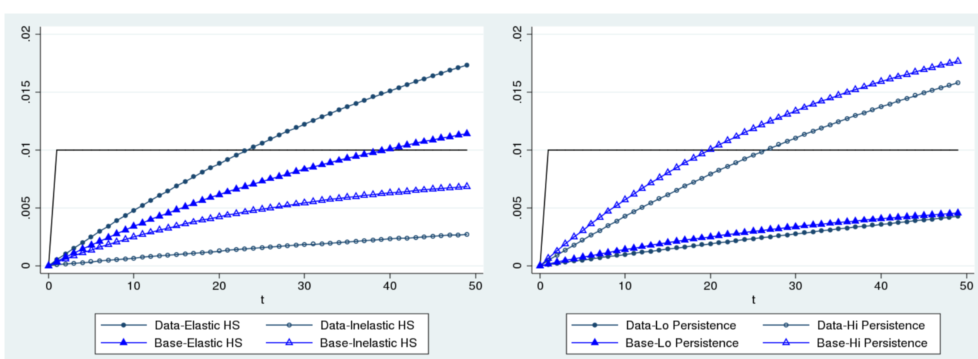

static environment instead of a dynamic state-dependent world, place-to-place mobility is absent and migration statistics do not apply.

A second counterfactual elevates the importance of the common part of utility, down-weighting the idiosyncratic portion of utility. Operationally, this is an increase in the γ_1 parameter, which makes the contribution of ε smaller relative to u. The change in parameter magnifies the utility contribution from wages, rents, and unemployment, but preserves the dynamic nature of the model.[46] The experiment reported uses $\gamma_1^{cf} = 3\gamma_1$ (I discuss the relative magnitude below). I refer to this as the "Less Idiosyncratic" scenario.

The Static and Less Idiosyncratic scenarios are reported in lieu of an experiment with "lower" moving costs.[47] Reducing the moving cost terms without adjustment to any other utility parameters by construction generates more idiosyncratic migration, uncorrelated with any market features. A static, free mobility scenario demonstrates full population adjustment, but "lower" moving costs is simply a poorly fitting dynamic model.

The third counterfactual is not a change of parameters but a change to the local market's equilibrium conditions. In the third scenario, housing demand is derived nationally as if housing were tradable across markets, not solely by local demand. There is still heterogeneity in price changes and levels, as housing supply varies across markets, but this scenario unlinks housing demand from local economic conditions. Housing prices no longer react to local changes in income or employment, but are determined by the interaction of local supply and national demand. I refer to this as the "National Housing Demand" scenario.

All scenarios leave unchanged the local market parameters governing housing supply elasticity and labor demand shock persistence or variance. However, the amenity terms were identified to replicate the size distribution of cities under the baseline prices/utility; alteration of these can distort the average size of a city. The purpose of the counterfactuals is to highlight the effects on population dynamics, not to make any city bigger or smaller on average, and therefore for each simulation I conduct a recalibration of the amenity terms to refit the average sizes under the new scenario. This is merely a rescaling to avoid unintended conflating factors, and does not disrupt the amenity rankings of the cities in the model.

For the sake of exposition, I report the effect of a one percent positive shock to both labor demand variables. The model is symmetric, so a negative shock produces the mirror image effect.[48] As above, I report aggregated reduced-form results for the two important moments generated by the model: the migration flows measured through a gravity model, and population changes measured through impulse responses.

5.1 Relaxing Frictions

Figures 7 and 8 compare, respectively, the changes to migration and the resultant population impulse response induced by the counterfactuals. Figure 7 graphically represents the marginal effects of labor

[46] Magnifying the contribution of consumption or decreasing the variance of idiosyncratic utility are conceptually equivalent.

[47] Results from such an experiment are available upon request.

[48] Thus, the model naturally generates the negative correlation in gross out- and in-migration rates in DFV's Figure 2.

Figure 7: Change in Net In-Migration Induced by 1% Change in Labor Demand

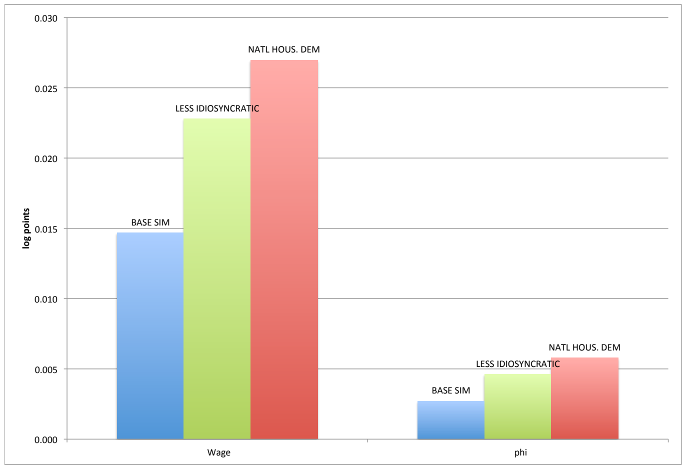

Table 6: Impulse Response Population Elasticities Across Simulations

	Migration Rate	Detrended Pop. Volatility (linear)	Detrended Pop. Volatility (HP filter)	Population Elasticity (%) 0.25 yr	1 yr	5 yrs	10 yrs	$\to \infty$
Baseline	1.02	1.34	0.28	0.03	0.13	0.54	0.89	1.48
Static	NA	3.55	1.95	1.06	1.06	1.06	1.06	1.06
Less Idiosyncratic	1.04	2.10	0.40	0.04	0.17	0.71	1.12	1.70
Nat'l Housing Dem.	1.03	1.89	0.37	0.04	0.15	0.63	1.04	1.75

demand differentials on place-to-place migration flows. (The point estimates that generate these are reported in Table 12 in the appendix.) The figures report that on average, a one percent change in relative wages will increase net migration in the direction of the wage-gaining location by 0.0147 log points in the baseline scenario. In the Less Idiosyncratic scenario, the same change generates a 0.023 log point change in flows, 56 percent larger. Under the National Housing Demand scenario, the relative wage increase generates a 0.027 log point increase, 83 percent larger. For the employment efficiency, the partial effects generate a similar pattern but are smaller in magnitude. Thus, migration is more sensitive to local labor demand conditions when the common portion of utility takes greater weight relative to idiosyncratic preferences and when housing demand is disconnected from these conditions.

How do these changes in place-to-place flows affect aggregated population responses? Figure 8 illustrates the impulse responses to a permanent one percent positive shock to the labor demands under each simulation. Table 6 reports the associated population elasticities at 1 quarter, 1 year, 5 years, 10 years, and in the limit. All three counterfactuals, in reducing frictions, generate greater population response than the baseline simulation. In the static scenario, population instantly and completely adjusts to the positive labor demand shock. The static scenario is, in a sense, the "gold standard" of population adjustment; indeed population *must* adjust to a positive shock, driving up unemployment risk and housing prices until utility is equalized to the other markets.[49]

The Less Idiosyncratic and National Housing Demand scenarios produce more population response than the baseline, but still under the moving cost constraints of the dynamic model. The responses are

[49] In general, it need not be of unit elasticity to accomplish this, as Figure 8, suggests; as will be shown later, this is a coincidence of reporting the average across markets.

Figure 8: Population Impulse Responses, Baseline and Counterfactual Simulations

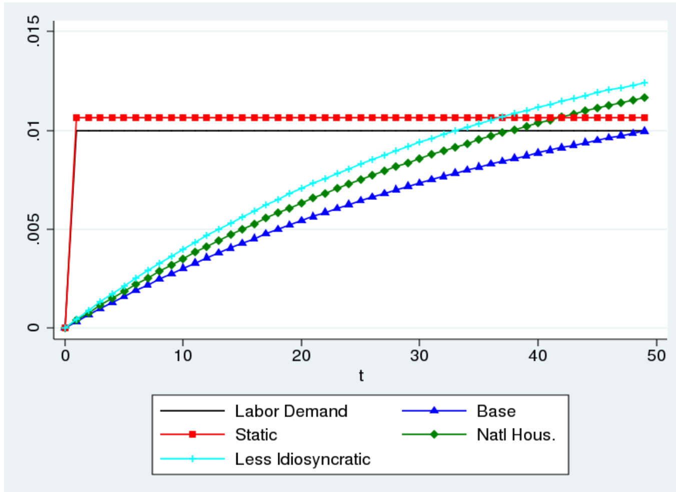

initially larger and, with the persistence of the positive shock, continue to pull away from the baseline impulse response as the effects manifest over time. The Less Idiosyncratic and National Housing Demand scenarios reach new steady states at 1.70 percent and 1.75 percent, respectively, above their initial levels. While the moving costs of the dynamic models make population adjustment slower, the persistence they generate implies a larger long run steady state response than the static scenario.

It is important to note that neither of the two dynamic experiments materially changes the *amount* of mobility, as shown in the first column of Table 6. That is, on average, people move at the same frequency.[50] The differences to the baseline simulation occur through the direction of mobility, so that the differences in Figure 7 are driving the effects exhibited in Figure 8. Under the counterfactuals, mobility is more *targeted*, occurring for market reasons. This non-obvious result is a key outcome of the simulations.

Both the Less Idiosyncratic and National Housing Demand simulations permit more population response to the shock than the baseline simulation, but the differing nature of the two scenarios means it is not instructive to make comparisons of magnitude from these results alone. Unlike the other counterfactual scenarios, Less Idiosyncratic is a matter of degree: *how much* weight is placed on u versus ε. More weight on γ_1 drives the results farther away from the baseline in the directions illustrated above; less weight on γ_1 is nearer the baseline. The choice of $3 \cdot \gamma_1$ is a (somewhat arbitrary) rescaling that produced results of similar magnitude to the National Housing Demand scenario.

