

Cities and Labor Market Dynamics

by

Kyle Douglas Mangum

Department of Economics
Duke University

Date: _____

Approved:

Patrick Bayer, Supervisor

Peter Arcidiacono

Christopher Timmins

Andrew Sweeting

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2012

ABSTRACT

Cities and Labor Market Dynamics

by

Kyle Douglas Mangum

Department of Economics
Duke University

Date: _____

Approved:

Patrick Bayer, Supervisor

Peter Arcidiacono

Christopher Timmins

Andrew Sweeting

An abstract of a dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2012

Copyright © 2012 by Kyle Douglas Mangum
All rights reserved

Abstract

People live and work in local markets spatially distinct from one another, yet space is absent from most economic models of the national labor market. Workers choose the markets in which they will participate, but there are costs to mobility. Furthermore, cities are heterogeneous in a number of dimensions, including their local labor market productivity, their housing supply, and their offering of amenities.

I examine the impact of these spatial considerations on the dynamics of local labor markets and the national market to which they aggregate. First I study the patterns of location choice through a gravity model of migration applied to rich panel data from the U.S. I find that location choices respond to temporal shocks to the labor market, but only after controlling for local heterogeneity. Next, with this result as motivation, I turn to development of a dynamic spatial equilibrium of the national labor market. I make a technical contribution to work in dynamic equilibrium modeling by empirically implementing an island economy model of worker mobility. I quantify the importance of worker mobility costs versus local housing prices for explaining spatial variation in the unemployment rate. I find that the link between the local housing market and the local labor market is important for explaining the spatial dispersion in unemployment, but mobility costs are not. Finally, I further exploit the dynamic equilibrium framework to examine the effect of local housing policy on labor market growth. I find that housing supply regulation is a constraint to growth, but is only binding on cities that are particularly desirable because of

their labor market opportunities or amenities. I find that some lightly regulated markets have a contingent of population that has been pushed out of more regulated markets by high housing prices.

Contents

Abstract	iv
List of Tables	viii
List of Figures	x
1 Migration Flows to Labor Markets	1
1.1 Introduction	1
1.2 Model	3
1.3 Empirical Application to Migration	5
1.3.1 Data	5
1.3.2 Estimation	8
1.4 Results	9
1.4.1 Baseline Regressions	9
1.4.2 Forward-looking Behavior and Migration Momentum	15
1.4.3 Housing Lock-in	15
1.4.4 Trends in Migration Rates	21
1.5 Conclusion	22
2 A Dynamic Equilibrium Model of Cities and Labor Market Dynamics	25
2.1 Introduction	25
2.2 Motivation	31
2.3 Model	36

2.3.1	The Setting	36
2.3.2	The Local Labor Market	37
2.3.3	Housing and Rents	41
2.3.4	Workers	43
2.3.5	Equilibrium	46
2.4	Equilibrium Approximation	50
2.5	Data and Estimation	54
2.5.1	Data	54
2.5.2	Estimation Stage 1: Local Data Series and Parameters	55
2.5.3	Stage 2: Migration Model	58
2.6	Results	62
2.6.1	First Stage Results: Location Heterogeneity	63
2.6.2	Second Stage Results: The Dynamic Model	65
2.6.3	Simulations	68
2.7	Conclusions and Directions for Future Work	80
3	Housing Regulation and Labor Market Growth	91
3.1	Introduction	91
3.2	Developing The Experiment	92
3.3	Results	95
3.3.1	Experiment 1: All Cities	95
3.3.2	Experiment 2: Reducing Regulation in Highly Regulated Cities Only	99
3.4	Conclusion	100
	References	103
	Author Biography	107

List of Tables

1.1	Baseline Results from the Gravity Model	11
1.2	Regressions Including Lagged Labor Market Growth	16
1.3	Summary of Lock-in Measures	19
1.4	Regressions Including Measures of Lock-in	20
1.5	Alternative Distance Functions	23
1.6	Results from the Gravity Model Using Migrants Only	24
2.1	R^2 from Regressions of Quarterly Local Labor Market Statistics on Time and MSA Dummies	36
2.2	Local Means and Volatilities of Employment State Variables	64
2.3	Housing Price Elasticities	66
2.4	Structural Parameters	67
2.5	Cross Sectional Unemployment Distribution, Data and Baseline Model	68
2.6	R^2 from Regressions of Unemployment Rates on Time and MSA Dum- mies, Data and Baseline Model	69
2.7	Mean Cross-Sectional Unemployment Dispersion Statistics from Sim- ulation Scenarios	72
2.8	Adjusted R^2 Statistics from Pooled Regressions of Unemployment on Time and Location Dummies	74
2.9	Mean Squared Error from National Unemployment Rate and National Unemployment Volatility	75
2.10	Cross-sectional Dispersion in Unemployment and Income Volatility . .	76
2.11	Mean Population Volatility by Simulation	76

2.12	Correlation of Population Differences to Baseline Simulation with Labor Market States	77
2.13	Correlations of Wage Premium/Population Difference Related to Locational Attributes	78
2.14	State Transition Processes	81
2.15	Correlation Across Locations in Shocks to Wages	82
2.16	Correlation Across Locations in Shocks to Employment Efficiency	83
2.17	Local Statistics From Baseline Simulation	84
2.18	Correlation of Population and Income Volatility Differences to Baseline Simulation with Rents and Amenities, By MSA	85
3.1	Conversion of Price Elasticities Using Saiz’s Estimates	95
3.2	Experiment 1: All Cities at Average Regulation Level—Mean Percentage Difference in Simulated Variables at End of Period	99
3.3	Experiment 2: Reducing Regulation in Highly Regulated Cities—Mean Percentage Difference in Simulated Variables at End of Period	101

List of Figures

1.1	Estimates of Time Dummies From Various Specifications	21
2.1	Unemployment Rates in MSAs, Quarterly	32
2.2	Dispersion in Cross-Sectional Distribution of Detrended Unemployment	33
2.3	Cross-Sectional Standard Deviation of Unemployment to Average Un- employment	33
2.4	Unemployment Rates in MSAs, Quarterly	34
2.5	Wages in 29 Largest MSAs, Local Difference to National Average, Quarterly, 1990-2010	35
2.6	Illustration of Local Heterogeneity in Housing Supply Elasticity . . .	42
2.7	Cross Sectional Unemployment Distribution in the Data and Baseline Simulation	68
2.8	Time Paths of Population and Unemployment in Detroit and Seattle	71
2.9	Cross Sectional Unemployment Distribution, Kernel-Density Estimates, by Simulation	73
2.10	Cross Sectional Standard Deviation in Unemployment Rate by Sce- nario, 1990-2010	74
2.11	Experiments' Differences in Dispersion Plotted Against Dispersion . .	79
2.12	Aggregate Statistics from Baseline Simulation	86
3.1	Price Elasticity Estimates Plotted Against Measures of Unusable Land and Regulation	96
3.2	Experiment 1: Population Change Plotted Against Price Elasticity Change	98

3.3	Experiment 2: Population Change Plotted Against Price Elasticity	
	Change	102

Migration Flows to Labor Markets

1.1 Introduction

There is an apparent disconnect in the literature on labor migration. On one hand, it has long been recognized that a person's location choice is made in a dynamic context, the result of forward-looking behavior. Therefore, the literature studying individual migration decisions, from Sjaastad (1962) to more recent structural work¹ is formulated in this dynamic setting. On the other hand, most studies of aggregate, interregional migration are framed, or at least empirically evaluated, in a static context or a very narrowly defined dynamic one;² either comparing migration flows between regions over a short period, or looking at before and after changes in region's populations. The limitation is often due to data availability.

Studies with relatively few periods of data are exploiting mainly (and perhaps only) spatial variation in location attributes, yet many attributes of labor and housing markets—salient topics of inquiry like wages, unemployment, and housing prices—are dynamically varying. This loses potential information from dynamic variation and

¹ See for example Kennan and Walker (2011), Bishop (2008), Gemici (2011), and Winkler (2011).

² See Greenwood (1985). Pissarides and Wadsworth (1989) and Dahl (2002) use individual data in estimation, but similarly in a narrowly defined dynamic context of one move decision.

it impairs the researcher's ability to study the effects of temporal shocks.

The lack of dynamic variability in local attributes could be especially problematic under the notion of a spatial equilibrium such as that proposed by Roback (1982). A feature of the disconnect in the literature is the issue of the importance of local labor market conditions on migration. Micro studies of migration (e.g. Kennan and Walker, 2011) find lifetime income prospects to be a key determinant of household migration choices; extrapolating from this one would expect local labor conditions to be importance to aggregate migration flows. But studies of aggregate migration do not find a clear positive relationship between population growth and local average wages (see, e.g., Blanchard and Katz, 1992; Greenwood, 1975, or Greenwood and Sexton, 2012). Under the notion of spatial equilibrium proposed by Roback (1982), prices in the labor and housing markets account for differences in local amenities. Thus, spatial variation in labor and housing prices tells us little about the underlying utility differentials that presumably are actually driving location choices. Thus, the lack of dynamic variability in local labor conditions may be part of the reason for this apparent contradiction.

This article seeks to narrow the gap in the literature by conducting a study of aggregate place-to-place migration under a long span of contiguous time. I utilize 19 consecutive years of migration flow data at a fine level of geography from a dataset compiled by the U.S. Internal Revenue Service (IRS). I apply a standard tool for the study of spatial interaction, the gravity model, but with the advantages afforded by the long time span of data.³ Estimation of the model is able to leverage temporal variation in local labor and housing market attributes in addition to spatial variation. Furthermore, I can account for unobserved location heterogeneity through the use of location or location pair fixed effects, which would not be possible without the

³ I also note in passing that the fine level of geography provided in the data allow me to specify the locations to be cities rather than states or regions. This is preferable for studying local labor and housing market attributes.

panel structure of the data. Results indicate that wages, unemployment, and housing prices have significant and intuitive effects on place-to-place migration flows, but only after conditioning on unobserved local heterogeneity. This is consistent with both the Roback notion of spatial equilibrium and the findings of recent structural studies of individual migration, resolving somewhat the contradiction between studies of individual and aggregate migration. Further, I find that recent local labor market growth is positively associated with migration flows, suggesting that migrants are forward-looking when choosing their destinations, which is consistent with the notion of migration being a dynamic decision problem.

The framework also allows me to examine the effect on migration of particular form of a temporal shock, the recent decline in housing prices. It has been suggested in the literature that such housing price declines impede labor mobility.⁴ After conditioning the gravity model on indicators of local housing market bust, I find at most weak evidence of so-called housing lock-in.

Finally, I find that the downward trend in migration rates noted by other researchers⁵ is present even after controlling for dynamic variation labor and housing market attributes.

The rest of the article proceeds as follows. The next section provides a brief review of the gravity model of spatial interaction. The third section describes the data and econometric specification. The fourth section discusses the empirical results. The fifth concludes.

1.2 Model

The gravity model, a mathematical expression of spatial interaction, borrows its name and intuition from the interaction of bodies in Newtonian physics. “Particles”

⁴ See, for example, Chan (2001) and Ferreira et al. (2010).

⁵ See, for example, Molloy et al. (2011).

(such as people or goods) float in space, with the size of the flows between points governed by the magnitude of the attraction and repulsion forces exerted by the points and the friction to movement imposed by distance.⁶ As in the gravitational pull of celestial bodies, the size or mass of the body affects its pull. The relationships can be expressed as:

$$M_{ij} = S_j S_i \frac{A_j R_i}{A_i R_j} \frac{1}{f(d)} \quad (1.1)$$

The size of the flow from i to j , M_{ij} depends positively on the sizes of i and j , S_i and S_j , the attraction forces of j , A_j , and the repulsive forces from i , R_i . Conversely, it depends negatively on the attraction of i , A_i , the repulsive forces of j , R_j , and by the distance friction, $f(d)$.

In logs, (1.1) becomes

$$m_{ij} = s_j + s_i + a_j - a_i + r_i - r_j - f(d) \quad (1.2)$$

where $\log(X) = x$. The linearity of equation (1.2) makes it easily applicable to a simple regression framework. The empirical application of the gravity model is then to find what constitutes attractive and repulsive forces.⁷

The gravity model has found use in transportation, hydrology, and telecommunications, though it is more commonly used in the social sciences, including economics, geography, and sociology (Kolaczyk [2009]). In economics, it has been used to study

⁶ Niedercorn and Bechdolt (1969) provide an economic derivation of the gravity model based on utility maximization constrained by time and/or travel costs. They also admonish the reader that analogy to physics can only be taken so far in a behavioral science and to recognize that the most useful feature of the gravity model is its empirical specification.

⁷ One could cast the gravity model completely in terms of attractive forces, depending on the definition of variables. For example, the attract/repel forces could be specified as pairwise differences, $x_{ij} = x_i - x_j$, then a positive difference either attracts to i or j . Ultimately it becomes an issue of mere semantics. The key is the sign of the variable in equation (1.2).

migration⁸ and trade.⁹

1.3 Empirical Application to Migration

This article applies a gravity model to U.S. domestic migration flows. I focus in particular on the impact of dynamically-varying attributes—prices in the local labor and housing markets—on migration flows. The point in space, or “body,” is defined as a metropolitan statistical area (MSA) because this is the area definition which most nearly corresponds to a labor market.

1.3.1 Data

The migration flow data come from data kept by the Statistics of Income Division of the Internal Revenue Service (IRS).¹⁰ The IRS compares a tax filer’s address from one year to the next to infer the filer’s migration status within the year: if the filer uses an address in A in year one, but an address of B in year two, she has apparently moved from A to B in the year; if she files at the same address, she has not moved.¹¹ The IRS then publishes the aggregated results in a place-to-place flow matrix. The data report gross flows.

The data are available at a county-to-county level from 1990-2008. Having multiple time periods, in successive years no less, is an advantage over most static gravity models.¹² The time dimension permits the use of time variation in addition to spatial variation in attractive and repulsive forces and the capture of unobserved location

⁸ See Zipf (1946) for a seminal example, and Greenwood (1975) for a survey.

⁹ See Carrere (2006) and citations therein.