If population is more able to adjust to local labor demand shocks, there may be implications for the aggregate output of the economy, to the extent there is "mismatch" in the sense of Shimer (2007). Table 7 reports the counterfactuals' changes relative to the baseline in aggregate employment and income.[51] On average, aggregate employment and income are about one-third of a percentage point higher under the counterfactuals, which corresponds to one-third of a standard deviation in cyclical

[50]Neither dynamic scenario alters the relative value of move/not move. The National Housing Demand scenario affects only how housing prices are calculated, not the moving costs or scale of utility. The Less Idiosyncratic scenario changes how certain factors of utility are weighted, but not the total expected utility relative to move costs. The effects are in where movers are going.

[51]"Aggregate" is defined as the collection of all cities in the sample.

Table 7: Differences in Aggregate Employment and Income

Employment		Difference From Baseline Simulation		
Stat	Units	Static	Less Idiosyncratic	National Housing
Mean	Pct pts.	0.24	0.30	0.30
Max	Pct pts.	2.07	0.86	0.92
Mean	Std Devs.	0.30	0.33	0.36
Max	Std Devs.	2.41	1.00	1.13
Income		Difference From Baseline Simulation		
Stat	Units	Static	Less Idiosyncratic	National Housing
Mean	Pct pts.	0.30	0.30	0.30
Max	Pct pts.	2.22	1.20	0.77
Mean	Std Devs.	0.13	0.09	0.11
Max	Std Devs.	0.96	0.40	0.26

quarterly aggregate employment and about one-tenth a standard deviation of quarterly aggregate income. At their maximum, aggregate employment and income improved by two percentage points in the static counterfactual, and just short of one percentage point in the two dynamic counterfactuals. Thus, there is some improvement in the matching of workers to markets, but the effect is relatively small in comparison to cyclical aggregate fluctuations. The modest degree of spatial mismatch is in line with the findings of Sahin et al. (2011).

5.2 Case Studies

The unique effects of the counterfactuals are illuminated further when examining case studies of particular markets. In this section, I detail the effects on: Detroit, Michigan, a market particularly hard-hit by the Great Recession; Phoenix, Arizona, a market with substantial growth in employment efficiency over the sample period; and San Francisco, California, a market exhibiting substantial income growth.[52] Analyses of specific, unique cities is a luxury of the model's rich locational heterogeneity.

Figure 9 displays several aspects of Detroit's dynamics under the baseline and counterfactual scenarios. The top left chart shows the time series of the metro area population in each simulation, with a plot of the wage and a vertical line for the timing of the precipitous drop in employment efficiency. In the baseline, Detroit's population shows modest growth during the stronger labor demand period of the 1990s before stagnating during the low labor demand of the latter 2000s. In the counterfactuals, population is more responsive to the cycle, with growth stronger in the 1990s and larger declines during the 2000s; the effects are most pronounced in the Static scenario. The top right and two center panel charts show how these population dynamics are generated through changes in migration rates (available for the dynamic simulations only). Migrations rates (reported three types–non-migration or "stayer" rate, out-migration, and in-migration) do show a response to dropping labor demand in the baseline simulation, but the dynamics are more pronounced in the counterfactual scenarios.

The bottom left chart shows the effect on the unemployment rates in the post-2007 period. While the labor demand shocks substantially raise unemployment in all scenarios, the increased population response in the counterfactuals reduces the spike in unemployment by about one percentage point in

[52]Recall that in all scenarios, the location fixed effects are refit so that average population shares are the same as the baseline, but the timing of population arrival/departure may vary.

the two dynamic scenarios and two percentage points in the static scenario. More population moves out, and there is less excess labor supply.

The bottom right chart shows the time series of housing prices in each simulation, focused in the post-2007 period. The differences in the scenarios are notable. In the Static and Less Idiosyncratic scenarios, population has responded more to the labor demand drop, so local housing demand drops even more than in the baseline, generating greater price declines. In contrast, in the National Housing Demand scenario, housing prices fall much less since they are determined by national demand instead of local. In this scenario, causality runs in the other direction: the *lack of drop* in housing prices is actually the mechanism pushing population out of the metro area. Thus Detroit is an illustration of the various forms of mobility frictions that can impede local market population adjustment.

Next I briefly discuss two growing markets, Phoenix and San Francisco. Phoenix exhibited a strong trend in employment efficiency, which generates substantial population growth in all scenarios as shown in Figure 10, but the Static and National Housing Demand scenarios deviate from the baseline simulation in contrasting ways at the end of the sample period. In the righthand panel of Figure 10, housing prices are higher when housing demand is determined locally, since employment in Phoenix is growing faster than the national average; therefore, more people enter the city when the local demand link is broken in the National Housing Demand scenario. In the Static scenario, in contrast, workers are more sensitive to local housing prices and, without forward-looking behavior, less responsive to the strong trend in employment efficiency growth, and population growth tapers off.

San Francisco exhibited strong growth in local incomes, especially since the late 1990s, in a market with high housing prices and a steep supply curve (see Table 4). As shown in Figure 11, population growth is evident in the baseline, but somewhat slow at one percent per year. Under the National Housing Demand scenario, however, population growth is much more responsive to the city's income growth, accelerating in the late 1990s/early 2000s. Notice, however, that housing prices (shown in the chart on the right) are not any lower on average than the other scenarios. How can this be? The answer highlights the importance of expectations in the dynamic model. The workers in the National Housing Demand scenario know that local housing prices in San Francisco will not be disproportionately impacted by the premium local income growth in current or future periods, and are therefore free to reap the benefits of above-average income growth without it being captured by land costs. In contrast, the Static and Less Idiosyncratic scenarios, by elevating the importance of flow utility, are actually more sensitive to the local house prices and therefore less sensitive to the income growth. San Francisco is evidently a market where the state of the housing market has great impact on population dynamics, and even relaxation of other frictions does not make population any more responsive to local labor demand.

5.3 Local Heterogeneity

The foregoing shows that population is made more responsive on average to local labor demand, but cities do not adjust in the same ways. In this section, I discuss how types of cities are differentially affected in the counterfactuals.

Figure 12 displays the changes to net migration induced by one percent changes to wages (left)

Figure 9: Case Study: Detroit Under Baseline and Counterfactual Simulations

NOTES: The vertical lines represent the timing of the precipitous drop in Detroit's employment efficiency variable.

Figure 10: Case Study: Phoenix Under Baseline and Counterfactual Simulations

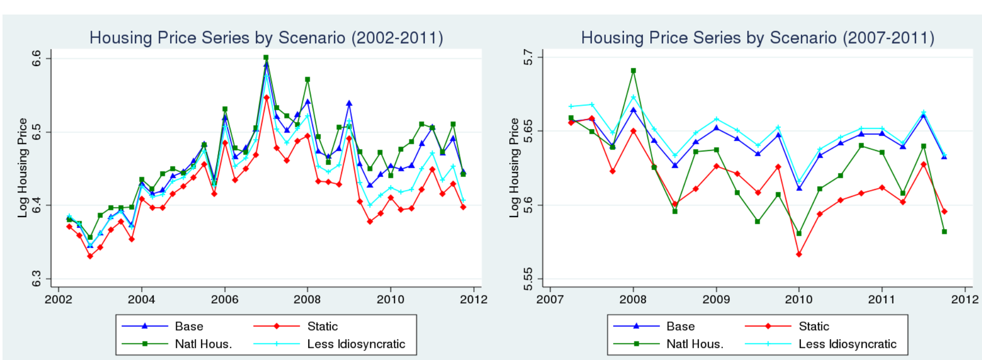

Figure 11: Case Study: San Francisco Under Baseline and Counterfactual Simulations

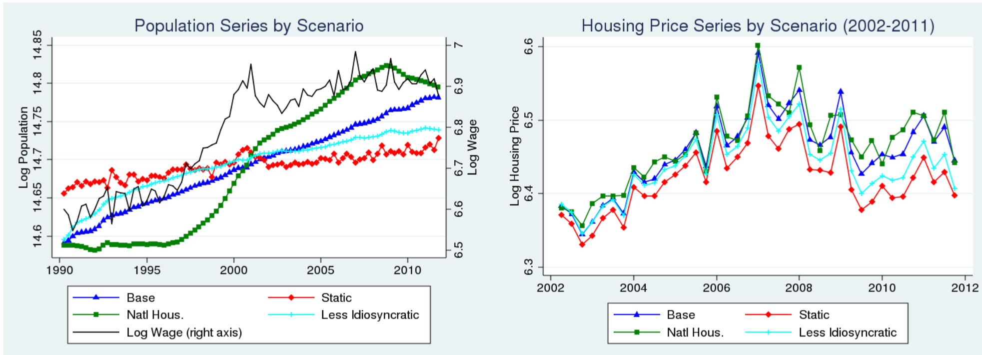

and employment efficiency (right), as measured by the gravity model. (Recall there is no place-to-place migration in the static model.) The top panel (A) distinguishes between cities with elastic and inelastic housing supply, and the bottom panel (B) distinguishes between cities with more and less persistent labor demand shocks. The reported effects are split by the types of cities, and the graph reads additively. The first block of a graphed bar is the effect between the least responsive cities, the next block is the additional effect between one elastic and one inelastic city, and the top block between the most elastic cities. Put another way, the second block is interaction effect if one of the cities was the responsive type, and the third block is the interaction effect when both are the elastic type.[53] For example, in the base simulation, a one percent increase in wage increases net migration by 1.1 percent in favor of the positively shocked city when both cities have inelastic housing supply. The effect is an additional 0.35 (for a total of 1.45 percent) when one of the cities has elastic housing supply, and another 0.35 (a total of 1.8 percent) effect when both cities are elastically supplied.