¹⁰ More information is available at <http://www.irs.gov/taxstats/article/0,,id=212683,00.html>. A simple internet search for “IRS SOI migration data” will also lead to the relevant web pages.

¹¹ There are reasons other than migration that an address may change; for instance, a person may maintain two homes. For now, such complications are ignored, as they are likely infrequent and orthogonal to the objectives of this study.

¹² For an application of gravity models to repeated cross-sectional data, see Greenwood and Sexton (2012)

heterogeneity through the use of location fixed effects. The data also allow for time trends in migration, a topic of recent interest (Molloy et al [2011]). Further, the IRS migration data is essentially a high-frequency census of migration flows, an advantage over even large samples such as Integrated Public Use Microdata Series (IPUMS) or the Current Population Survey (CPS). Of course, it is a census of *taxpayers* only, which may impose some sample selection issues. While this is likely a small disadvantage for purposes of studying local labor markets, the reader should bear the disclaimer in mind.

Because the location status is derived from the individual’s tax return, which is filed during the first quarter of the year, the “migration year” does not exactly correspond to the calendar year, but leads one quarter. When I refer to year 1990, for instance, I am referring to location decisions made from the second quarter of 1990 to the first quarter of 1991.¹³ The labor and housing market data are adjusted to correspond to the migration year.

For each place-to-place cell, the data report the number of tax returns (essentially the number of households), the number of exemptions (the number of persons), and the taxable income of the filers in the cell. In order to keep the flow matrix of a manageable size, cells with fewer than ten returns are not reported. While this truncation and imputation is much less of an issue for the large urban counties used in this analysis, it does have a few practical consequences for the current study. First, I use the returns/households level data instead of the population data. Second, while the flow matrix may not be symmetric, it is not known whether a particular missing place-to-place pair actually has zero migrants, or just a number so small as to be unobserved. However, the total number of out-migrants are reported for each origin, and the total number of in-migrants are reported for each destination,

¹³ This is my own convention for ease of exposition. The IRS would refer to the same timeframe as the 1990-1991 data—from the tax return filed in 1990 (for income earned in 1989) to the tax return filed in 1991.

so the aggregate mobility is known. To mitigate the impact of the truncation, I impute unobserved cells by allocating the unaccounted for migrants according to a score based on the frequency of migration in other (observed) years. This is designed to correct for the cell which, by way of example, averages a steady flow of 15-20 migrating households per year, but for one year dips below the threshold of ten and is unobserved. This ignores the other kind of cell which, again by way of example, rarely has any migrating households, but happens to have one in a particular year; there is much less information to impute such a cell, and these cells are likely to be of little consequence anyway. If a pair still has zero migrants, I exclude the pair from the analysis rather than enter a zero (which when logged may skew a regression analysis).¹⁴

To examine attractive and repulsive forces, I also employ data on local labor and housing markets. The wage data come from the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics (BLS). This is a quarterly employer-level census of payroll size, in number of employed persons and total dollar amount. The average wage is derived directly from this. The unemployment rate data come from the Local Area Unemployment Statistics (LAUS), also from the BLS. The housing price data come from two sources. First, the price dynamics come from a metro level housing price index calculated by the Federal Housing Finance Agency (FHFA); this is pegged to a dollar amount using the median house value from the 2000 Census. All of these data sources were selected because of their fine level of geographic detail (county level for the BLS products, metro level for the FHFA) and relative frequency (quarterly or monthly for LAUS). The frequency allowed me to fit the data to the staggered migration year, and use the staggering to construct instruments. All data were aggregated to the MSA level.

¹⁴ Entering a small value, say, two, for unobserved pairs would cause an excessive number of total migrants.

1.3.2 Estimation

The econometric specification is the application of (1.2) to the migration flow and location attribute data. The estimating equation is:

$$m_{ijt} = \theta_0 + \theta_1 s_i + \theta_2 s_j + \alpha_i X_{it} + \delta_j Z_{jt} + f(d_{ij}, \mu) \quad (1.3)$$

The log migration flow from origin i to destination j at time t depends linearly on the log size of the origin s_i and destination s_j , the attributes of the origin, X_{it} , and destination Z_{jt} , and a linear in parameters function of the distance between the origin and destination, $f(d_{ij}, \mu)$. Some specifications contain time and or origin and destination fixed effects.

Note that the attributes of a location are indexed by t . The data are essentially an (unbalanced) panel, and thus there is time variation in addition to spatial variation in the attribute variables.¹⁵ The attributes employed are the local wages, unemployment rates, housing prices, and in some specifications, indicators of labor market growth and/or housing market shocks; all of which are varying dynamically. Time invariant attributes, such as climate, are either excluded or absorbed by the origin and destination fixed effects, depending on the specification. The parameters to estimate are $\theta, \alpha, \delta, \mu$. The signs and magnitudes of these parameters—with primary interest in α, δ —allow for interpretation of the attributes as attractive or repulsive forces.

The inclusion of origin and destination sizes is important. On the origin side, it accounts for larger flows resulting from a larger at-risk population; in the context of location choice, it turns the migration flow into a choice probability. The “at-risk” intuition is similar on the destination side. If one were simply drawing at random, a large city would be more likely than a small city to produce a new job, be home to a

¹⁵ Carrere (2006) applies a gravity model to panel data on trade flows.

family member, and the like, and hence be the destination. Also, a large destination is more likely to have good amenities and therefore be a desirable location.¹⁶

One could work with origin/destination differences—pairwise differences instead of origin and destination attributes—and the model is effectively the same, but I prefer to split out the origin and destination variables to maintain the attract/repel intuition. In what follows, the standard errors are clustered around the time-origin group to account for potential correlation in destination flow outcomes. For example, if there are idiosyncratically more migrants to destination A from origin X, then there will be fewer migrants in origin X available to choose all other destinations B, C, and so on.¹⁷

1.4 Results

This section presents the results from the application of the gravity model to the IRS migration data.

1.4.1 Baseline Regressions

Table 1.4.1 presents the main results. The distance friction employed here and in subsequent specifications is a dummy for whether the destination was a move, plus a coefficient times the log distance in miles between the origin and destination. The dummy accounts for the discontinuous drop in the level of flows between moves and non-moves, while the continuous variable allows the friction to change with distance in a concave way. Table 1.5 at the end of the article reports results for other specifications of the distance friction, including the move dummy alone, and linear

¹⁶ This statement derives intuition from a combination of Roback (1982) and Berry (1994).

¹⁷ Technically, for asymptotics on a true cluster sample, the econometrician would want the number of clusters to tend to infinity while the size of each cluster would remain fixed (Wooldridge, 2002). In the gravity model, adding more locations means the size of the cluster will also increase, suggesting that the cluster group should be only the time period. The results are not substantially affected by using the time as the only cluster group, but I prefer the finer cluster grouping of time-origin to account for the intuition explained in the main text.

and polynomial distances with and without the move dummy. While the distance specification does little to alter the attract/repulse results, the dummy plus log miles specification proved the best fit to the data. Naturally, both elements of the distance friction are negative and highly significant in all specifications, evidence that not only do migration flows decline with physical distance, but also that inertia is great: moving at all is a nontrivial decision.

Recall that the gravity model specifies that the origin and destination sizes be included as right-hand side variables, as sheer mass creates more flows in and out. In all specifications, these are positive and highly significant, in line with intuition.

The other regressors—the candidate attraction and repulsion forces—are the origin and destination unemployment rates, wages, and housing prices (all in logs). All attraction/repulsion variables are expressed as the difference from the contemporaneous national average.¹⁸ Note that a positive coefficient indicates the variable increases flows, which makes the variable a repulsive force when associated with the origin and an attractive force when associated with the destination. A negative coefficient is then indicative of an attractive force at the origin and a repulsive force at the destination.

Column 1 of Table 1.4.1 is the basic gravity model with origin and destination sizes, the distance friction, and the candidate attraction and repulsion forces. The origin and destination unemployment rates have intuitive signs, with unemployment acting as a repulsion from the origin and destination.¹⁹ Wages and housing prices, however, do not tell a consistent attract/repel story in this specification, as the origin and destination have the same sign. These results say that high wage locations have lower flows between them (conditional on size), and high house price locations have higher flows between.

¹⁸ Though the results are little different without this detrending, especially when time dummies are included, as the detrending amounts to the same model as that of time dummies.

¹⁹ Pissarides and Wadsworth (1989) found similar results regarding unemployment using British data from the 1970s and 1980s.

Table 1.1: Baseline Results from the Gravity Model

Variable	OLS						IV											
	1		2		3		4		Local Premia		5		6		7		8	
	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
Cons	-5.615	0.128	-5.980	0.130	-2.169	0.738	-12.463	0.709	-3.587	0.114	-5.870	0.131	-6.381	0.133	-2.666	0.761	-2.461	0.025
Move dummy	-4.429	0.032	-4.396	0.032	-2.464	0.024			-4.955	0.032	-4.436	0.033	-4.391	0.033	-2.461	0.025		
Log mi.	-0.804	0.005	-0.811	0.005	-1.170	0.005			-0.711	0.005	-0.804	0.005	-0.813	0.005	-1.170	0.005		
Orig size	0.734	0.008	0.756	0.008	0.684	0.056		0.672	0.671	0.007	0.741	0.008	0.769	0.008	0.685	0.058		
Dest size	0.714	0.005	0.737	0.005	0.577	0.040		0.594	0.638	0.005	0.726	0.005	0.755	0.005	0.621	0.041		
Orig Unemp Rate	0.041	0.018	0.059	0.017	0.180	0.009		0.181	0.201	0.043	-0.036	0.020	-0.026	0.019	0.153	0.009		
Dest Unemp Rate	-0.126	0.006	-0.105	0.006	-0.179	0.014		-0.178	-0.176	0.016	-0.207	0.007	-0.195	0.007	-0.175	0.016		
Orig Wage	-1.460	0.062	-1.597	0.065	-0.322	0.065		-0.291	-0.743	0.234	-1.523	0.061	-1.734	0.066	-0.350	0.057		
Dest Wage	-1.479	0.022	-1.614	0.020	0.579	0.085		0.614	0.104	0.132	-1.579	0.026	-1.791	0.021	0.824	0.104		
Orig House Price	0.668	0.020	0.671	0.019	0.233	0.014		0.232	0.255	0.085	0.680	0.020	0.696	0.020	0.258	0.016		
Dest House Price	0.627	0.008	0.632	0.008	-0.142	0.028		-0.145	-0.093	0.043	0.630	0.008	0.648	0.008	-0.255	0.030		
Dummies:																		
Time?			y		y		y			y					y			
Orig?					y										y			
Dest?					y										y			
Orig-Dest?					y		y								y			

Both wages and housing prices have strong time trends in many locations, even relative to the national average, and it is already known that migration rates have been gradually declining in the U.S. over the last few decades (see Molloy et al [2011]). Perhaps the negative association of wages to flows is due to the inverse patterns in their respective trends. To test this, column 2 adds time dummies to the gravity model to soak up trends in migration rates. The results are essentially unchanged. The time dummies do pick up the downward trend in mobility, exhibited in Figure 1.4.4, which will be addressed in more detail in a later section.

Column 3 of Table 1.4.1 adds origin and destination fixed effects to the specification. The fixed effects are meant to capture unobserved heterogeneity in the MSAs that, over the period of data, is causing them to be net senders or receivers of population. Note that within a study of mobility, the unobserved heterogeneity is a dynamic heterogeneity; this is conceptually different from some permanent unobserved “quality of life” that causes a metropolitan area to be bigger or smaller than expected conditional on its wages, unemployment risk, housing prices, etc. For example, a favorable climate may allow a city to have a larger market share of population than another city with an identical labor and housing market but an unfavorable climate; but that example is the result some long run spatial equilibrium. The key in this context is that, for the fixed effect to matter, there is something changing which is allowing, e.g., the favorable climate city to continue to absorb more people from the unfavorable climate city. This could include changing amenities, such as an improvement in local schools or development of the arts. The fixed effects will also pick up if a city is more or less mobile than average, for reasons such as industry or occupation composition and potential differences in turnover, a city being a retirement, schooling, early career or immigrant destination, and perhaps many others.

The addition of the fixed effects affects the coefficients of the attract/repel variables, especially the wages and housing prices. These now tell more of a push/pull story: higher origin wages reduce flows out while higher destination wages increase flows in, and higher origin housing prices increase flows out and higher destination

housing prices reduce flows in. Interestingly, the origin and destination unemployment rate estimates are not drastically affected. Comparing column 3 to 2 or 1 suggests that some high wage locations have been net senders, while some low wage locations have been net receivers. Wages and housing prices do seem to affect location decisions, but there has not been a widespread movement of workers from low wage to high wage locations. Clearly, the unobserved heterogeneity in cities as destinations or origins has an important effect on how migrants respond to labor and housing market features of cities.

Column 4 carries the fixed effects strategy one step further, using origin-destination pair dummies. These fixed effects pick up an even greater degree of observed and unobserved heterogeneity (including the distance between two locations). The results on the attract/repel variables are much the same.

In a similar vein, column 5 uses the locally demeaned wages, housing prices and unemployment as the attract/repel regressors. This purges the regression of the more permanent differences in local labor and housing markets—the long run spatial equilibrium—and focuses on the effects of the temporal, cyclical differences. Even without origin and destination fixed effects, the attract/repel relationship of unemployment, wage, and house price *premia* is evident, especially from the origin side. That is, migrants appear to be responding to cyclical variation in these variables, but not simply leaving low wage areas for high. This is further evidence that a Roback-style (1982) spatial equilibrium holds in the long run, but at higher frequency, workers at the margin are responding to labor and housing market dynamics.

The discussion so far has ignored the potential for simultaneity bias in the estimates. One might worry that the gravity model amounts to a classic endogenous regressors error: the regression of labor supply on wage and/or unemployment rates; the concern is similar for housing demand. Fortunately, a simple remedy is available in the usage of lagged values as instruments. The working assumption of the this instrumental variables (IV) strategy is that lagged values are uncorrelated with current unobserved migration; this means that, e.g., wages were not set in anticipation of future migration flows. Columns 6, 7, and 8 repeat the specifications of

columns, 1, 2, and 3, respectively, but using lagged values of origin and destination unemployment rates, wages, and housing prices as instruments for the current values.²⁰ The coefficients rise in magnitude for destination wages and housing prices and origin wages, though not significantly so, and the effects on the unemployment rate coefficients and origin house prices are negligible. Still, the regression in column 8 is preferred in principle.