Panel (A) displays results by housing supply slope. Both the Less Idiosyncratic and the National Housing Demand scenarios exhibit greater migration response in favor of higher wage and higher employment efficiency cities. For each, the total effect for elastically supplied city pairs–the top of the bars–is higher than in the baseline scenario. However, the effect for National Housing Demand is especially concentrated between inelastically supplied markets: notice that the change induced by this scenario is greatest in the first block of the bar, whereas the effect for Less Idiosyncratic is proportional in all types of cities. That is, the effect of the housing demand channel is greatest in housing-constrained markets, which is intuitive manifestation of supply elasticity. (Recall that the experiment merely unlinks housing demand from local conditions, and does not change the local supply function.) Thus, there is heterogeneity in response to the type of friction relaxed.

The lower panel (B) distinguishes cities by the relative persistence of their shocks. As with panel A, the effects under the Less Idiosyncratic scenario are spread out among the different types of cities, but the effects under the National Housing Demand are concentrated among cities with persistent shocks, where a local demand shock would affect housing costs for a longer period.

Figure 13 plots the impulse responses for each counterfactual simulation, by each form of city

[53] The effect of a labor demand change to net migration could be written as: $\Delta net\_migration_{i,k} = \beta_0 + \beta_1 I[i = elastic\_type] + \beta_2 [k = elastic\_type], \;\; \beta_0, \beta_1, \beta_2 > 0.$

heterogeneity, and Table 8 lists the associated population elasticities. The impulse responses from the baseline simulation are included for reference. Starting with the Static scenario in panel A, Figure 13 shows (on the left) that inelastic housing supply chokes off population response more severely. With inelastic supply, housing prices more quickly congest the positively shocked market and have it reach utility equalization with less population reallocation. The population response is twice as large, on average, for elastically supplied markets to reach utility equalization. On the right, there is little distinction between high and low shock persistence markets, since in the static model horizon length is irrelevant.

As shown in Panel B, the Less Idiosyncratic scenario induces more population response in proportional ways for each type of city.[54] These results are an interesting contrast to Panel C, which shows the impulse responses from the National Housing Demand scenario. The inelastically supplied housing markets are much more responsive than in the baseline, especially in short horizons. There is effectively no change to elastically supplied markets. The population responses from the elastic and inelastic cities are driven closer together, as if the type of housing supply little matters.[55] Moreover, high persistence markets are markedly more responsive, while low persistence markets are scarcely affected.[56] Thus, the examination by city type reveals that the Less Idiosyncratic scenario replicates underlying heterogeneity, but the National Housing Demand scenario's effects are pronounced in the most housing-constrained markets.

The conclusion is that the degree of idiosyncratic preference affects population elasticities across the board, whereas housing frictions are most important in the most constrained markets. Unlinking housing demand from local labor conditions unleashes population response in inelastically supplied markets and in those with persistent shocks.

Recall the under both dynamic counterfactuals, mobility is more *targeted* towards positively shocked markets, but that mobility does not rise on average. The difference between the two dynamic counterfactuals is the spatial heterogeneity present in one: for the National Housing Demand scenario, the targeting effect is especially pronounced in inelastic housing markets and in those with longer shock persistence.

Before moving to a discussion of implications, it is interesting to compare these results to those of Davis et al. (2013, DFV), a paper with similar motivations but a very different approach. My results and theirs broadly agree on the importance of idiosyncratic migration as a factor inhibiting population adjustment; in the language of DFV, "guided trips" produce the labor-market targeted migration that the Less Idiosyncratic scenario generates. However, I find an important role for housing, whereas DFV do not. There are several possible reasons for this disagreement. First, I point out that in my results the importance of housing varied greatly across markets, and DFV do not allow for locational

[54]Also, the correlation of city-by-city labor demand elasticities and housing supply and persistence is little changed from the baseline.

[55]In city-by-city estimates of impulse response (Table 13 in the appendix), the correlation between housing supply elasticity and wage drops from -0.49 in the baseline to -0.35; the correlation between housing supply elasticity and employment efficiency elasticity moves from -0.24 to 0.09.

[56]In city-by-city estimates of impulse response (Table 13), the correlation between a city's wage persistence and its population elasticity grows from 0.69 in the baseline to 0.85; the for employment efficiency, the correlation between persistence and its population elasticity moves slightly from 0.49 to 0.53.

Figure 12: Changes in Net In-Migration in Counterfactuals, by City Type

Panel A: By Housing Supply

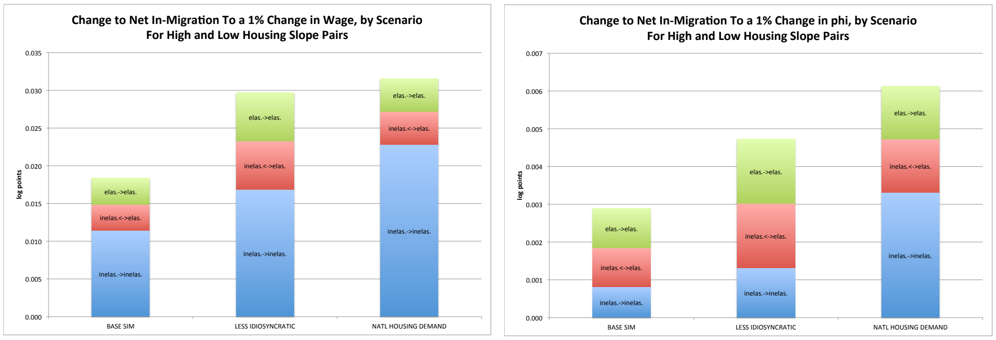

Panel B: By Shock Persistence

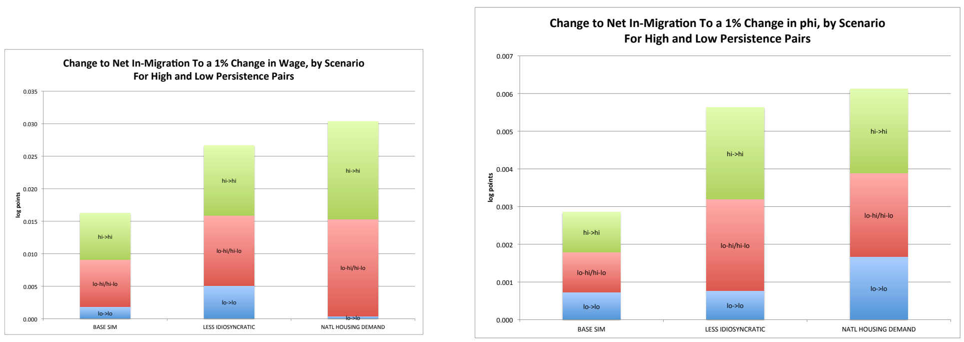

heterogeneity in housing supply. In the large cities in my sample, housing supply is likely more limiting than in the broader sample of small metros that DFV include. Second, the housing mechanisms of the two papers are different. My model contains a house price effect whereby local shocks and population inflows are capitalized, and housing per person is inelastic. The congestion effect of DFV is through declining housing per person, without a price capitalization effect on utility. The price effect is clearly important, and seemed sensible for measuring relatively short run fluctuations in housing demand. But in light of DFV's results, future work should consider relaxing the inelasticity in housing services consumed.

6 Implications for Policy

The simulation results led to three main conclusions. First, a dynamic model more accurately describes population movements across local labor markets. Second, the reallocation of labor supply over markets is hampered by at least two mechanisms: the large degree of idiosyncratic preferences for locations, and the capitalization of local labor demand shocks into housing costs. These prevent labor market-directed mobility. Migration can be made more targeted without changing in volume. Third, heterogeneous local labor markets adjust at different rates; in particular, markets with elastic housing adjust more quickly and to a greater degree, as do markets with perceived persistent demand shocks.

Figure 13: Heterogeneity in Population Impulse Responses in Counterfactuals, by City Type

Panel A: Static

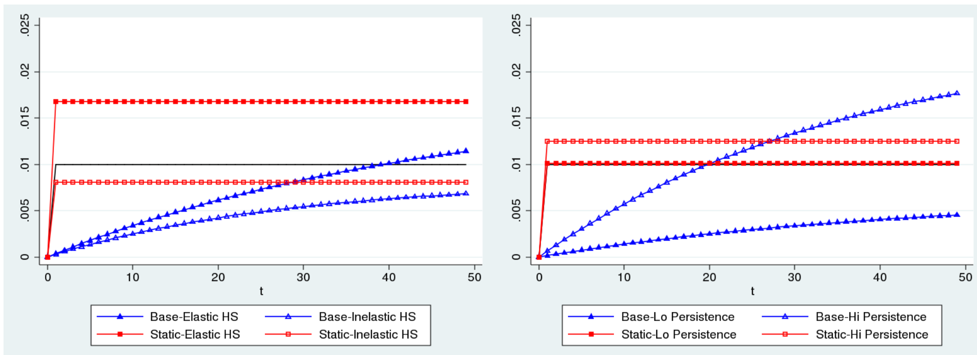

Panel B: Less Idiosyncratic

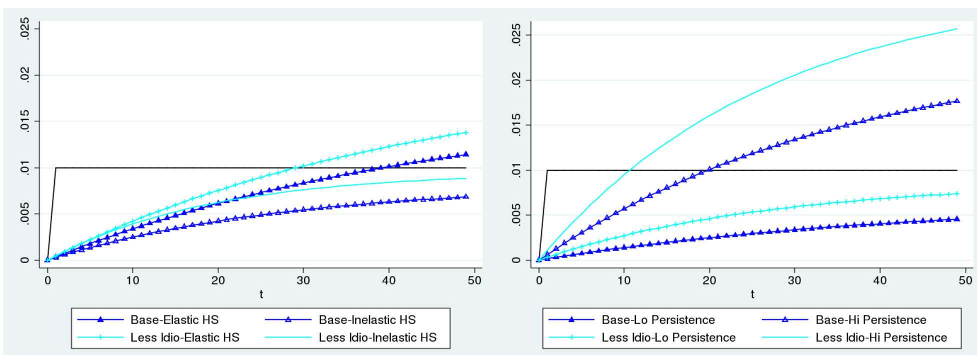

Panel C: National Housing Demand

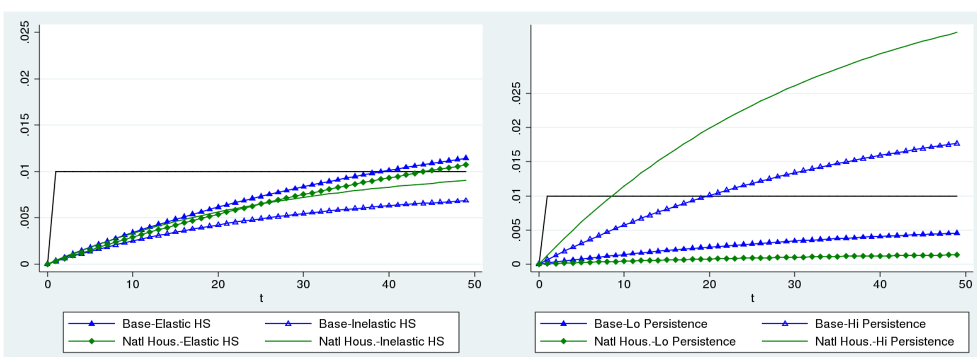

Table 8: Impulse Response Population Elasticities Across Simulations, by City Type