Table 1.5 at the end of the article repeats Table 1.4.1 using only migrant observations; that is, excluding non-movers from the data. The results suggest that the patterns in Table 1.4.1 are not driven by non-movers. Table 1.5 also includes an additional column of 3.1, which includes non-movers in the data, but allows the destination variables to enter the regression separately for movers through the interaction with the move indicator. Little intuition is gained through this, except to see the apparent endogeneity of the destination house prices.

Before proceeding, consider the magnitudes of the coefficients in Table 1.4.1. While the coefficients on the unemployment rates are smaller, it would be incorrect to read the results as indicating that unemployment rates have only relatively minor effects on migration patterns. As displayed, it appears that wages are the strongest force for attraction or repulsion, followed by housing prices and then unemployment, and in the most straightforward sense, this is so. But also consider the feasibility of large changes in the underlying variables, one of which is a *rate*. Recall that the coefficients in a log-log specification like a gravity model are elasticities. A 100 percent increase in wages—a doubling of incomes—would be a massive change to a labor market, while a 100 percent increase in the unemployment rate (from say, 4 percent to 8 percent) would be a significant but considerably more modest change. Thus, while the coefficient is smaller, in reality the impact of the variable is not.

²⁰ Recalling the timing of the migration data, I use the first quarter values of the variables as instruments; for example, wages in the first quarter of 1990 serve as the instrument for the 1990:q2-1991:q1 annual variable. This allows me to preserve the first year of migration flows as data in the regression.

1.4.2 *Forward-looking Behavior and Migration Momentum*

If moving between cities is a costly and infrequent occurrence, as the rate of migration suggests, it is reasonable to expect agents to be forward-looking when making a migration decision. I examine this hypothesis in the context of the gravity model. We would expect predictors of future growth to be an attractive force for in-migration, conditional on the current attract/repel forces. This is indeed what I find.

Metro areas demonstrate persistence in employment shocks, as evidenced by the persistence of wage and employment growth at the metro level. In the data,²¹ the autocorrelation of employment growth is 0.46, of wages it is 0.88, and of wage growth it is 0.09. Thus wage and employment growth make good candidates for predictors of future labor market conditions.²² In Table 1.4.2, I extend the regressions from Table 1.4.1, columns 3 and 8 (which were, respectively, the OLS and IV models including origin and destination fixed effects) by adding lagged wage and employment growth as attract/repel variables; the use of the lag also sidesteps the issue of simultaneity.

In all specifications, the wage and employment growth variables exhibit attractiveness, especially at the destination. Note that this attractiveness is exhibited on top of the attractiveness of current wages. The coefficients on the destination growth are larger in magnitude, and in the specification of column 5 using migrants only, the coefficients are larger still. This suggests that growth predictors are not merely a blanket attractive force, but are especially important to migrants in choosing their particular destination. The fact that past employment growth is positively related to current migration shows there is some momentum in city growth.

1.4.3 *Housing Lock-in*

In this section I use the gravity model to examine whether declining housing prices are associated with reduced migration, as some literature has suggested. The intuition

²¹ Again, the data are expressed relative to the national average.

²² Even if the autocorrelation of wage growth is low, wage growth can be attractive to potential migrants because the autocorrelation of wages is high, and thus growth that has already occurred will translate to persistently higher wages.

Table 1.2: Regressions Including Lagged Labor Market Growth

Variable	OLS						IV					
	1		2		3		4		5		6	
	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
Cons	-2.003	0.420	-1.882	0.420	-1.911	0.417	-2.524	0.904	-2.443	0.910	-2.442	0.911
Move dummy	-2.475	0.015	-2.475	0.015	-2.475	0.015	-2.474	0.026	-2.474	0.026	-2.474	0.026
Log mi.	-1.170	0.002	-1.170	0.002	-1.170	0.002	-1.170	0.005	-1.170	0.005	-1.170	0.005
Orig size	0.643	0.027	0.649	0.027	0.646	0.027	0.644	0.067	0.654	0.068	0.649	0.068
Dest size	0.595	0.027	0.580	0.027	0.585	0.027	0.644	0.051	0.628	0.051	0.632	0.051
Orig Unempl Rate	0.176	0.008	0.177	0.009	0.172	0.009	0.147	0.010	0.145	0.009	0.136	0.009
Dest Unempl Rate	-0.174	0.008	-0.160	0.009	-0.148	0.009	-0.175	0.017	-0.148	0.018	-0.137	0.018
Orig Wage	-0.235	0.048	-0.262	0.046	-0.235	0.048	-0.215	0.072	-0.281	0.062	-0.211	0.071
Dest Wage	0.456	0.048	0.539	0.046	0.462	0.048	0.751	0.126	0.808	0.115	0.724	0.125
Orig House Price	0.233	0.014	0.235	0.014	0.232	0.014	0.247	0.016	0.252	0.017	0.243	0.017
Dest House Price	-0.142	0.014	-0.144	0.014	-0.135	0.014	-0.261	0.032	-0.253	0.032	-0.241	0.032
Orig Lag Empl Growth			-0.102	0.084	-0.094	0.084			-0.221	0.097	-0.214	0.097
Orig Lag Wage Growth	-0.189	0.086			-0.185	0.086	-0.312	0.084			-0.311	0.084
Dest Lag Empl Growth			0.853	0.085	0.833	0.085			0.881	0.119	0.873	0.119
Dest Lag Wage Growth	0.523	0.086			0.489	0.086	0.370	0.143			0.376	0.143
Dummies:												
Time, Orig, Dest	y		y		y		y		y		y	

of lock-in is simple: housing price declines give homeowners disincentive to move because they must take an asset loss to sell their home. An extended time-on-market for a home (presumably because the price needs to come down even more to find a buyer) functions the same. For leveraged homeowners, the disincentive may be even greater if they owe more on their mortgage than they would receive in a home sale; this is commonly referred to as being “underwater” on the property.

Recently, housing lock-in has received considerable attention in the literature.²³ Much of the concern about lock-in is whether it causes widespread immobility that leads to an inefficient allocation of labor, or spatial mismatch. The data and model of this article are well-suited to examine whether housing price declines have reduced aggregate mobility and created a widespread lock-in effect. The following is not an examination of whether housing price declines constrain an individual’s mobility per se, but whether the recent housing market price bust has diminished the rate of aggregate migration.

The technique is much the same as previous subsections: include variables indicative of lock-in and see whether they act as attract or repel variables in the gravity model. Specifically, an origin experiencing a housing price decline would see less outflow, conditional on the other attract/repel variables in the model. First, there is the matter of operationalizing the lock-in idea in the context of the gravity model. Because much of the concern over lock-in has been driven by the recent housing bust, I use measures that emphasize the period after a location’s housing price has peaked. I use three simple approaches. One is a dummy variable for whether a location is below its peak housing price. Table 1.4.3 below shows that virtually all locations are post their peak price in the late 2000s, though there is some useful variation in precisely when. But since this discrete variable may lump some slightly affected locations in with the critically-affected, the other two variables are continuous measures of the size of the decline: one in percentage from the peak, and the other in the magnitude of the decline in (log) dollars. These latter two measures are also

²³ See, for example, Chan (2001); Ferreira et al. (2010); Karahan and Rhee (2011); Winkler (2011). Conversely, Molloy et al. (2011) find lock-in to have little impact on mobility.

summarized in Table 1.4.3.

Table 1.4.3 reports the results from including these lock-in measures as attract/repel variables in the gravity model. The baseline specification is from Table 1.4.1, column 6—time, origin, and destination dummies; instrumented wages, unemployment, and housing prices; and lagged wage and employment growth. Table 1.4.3, column 1 adds to the baseline the post peak dummy the origin side. The coefficient on the post-peak dummy is actually positive, indicating no diminished mobility resulting from starting in a location with a housing price decline. Columns 2 and 3 repeat the specification with the continuous measures of housing price decline. Here again, there is no indication of lock-in.

Columns 4, 5 and 6 use the three measures of lock-in on both the origin and destination sides. Interestingly, each of the destination side measures exhibit a negative sign; that is, locations experiencing a housing price decline are less attractive destinations. This indicates some impact on mobility, though not in a way directly consistent with a lock-in story, and thus the intuition behind the result is not obvious. The result may indicate that a housing bust has some impact on the expectations of forward-looking migrants, or it is perhaps even the result of some uncorrected simultaneity. Columns 7, 8 and 9 repeat the analysis of columns 4-6, but using only migrants (non-movers excluded from the data). The limiting of the sample has virtually no effect on the results.

The final column of Table 1.4.3 includes both the post peak dummy and the log magnitude of the decline. It is only here that the results indicate some immobility caused by a housing price decline in the origin location: while the post peak dummy is still positive, the coefficient on the price decline is negative. However, for a net negative effect on mobility at these point estimates, the origin must be experiencing a price loss of 11.14 log points, or about \$69,000. The coefficient on the percentage decline measure (not reported) was essentially zero.

Thus, the model, data, and empirical strategy employed here provide scant evidence of widespread diminished mobility, or lock-in, due to housing price declines.

Table 1.3: Summary of Lock-in Measures

Migration Year	Cities Peaked	Cities Post Peak	Percent Decline Post Peak			Dollar Decline Post Peak			
			Mean	SD	Max	Mean	SD	Max	
1990-2002	6	6	0.13	0.08	0.01	12,576	7,240	649	23,026
2003	10	16	0.05	0.06	0.01	5,357	6,063	483	12,239
2004	21	37	0.03	0.03	0.01	2,744	2,764	790	10,584
2005	31	68	0.03	0.02	0.01	2,708	1,762	640	10,339
2006	131	199	0.04	0.02	0.01	7,377	5,699	1,315	29,593
2007	9	208	0.06	0.05	0.01	12,889	18,703	583	122,774
2008	63	271	0.14	0.14	0.03	28,422	40,662	2,558	238,558

Table 1.4: Regressions Including Measures of Lock-in

Variable	1			2			3			4			5			6			7			8			9			10						
	Est	SE		Est	SE		Est	SE		Est	SE		Est	SE		Est	SE		Est	SE		Est	SE		Est	SE		Est	SE					
Cons	-2.50	0.92		-2.43	0.91		-2.49	0.92		-2.29	0.92		-2.49	0.91		-2.25	0.92		-4.72	0.94		-4.92	0.94		-4.68	0.94		-2.23	0.92		-2.47	0.94		
Move dummy	-2.47	0.03		-2.47	0.03		-2.47	0.03		-2.47	0.03		-2.47	0.03		-2.47	0.03		-1.19	0.01		-1.19	0.01		-1.19	0.01		-1.17	0.01		-1.17	0.01		
Log mi.	-1.17	0.01		-1.17	0.01		-1.17	0.01		-1.17	0.01		-1.17	0.01		-1.17	0.01		0.66	0.07		0.65	0.07		0.66	0.07		0.66	0.07		0.66	0.07		
Orig size	0.66	0.07		0.65	0.07		0.65	0.07		0.66	0.07		0.65	0.07		0.66	0.07		0.66	0.07		0.65	0.07		0.66	0.07		0.66	0.07		0.66	0.07		
Dest size	0.63	0.05		0.63	0.05		0.63	0.05		0.61	0.05		0.64	0.05		0.61	0.05		0.62	0.05		0.64	0.05		0.61	0.05		0.61	0.05		0.61	0.05		
Orig Unempl Rate	0.13	0.01		0.14	0.01		0.13	0.01		0.13	0.01		0.13	0.01		0.13	0.01		0.13	0.01		0.14	0.01		0.13	0.01		0.13	0.01		0.13	0.01		
Dest Unempl Rate	-0.14	0.02		-0.14	0.02		-0.14	0.02		-0.13	0.02		-0.13	0.02		-0.13	0.02		-0.12	0.02		-0.13	0.02		-0.12	0.02		-0.12	0.02		-0.12	0.02		
Orig Wage	-0.20	0.07		-0.21	0.07		-0.20	0.07		-0.20	0.07		-0.21	0.07		-0.20	0.07		-0.19	0.07		-0.20	0.07		-0.19	0.07		-0.20	0.07		-0.20	0.07		
Dest Wage	0.73	0.12		0.72	0.12		0.73	0.12		0.68	0.13		0.68	0.13		0.68	0.13		0.70	0.13		0.70	0.13		0.70	0.13		0.68	0.13		0.68	0.13		
Orig House Price	0.23	0.02		0.24	0.02		0.23	0.02		0.23	0.02		0.24	0.02		0.23	0.02		0.23	0.02		0.24	0.02		0.23	0.02		0.25	0.02		0.25	0.02		
Dest House Price	-0.24	0.03		-0.24	0.03		-0.24	0.03		-0.21	0.03		-0.21	0.03		-0.21	0.03		-0.21	0.03		-0.21	0.03		-0.20	0.03		-0.20	0.03		-0.20	0.03		
Orig LEGR	0.87	0.12		0.87	0.12		0.87	0.12		0.88	0.12		0.83	0.12		0.87	0.12		0.89	0.12		0.84	0.12		0.88	0.12		0.87	0.12		0.87	0.12		
Orig LWGR	0.38	0.14		0.38	0.14		0.38	0.14		0.36	0.14		0.36	0.14		0.36	0.14		0.36	0.15		0.35	0.14		0.35	0.14		0.35	0.14		0.35	0.14		
Dest LEGR	-0.22	0.10		-0.21	0.10		-0.21	0.10		-0.21	0.10		-0.21	0.10		-0.21	0.10		-0.20	0.10		-0.20	0.10		-0.20	0.10		-0.22	0.10		-0.22	0.10		
Dest LWGR	-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.31	0.08		-0.32	0.08		-0.32	0.08		
Orig Post. Peak	0.02	0.01					0.02	0.01		-0.05	0.01					-0.05	0.01		-0.05	0.01								0.04	0.01		0.04	0.01		
Dest Post. Peak				0.02	0.04																	0.03	0.04											
Orig Pct HP Loss																																		
Dest Pct HP Loss																																		
Dest HP Loss																																		
Dummies:																																		
Time, Orig, Dest	y			y			y			y			y			y			y			y			y			y			y			

²⁴LEGR=lagged employment growth; LWGR=lage wage growth




FIGURE 1.1: Estimates of Time Dummies From Various Specifications

1.4.4 Trends in Migration Rates

Other researchers have noted that migration rates in the U.S. have declined over the last few decades without a satisfying explanation (Molloy et al. [2011]). The empirical analysis of this article arrives at a similar conclusion. Figure 1.4.4 plots the time dummies from several specifications of the gravity model (normalized to the same value in 1991). The downward trend in migration rates is evident. What is notable is that the trend is present even after controlling for the each location's time variation in labor and housing markets, including measures of expectations of labor market growth and potential housing lock-in. Thus, the analysis herein can only continue to offer the negative result that the secular trend remains unsatisfactorily explained.

I also note that accounting for the downward trend in migration was critical in the previous analysis of housing lock-in. Because virtually all locations' housing prices peaks and declined in the latter years of the data, there is substantial correlation between the lock-in measures and a time trend. Failing to account for the time trend would yield results indicating the presence of housing lock-in. Thus it is even more

remarkable that the trend in migration rates seems independent from the housing bust.

1.5 Conclusion

This article applied a version of the gravity model of population migration to migration flow data set including temporal variation. Therefore, the empirical strategy leveraged both temporal and spatial variation in location attributes. Results indicate that wages, wage growth, and employment growth act as attractive forces, while unemployment and high housing prices act as repulsive forces, but the result is clear only after controlling for unobserved location heterogeneity, which is made possible by the temporal variation in the flow data. The attraction of wage and employment growth suggests forward-looking behavior on the part of migrants.

A further advantage of the time-varying data is the ability to evaluate the importance of temporal shocks to locations. In particular, I examined whether housing price declines constrain mobility. The data provide only weak evidence of diminished mobility, and if present at all, only in extreme declines.

In future work, the gravity model of this paper could be extended to further exploit the time variation in the flow data. Potential extensions include time-varying attract/repel coefficients and an analysis of network effects—whether past migration flows encourage future flows along the same origin-destination path. Other extensions could include more location attributes, such as time varying amenities, and more sophisticated specifications of distance between locations that account for industry, occupational, or amenity space in addition to actual physical space.

Finally, the importance of locational fixed effects and the indications of forward-looking behavior together suggest that notions of spatial equilibrium are extant, yet perhaps in a dynamic setting. I began this article by noting the disconnect between structural models of individual migration and inter-regional studies of aggregate migration. A structural equilibrium model of migration would further close the gap in these literatures.

Table 1.5: Alternative Distance Functions

Variable	1		2		3		4		5		6		7		8		9		10	
	Est.	SE	Est.	SE	Est.	SE	Est.	SE	Est.	SE	Est.	SE	Est.	SE	Est.	SE	Est.	SE	Est.	SE
Cons dummy	-1.66	0.05	-11.34	0.05	-3.90	0.04	-12.83	0.05	-5.38	0.04	-2.19	0.52	-10.90	0.59	-2.44	0.43	-9.40	0.52	-2.12	0.39
Miles	-9.41	0.02	0.00	0.00	0.00	0.00	0.00	0.00	-8.22	0.02	-9.46	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Miles Sqrd	0.58	0.00	0.64	0.00	0.68	0.00	0.72	0.00	0.73	0.00	0.68	0.03	0.69	0.04	0.69	0.03	0.69	0.03	0.68	0.03
Orig size	0.55	0.00	0.60	0.00	0.65	0.00	0.70	0.00	0.71	0.00	0.57	0.03	0.58	0.04	0.58	0.03	0.57	0.03	0.57	0.03
Dest size	0.03	0.01	0.16	0.01	0.12	0.01	0.06	0.01	0.05	0.01	0.18	0.01	0.18	0.01	0.18	0.01	0.18	0.01	0.18	0.01
Orig Unempl Rate	-0.10	0.01	0.02	0.01	-0.03	0.01	-0.08	0.01	-0.10	0.01	-0.18	0.01	-0.18	0.01	-0.18	0.01	-0.18	0.01	-0.18	0.01
Dest Unempl Rate	-0.82	0.02	-1.17	0.02	-1.23	0.02	-1.58	0.02	-1.50	0.02	-0.32	0.06	-0.32	0.07	-0.32	0.05	-0.32	0.06	-0.32	0.04
Orig Wage	-0.76	0.02	-1.10	0.02	-1.19	0.02	-1.57	0.02	-1.51	0.02	0.58	0.06	0.58	0.07	0.58	0.05	0.58	0.06	0.58	0.04
Dest Wage	0.30	0.01	0.73	0.01	0.68	0.01	0.59	0.01	0.59	0.01	0.23	0.02	0.23	0.02	0.23	0.02	0.23	0.02	0.23	0.01
Orig House Price	0.26	0.01	0.70	0.01	0.64	0.01	0.54	0.01	0.54	0.01	-0.14	0.02	-0.14	0.02	-0.14	0.02	-0.14	0.02	-0.14	0.01
Dest House Price																				
Dummies:																				
Time?	y		y		y		y		y		y		y		y		y		y	
Orig?																				
Dest?																				

Table 1.6: Results from the Gravity Model Using Migrants Only

Variable	OLS						IV											
	1		2		3		4		5		6		7		8		3.1	
	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE	Est	SE
Cons	-10.205	0.126	-10.550	0.128	-4.576	0.758	-12.672	0.723	-8.622	0.111	-10.474	0.129	-10.957	0.131	-5.072	0.781	-2.425	0.350
Log mi.	-0.809	0.005	-0.816	0.005	-1.189	0.005	0.000	0.000	-0.713	0.005	-0.809	0.005	-0.819	0.005	-1.190	0.005	-1.175	0.002
Orig size	0.741	0.008	0.764	0.008	0.687	0.057	0.676	0.057	0.675	0.007	0.748	0.008	0.777	0.008	0.688	0.059	0.686	0.023
Dest size	0.722	0.005	0.746	0.005	0.581	0.041	0.599	0.011	0.642	0.005	0.734	0.005	0.764	0.005	0.625	0.042	0.579	0.023
Orig Unempl Rate	0.039	0.019	0.057	0.017	0.180	0.009	0.181	0.009	0.202	0.044	-0.041	0.021	-0.030	0.019	0.153	0.009	0.179	0.008
Dest Unempl Rate	-0.130	0.007	-0.108	0.006	-0.179	0.014	-0.178	0.003	-0.175	0.016	-0.213	0.007	-0.200	0.007	-0.175	0.016	-0.099	0.035
Orig Wage	-1.504	0.063	-1.645	0.066	-0.315	0.066	-0.282	0.067	-0.747	0.237	-1.569	0.063	-1.786	0.068	-0.339	0.058	-0.320	0.040
Dest Wage	-1.526	0.022	-1.666	0.020	0.586	0.086	0.624	0.020	0.102	0.134	-1.628	0.026	-1.847	0.021	0.836	0.105	0.715	0.096
Orig House Price	0.684	0.020	0.688	0.019	0.233	0.014	0.233	0.014	0.257	0.086	0.697	0.021	0.714	0.020	0.257	0.016	0.249	0.013
Dest House Price	0.644	0.008	0.649	0.008	-0.141	0.029	-0.144	0.006	-0.091	0.044	0.647	0.008	0.666	0.008	-0.255	0.030	-1.876	0.037
Dummies:																		
Time?			y		y		y		y						y		y	
Orig?					y										y		y	
Dest?					y										y		y	
Orig-Dest?																		
Move Dummy																		
Mig X Dest UR																	-2.217	0.019
Mig X Dest Wage																	-0.081	0.034
Mig X Dest HP																	-0.135	0.088
																	1.751	0.035

A Dynamic Equilibrium Model of Cities and Labor Market Dynamics

2.1 Introduction

Most economic models of the labor market abstract from geographic space. In reality the national labor market consists of many local labor markets, or cities. While there is a significant national component to business cycles, much of the variation in labor market shocks is local. When cities are hit by heterogeneous shocks, the efficient matching of workers to jobs requires the reallocation of people across space, but spatial arbitrage is far from complete. At any point in time, the dispersion in unemployment rates across cities is considerable—among the 30 largest U.S. cities the difference between the minimum and maximum unemployment rate has averaged 4.8 percentage points over the past twenty years—and this dispersion is especially high in recessions. In 2009, for example, unemployment differences among the 30 largest cities was as high as 9.6 percentage points.

This paper seeks to understand how the spatial organization of the economy in geographically distinct markets affects labor market dynamics. The main goal

of the analysis is to quantify the impact of frictions that hinder the reallocation of workers across cities on the employment dynamics of individual cities and the national economy. There are a number of reasons that workers collectively might not fully arbitrage away spatial differences in labor market opportunities. First, migrating to a new city is costly in financial, psychological, and social terms. For migration to be optimal to workers, labor market differentials would need to be substantial enough to outweigh these costs. Bearing these migration costs may simply not be worth the trouble, especially if the labor market differentials are not expected to persist. Similarly, strong preferences for the many other features and amenities of their city might keep people from moving despite the availability of better labor market conditions elsewhere.

A second potential source of friction is the fact that the prices of housing and many other goods and services depend on local demand. The actual impact of a labor market shock on consumption may be mitigated to the extent that these prices are responsive to local labor market conditions. For example, the hit that workers take to consumption following a negative labor market shock might not be nearly as large as the hit to their wages. How responsive local prices are to labor market shocks will generally depend on how elastically these goods are supplied in the city, a feature which can potentially vary significantly across locations. Furthermore, because housing is durable and immobile across cities, a drop in local demand will generally result in a sharp decline in the price (rather than quantity) of housing. More generally, we would expect the prices of any local goods that depend on durable and immobile capital to be quite responsive to negative shocks to local labor market conditions. In this way, the price of housing and other local goods works as a sort of “pro-cyclical” migration cost, reducing the migratory response of workers to geographic variation in labor market opportunities. Thus I identify two types of friction—one that is a direct result of geographic distance, and one that is an indirect

result of distance causing some goods to be exclusively locally consumed.

The spatial fragmentation of the labor market is particularly interesting vis-a-vis other segmentations of the labor market (e.g., industries and occupations) because many of the impediments to spatial reallocation are not directly related to the labor market itself. Instead, the frictions described above affect how workers choose the local labor market in which to participate. Therefore, the term space should be interpreted broadly. At its core, the concept of space is merely the organization of the economy into distinct units, with local markets receiving their own productivity shocks and clearing according to local equilibrium conditions. Underlying the separation of markets, however, is that utility-maximizing agents face these other non-labor market inducements which, while rational from the point of view of the agent, can amount to frictions to the labor market. Otherwise, space would not matter, and the economy would be one unified market.

Whether directly tied to the labor market or not, impediments to mobility could have important implications for macroeconomic adjustment. Blanchard and Katz (1992) argue that labor mobility—not firm mobility or firm entry and exit—is the primary mechanism for adjustment to regional business cycles. Similarly, Topel (1986) finds worker mobility in response to local labor demand shocks to be significant in explaining wage differentials.

The particular goal of this paper is to quantify and compare the importance of the two main forms of friction highlighted above, the migration costs/location preferences versus the prices of local goods. How much of the spatial dispersion in unemployment is owed to these frictions?

To examine this question, one needs a model of how workers choose the markets in which they will supply labor, subject to the non-labor market incentives already described. This paper builds a dynamic spatial equilibrium model of the national labor market. In the model, workers choose one market at a time, thereby subjecting

themselves to the local labor demand and housing market conditions at the time. The model setting is a system of a finite number of local labor markets (in practice, MSAs), with each “island” defined by its own attributes and its point in space, receiving its own process of productivity shocks. Workers face relocation costs, so their optimal location decision is a dynamic problem. Workers also face congestive interactions through competition in the labor and housing markets and therefore are forced to take into account the equilibrium actions of other workers. Local heterogeneity is modeled by variation across locations in worker productivity, firm-worker matching efficiency, and the elasticity of housing supply. I structurally model the worker’s location decision, accounting for relocation costs, a worker’s expectations of the future paths of location attributes, and congestive interactions in the labor and housing market. Employing a dynamic discrete choice framework allows me to model the spatial cross-section and permits a rich characterization of location heterogeneity to be linked directly with city-level data.¹ A novel equilibrium computation routine allows for feasible estimation and simulation of the model. The model is then simulated under alternative scenarios in which relocation costs are eliminated and which the housing market is untied from the labor market.

The modeling antecedents of this paper are the island economy model of Lucas and Prescott (1974) and the canonical work in spatial equilibrium, Roback (1982). Importantly, Roback’s model is a static equilibrium in which workers face no costs to relocation. The literature on migration (such as Kennan and Walker (2011) and Bishop (2008)) has found that while labor market outcomes and amenities affect location decisions, relocation comes at significant cost in terms of utility. This is needed to rationalize why inter-city moves are observed so infrequently, with only a small fraction of households relocating within a year. This paper extends the

¹ This framework is much richer than a classic urban economics “toy” two-location model of a city and reservation locale.

spatial equilibrium model of Roback (1982) to a dynamic context where workers face adjustment costs to relocation.

An extension this paper makes over the individual migration literature, however, is the congestive interactions of agents in the equilibrium framework. This allows me to distinguish between the simple inertia effect on moving costs—that people face psychic or monetary costs to changing location—from the equilibrium feedback, such as a locally priced good, that mitigates locational differentials. If workers face competition in the labor or housing markets, then the locational differential is smaller in general equilibrium than in partial equilibrium. While the partial equilibrium may be sufficient for individual migration studies, a study of aggregate population flows such as this must take the equilibrium seriously. To my knowledge, there is no other equilibrium model of costly migration.