		Population Elasticity				
	City Types	0.25 years	1 year	5 years	10 years	$\rightarrow \infty$
Baseline	Elastic Housing	0.04	0.14	0.61	1.01	1.75
	Inelastic Housing	0.03	0.11	0.42	0.63	0.82
	Less Persistent Shocks	0.02	0.06	0.25	0.41	0.66
	Persistent Shocks	0.06	0.25	1.01	1.59	2.40
Static	Elastic Housing	1.68	1.68	1.68	1.68	1.68
	Inelastic Housing	0.81	0.81	0.81	0.81	0.81
	Less Persistent Shocks	1.01	1.01	1.01	1.01	1.01
	Persistent Shocks	1.25	1.25	1.25	1.25	1.25
Less Idiosyncratic	Elastic Housing	0.05	0.18	0.75	1.23	2.04
	Inelastic Housing	0.05	0.18	0.62	0.84	0.96
	Less Persistent Shocks	0.03	0.12	0.46	0.68	0.88
	Persistent Shocks	0.11	0.42	1.60	2.37	3.07
National Housing Dem.	Elastic Housing	0.03	0.12	0.54	0.93	1.98
	Inelastic Housing	0.04	0.15	0.56	0.83	1.08
	Less Persistent Shocks	0.00	0.02	0.08	0.12	0.19
	Persistent Shocks	0.13	0.50	1.99	3.08	4.41

The capitalization of local labor demand shocks into housing prices is an especially salient friction in locations with persistent shocks.

Of course, these thought experiments are tests of mechanisms, not particular policies. What are the implications for policy, if any?

6.1 The Importance of Idiosyncratic Preference

One implication is that simply incentivizing mobility need not lead to better allocation of workers to markets. The importance of idiosyncratic preferences shows that is not the total quantity of mobility per se, but the degree to which it is directed by common factors.[57] The empirical exercises above show that horizontal preferences largely trump vertical preferences, and gross migration dwarfs net migration. Magnifying the import of the vertical elements leads to more worker reallocation.

What are "idiosyncratic preferences?" In the model, they act a stand in for household heterogeneity of many types.[58] If the result stems from heterogeneity in returns to mobility, then there may be no inefficiency to correct. For example, migration idiosyncratically motivated by family is not an issue not regional policy. Or, if the extent of gross migration is due to individual heterogeneity in preferences (say, beach-lovers born in Denver and snow-skiers born in Miami), then gross migration is the natural process of sorting. Moreover, even if heterogeneity in the impact of local shocks on different types of workers causes out-migration from good markets and in-migration to bad, the question becomes one of inequality in the labor market, not regional policy.

However, if such "two-way traffic" is due to a lack of information, then there is a standard form of inefficiency, and some policy mechanism to correct misinformation would be in order. If workers in Detroit are simply not aware of better opportunities elsewhere, this would appear as an outsized idiosyncratic preference for Detroit. If workers perceive a city's labor demand to be higher than it

[57] This suggests that concern over declining mobility (as described in Molloy et al. (2011)) may be misplaced.

[58] Permanent heterogeneity in markets' industrial composition is picked up in the moving cost function.

actually is–because of, say, adaptive expectations or even irrational exuberance for a growing market–then they would be in-migrants with an outsized idiosyncratic preference for the location.

The good news is that, the more appropriate policy is, the more readily it is available. That is, if misinformation were even part of the cause of idiosyncratic preferences, then policies fostering the collection and dissemination of accurate information would enhance workers' ability to reallocate. For example, the government could offer job search assistance (but not generic relocation assistance, which would merely increase total mobility). Accurate forecasts of local labor demand, such as those provided in the U.S. by federal agencies or regional Federal Reserve banks would enhance targeted migration. This may include the unpalatable task of projecting the decline of a city, something governments are not eager to do.

The bad news is that it is not clear that idiosyncratic migration is actually due to misinformation. The result that markets with persistent shocks are more responsive, even in the short run, is prima facie evidence that misinformation is unlikely to be the major cause of idiosyncratic migration, and heterogeneity in returns or preferences seems plausible. However, a complete analysis of the causes of idiosyncratic migration is beyond the scope of the current article. But since the appropriateness of policy depends on a decomposition of the causes, further research is warranted.

6.2 The Link Between Local Labor Demand and Prices of Nontradables

There is little reason to expect the capitalization of local labor demand shocks into housing costs is welfare or productivity enhancing. While positive shocks may benefit current landowners, this comes at the expense of future migrants (or more perhaps accurately, those never-realized), as noted by Moretti (2011).[59] Even so, it is not obvious that landowners are better off in general, as they are subject to capital risk.

This is not an argument for or against homeownership, as the model in this paper did not use tenure choice to generate substantial differences between markets. Neither is it a naive bemoaning a high cost of living, as high land costs can result from high amenities. Rather, the results of this paper illuminate the joint dynamics of labor demand and housing cost. The link between labor demand and nontradable consumption dampens labor supply responses to labor demand shocks both positive (Saks (2008)) and negative (Glaeser and Gyourko (2005)). In this way, the simulations contain implications for policy.

Housing supply restrictions should be weighed carefully. Some restrictions may be welfare enhancing, but the effect on supply elasticity can indirectly have consequences for local labor market adjustment. Moreover, local restrictions could be a case of NIMBY (not in my backyard), in which the benefits are private, but the costs public.

The counterfactual experiment in this paper was to unlink local housing costs from local demand without changing supply, and yet the effects were large and predictably heterogeneous across markets of different supply elasticity. Therefore, a second implication is that a policy or financial instrument

[59] It is interesting that while a dynamic model is more appropriate for studying labor adjustment, the basic mechanism of a static model, as in Moretti (2011) looms large. Indeed, the prospect of a change to housing payments over a long horizon only exacerbates the effect.

which diversifies a household exposure to housing prices could affect regional dynamics. Examples may include alternative mortgage products (Kung (2012)), down payment assistance, insurance against capital loss (Shiller (2009)), holdings of other locations in a household's portfolio (Hizmo (2010)), and a readiness by governments local or national to remove unwanted housing stock (Glaeser and Gyourko (2005)).

7 Conclusion

This paper has argued that a dynamic spatial equilibrium model is most appropriate for modeling labor mobility since migration is costly for workers and locations are heterogeneous. The paper proposed, estimated, and simulated such a model using quarterly data on large U.S. metropolitan areas. Simulations conducted thought experiments designed to test several mechanisms that impede labor supply adjustment across locations.

Clearly a static, free mobility world would have the most labor supply adjustment over markets. But within the context of a realistic model featuring adjustment costs, more mobility does not necessarily mean more labor supply response. The large degree of idiosyncratic migration and the capitalization of local labor demand shocks into housing prices generate frictions to the targeting of migration, even if they do not affect aggregate volume of mobility.

The simulations also showed how different types of markets responded to changes in the economy. Locational heterogeneity did more than add realism to the model; it facilitated tests of the mechanisms. For example, differences in shock persistence are likely generalizable to within-market variation as well as between-market.

On the methodological side, the model offers a strategy for empirically implementing dynamic equilibrium models of segmented heterogeneous markets, which could open new opportunities for dynamic spatial analysis. Segmented markets are common, and this paper offers a blending of micro- and macro-economists' toolkits which may be useful in other applications, possibly including occupational and industry mobility, trade, and oligopolistic or monopolistic competition.

However, the model's complexity meant it was still limited in its reach. Future work could relax assumptions about worker homogeneity, permit durable capital and housing, and allow for directed search, to name a few. The difficulty in this paper was in large part due to the need to estimate many parameters (especially the scale parameter and location amenities) within context of the model. Contexts in which full estimation can be avoided may permit even greater richness.

References

Victor Aguirregabiria and Pedro Mira. Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53, 2007.

David Albouy. Are big cities bad places to live? estimating quality of life across metropolitan areas. Technical report, National Bureau of Economic Research, 2008.

David Albouy. The unequal burden of federal taxation. *Journal of Political Economy*, 117(4):635–667, 2009.

David Albouy and Gabriel Ehrlich. Metropolitan land values and housing productivity. University of Michigan, 2012.

David Albouy and Bryan Stuart. Urban population and amenities. Technical report, National Bureau of Economic Research, 2014.

Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297, 1991.

Erhan Artuc, Shubham Chaudhuri, and John McLaren. Trade shocks and labor adjustment: A structural empirical approach. *American Economic Review*, 100(3), 2010.

Patrick Bajari, C. Lanier Benkard, and Jonathon Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007.

Timothy J Bartik. Who benefits from state and local economic development policies? WE Upjohn Institute for Employment Research, 1991.