These two features together, the moving costs and general equilibrium setting, allow me to compare two impediments to labor reallocation: generic migration costs and the link between labor and housing markets. After building and estimating the model, the main exercise of this paper is to compare two counterfactual scenarios, one in which workers face no migration costs, and another in which the link of the housing market to the labor market is relaxed. I find that, in equilibrium, the housing market forms more of an impediment to labor allocation than generic migration costs. When the housing market link is relaxed, labor supply is more responsive to local labor market conditions, unemployment differences between cities are reduced, and cities more closely follow the national cycle. Spatial arbitrage of labor is increased by about one-third. Reducing migration costs has a smaller impact, but primarily affecting the tails of the unemployment distribution, with no broad effects on the spatial arbitrage of labor. That is, the reduction in migration costs moves the outlying high and low unemployment cities toward the national average, but does not induce broader convergence in unemployment rates.

The key difference between the two experiments is that housing prices form a pro-cyclical moving cost.² When a location receives a positive labor market shock, housing prices increase, reducing the relative differential in local real wages; when there is a negative shock, the decline in local housing prices mitigates the effect on utility. Thus relaxing the housing market link provides margin for reallocation in any state of the world, not just when a few locations are outliers.

This finding relates to the growing literature on the effect of housing on labor allocation, much of which has focused on the “lock-in” effect of underwater mortgages (see, for example, Chan (2001), Ferreira, Gyourko, and Tracy(2010), Davis, Fisher, and Veracierto (2010), Karahan and Rhee (2011)). The housing market friction of this paper is more general, as the effects of lock-in can occur (though not explicitly through a negative housing equity channel), but “lock-out” can also occur as rising housing prices in growing markets form a barrier to entry. The lock-out feature of housing in this model is consistent with the findings of Saks (2008), Sinai and Souleles (2009), Gyourko, Mayer, and Sinai (2005), and Van Nieuwerburgh and Weill (2010). Saks (2008) finds that MSAs with higher degrees of land use regulation exhibit smaller employment changes in response to labor demand shocks. Sinai and Souleles (2009) find that when households do relocate, they tend to move between correlated housing markets, which is consistent with the structure of the model of this paper, where workers evaluate locations in terms of consumption differentials. The latter two papers relate increasing cross-city dispersion in housing prices to increasing dispersion in the labor market. A key result of this paper is that if housing was unlinked from the labor market, the cross-sectional dispersion in unemployment rates would decrease.

That housing appears like a pro-cyclical moving cost also relates to the insights

² As Chan (2001) aptly puts it in the context of homeownership, “A great degree of risk is introduced to the homeowner’s portfolio as there is positive covariance between its most important assets: human capital and housing capital.”

of Glaeser and Gyourko (2005) and Notowidigdo (2010). Glaeser and Gyourko (2005) suggest that the irreversibility of housing contributes to urban decline, as the “kinked” housing supply curve, due to the fixed and durable nature of housing, results in large price effects in periods of declining demand. Notowidigdo (2010) extends this, finding that falling housing prices are part of the apparent high moving costs to low-skill workers in declining labor markets. The simulation exercises of this paper provide support for the intuition of these studies, suggesting that the immobile and durable nature of infrastructure is an important part of local growth and especially decline.

Ultimately, the goal of this line of research is to better understand what constitutes a friction to the economy. If different types of spatial frictions affect the economy in different ways, conclusions about the efficiency of the economy and prescriptions for policy will depend on the nature of the friction. This paper is a first step in classifying and quantifying spatial frictions in the labor market.

The rest of the paper proceeds as follows. Section 2 briefly characterizes unemployment dispersion in the U.S. Section 3 describes the model, starting with the equilibrium conditions within an arbitrary market, and then describing the worker’s optimal location problem. A section briefly describing the computational routine follows. Section 5 describes the empirical implementation of the model. Section 6 conducts the model’s simulation under alternative scenarios and discusses the results. Section 7 concludes. An appendix provides further details on the computational routine.

2.2 Motivation

This section characterizes the degree of spatial dispersion in unemployment in the U.S. The data, explained in more detail in section 5, are quarterly observations of metropolitan statistical area (MSA)-level unemployment rates for the period 1990-




FIGURE 2.1: Unemployment Rates in MSAs, Quarterly

2010, obtained from the Bureau of Labor Statistics. I focus on the 29 largest MSAs.

Figure (2.1) is a scatterplot of quarterly unemployment rates in the 29 largest MSAs; each blue point is a city-quarter observation, and the black line is the national average. Clearly, there is a national component to the movement of MSA-level unemployment rates over time, with the points moving in swarm with the national average. But the thickness of the band indicates there is also a significant degree of dispersion at all points in the business cycle.

Figure (2.2) displays the quartiles of the cross-sectional distribution of the detrended unemployment rate (local unemployment less the national average). While the dispersion clearly varies over the business cycle, observe that the interquartile range is as high as 2.5 percentage points and rarely smaller than 1 percentage point.

The next figure highlights the dispersion/business cycle relationship alluded to earlier. Figure (2.3) plots the cross sectional standard deviation in unemployment to the national average unemployment rate. Periods of higher national unemployment are associated with more dispersion between cities. Comparing the dispersion of the locally demeaned unemployment rate retains this pattern. This suggests that national downturns in employment impact some markets more than others.

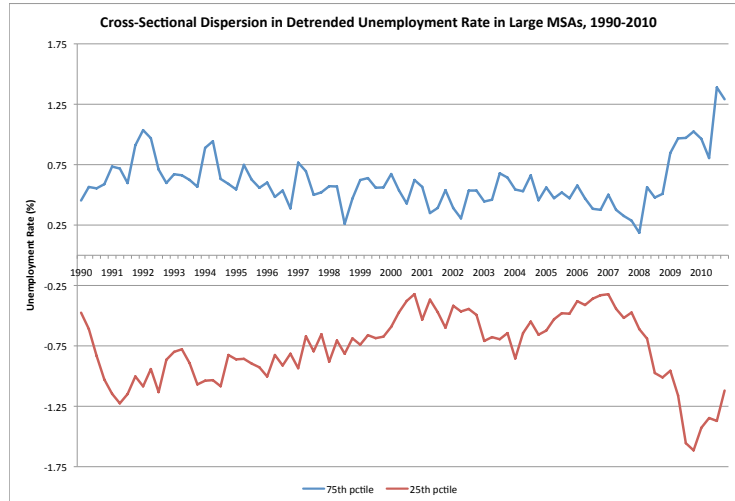


FIGURE 2.2: Dispersion in Cross-Sectional Distribution of Detrended Unemployment

The previous two figures demonstrated that the magnitude of unemployment dispersion varies over time. The next question is whether this dispersion is the result of transitory or permanent differences in local unemployment rates. A plausible argument derived from Roback (1982) could be that permanent local differences in unemployment are necessary (for whatever reason), and that wages, amenities, or

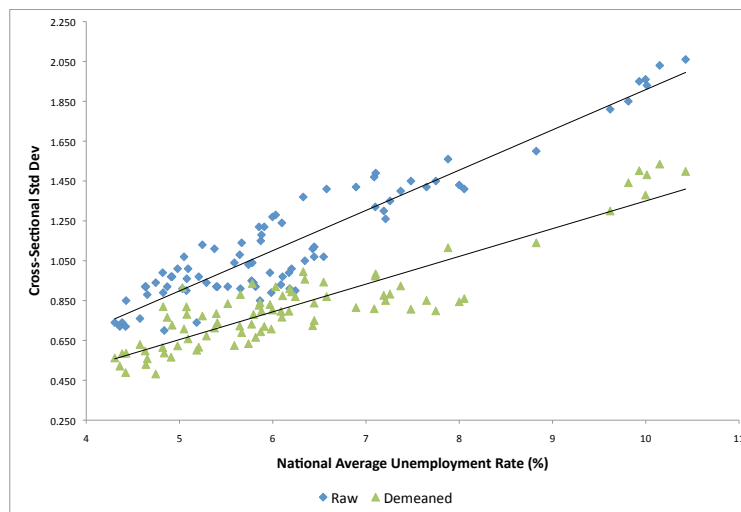


FIGURE 2.3: Cross-Sectional Standard Deviation of Unemployment to Average Unemployment




FIGURE 2.4: Unemployment Rates in MSAs, Quarterly

housing prices compensate workers for this labor market disamenity. However, the next figure illustrates that there are both permanent and transitory components. Figure (2.4) plots the detrended unemployment rate for selected MSAs. Observe that some locations, such as Detroit and Los Angeles, are near or below the national rate at some points while well above it at others. Some cities, such as Chicago and Pittsburgh, oscillate around the national average, and some, most notably Washington DC, are consistently below the national rate. While there are likely varied explanations for the behavior of the individual cities, Figure (2.4) suggests that some local cycling is present, and the dispersion is not entirely due to permanent local differences in unemployment rate. Therefore, the posited “compensating differential” explanation is not especially persuasive.

Permanent differences between cities are much more visible in wages, however, suggesting that the compensating differential argument is extant in some dimensions of the labor market. Figure (2.5) plots the relative wage series for the 29 largest MSAs. The stratification among cities is evident, and there is little turnover. Perhaps it is wages that compensate workers for permanent differences in amenities, while unemployment plays the role of shorter-run adjustment mechanism.




FIGURE 2.5: Wages in 29 Largest MSAs, Local Difference to National Average, Quarterly, 1990-2010

The final exhibit in this section summarizes the above findings and lays the foundational motivation for this paper. Table (2.1) reports the R^2 statistics from pooled regressions of local unemployment rates or wages on sets of dummies. Times dummies pick up variation that is due to economy-wide effects—the national business cycle. The first row of Table (2.1) shows that this type of variation is more prevalent in unemployment than wages. MSA dummies pick up variation that is due to permanent differences between cities. The second row of Table (2.1) shows that this type of variation is more prevalent in wages than unemployment. When both sets of dummies are included, the remaining variation is MSA by time, or local non-permanent variation. It is this local/temporal variation that is of primary interest for this paper. Under complete spatial arbitrage of labor supply, this variation would not exist. The last row of Table (2.1) demonstrates that a substantial amount of local/temporal variation remains for unemployment, though not for wages.

The preceding exhibits have illustrated the spatial dispersion in the unemployment rate, that it varies over time, and that at least some of the local variation is transitory. I now turn to an examination of why labor mobility does not arbitrage

Table 2.1: R^2 from Regressions of Quarterly Local Labor Market Statistics on Time and MSA Dummies

Left-side Variable:	Unemployment	Wage
Time dummies	.64	.21
MSA dummies	.16	.73
Time and MSA dummies	.82	.96

away the spatial variation in unemployment. The next section describes a spatial equilibrium model of the national labor market.

2.3 Model

2.3.1 *The Setting*

Time is discrete and the horizon infinite. The economy consists of a finite number of “islands,” or local markets. There are J locations, with j referring to an arbitrary single location. A labor market is defined as the area in which a person is constrained to live and work within one period of time.³ There is a continuum of identical, infinitely-lived, risk-averse workers distributed about the local markets, with a national labor force population of measure \mathcal{P} .⁴ The labor force population at a location j is denoted p_j ; $\sum_j p_j = \mathcal{P}$. A worker lives and supplies labor inelastically in one market in each time period t . Each market has its own exogenous processes for productivity which govern wages and employment rates.

The worker’s objective is to maximize her expected lifetime utility through optimal choice of labor market. When choosing a location, all workers have full information about the present state of the economy, which is described by the labor force distribution from the end of last period and the current productivity states

³ As is standard (see, e.g., Van Nieuwerburgh and Weill (2010)), the empirical counterpart is a Metropolitan Statistical Area (MSA).

⁴ In this model, population and labor force are equivalent.

at each location. Workers know the distributions of all shocks to the economy and can form expectations of the values of future state variables. Each market exhibits congestion through prices; that is; workers prefer a market with fewer other workers, all else equal. The population, employment rate, and housing costs of each location are determined endogenously.

The model's timing is as follows. At the beginning of each period, the labor demand of each location is revealed. Workers then simultaneously choose the market in which they will live and work for the period, paying a moving cost if they relocate. Then in their chosen location, workers meet firms through a matching process, and the employed produce output and receive wages. All pay the local housing cost and consume the remainder of their wage. The timing then repeats.

2.3.2 The Local Labor Market

Within a location in each period, workers are matched to jobs created by firms through a one-shot matching process, as in Ahn, Arcidiacono, and Wessels (2011). This matching process which I now detail determines the employment and wage rates for a local labor market in a given time period.

Firms

In each location there are a large number of identical firms. The production technology is linear in the sole input, labor, and output is sold on a perfectly competitive world market. Production is either consumed within the economy (with a price normalized to one), used to construct housing, or sold outside the economy. The output of each worker is y , an exogenously determined worker productivity parameter.

To hire labor, firms enter the matching process by posting vacancies. In order to enter, firms must pay an employment cost, specified as a per-vacancy cost of $c(v)$, a function which is increasing in the total number of vacancies, $c'(v) > 0$. The

employment cost is a search friction meant to capture all non-wage costs to production, such as the recruitment, training, and management of workers, production fixed costs, worker benefits, and local regulatory compliance. That the per-vacancy cost is increasing in the number of employed workers reflects convexity in non-wage costs. Thus, even with constant returns to scale production and matching technology, firm profitability exhibits decreasing returns. The employment cost function is constant over time,⁵ and within a location, the same for all firms. The employment cost is determined by the equilibrium number of posted vacancies.

The purpose of the convex employment cost feature is to pin down city size. It is a necessary feature for congestion in the labor market, because without some form of decreasing returns to scale,⁶ the employment rate is nondecreasing in the number of workers, and the size of a city is unbounded. The necessity of this feature distinguishes this model of matching in multiple local markets from a national matching model.⁷ Secondly, it provides a simple rationalization for the significant variation in the average size of labor markets (e.g. the New York market employs many more than Boston).

The two costs to the firm are the employment cost and the wage it pays to the workers with whom it matches. The firm finds a match with probability ρ_f . When a firm matches with a worker, the firm receives the output from the match, y , and pays the worker a wage w . The expected profit of a firm posting a vacancy is then:

$$\pi = \rho_f(y - w) - c(v) \tag{2.1}$$

With a large number of potential entrants, firms will enter the market by posting vacancies, decreasing ρ_f and increasing $c(v)$ until expected profit is driven to zero.

⁵ Temporal variation in the efficiency of the matching process will be introduced in another channel.

⁶ Recall that I assumed the production technology is linear in labor input.

⁷ For a survey, see Petrongolo and Pissarides (2001).

This zero profit condition will factor importantly into the empirical implementation of the local market's matching process.

Workers and Wages

A worker enters the matching process by locating in a market. The total number of available workers in a location is p , the local labor force population. The worker's value from unemployment is normalized to 0, so a worker finding a match will accept any strictly positive wage. As is common in the search literature,⁸ wages are determined according to a Nash bargaining rule in which the worker receives a share $\delta \in (0, 1)$ of her output, or

$$w = \delta y \tag{2.2}$$

All workers are equally productive, so $w = \delta y$ for every match. The value to unemployment is normalized to 0 for all workers, locations, and time periods, so workers will accept any match they receive from the search process.

Matching

The matching technology is specified as a constant returns-to-scale function with labor supply elasticity parameter α :

$$n = m(v, p, \alpha) \tag{2.3}$$

The number of employed (matched) workers, n , is increasing in both the available vacancies, v , and the available workers, p , with the elasticity of employment to each argument determined by α . Workers not finding a match are left unemployed for the period. The resulting employment rate (the worker's match probability) is $er = \frac{n}{p} = \frac{m(v,p,\alpha)}{p}$. The firm's match probability is $\rho_f = \frac{m(v,p,\alpha)}{v}$.

⁸ See, for example, Pissarides (1985).

I will refer to the parameter α as the “matching efficiency” because it will determine how responsive the number of matches is to the number of available workers. This is a key determinant of the employment rate, and it governs how quickly the labor market congests. A better matching state is one in which an available worker very likely translates into a job match; hence, a more efficient matching state (high α , or greater elasticity with respect to p) is one in which the number of job matches is determined mostly by the number of available workers. The matching efficiency may vary between locations and over time within a location.⁹

Locational Heterogeneity

The two types of labor market states, labor productivity and the matching efficiency, are the sources of temporal variation in the labor market. They are also the sources of locational heterogeneity in the model. Aggregate demand may drive temporal changes in labor productivity, but cross-sectional variation in labor productivity due to, for instance, differences in the local industry mix, will result in varying wage movements in different locations. Furthermore, differences in the matching efficiency can yield heterogeneous employment responses to labor productivity shocks. A location with lower matching efficiency, reflecting for instance, search frictions or local labor regulations, will exhibit less employment growth for a given level of labor productivity.¹⁰ Some locations may experience variation in labor productivity but not match efficiency and will therefore exhibit more wage variation than employment variation; in other locations, the opposite may occur. The intuition is that two dimensions of labor states allow locations to exhibit wage volatility but not

⁹ Recall that the employment cost function may also differ between locations but in contrast to the matching efficiency, it is assumed constant over time.

¹⁰ One can see this by looking at the zero profit condition of the firm. Less efficient matching will drive down the ρ term more quickly, so for a given level of match profitability ($y - w$), firms will be less willing to open vacancies.

employment volatility, and vice versa.¹¹

2.3.3 *Housing and Rents*

The local market is defined as the area in which a person lives and works in a period. Each member of the labor force population in a market area consumes one unit of housing for the period, paying the local rental rate. Housing quality is assumed constant over time and space.

Housing demand is given by the inverse demand function $h_d = D(q_d, I)$, where the price h_d depends on the quantity demanded and the income of the local population, I . Thus, demand shifts as population flows in or out of the location and as workers in a location have more or less income to bid for housing. Each location has a large number of real estate firms that supply housing units of constant quality. Each period, firms construct housing units using firm production as input and rent the units to the local population. The market is perfectly competitive, so the real estate firms earn zero profits, and the housing supply curve is the marginal cost curve of housing production. Because land in a location is in fixed supply, the supply curve is assumed not to shift. The inverse supply function is $h_s = S(q_s, \tau_j)$, where τ is a location-specific parameter that determines the slope of the housing supply function.¹² The market clearing condition is $q_d = q_s$. Combining the supply¹³ and demand functions yields $h_d = D(S^{-1}(h_s, \tau_j), I)$, which yields the equilibrium price function $h = H(I, \tau_j)$.

Each location has its own supply curve as represented by τ_j , and differences in the slope of the supply curve are another form of location heterogeneity in the model.

¹¹ Shimer (2005) convincingly documents the failure of a single labor productivity dimension in rationalizing employment volatility in the U.S.

¹² One can think of this as how much of the consumption good is needed to produce additional housing in the location, which is not explicit but reflected in the price of housing.

¹³ The supply function is assumed continuous and monotonic and is thus invertible.

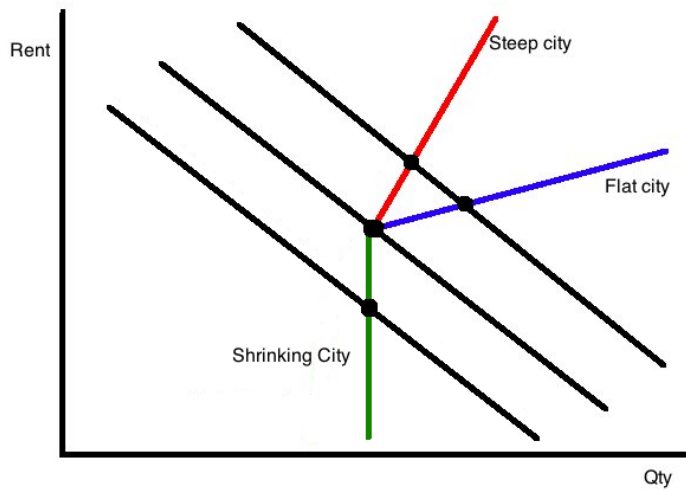


FIGURE 2.6: Illustration of Local Heterogeneity in Housing Supply Elasticity

As the literature on housing supply has shown,¹⁴ supply elasticity of housing can vary greatly between cities because of differences in local regulation, topography, and available land. Furthermore, because housing is a highly durable, locationally-fixed good, it cannot easily be removed from supply, a point argued by Glaeser and Gyourko (2005). Thus, cities below their housing stock capacity are on a “vertical” part of the supply curve.

Heterogeneity in housing supply curves is illustrated in Figure (2.6) which depicts hypothetical demand and supply curves for three types of cities. In the illustration, an identical upward demand shift will increase prices more in the ‘steep’ city and quantity more in the ‘flat’ city; a downward demand shift will cause predominantly price effects in the ‘shrinking’ city. This is the intuition behind allowing heterogeneous supply curves, and why differing price-demand relationships appear in the data, though the steep/flat/shrinking classification is merely for illustrative purposes.

¹⁴ See, for example, Glaeser, Gyourko, and Saks (2006), Saiz (2010), Paciorek (2010).

2.3.4 Workers

Having established the conditions of the local market within a period, I now turn to the sorting of the labor force over locations. First I will define the worker's utility function.

Utility Specification

Workers are risk averse and have utility over consumption, which is the remainder of their local wage after paying the local housing price, $c_{jt} = w_{jt} - h_{jt}$.¹⁵ Recall that all residents pay the housing price, but only matched workers receive a wage. To account for the risk of unemployment, I specify utility as the employment probability-weighted utility of consumption:

$$u(w_{jt}, \alpha_{jt}, p_{jt}) = E(u(c_{jt})) = er_{jt} \cdot u(w_{jt} - h_{jt}) + (1 - er_{jt}) \cdot u(-h_{jt}) \quad (2.4)$$

where er_{jt} , w_{jt} , h_{jt} are, respectively, the employment probability, the wage, and the rent in j at time t . Utility varies by location and over time according to the changes in wages, employment rates, and housing prices.

Note that unemployment directly enters the utility function without specifically affecting any individual, in a specification equivalent to a full local unemployment insurance system.¹⁶ This accounts for unemployment as a “bad” while maintaining tractability in the state space of the problem, and maintains consistency with the timing of the location decision before the labor matching process; that is, the worker chooses a location before she knows whether she will match to a job.

Also, all workers in j receive a flow amenity value of μ_j , which is assumed to be constant over time for each location.

¹⁵ Note that this relationship is implied by a simple household budget constraint.

¹⁶ The expected utility specification is an extension of Harris and Todaro's (1970) “expected earnings.”

Location Choices

Workers choose the local market to which they will supply their labor and in which they will reside for the period in order to maximize their lifetime expected utility. The workers' problem is dynamic because: 1) they face adjustment costs to changing locations, and 2) the states of the labor market in each location change over time.¹⁷ The workers choose an optimal location j according to current and expected future utility flows. The processes governing local labor productivity and matching efficiency are assumed exogenous, and each process has its own location-specific variance and degree of persistence. I assume workers know the law of motion for the labor productivity and matching efficiency in each location and can therefore form expectations about the future state of the labor market in terms of wages and employment.

At the start of the period, productivity and cost shocks to local markets are revealed. Workers know the wage and match efficiency in each local market, and therefore know the $\{w, er\}$ pair for an arbitrary value of the labor force population. The employment and wage will also determine the rental price of housing, h . Thus, workers know the utility provided by each location conditional on its states; the location's labor force population is the endogenous state to be determined by a sorting equilibrium.

The worker faces moving costs if she chooses a location other than her current one. The moving costs are a function of distance and are therefore symmetric for a location pair. I assume they are time-invariant and are the same for all workers choosing between k and j . Population new to the system does not face a moving cost when selecting a location.¹⁸

¹⁷ This reasoning goes back to at least Sjaastad (1962).

¹⁸ This is to represent, e.g. immigrants and young people new to the labor force who are choosing an initial location.

Workers then receive an idiosyncratic location shock, ε_{ijt} , and decide where to locate. Because the utility of a location depends on its population (the number of workers making the choice to live there), the workers' equilibrium decision rules are determined simultaneously.

The state variables in the workers' problem are the labor productivity, y_{jt} , and match efficiency, α_{jt} in each local market, and the labor force population of the local market that resulted from location choices in the previous period, p_{t-1} . The vector of all locations' productivity states are collected in W , the matching efficiency states in A , and the vector of the initial population distribution over locations is P . The other state variable is k , the worker's current location. Location enters as a state variable because the moving cost to a destination j varies according to origin k . Assuming a stationary Markov structure on the transitions of the productivity and matching efficiency states, the workers' problem can be represented as a discrete choice problem in Bellman equation form:

$$V(W_t, A_t, P_{t-1}, k) = \max_j [u(w_{jt}, \alpha_{jt}, p_{jt}) + \mu_j + mc_{kj} + \varepsilon_{ijt} + \beta EV(W_{t+1}, A_{t+1}, P_t, j)] \quad (2.5)$$

with β the worker's discount rate. The expectation is taken with respect to the future movements in local productivities and match efficiencies as well as the future idiosyncratic draws; the population states are determined by the equilibrium in the location choice problem. As written, W and A appear to lead the population state variable, but they are contemporaneous with respect to the migration decision: given the timing of the model, $\{W_t, A_t\}$ and P_{t-1} are both revealed at the time of worker location decision. They appear as state variables at time t because they compose the information set at the time the worker makes a location decision.

When valuing a location, the worker cares about *all locations'* productivities, match efficiencies, and populations—not only the location evaluated as a potential

destination—because each of these affect how likely it is that other workers choose the same market as she.

Theoretical island economies models such as Lucas and Prescott (1974) often assume an infinite number of islands, and use the massless nature of the islands to specify an equilibrium. An empirical application such as this needs to use a finite number of islands as the islands correspond to actual markets. Using a finite number of locations adds computational complexity beyond that of standard island economies models, as the identities of the islands matter, and therefore the discrete choice problem carries a large state space. The advantage of the discrete choice setting, however, is the direct link between model and data—each island corresponds to an actual city for which data series exist. The heterogeneity of locations is explicitly accounted for, and estimation of the model can use the empirical distributions of wages, employment, and housing prices of actual MSAs. To my knowledge, there is no other literature that solves an island economies equilibrium through a discrete choice optimization.

2.3.5 Equilibrium

I define the equilibrium in the model as a rational expectations equilibrium: workers choose locations conditional of their expectations of the location choices of other in the system, and in equilibrium, workers' expectations of others' strategies are correct. Because of the differential moving costs affect where workers move conditional on their origins, the current and future distribution of population is dependent on the history of population sorting. I am only interested in an equilibrium where, given vectors of labor market states and the current population distribution, agents can rationally make location choices for the next period. The implicit assumption in writing the workers' problem in (2.5) is that the geographic distribution of population is sufficient for summarizing both the path of play and agents' expectations about

the future.¹⁹ That is, P is a sufficient state variable in a recursive problem.

An equilibrium in this model is now defined. Capital letters again denote the vectors collecting the all locations' values of a variable.

Equilibrium. Given productivity and efficiency states, S , and last period's distribution of population over locations, P , a rational expectations equilibrium is a set of employment rates ER , housing rental prices, H , an allocation of population P' , such that:

1. Each local labor market matches according to $n = m(v, p, \alpha)$ and $\pi = 0, \forall j$.
2. Each local housing market clears according to $h = H(I, \tau_j) \forall j$.
3. Worker's locations are optimal.

A few technical assumptions on the value functions are needed to ensure an equilibrium exists and is reasonable.²⁰ The underlying state variables—most notably, wages and national population—may have trends which could cause the value of a particular location to become unbounded, in which case workers will not sort over markets but instead all enter the same location. I rule out by assumption any equilibrium in which any city absorbs the entire national population, or one in which a city is completely emptied.²¹

¹⁹ The equilibrium is similar to Lee and Wolpin (2006), Lee (2005), and Krusell and Smith (1998) in that I have an rational expectations equilibrium process, except that rather than assuming a particular form for a forecasting rule for the endogenous state and then iterating on it to find parameters consistent with a rational expectations equilibrium, my forecasting rule is degenerate: population distribution over the islands is all the information agents need about the endogenous states. The population distribution itself is the equilibrium object.

²⁰ While the model has only congestive interactions and not the agglomerative forces that often lead to multiple equilibria, I cannot say that in general an equilibrium is unique. In practice, I keep the equilibrium in the model as close to that of the data as possible.