Patrick Bayer and Christopher Timmins. On the equilibrium properties of locational sorting models. *Journal of Urban Economics*, 57(3):462–477, 2005.

Patrick Bayer, Fernando Ferreira, and Robert McMillan. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638, 2007.

Patrick Bayer, Robert McMillan, Alvin Murphy, and Christopher Timmins. A dynamic model of demand for houses and neighborhoods. Working Paper 17250, National Bureau of Economic Research, July 2011. URL `http://www.nber.org/papers/w17250`.

Paul Beaudry, David A. Green, and Benjamin M. Sand. Does industrial composition matter for wages? an empirical evaluation based on search and bargaining theory. *Econometrica*, 80(3):1063–1104, 2012.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, pages 841–890, 1995.

Steven T. Berry. Estimating discrete choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262, 1994.

Kelly C. Bishop. A dynamic model of location choice and hedonic valuation. Olin School of Business, Washington University in Saint Louis, 2008.

Oliver Jean Blanchard and Lawrence F. Katz. Regional evolutions. *Brookings Papers on Economic Activity*, 1:1–37, 1992.

John Bound and Harry J Holzer. Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, 18(1):20–54, 2000.

Jeffrey C. Brinkman, Daniele Coen-Pirani, and Holger Sieg. Estimating a dynamic equilibrium model of firm location choices in an urban economy. Federal Reserve Bank of Philadelphia, Working Paper 12-26 2012.

Daniele Coen-Pirani. Understanding gross workers flows across u.s. states. *Journal of Monetary Economics*, 57:769–784, 2010.

Morris A. Davis, Jonas D.M. Fisher, and Marcelo Veracierto. Gross migration, housing and urban population dynamics. Federal Reserve Bank of Chicago, 2013.

Rebecca Diamond. The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. mimeo, Stanford University, December 2013.

Ahu Gemici. Family migration and labor market outcomes. *Manuscript, New York University*, 2011.

Edward L. Glaeser and Joseph Gyourko. Urban decline and durable housing. *Journal of Political Economy*, 113(2):345–365, 2005.

Michael J. Greenwood. Chapter 12 internal migration in developed countries. In Mark R. Rosenzweig and Oded Stark, editors, *Handbook of Population and Family Economics*, volume 1, Part B of *Handbook of Population and Family Economics*, pages 647 – 720. Elsevier, 1997.

Michal J. Greenwood. Human migration: Theory, models, and empirical studies. *Journal of Regional Science*, 25(4):521–544, 1985.

Aurel Hizmo. Risk in housing markets: An equilibrium approach, 2010.

V Joseph Hotz and Robert A Miller. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529, 1993.

Kenneth Judd. *Numerical Methods in Economics*. MIT Press, Cambridge, MA, 1998.

Fatih Karahan and Serena Rhee. Housing and the labor market: The role of migration on aggregate unemployment. University of Pennsylvania, 2011.

John Kennan and John R. Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79:211–251, 2011.

Patrick Kline. Understanding sectoral labor market dynamics: An equilibrium analysis of the oil and gas field services industry. mimeo., University of California-Berkeley, 2008.

Patrick Kline and Enrico Moretti. Place based policies with unemployment. *American Economic Review Papers and Proceedings*, 103(3):238–43, May 2013.

Per Krusell and Anthony A. Smith. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998.

Edward Kung. Continuous workout mortgages in a structural model of housing and mortgage markets. mimeo, UCLA, 2012.

Somik V Lall, Christopher Timmins, and Shouyue Yu. Connecting lagging and leading regions: The role of labor mobility. *Brookings-Wharton Papers on Urban Affairs*, pages 151–174, 2009.

Donghoon Lee and Kenneth I. Wolpin. Intersectoral labor mobility and the growth of the service sector. *Econometrica*, 74:1–46, 2006.

Robert E. Lucas and Edward C. Prescott. Equilibrium search and unemployment. *Journal of Economic Theory*, 7:188–209, 1974.

Mario J. Miranda and Paul L. Fackler. *Applied Computational Economics and Finance.* MIT Press, Cambridge, MA, 2002.

Ian Molho. Migrant inertia, accessibility, and local unemployment. *Economica*, 62:123–132, 1995.

Raven S. Molloy, Christopher L. Smith, and Abigail Wozniak. Internal migration in the united states. *Journal of Economic Perspectives*, 25(3):173–196, 2011.

Enrico Moretti. Local labor markets. *Handbook of labor economics*, 4:1237–1313, 2011.

Plamen Nenov. Labor market and regional reallocation effects of housing busts. Norwegian Business School, 2012.

Stign Van Nieuwerburgh and Pierre-Olivier Weill. Why has house price dispersion gone up? *Review of Economic Studies*, 77(4):1567–1606, 2010.

Mark D Partridge and Dan S Rickman. The waxing and waning of regional economies: The chicken–egg question of jobs versus people. *Journal of Urban Economics*, 53(1):76–97, 2003.

Martin Pesendorfer and Philip Schmidt-Dengler. Asymptotic least squares estimators for dynamic games. *Review of Economic Studies*, 75:901–928, 2008.

James Poterba. Taxation and housing: Old questions, new answers. *AEA Papers and Proceedings*, 82 (2):237–242, 1992.

Jennifer Roback. Wages, rents, and quality of life. *Journal of Political Economy*, 90(6):1257–1278, 1982.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, pages 999–1033, 1987.

Aysegul Sahin, Joseph Song, Giorgio Topa, and Giovanni L. Violante. Measuring mismatch in the u.s. labor market. Federal Reserve Bank of New York, 2011.

Albert Saiz. The geographic determinants of housing supply. *Quarterly Journal of Economics*, 125(3): 1253–1296, 2010.

Raven E. Saks. Job creation and housing construction: Constraints on metropolitan area employment growth. *Journal of Urban Economics*, 64:178–195, 2008.

Sam Schulhofer-Wohl and Greg Kaplan. Understanding the long-run decline in interstate migration. NBER Working paper No. 18507, 2012.

Robert J Shiller. *The new financial order: Risk in the 21st century*. Princeton University Press, 2009.

Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49, 2005.

Robert Shimer. Mismatch. *American Economic Review*, 97(4):1074–1101, 2007.

Larry A. Sjaastad. The costs and returns of human migration. *Journal of Political Economy*, 70(5): 80–93, 1962.

Nancy L Stokey, Robert E Lucas Jr, and Edward C. Prescott. *Recursive methods in economic dynamics*. Harvard University Press, 1989.

Andrew Sweeting. Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica*, 81(5):1763–1803, 2013.

Robert Topel. Local labor markets. *Journal of Political Economy*, 94(3):111–143, 1986.

Gabirel Y. Weintraub, C. Lanier Benkard, and Benjamin VanRoy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.

Chamna Yoon. The decline of the rust belt: a dynamic spatial equilibrium analysis. mimeo., University of Pennsylvania, 2013.

A Additional Tables

Table 9: Labor Demand Persistence and Shock Covariances

MSA	Persistence		Labor Demand Shock Covariances (in 10^{-3})				
	δ_w	δ_ϕ	σ_w	σ_ϕ	$\sigma_{w,\phi}$	$mean(\sigma_{w,jk})$	$mean(\sigma_{\phi,jk})$
Outside_option	0.584	1.000	0.404	0.044	-0.099	0.445	0.445
Atlanta_GA	1.000	1.000	0.733	0.076	-0.185	0.504	0.492
Baltimore_MD	0.573	1.000	0.517	0.059	-0.127	0.448	0.445
Boston_MA	1.000	1.000	1.229	0.100	-0.311	0.671	0.603
Chicago_Gary_IL	0.661	1.000	0.850	0.081	-0.200	0.639	0.569
Cincinnati_OH	0.724	1.000	0.688	0.069	-0.175	0.550	0.528
Cleveland_OH	0.575	1.000	0.574	0.053	-0.136	0.504	0.463
Dallas_FW_TX	1.000	1.000	0.713	0.060	-0.170	0.560	0.516
Denver_Boulder_CO	1.000	1.000	0.757	0.067	-0.173	0.499	0.440
Detroit_MI	0.712	1.000	1.205	0.157	-0.353	0.509	0.554
Houston_Braz_TX	0.351	0.654	0.874	0.082	-0.180	0.552	0.478
Indianapolis_IN	0.567	0.686	1.264	0.119	-0.295	0.674	0.571
Kansas_City_MO	0.636	0.663	0.869	0.078	-0.200	0.507	0.426
Los_Angeles_CA	0.555	1.000	0.584	0.070	-0.153	0.512	0.554
Miami_Hialea_FL	0.655	1.000	0.684	0.083	-0.179	0.546	0.550
Minneapolis_MN	0.729	1.000	0.920	0.081	-0.227	0.638	0.547
New_York_NJ_NY	0.288	0.368	2.743	0.216	-0.472	0.935	0.641
Orlando_FL	1.000	1.000	0.693	0.072	-0.165	0.492	0.496
Philadelphia_PA	0.479	1.000	0.703	0.076	-0.177	0.570	0.576
Phoenix_AZ	1.000	1.000	0.742	0.068	-0.180	0.561	0.521
Pittsburgh_B_PA	0.206	1.000	0.457	0.053	-0.107	0.440	0.477
Portland_Van_OR	1.000	1.000	0.513	0.065	-0.131	0.452	0.500
Riverside_Sa_CA	1.000	1.000	0.418	0.071	-0.119	0.370	0.461
Sacramento_CA	0.707	1.000	0.555	0.077	-0.158	0.394	0.427
St_Louis_MO	0.550	1.000	0.671	0.070	-0.176	0.491	0.459
San_Diego_CA	0.752	1.000	0.703	0.072	-0.182	0.556	0.542
San_Francisc_CA	1.000	1.000	1.322	0.097	-0.276	0.583	0.596
San_Jose_CA	1.000	1.000	2.041	0.184	-0.531	0.671	0.677
Seattle_Ever_WA	1.000	1.000	1.422	0.122	-0.364	0.516	0.515
Tampa_St_Pe_FL	0.591	1.000	0.517	0.072	-0.141	0.471	0.538
Washington_DC	0.661	0.711	0.588	0.050	-0.136	0.458	0.400

NOTES: The estimates are conducted city-by-city using an AR1 structure of the log of the series. A δ value of 1 indicates a Dickey-Fuller test rejected the null of no unit-root.