²¹ The issue is not whether a single location may someday be so desirable as to absorb all population at that date. The practical problem is that without boundedness assumptions on the value functions, expectations of future trends may cause that 'supercity' to be realized today.

Therefore, I assume $|V(S_t, P_{t-1}, k)| < \infty \forall k, t$.²² A reasonable assumption which rules out $V = -\infty$ is that $w_{jt} > 0 \forall j, t$. Then to rule out $V = \infty$, I make the assumption that national population growth will always be large enough to congest locations (through e.g., house prices) with positive trends in wages, ensuring consumption is finite.²³

The last equilibrium condition, the worker's utility maximization condition, is the key to determining the equilibrium. The model specifies the equilibrium relationship in the local labor and housing markets, so the first two equilibrium conditions follow trivially when the last is satisfied. That is, all the endogenous variables are pinned down by the sorting of labor across locations. An equilibrium is then a fixed point in the spatial distribution of population.

$$P \Sigma(S, P') = P' \tag{2.6}$$

where

$$P = (p_1 \quad p_2 \quad \cdots \quad p_J)$$

and

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1J} \\ \sigma_{21} & \sigma_{22} & \cdots & \sigma_{2J} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{J1} & \sigma_{J2} & \cdots & \sigma_{JJ} \end{pmatrix}$$

P is the spatial distribution of population from the end of the previous period, P' is the new population distribution, and Σ represents the migration probabilities of workers as a function of the vector of local productivities and the worker's expectation of the population distribution; σ_{kj} represents the probability of residents in k choosing location j .

²² Bajari, Benkard, and Levin (2007) make the same assumption.

²³ A convenient but less general assumption is that utility is bounded above. For example, specifying CARA utility of $u(c) = 1 - \exp(-c)$ ensures that $\sup[u] = 1$ and the maximal discounted value of flow utility is simply $\frac{1}{1-\beta}$, which is finite.

Condition (2.6), which represents how workers behave optimally given the actions of other workers, will hold for each period t . Intuitively, population will flow into desirable locations and away from undesirable ones until choice value differences are arbitrated away by the congestion in prices (the rental price of housing and the employment probability).

However, this is not to say that *utilities* across locations are equilibrated, as in Roback(1982). In contrast to Roback’s static model, the model has costs to mobility, so that choice values—which incorporate moving costs—may be in equilibrium without convergence in flow utilities. Furthermore, the workers are maximizing their lifetime utility, and may trade off current and future consumption according to their expectations of the labor market states. The richer dynamic equilibrium of this model uses the value function to account for moving costs, the expectations of future states, and the expected future actions of other workers in the system.

The full distribution of population is relevant to each worker’s choice, as workers in comparatively good markets are less likely to move than workers in comparatively bad markets, and those “far away” (in moving cost terms) are less likely to make the same location decision as those “nearby.” The dynamic sorting problem thus carries the entire spatial distribution of population as state variables. The assumption of atomistic agents implies that choice probabilities are equivalent to moving shares, so the optimization problem need not be infinite dimensional. But even so, the heterogeneity of J locations makes the problem akin to a game of J players. Shortly I will discuss the numerical technique for finding the equilibrium, but for now, merely note that the spatial element of the model means state space grows exponentially fast in the number of locations, and so for any “large” number of locations (more than about 4), the state space is extremely large.²⁴

²⁴ With 30 locations, three states each, a discretization method for value function iteration with N discretization points would involve N^{90} points in the state space.

2.4 Equilibrium Approximation

I now describe the technique for numerically finding the equilibrium of the model. More details are available in the Appendix.

Finding the equilibrium of this model for a given set of parameters requires finding a fixed point in the spatial distribution of population. The values of the other endogenous variables, the local housing prices and employment rates, are pinned down by inserting the population into the equilibrium conditions for the other local variables.

The worker's migration decision is specified as a discrete choice optimization problem with social interactions, the interactions being the endogenous congestion. In equilibrium, a worker making a migration decision should accurately predict the migration decisions of all other workers. The assumption of atomistic agents implies that choice probabilities are equivalent to moving shares, so an equilibrium in population is equivalent to an equilibrium in choice probabilities:

$$\sigma_{kj}(\sigma_{-k,j}^*) = \sigma_{kj}^* \quad \forall k, j \quad (2.7)$$

where σ_{kj} is the probability of choosing to locate in j conditional on currently residing in k , and σ^* denotes the equilibrium value. I will assume the idiosyncratic component of utility, ε_{ijt} , is distributed iid Type 1 extreme value, meaning the value differences are distributed logistically, and the resulting choice probabilities are:

$$Pr(\text{locate in } j | \text{origin } k) = \sigma_{kj} = \frac{v_{kj}(X)}{\sum_i v_{ki}(X)} \quad (2.8)$$

where $v_{kj}(X)$ is the non-idiosyncratic relative values of locating in market j conditional on the state space X and origin k .

Finding the equilibrium of this system of location decisions amounts to a system of J^2 nonlinear equations. It is similar to finding the equilibrium strategies of a J -player game of private information, except that the "players" (locations) can switch

identities by migrating. Numerical methods for finding the equilibrium of such a dynamic game (without the identity switching) have been established in the industrial organization literature, beginning with Pakes and McGuire (1994). A numerical solution algorithm is as follows:

1. Initialize a guess of the choice probabilities, Σ^0 (the policy function) for workers at all locations J .
2. Solve the value function for each location, V_j^0 , using the guess Σ^0 .
3. Update the choice probabilities, Σ^1 , using the converged values of V_j^0 in (2.8).
4. Repeat steps 2-3 until convergence.

What complicates the direct application of these methods to the present model, however, is the sheer size of the state space. Each location has only three state variables, the productivity, y , the employment efficiency, α , and the population, p , but the spatial considerations cause the states of *all* locations to affect the equilibrium probability of selecting any location. A standard value/policy iteration involving discretization of the state space into N points quickly becomes infeasible. In the full solution method, where each choice probability depends on all locations states, the value function would be evaluated at N^{3J} points. For any interesting number of locations, the state space is prohibitively large. On top of that, such a value function would need to be evaluated for each location, with all functions converging jointly.

To make the equilibrium computation feasible, I simplify the computation through two approximation techniques. The first form of approximation is state aggregation. The separate tracking of each local state is what causes the explosive growth in the state space in the full solution. When valuing a particular location, workers care how many others are likely to locate there, but exactly where the congestion comes from should not matter. Thus a natural approximation of relative value differences

is to summarize the value of a location using its own states and a cross-sectional summary (such as an average) of the states of other locations in the system.

The technique used here has conceptual antecedents in Weintraub, Benkard and Van Roy (2008) and Krusell and Smith (1998). Weintraub et. al. suggested using a technique they called “oblivious equilibrium” (OE) for games with large numbers of players. Rather than a player keeping track of all rivals’ strategies, the player behaves “obliviously” regarding the individual players and instead summarizes the aggregate long-run industry state into a single state variable. Krusell and Smith (1998) applied similar intuition in a different setting. In their model, a continuum of heterogenous agents make savings decisions which depend on the distribution of wealth in the economy; the full solution involves an agent considering the wealth of all the other agents as states, an infinitive-dimensional problem. Their solution is to assume agents summarize information about the wealth distribution into a finite set of moments. In practice, they find that an equilibrium using only the mean wealth as a state was sufficient for matching the empirical distribution of wealth in the U.S. The approach employed in the current setting is not exactly like either one of these, but the intuition is similar. The routine inherits the “myself, and everyone else” categorization of states from OE, but through the use of summary statistics as in Krusell and Smith.

The approach grows more attractive as the number of locations grows large—not only for its relative computational gains, but also because it should become a more accurate approximation as the locations approach an economy of infinite islands. And while not exploited in this paper, it also naturally scales into regional classification, such as having the states of nearby locations enter the value function separately from faraway locations.

Even with state aggregation, however, the state space is large enough to make computation of the model slow. To speed up the solution of the value functions, I

follow Benitez-Silva et. al (2000) and Sweeting (2010) in applying Parameterized Policy Iteration (PPI) to find the fixed point of the joint worker optimization problem. Instead of discretizing the state space, I approximate the value function by specifying it to be a Q -termed, linearly parametrized combination of basis functions the states X :

$$v(X) \approx \hat{v}(X) = \sum_q^K \lambda_q g_q(x)$$

The advantage of the approximation is that it reduces the Bellman equation in the worker's problem, (2.5), to a more easily-solved linear problem. The parameters λ are chosen so that they most nearly fit the contraction in the worker's problem:

$$\hat{\lambda} = \underset{\lambda}{\operatorname{argmin}} \left\| \left[\sum_q \lambda_q g_q(x) - u(X') - \beta E \left(\sum_q \lambda_q g_q(x') \right) \right] \right\| \quad (2.9)$$

Using this approximation, the algorithm is now:

1. Form the functions $g_q(x)$.
2. Guess a vector of parameters for the value function approximation, $\hat{\lambda}^0$.
3. Initialize a guess of the choice probabilities, Σ^0 (the policy function) for workers at all locations J .
4. Form the functions $E(g_q(X'))$, where X' is the new states implied by Σ^0 .
5. Find λ according to (2.9). This updates the guess of $\hat{v}(X)$.
6. Update the choice probabilities, Σ^1 , using the new guess of $\hat{v}(X)$ in (2.8).
7. Repeat steps 4-6 until convergence in Σ .

A further advantage of PPI is that the state space need not grow exponentially in the number of players/locations. I simply need a “good basis” of functions $g_q(x)$

in (2.9), and then enough points to reliably find λ in (2.9); at a minimum, $N > Q$ is required for identification. In practice, I use several hundred (carefully selected) points per location.²⁵

Applying both methods of approximation yield significant computational gains. For example, the approximation routine converges in about one to five minutes (depending on starting values) on a desktop machine when $J = 30$. In comparison, “full solution” routines can take hours to converge with just $J=3$. Importantly, I have found the approximation is fast enough to nest inside an estimation routine.

2.5 Data and Estimation

This section describes the empirical implementation of the model. First I will describe the data, and then go on to describe the two estimation stages.

2.5.1 Data

Estimation of the model requires geographically detailed data on local labor and housing markets. The data employed are taken from several public sources:

- Local Area Unemployment Statistics (LAUS) from the Bureau of Labor Statistics (BLS). This dataset contains monthly employment and labor force size, at the county level county, from 1990-2010. The data are aggregated up to the MSA level.
- Quarterly Census of Employment and Wages (QCEW) from the BLS. This dataset contains monthly employment and quarterly wages at the county level, by industry, for the period 1990-2010 (through quarter 2). Its exhaustive information makes the QCEW valuable. Firms are subject to this census (not survey) where, each quarter, they are asked their payroll in the last month of

²⁵ See the appendix for more explanation of the choice of basis points.

the quarter and how many paid employees they had each month of the quarter. The data were aggregated up to one industry at the MSA level.

- County-to-County Migration from the Internal Revenue Service (IRS). This dataset contains the annual county-to-county migration flows as calculated by address changes on tax returns for the period 1990-2009. This dataset is valuable because it is essentially a census of taxpayers migration decisions at a fine level of geography and at a relatively high frequency.
- City Housing Price Indices (HPI) from the Federal Housing Finance Agency (FHFA, formerly OFHEO). The HPI are weighted repeat-sales indices of the value of single-family homes at the MSA level. These data are converted to dollar values using the county level median home price from the 2000 Census,²⁶ then aggregated back to the MSA level using the labor force population as weights.
- Consumer Price Index (CPI) from the BLS. The “All items, less housing” series²⁷ was used to deflate all dollar values to 2000 dollars.

The LAUS, QCEW, and HPI data are seasonally adjusted for each MSA by regressing the raw time series of data on quarterly dummies, and then taking the residuals from this regression to be the seasonally adjusted data. This is done separately for each MSA to allow for different seasonality in different locations.

2.5.2 Estimation Stage 1: Local Data Series and Parameters

Each of the following types of parameters are estimated or derived from time series data for each MSA. Allowing the parameters to be specific to each MSA is important for incorporating local heterogeneity.

²⁶ Downloaded from the NHGIS, www.nhgis.org.

²⁷ Series ID: CUSR0000SA0L2

Labor Matching

This section explains how I use the matching model described above to derive the variables describing the state of each local labor market. The available data contain employment, labor force population, and wages for each market over several periods of time. Importantly, vacancies and labor productivity are unobserved. However, the structure of the matching model allows me to derive the two key state variables, conditional on further assumptions about the δ and c parameters.

The first important feature of the labor market is labor productivity, or the output produced when a worker matches with a firm, y . I assume the parameter δ is constant over time, and with wages assumed to be a constant fraction of the output in (2.2), the wage can be taken to fully account for all movements in labor productivity. Thus each local market's wage series is used directly as a state variable.

The second important state variable is the matching efficiency, α . I use a constant returns to scale Cobb-Douglas for the matching technology, $n = v^{1-\alpha}p^\alpha$, $\alpha \in [0, 1]$, and set a linear form for the vacancy cost, $c(v) = cv$. The Cobb-Douglas form makes it apparent why I refer to α as the matching efficiency: it is the elasticity of employment with respect to labor supply. A higher value of α translates to less congestion in the labor market, as adding more available workers increases the number of matches at a higher rate.

The matching model is under-identified with the available data. I will therefore use the available data to derive an implied value of the match efficiency. Doing so leverages the firm's zero expect profit condition. Setting expect profits to 0, substituting the match function in for the unobserved vacancies, and solving in logs, I obtain:

$$\log(n) = \frac{2\alpha}{1+\alpha}\log(p) + \frac{1-\alpha}{1+\alpha}\left[\log(w) + \log\left(\frac{1-\delta}{\delta}\right) - \log(c)\right]$$

Recall that the vacancy cost function introduces the decreasing returns to scale that set city size. I will therefore set $\phi = \log(\frac{1-\delta}{\delta}) - \log(c)$ separately for each city to pin down its average size. I set ϕ so that it produces a value of $\alpha = 0.72$ at the average values of wages, employment, and population.²⁸ Then with ϕ set, I use the observed employment, population, and wages to derive implied match efficiency:

$$\alpha = \frac{\log(n) - \log(w) + \phi}{\log(p) - \log(\frac{n}{p}) - \log(w) + \phi} \quad (2.10)$$

The intuition of using (2.10) to derive the implied match efficiency is simple. In periods with high unemployment (low n relative to p), match efficiency must have been poor; conversely, low unemployment corresponds to higher values of α . The productivity of workers, captured through the wage, also provides information on the match efficiency, as firms should be willing to creating more vacancies when workers are more productive; if they do not and employment does not grow, then match efficiency must have been poor.