Table 10: Amenity Term μ_j: Baseline Estimates and Recalibrations

MSA	Baseline Amenity Estimate	Counterfactuals: Recalibration Static	Natl Housing	Less Idiosyncratic
Outside_option	0.000	0.000	0.000	0.000
Atlanta_GA	-0.070	-0.014	-0.073	-0.102
Baltimore_MD	-0.071	-0.007	-0.069	-0.083
Boston_MA	-0.078	-0.019	-0.082	-0.121
Chicago_Gary_IL	-0.052	-0.011	-0.051	-0.070
Cincinnati_OH	-0.077	-0.010	-0.077	-0.092
Cleveland_OH	-0.081	-0.011	-0.081	-0.101
Dallas_FW_TX	-0.072	-0.019	-0.076	-0.116
Denver_Boulder_CO	-0.072	-0.009	-0.079	-0.095
Detroit_MI	-0.072	-0.016	-0.073	-0.106
Houston_Braz_TX	-0.065	-0.023	-0.065	-0.093
Indianapolis_IN	-0.085	-0.013	-0.085	-0.108
Kansas_City_MO	-0.080	-0.011	-0.081	-0.100
Los_Angeles_CA	-0.034	0.009	-0.029	-0.027
Miami_Hialea_FL	-0.075	-0.015	-0.076	-0.100
Minneapolis_MN	-0.069	-0.013	-0.069	-0.092
New_York_NJ_NY	-0.044	-0.013	-0.041	-0.068
Orlando_FL	-0.075	-0.002	-0.079	-0.085
Philadelphia_PA	-0.064	-0.016	-0.063	-0.091
Phoenix_AZ	-0.066	-0.006	-0.072	-0.082
Pittsburgh_B_PA	-0.079	-0.013	-0.079	-0.100
Portland_Van_OR	-0.071	-0.001	-0.077	-0.077
Riverside_Sa_CA	-0.047	0.014	-0.052	-0.024
Sacramento_CA	-0.066	0.001	-0.066	-0.060
St_Louis_MO	-0.076	-0.012	-0.076	-0.098
San_Diego_CA	-0.053	0.011	-0.050	-0.032
San_Francisc_CA	-0.057	0.000	-0.068	-0.060
San_Jose_CA	-0.094	-0.017	-0.101	-0.132
Seattle_Ever_WA	-0.071	-0.007	-0.076	-0.090
Tampa_St_Pe_FL	-0.069	-0.005	-0.069	-0.075
Washington_DC	-0.061	-0.016	-0.060	-0.084

Table 11: Estimates of Gravity Model, Data and Baseline Simulation

		Data	Baseline
orig wage	Est	-0.867	-0.834
	se	0.079	0.011
dest wage	Est	0.514	0.843
	se	0.079	0.011
orig ϕ	Est	-0.187	-0.143
	se	0.020	0.003
dest ϕ	Est	0.316	0.147
	se	0.020	0.003

NOTES: Data are annual frequency.

Table 12: Point Estimates of Gravity Model, by Simulation Scenario and City-type Grouping

Avg. for Group:	Labor Demand	Baseline	Natl Housing Demand	Less Idiosyncratic
All	Wage	0.735	1.350	1.139
	ϕ	0.135	0.290	0.231
Inelastic HS	Wage	0.571	1.142	0.842
	ϕ	0.041	0.166	0.067
Elastic HS	Wage	0.920	1.579	1.484
	ϕ	0.145	0.308	0.238
Persistent Shock	Wage	0.818	1.523	1.337
	ϕ	0.143	0.307	0.283
Less Persistent Shock	Wage	0.092	0.018	0.253
	ϕ	0.037	0.084	0.040

NOTES: The estimates from these specifications generate Figures 7 and 12. Data are at quarterly frequency. Point estimate is for destination attribute for incumbent workers; origin point estimates and those for new workers are essentially the same. Standard errors are suppressed.

Table 13: Point Estimates of Population Autoregression, Data and Baseline Simulation, by City-type Grouping

Averaged					DATA			BASE SIM		
				Grouping	lag p	wage	ϕ	lag p	wage	ϕ
				All	0.989	0.004	0.021	0.977	0.013	0.021
				Inelastic HS	0.991	0.000	0.007	0.965	0.015	0.015
				Elastic HS	0.983	0.013	0.039	0.979	0.016	0.021
				Persistent Shocks	0.985	0.009	0.037	0.973	0.029	0.035
				Less Persistent Shocks	0.994	0.001	0.009	0.976	0.003	0.012
MSA-by-MSA	MSA attributes				DATA			BASE SIM		
MSA Name	HS slope	wage persist	ϕ persist	Inelas.,Pers.	lag p	wage	ϕ	lag p	wage	ϕ
Atlanta_GA	0.40	1.00	1.00	0,1	0.976	0.003	0.120	0.974	0.048	0.049
Baltimore_MD	1.35	0.57	1.00	1,0	0.976	0.006	-0.007	0.983	0.002	0.030
Boston_MA	1.18	1.00	1.00	1,1	1.022	-0.003	0.018	0.979	0.030	0.050
Chicago_Gary_IL	1.04	0.66	1.00	1,0	0.995	-0.007	0.004	0.995	-0.003	0.025
Cincinnati_OH	0.48	0.72	1.00	0,0	0.980	0.004	-0.003	0.989	0.004	0.030
Cleveland_OH	1.13	0.57	1.00	1,0	1.008	-0.017	-0.009	0.975	0.009	0.012
Dallas_FW_TX	0.31	1.00	1.00	0,1	0.987	0.007	0.066	0.972	0.053	0.041
Denver_Bould_CO	0.91	1.00	1.00	0,1	0.987	0.004	0.062	0.972	0.029	0.049
Detroit_MI	1.79	0.71	1.00	1,0	0.916	0.011	0.030	0.998	0.000	0.013
Houston_Braz_TX	0.44	0.35	0.65	0,0	0.973	0.027	0.077	0.975	0.006	0.015
Indianapolis_IN	0.43	0.57	0.69	0,0	0.993	0.008	0.001	0.980	0.005	-0.015
Kansas_City_MO	0.80	0.64	0.66	0,0	0.990	0.005	-0.003	0.988	-0.001	0.003
Los_Angeles_CA	2.34	0.56	1.00	1,0	0.978	0.007	0.005	0.962	0.000	-0.001
Miami_Hialea_FL	1.45	0.66	1.00	1,0	1.005	-0.025	-0.001	0.981	-0.001	0.018
Minneapolis_MN	1.08	0.73	1.00	1,0	0.990	0.004	0.020	0.982	-0.001	0.027
New_York_NJ_NY	1.54	0.29	0.37	1,0	0.996	-0.005	-0.017	0.985	-0.002	-0.003
Orlando_FL	0.62	1.00	1.00	0,1	0.981	0.019	0.050	0.970	0.040	0.037
Philadelphia_PA	1.15	0.48	1.00	1,0	0.981	0.005	-0.010	1.014	-0.003	0.024
Phoenix_AZ	0.65	1.00	1.00	0,1	0.977	0.012	0.071	0.973	0.035	0.038
Pittsburgh_B_PA	0.70	0.21	1.00	0,0	0.981	-0.021	-0.039	0.929	0.009	0.019
Portland_Van_OR	1.09	1.00	1.00	1,1	0.995	-0.010	0.018	0.974	0.009	0.023
Riverside_Sa_CA	0.74	1.00	1.00	0,1	0.961	0.090	0.018	0.967	0.024	0.013
Sacramento_CA	0.78	0.71	1.00	0,0	0.973	0.031	0.024	0.970	0.002	0.017
St_Louis_MO	1.21	0.55	1.00	1,0	0.989	0.002	0.005	0.970	0.009	0.070
San_Diego_CA	1.08	0.75	1.00	1,0	0.988	0.007	0.000	0.976	-0.002	-0.001
San_Francisc_CA	1.23	1.00	1.00	1,1	0.995	-0.001	0.003	0.982	0.004	-0.008
San_Jose_CA	0.94	1.00	1.00	1,1	1.037	-0.004	0.031	0.953	0.042	0.063
Seattle_Ever_WA	0.82	1.00	1.00	0,1	0.994	0.002	0.015	0.966	0.039	0.047
Tampa_St_Pe_FL	0.97	0.59	1.00	1,0	0.992	0.009	0.032	0.982	-0.002	0.023
Washington_DC	0.92	0.66	0.71	0,0	0.988	0.011	-0.003	0.977	0.003	0.003
				Correlation b/tw est & HS slope		-0.238	-0.426		-0.487	-0.242
				Correlation b/tw est & persistence		0.226	0.152		0.694	0.478

NOTES: The estimates from these specifications generate Figures 1, 2,4, 6. Data are at quarterly frequency. Standard errors are suppressed.

Table 14: Point Estimates of Population Autoregression, by Simulation Scenario and City-type Grouping

	BASE SIM		
Grouping	lag p	wage	ϕ
All	0.977	0.013	0.021
Inelastic HS	0.965	0.015	0.015
Elastic HS	0.979	0.016	0.021
Persistent Shocks	0.973	0.029	0.035
Less Persistent Shocks	0.976	0.003	0.012
	STATIC		
Grouping	lag p	wage	ϕ
All	0.000	0.608	0.456
Inelastic HS	0.000	0.319	0.491
Elastic HS	0.000	1.106	0.574
Persistent Shocks	0.000	0.653	0.597
Less Persistent Shocks	0.000	0.594	0.417
	NATL. HOUSING DEMAND		
Grouping	lag p	wage	ϕ
All	0.978	0.016	0.022
Inelastic HS	0.964	0.018	0.021
Elastic HS	0.984	0.017	0.014
Persistent Shocks	0.970	0.052	0.078
Less Persistent Shocks	0.975	-0.005	0.010
	LESS IDIOSYNCRATIC		
Grouping	lag p	wage	ϕ
All	0.973	0.013	0.032
Inelastic HS	0.949	0.016	0.034
Elastic HS	0.977	0.019	0.028
Persistent Shocks	0.964	0.046	0.065
Less Persistent Shocks	0.963	0.000	0.032

NOTES: The estimates from these specifications generate Figures 8, 13. Data are at quarterly frequency. Standard errors are suppressed.

B Existence and Uniqueness of Approximated Equilibrium

The endogenous state variables are the populations in each location j. In the full solution to (10), the population of each j enters as a separate state because of their potential impact on future flows of utility as migrants of those locations may choose i in the future. The simplification made to the approximated solution is to ignore the individual states, including populations, of other locations; the approximated problem uses only the population state immediately relevant to current and future flow utility.

Let $p_i^{'}$ be the population that results from the migration probabilities to location i, σ_{ij}, of the population currently residing in locations $j \in \{1, 2, ..., J\}$.

$$p_i^{'} = \sum_j \sigma_{ij} p_{ij}$$

The migration probabilities depend on what the population $p_i^{'}$ is expected to be, in addition to labor market demand states summarized by X. Let p_i^* represent the expected population. In equilibrium,

$$p_i^{'}(X, p_i^*) = p_i^* = \sum_j \sigma_{ij}(X, p_i^*) p_{ij} \tag{16}$$

The population that obtains is a result of the summation of location demand functions σ. In order to show that the equilibrium is unique, one must show that the demand function, (16), has one fixed point, or equivalently, that it crosses the 45 degree line once and only once. This amounts to showing that the demand function, (16), is strictly positive at 0 and monotonically decreasing.

Theorem. Let the demand for a location as a function of expected population be given by (16), with the associated choice probabilities given by (11) and value functions given by (10) under the summarized state space. The demand function has exactly one fixed point.

Proof. First, to achieve existence, one must show that demand for the location is strictly positive at $p_i = 0$; that is, that $\sum_j \sigma_{ij}(X, 0) p_{ij} > 0$

Recall that the choice probability is

$$\sigma(X, p_i) = \frac{exp(v_{ij}(X_i, p_i))}{\sum_k exp(v_{kj}(X_k, p_k))}$$

Given the form of the choice probabilities, the result is trivial, and I explain it without a lengthy derivation. Demand is strictly positive demand as long as $v_{ij} \neq -\infty$. If $p_i = 0$, then $er = 1$ and $r = 0$, so flow utility is strictly positive. I rule out by assumption expectations so dystopian as to render the continuation value, $E[v_{ij}^{'}]$ to be negative infinity. Infinitely negative amenity values, μ_i, are also ignored. Hence, $\sigma(X, 0) > 0$.

Next, to achieve uniqueness, one must show that the demand function is monotonically decreasing in p_i, $\frac{\partial p_i^{'}}{p_i} < 0$. I now show that the derivative is negative:

$$\frac{\partial p_i'}{p_i} = \frac{\partial \sum_j \sigma_{ij}(X, p_i) p_{ij}}{p_i} < 0$$

The value function for the destination location varies with the origin, but only in the ij-specific moving cost. The moving cost is additive and not dependent on p_i, so it is not present in the derivative, meaning that it is sufficient to show that the choice probability is monotonically decreasing in p_i. (The choice probabilities from different origins j are different in levels due to the moving costs, but not slopes.) Therefore I show the derivations for an arbitrary origin j.

$$\frac{\partial \sigma(X, p_i)}{p_i} = \frac{\partial \frac{exp(v_{ij}(X_i, p_i))}{\sum_k exp(v_{kj}(X_k, p_k))}}{p_i} < 0$$

For ease of notation, I refer to the numerator of the choice probability as N, and the denominator as D:

$$N = exp(v_{ij}(X_i, p_i))$$

$$D = \sum_k exp(v_{kj}(X_k, p_k))$$

The derivative of the choice probability is

$$\frac{\partial \sigma(X, p_i)}{p_i} = \frac{\frac{\partial N}{p_i} D + N \frac{\partial D}{p_i}}{D^2}$$

Clearly, N, D, and D^2 are positive. The result will obtain if $\frac{\partial N}{p_k} < 0$, $\frac{\partial D}{p_k} < 0$.

$$\frac{\partial N}{p_k} = exp(v_{ij}(X_i, p_i)) \frac{\partial v_{ij}}{p_k} = exp(v_{ij}) \left(\frac{\partial u(er, w, r)}{p_k} + \beta \frac{\partial E[v_{ij}]}{p_k} \right)$$

What remains to be shown is that $\frac{\partial u(er,w,r)}{p_k} < 0$, as $\frac{\partial E[v_{ij}]}{p_k} < 0$ if $\frac{\partial u(er,w,r)}{p_k} < 0$ by Theorem 9.10 of Stokey et al. (1989). This will be true so long as the interactions between workers in the model are competitive, in the sense that more workers in one's location means lower utility, ceteris paribus. These conditions are satisfied in the model of this paper, as the workers compete for jobs and housing.

By way of example, I derive this result for the specific utility function used above. Fully expanded to see all the components, (8) is

$$u(er, w, r) = er \cdot u(w, r) + (1 - er) \cdot u(0, r)$$

$$\phi w^{1-\alpha} p^{\alpha-1} (u_0 - exp(-\gamma_2(w - \rho_0 - \rho_1 \phi w^{2-\alpha} p^{\alpha})) + (1 - \phi w^{1-\alpha} p^{\alpha-1})(u_0 - exp(-\gamma_2(0 - \rho_0 - \rho_1 \phi w^{2-\alpha} p^{\alpha}))$$

where u_0 is an arbitrary constant in the CARA utility specification.

Its derivative with respect to p is

$$\frac{\partial u(er,w,r)}{p} = (\alpha-1)\phi w^{1-\alpha}(u(w,r)-u(0,r)) + \gamma_2\rho_1\phi w^{2-\alpha}p^{\alpha}(er(u(w,r)-u_0)+(1-er)(u(0,r)-u_0))$$

The employment congestion term, $(\alpha-1)\phi w^{1-\alpha}(u(w,r)-u(0,r))$, is negative because $\alpha < 1$, and $\phi w^{1-\alpha} > 0, (u(w,r)-u(0,r)) > 0$ for $w > 0$ (i.e., wages are strictly positive).

The rent increase term, $\gamma_2\rho_1\phi w^{2-\alpha}p^{\alpha}(er(u(w,r)-u_0)+(1-er)(u(0,r)-u_0))$, is negative because $\gamma_2\rho_1\phi w^{2-\alpha}p^{\alpha} > 0, 0 \leq er \leq 1$, and $u_0 > u(w,r), u_0 > u(0,r)$. Hence, each term is negative, so $\frac{\partial u(er,w,r)}{p_k} < 0$, as was to be shown.

The denominator D is the summation of the numerator and similar $exp(\cdot)$ expressions for the other locations, one of which is affected by the population of i because of the adding-up constraint: without loss of generality, any population added to i is subtracted from location $j = 0$. So

$$\frac{\partial D}{p_k} = \frac{\partial N}{p_k} + \frac{\partial exp(v_{0j}(X_i, P - \sum_{k\neq 0} p_k))}{\partial p_i}$$

By a similar but converse argument as above, the value function for the outside option is increasing as the population of the location i increases–there is less congestion in $j = 0$ as the population in i increases. Hence by the chain rule,

$$-\frac{\partial exp(v_{0j}(X_i, P - \sum_{k\neq 0} p_k))}{\partial p_i} < 0$$

So then

$$\frac{\partial D}{p_k} = \underbrace{\frac{\partial N}{p_k}}_{-} - \underbrace{\frac{\partial exp(v_{0j}(X_i, P - \sum_{k\neq 0} p_k))}{\partial p_i}}_{+} < 0$$

Thus,

$$\frac{\underbrace{\frac{\partial N}{p_i}}_{-}\underbrace{D}_{+} + \underbrace{N}_{+}\underbrace{\frac{\partial D}{p_i}}_{-}}{\underbrace{D^2}_{+}} < 0$$

and the result obtains.

■

The proof was for an arbitrary location, and hence applies to all locations in the system.

While the derivation was cumbersome, the basis of the proof is quite simple. As long as the interactions in the model are congestive, the equilibrium of the approximated problem is unique, a result analogous to the static model of Bayer and Timmins (2005).

C Equilibrium Approximation Method

The following appendix section describes the method of equilibrium approximation in more detail. It includes tips for practitioners.

C.1 The PPI Algorithm

From (14), the value function is approximated as

$$v_j(X) \approx \hat{v}_j(X) = \sum_q^Q \lambda_{j,q} g_q(x_j)$$

where j superscripts indicate that these are the choice-specific value functions (i.e. one for each location). The probability of a worker choosing a particular location is shown in (11). The parameters λ are chosen so as to fit the contraction implied by the problem as nearly as possible:

$$\hat{\lambda} = argmin_\lambda \| [\sum_q \lambda_q g_q(x) - u(X') - \beta E(\sum_q \lambda_q g_q(x')] \| \tag{17}$$

The PPI algorithm is

(I) Initialize

1. Select X by state aggregation
2. Choose basis points in the state space at which to evaluate the problem
3. Form the functions $g(X), E(g(X'))$.
4. Guess an initial (i.e. iteration 0) vector of parameters for the approximation, λ^0.

(II) Iterate

1. Form the choice probabilities $\Sigma(\lambda_0)$ according to (11), using the current guess of λ in (14)
2. Using $\Sigma(\lambda^0)$, update $E(g(X'))$
3. Find λ^1 according to (17). This updates the guess of each $\hat{v}_j(X)$, yielding updated choice probabilities, $\Sigma(\lambda_1)$, returning to iteration step 1.
4. Repeat iteration steps II-1,2,3 until convergence in $\Sigma(\lambda^r)$ by a chosen metric such as $sup||\Sigma(\lambda^r) - \Sigma(\lambda^{r-1})|| < \epsilon$. (Convergence in $\lambda, \hat{v}_j(X)$) should follow.).

Except for the state aggregation step I-1, the steps are specific to the PPI application of projection methods. Comparable value function solution methods may be substituted; other methods will differ in specifics of implementation but follow a similar iterative structure. I next describe the steps in more detail.

C.2 State Aggregation

How to aggregate states is a matter of researcher judgment in the context of the specific problem. In the implementation above, the states used for approximating the conditional value function are the location's own wage, employment efficiency ϕ, and population, the average wage and employment efficiency of all other locations, and the economy's total population, for a total of six states. In the model of this paper, population level is important for utility, but the source of competing population is irrelevant, so the anonymous "own/other" structure seemed appropriate. (Experimentation with other definitions of states yielded no better summarizations.) In some applications, states describing individual markets (e.g. the productivity of a near neighbor) make be appropriate.

C.3 Choosing the Basis Points

Projection methods require the functional equation to be evaluated at sampled points in the state space. The λ parameters on the approximating function depend on the state points at which the value function is evaluated. Choosing a "good basis" therefore means not only specifying the parametric form of the approximating function, but also selecting "good data" on which to evaluate the approximating function. As in any curve-fitting, the more data points the better, but these basis points come at some computational cost. (For a choice of N basis points, the matrix in the policy improvement step has NJ rows.) Some judgment on the part of the researcher is necessary to parsimoniously choose points that yield a good approximation to the underlying value function and good performance in the numerical routine.

I use $N = 1,000$ basis points, drawn in three ways. First, I use states that actually occur in the data, but in uncorrelated manner, randomly choosing the points from the empirical distribution from each city. This draw is intended to give the approximating function relevance for the equilibrium played in the data. Second, I randomly draw points from a range slightly larger than the empirical distribution. For example, if the population of a city in the data ranges from 1 to 1.2 million, I draw from a range such as 0.9 to 1.3 million. This draw is intended to add information to the approximating function for counterfactual points in the state space, while adding density to the data-relevant range. Third, I draw points imposing some negative correlation between the three types of states for each city. This third draw is intended to add information for potential counterfactual scenarios as well as improve the performance of the matrix operations that find the approximating parameters. Within this draw, there are four ways to reverse the correlation in the data: three pairwise restricted (wage and population, wage and ϕ, population and ϕ), with the third state drawn as uncorrelated to the two restricted, and one with a restriction on the full 3-by-3 covariance matrix.

The same set of basis points is used for all evaluations of the approximating function, both within estimation and for counterfactuals, to prevent simulation error from affecting results.

C.4 Forming the Approximating Polynomial Basis Functions $g(\cdot)$

Here again, some researcher judgment is necessary, and in principle many sets of polynomials could work. Numerical performance is improved when the matrix $[g_1(x) \cdots g_Q(x)]$ is well-conditioned (e.g.

when the columns are orthogonalized).[60] It is often useful that the variable first be normalized; this prevents number scaling issues from affecting numerical performance, especially in early iterations when the problem is relatively far from its solution. For each state i, I used $z = \frac{x_i - \bar{x}_i}{\sigma_i}$, where the means and standard deviations are taken from the sample data. This expresses each point in the basis space relative to the sample mean.

After much experimentation, I settled on a fully interacted polynomial in the six states:

$$Z = [1 \; z_1 \; z_2 \; z_3 \; z_4 \; z_5 \; z_6 \; z_1z_2 \; z_1z_3 \ldots z_2z_3 \; z_2z_4 \ldots z_1z_2z_3 \ldots z_1z_2z_3z_4z_5z_6]$$

A constant was included for each city to allow the mean value to vary across locations. The higher order terms permitted the value function to have nonlinear shape; the highest order term was a full interaction of all the states. Higher order terms besides interactions (such as squares, cubics, etc) did not improve performance and were omitted.

The matrix of expected future states, $E(g(X'))$, is important to the problem because it reflects different transition processes in the state variables between locations. For example, an above-mean value of x indicates an above-mean value of x' in locations with more persistent shocks, and hence this state may be judged to be more valuable than in a location with states that more quickly mean revert. To form the $E(g(X'))$ matrix, standard integration methods may be used. The optimal choice of method depends on how often the expected future state matrix needs to be evaluated. As explained below, the $E(g(X'))$ could be precomputed under the structure of this problem, and simple Monte Carlo integration sufficed. Simulation draws use each location/state AR(1) transition function, $x'_r = \delta x + \epsilon_r$, where δ is the AR(1) parameter and ϵ_r is a particular draw from the shock distribution. The Monte Carlo integration uses the full $2J$-dimensional empirical shock covariance matrix with 15,000 error draws. Note that with nonlinear approximating functions, $E(g(X')) \neq g(E(X'))$; hence, the states must be drawn first, the matrix Z formed, and $E(g(X')) = \frac{1}{R}\sum_r^R Z_r$.

Recall that because agents are assumed atomistic, the policy predicts the population shares exactly. Thus, conditional of the policy, there is no uncertainty regarding the endogenous states. This turns out to be a very convenient feature. Since the expectation of future states is taken over the exogenous variables only, the $E(Z)$ matrix can be precomputed outside the PPI routine. $E(Z|P)$ is then converted to $E(Z|P')$ using a simple linear operation within the PPI routine. The computational time savings are substantial.

C.5 Iteration

As illustration of the iterative structure of the approximation method, consider a two-location example. The choice values can be written as

$$v_1 = \pi_1 + \beta[\sigma_{1|1}E(v_1) + \sigma_{2|1}E(v_2)]$$

$$v_2 = \pi_2 + \beta[\sigma_{1|2}E(v_1) + \sigma_{2|2}E(v_2)]$$

[60] Drawing uncorrelated basis points also improved the performance of matrix operations.

where $\pi_j = \sigma_{1|j}(u_1 + m_{1|j} + \gamma - log(\sigma_{1|j})) + \sigma_{2|j}(u_2 + m_{2|j} + \gamma - log(\sigma_{2|j}))$ is the choice-weighted expected value of flow utility less moving costs, conditional on being in j. States X are suppressed in notation. γ is Euler's constant, the mean of the T1EV distribution. One can think of π as the expected value of flow utility of the optimal choice before the idiosyncratic ε shocks are revealed. Putting the expected value of the ε in the flow term will preserve linearity in the discounted expected value function terms. This expression of the value function follows Aguirregabiria and Mira (2007) and Sweeting (2013).

Step II-1. The choice probabilities are formed using the current iteration's guess of λ_s. Find

$$v_{j|k} = u_j + m_{j|k} + \beta \hat{Z}'_j \lambda_s$$

and then insert into (11) to make $\Sigma(\lambda_s)$.

Step II-2. The expected future approximating functions $E(g(X'))$ are formed/updated, conditional on $\Sigma(\lambda_s)$.

Step II-3. The routine leverages the linearity in the approximation. Let $Z = [g_1(x) \cdots g_Q(x)]$ and $\hat{Z} = E([g_1(x') \cdots g_Q(x')])$. The J approximations of (14) yields a system of equations, represented in matrix form as:

$$\begin{pmatrix} Z_1 & 0 \\ 0 & Z_2 \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} \pi_1 \\ \pi_2 \end{pmatrix} + \beta \begin{pmatrix} \sigma_{1|1}\hat{Z}'_1 & \sigma_{2|1}\hat{Z}'_2 \\ \sigma_{1|2}\hat{Z}'_1 & \sigma_{2|2}\hat{Z}'_2 \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix}$$

Each location j has N rows (the basis points), so the matrices are of size $JN \times Q$ and vectors $JN \times 1$. The lefthand side matrix is block diagonal of the row vectors Z_j, representing the value in the origins 1 and 2. On the righthand side, the expected states function $\hat{Z}$ is weighted by the choice probabilities, which represents that the choice values incorporate the probabilities of transitioning to other markets from j. Note that these probabilities are heterogeneous by origin, since location-to-location moving costs vary. This complication is unique to a segmented markets model, and is unlike, for instance, players in a dynamic game who maintain consistent identity. Because of the joint dependence, it is useful to solve the functional equation as a multi-equation matrix problem.

The above set of linear equations can be rearranged as

$$\underbrace{\begin{pmatrix} Z_1 & 0 \\ 0 & Z_2 \end{pmatrix} - \beta \begin{pmatrix} \sigma_{1|1}\hat{Z}'_1 & \sigma_{2|1}\hat{Z}'_2 \\ \sigma_{1|2}\hat{Z}'_1 & \sigma_{2|2}\hat{Z}'_2 \end{pmatrix}}_{C} \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix} = \underbrace{\begin{pmatrix} \pi_1 \\ \pi_2 \end{pmatrix}}_{\pi}$$

The terms λ are chosen to minimize the distance between the two sides of the approximated contraction, the policy improvement step of PPI. Since the λ terms enter linearly, the distance can be minimized and the λ updated by, inter alia, a simple ordinary least squares formula:

$$\hat{\lambda} = (C'C)^{-1}(C'\pi)$$

The update returns the algorithm to step II-1. This routine is iterated until convergence in the policy function (the migration share matrix) by some metric such as the $sup\,norm$. Convergence in the approximating parameters λ occurs concomitantly.