I back of the time series of values of α for each city. With the time series for w and α in hand, I estimate the process for transitions in the states. The estimates of these processes are used to form workers' expectations of future states in the dynamic rational expectations equilibrium. I use a simple AR(1) process with trend.

$$w_{t+1} = \psi_0 + \psi_1 w_t + \psi_2 t + \omega_t \quad (2.11)$$

$$\alpha_{t+1} = a_0 + a_1 \alpha_t + a_2 t + \nu_t \quad (2.12)$$

Before estimating (2.11) or (2.12), I test the data for a unit root. If the data fail to reject the null in a Dickey-Fuller test of a unit root with trend, I assign $\psi_1 = 1$ or $a_1 = 1$, and I regress the log differences on a constant to check for a trend in the process.

²⁸ Shimer (2005) sets $\alpha = 0.72$ in a Cobb-Douglas matching technology for the aggregate labor market.

Housing Prices

The model specified a relationship between demand the price of housing, with the elasticity coming from the slope of the housing supply function. Estimation will recover the extent to which housing prices change as demand shifts within in each MSA. As with the wage and employment cost processes, estimating each location separately is important for recovering location heterogeneity.

Demand can shift with changes in population size or income. To reflect both, I use local total income—the sum of all wages paid to workers in the period—as the demand statistic; I denote this measure I . I then use a log-linear specification to recover the housing price elasticity.

$$h_t = \tau_0 + \tau_1 I_t + \xi \quad (2.13)$$

Note that I endogenous because it contains current population, so that $E(I\xi) \neq 0$. If housing prices are, for example, temporarily higher due to an unobserved cost shock, fewer people will choose the location. To account for this, I instrument for I with its lagged value.

2.5.3 Stage 2: Migration Model

With first-stage estimates in hand, I now turn to the estimation of the dynamic equilibrium model of worker location.

Utility Specification

I parameterize utility over consumption as CARA with scaling parameter γ_1 and risk aversion parameter γ_2 .

$$u(er, w, h) = \gamma_1 [er \cdot (1 - \exp(-\gamma_2(w - h))) + (1 - er_{jt}) \cdot (1 - \exp(-\gamma_2(-h)))]$$

Moving costs are parameterized as a linear function of the log distance between two cities in miles. This is parsimonious, while establishing an intuitive concave relationship between moving costs and distance. The exception is migration to or from the aggregated outside option, $j = 0$, which has no definable distance. In these cases the intercept is permitted to vary but distance is not assigned.

$$mc = mc_0 + mc_1(k, j = 0) + mc_2(k = 0, j) + mc_3 \log(\text{miles})$$

The structural parameter vector of interest in this stage is then $\theta = [\gamma_1; \gamma_2; mc']$, the scaling and risk aversion parameters, and the vector of moving costs parameters. The amenity parameters μ are backed out for each candidate value of θ to make average population shares fit their observed shares, similar to Berry (1994).

The Estimation Algorithm

The estimation routine is as follows. First, for each candidate vector of parameters, the dynamic equilibrium approximation routine is conducted to estimate how agents value the future. To do so, I began with the basis states and form the approximating polynomial terms, $g_q(x)$. I then use the parameters and residuals from the estimation of the evolution of wages and employment cost, (2.11) and (2.12), to form the expected future states conditional on the basis, $E(g_q(x'|x))$. I then find the fixed point of the equilibrium sorting problem through PPI, using the full set of basis states to estimate the approximation parameters, λ . Within the PPI routine, utility is calculated using the elasticity of housing price estimated in the first stage and the structural parameters. Because each location has its own housing price elasticity and own process for evolution of the labor market states, there are separate approximation parameters for each location.

The output of the equilibrium approximation is the value agents place on an arbitrary vector of states, given the structural parameters, the expected evolution of

states, and the equilibrium actions of other agents. The next step is to apply this estimate of the value judgment encapsulated in λ to the observed location decisions of workers. To do so, I feed the actual wage and match efficiency data into a simulation of the model, with the value of each choice expressed as:

$$\hat{v}_{kj} = u_j + m_{kj} + \mu_j + \beta E\left(\sum_q \hat{\lambda}_q g_q(x')\right) \quad (2.14)$$

The choice values (2.14) are plugged into (2.8) to get the predicted migration flow matrix $\hat{\Sigma}$. Further nested within the fixed point problem is the Berry-style contraction to make each city's average population shares fit their observed shares for each candidate value of θ .

The output of the simulation is the predicted flow matrix and the population size of each location in each period. The model is simulated at a quarterly frequency, but the migration flow data is at an annual frequency, so the quarterly migration matrices are aggregated to a predicted annual flow matrix.

The matrix equation representing this for the r^{th} guess at θ is:

$$P_{t-1} \hat{\Sigma}(X_t, \theta^r) = \hat{P}_t \quad (2.15)$$

where

$$\hat{\Sigma}(X_t, \theta^r) = \begin{pmatrix} \hat{\sigma}_{11}(X_t, \theta^r) & \hat{\sigma}_{12}(X_t, \theta^r) & \cdots & \hat{\sigma}_{1J}(X_t, \theta^r) \\ \hat{\sigma}_{21}(X_t, \theta^r) & \hat{\sigma}_{22}(X_t, \theta^r) & \cdots & \hat{\sigma}_{2J}(X_t, \theta^r) \\ \vdots & \vdots & \ddots & \vdots \\ \hat{\sigma}_{J1}(X_t, \theta^r) & \hat{\sigma}_{J2}(X_t, \theta^r) & \cdots & \hat{\sigma}_{JJ}(X_t, \theta^r) \end{pmatrix}$$

The procedure predicts the spatial distribution of population and the migration flows in each time period for each candidate value of θ . Values for θ can be chosen to minimize the distance between the prediction population distribution and its values in the data. Following Pesendorfer and Schmidt-Dengler (2008), the moment conditions are the system of equations characterized by the (approximated) equilibrium

of the model. With labor force population at a quarterly frequency and migration shares at an annual frequency, the moment conditions in use are:

$$E(\hat{P}_t(\theta) - P_t) = 0, \quad E(\hat{\Sigma}_t(\theta) - \Sigma_t) = 0, \quad \forall t \quad (2.16)$$

Let M be the collected moment conditions. This suggests the weighted least squares estimator

$$\hat{\theta} = \underset{\Theta}{\operatorname{argmin}} \sum_t^T \sum_j^J W_{jt} (\hat{M}_{jt}(\theta) - M_{jt})^2 \quad (2.17)$$

where W is the weight.²⁹ A search over candidate values of θ is conducted, with the equilibrium approximation technique nested inside the routine.³⁰

Identification

The risk aversion parameter γ_2 governs how much agents are willing to trade off higher wages versus the risk of being unemployed. Thus, the parameter is identified by the average level of congestion in locations (i.e. how much unemployment risk the agents accept to have a chance at a wage/housing price pair in a location). Clearly, the migration flow data identify the moving cost parameters. Less obviously, however, the scaling parameter γ_1 is identified from migration as well. The parameter γ_1 captures how important in the location choice problem utility from consumption is relative to idiosyncratic location shocks, so the level of idiosyncratic migration (i.e. flows driven opposite of consumption differentials) helps identify this parameter.³¹

Also not obvious is that moving costs are helpful for predicting population size as

²⁹ The weighting metric I use is the inverse of the magnitude of the moment, treating the errors as percentage differences; this prevents the large cities from dominating the objective function. Work in progress incorporates the suggestions of Pesendorfer and Schmidt-Dengler (2008) for the efficient weight matrix.

³⁰ As the parameters near their converged values, using the policies from the previous iteration as starting values speeds convergence.

³¹ The variance of the shocks was implicitly normalized to 1, but estimation of the scaling parameter is isomorphic to estimation of the variance of the idiosyncratic shocks.

well. This is because the persistence and variance of the states is heterogeneous between locations, and since the moving costs affect how easily workers change locations, they affect how willing workers are to choose, for instance, a location with a large but short-lived shock to the local labor market. This of course also depends on γ_1 and γ_2 , so these parameters must be jointly estimated in a way that can handle the dynamic implications of location choice, such as the proposed method of this paper.

The data contain sufficient moments for identification. For each location, there are data for four quarters over 21 years, comprising $84J$ population moments to match. The IRS migration flows data come at only an annual frequency, but there are J^2 moments for 19 years, for a total of $84J + 19J^2$ moments. This lightly parametrized model has only six parameters.

Normalization

As is typically the case in discrete choice models, only the difference in choice values is identified, so some scale normalization is in order. After selecting $J - 1$ MSAs of interest, the rest of the country is collected into one aggregate location, taken to be the “outside option” (or the “reservation locale” as it is frequently termed in the urban economics literature). I will normalize by assuming the outside option has a utility of 0 in all states. The utility in the MSAs of interest is then identified relative to the outside option, and the moving costs in relation to utility.

2.6 Results

This section presents results of estimation and simulation of the model for an economy of 30 locations, the 29 largest MSAs in the U.S. plus an outside option aggregating the rest of the contiguous 48 states. The specified MSAs constitute slightly less than half the workforce population.

2.6.1 First Stage Results: Location Heterogeneity

The Labor Market States

Table (2.6.1) contains some descriptive statistics of the labor market states of each location. Presented are the means and standard deviations of the local wage, match efficiency, and unemployment rates. The cities exhibit significant heterogeneity in means and volatilities. In wages, there are high mean, high volatility locations (e.g. New York, San Jose) and low mean, low volatility locations (e.g. Tampa and Pittsburgh), but also high mean, low volatility (Washington D.C.), and low mean, high volatility (Portland). For match efficiency and unemployment rate, a similar picture emerges, but offering another dimension of heterogeneity. Cities with high volatility in wages are not necessarily the same as those with high volatility in match efficiency.

Table (2.14) (placed at the end of the article) presents the results from the estimation of (2.11) and (2.12), the transition processes of the labor market states. The heterogeneity between cities is especially evident in the persistence of the wage processes, ranging from cities with low persistence (Houston, St. Louis, Philadelphia) to others having full persistence (San Francisco, San Jose, Denver, Dallas). Most cities exhibit a unit root in the employment efficiency process.

Tables (2.7) and (2.7) (end of article) present the full locational matrix of contemporaneous correlation in shocks to the labor market states. Shocks are defined as the residuals from the processes estimated in (2.14). On the whole, while far from being perfectly correlated, the large degree of correlation in shocks to the local labor market states indicates a strong national component to shocks, with the local heterogeneity captured more by the persistence of the processes. Note, however, the differences in pairwise correlation can be substantial. Also, the mean correlation statistics show that some places are more “unique” in their labor market. For example, wage shocks in Detroit, San Jose, and Seattle are on average less correlated with

Table 2.2: Local Means and Volatilities of Employment State Variables

MSA	Mean Wage, dollars/wk	Mean Log Wage	StDev of detranded log wage	Mean α	StDev α	ϕ Avg empl cost	Mean Unempl Rate	Std Dev. Unempl Rate
Rural/O.O.	552	6.3112	0.0215	0.6614	0.0655	-11.51	6.44	1.44
Atlanta GA	720	6.5753	0.0424	0.5784	0.1585	-7.85	5.04	1.84
Baltimore MD	678	6.5153	0.0247	0.6688	0.0806	-7.29	5.50	1.29
Boston MA	850	6.7377	0.0453	0.7295	0.0860	-7.56	5.42	1.74
Chicago Gary IL	762	6.6330	0.0338	0.7416	0.0615	-8.25	6.55	1.71
Cincinnati OH	639	6.4564	0.0351	0.6811	0.0784	-7.09	5.57	1.66
Cleveland OH	654	6.4811	0.0273	0.7142	0.0657	-7.00	6.48	1.52
Dallas Fort TX	721	6.5773	0.0376	0.6676	0.0755	-7.96	5.38	1.43
Denver Boul'd CO	719	6.5732	0.0432	0.6835	0.0913	-7.41	4.98	1.48
Detroit MI	768	6.6428	0.0456	0.7027	0.1066	-7.59	7.40	2.89
Houston Braz TX	754	6.6205	0.0302	0.6942	0.0475	-7.74	5.79	1.20
Kansas City MO	632	6.4463	0.0331	0.7352	0.0744	-6.98	5.39	1.43
Los Angeles CA	740	6.6051	0.0268	0.7254	0.0742	-8.57	6.83	2.06
Miami Hialea FL	622	6.4301	0.0278	0.5592	0.1858	-7.46	5.99	2.15
Minneapolis MN	712	6.5638	0.0381	0.6329	0.1055	-7.63	4.50	1.38
New York Nor NY	993	6.8951	0.0509	0.6804	0.0650	-8.73	6.63	1.53
Orlando FL	562	6.3291	0.0313	0.6496	0.1837	-7.07	5.31	2.39
Philadelphia PA	750	6.6163	0.0258	0.6791	0.0770	-7.99	5.76	1.45
Phoenix AZ	624	6.4323	0.0328	0.6483	0.1423	-7.63	4.74	1.62
Pittsburgh B PA	636	6.4527	0.0210	0.7179	0.0508	-7.13	6.35	1.12
Portland Van OR	653	6.4786	0.0345	0.7387	0.0684	-6.94	6.24	1.99
Riverside Sa CA	561	6.3289	0.0267	0.5961	0.1638	-7.54	7.55	2.76
Sacramento CA	668	6.5011	0.0291	0.6333	0.0925	-6.93	7.35	2.24
St Louis MO	651	6.4761	0.0294	0.7431	0.0727	-7.25	5.96	1.68
San Diego CA	676	6.5106	0.0379	0.6650	0.0916	-7.33	5.57	1.96
San Francisc CA	889	6.7812	0.0489	0.7437	0.0987	-7.51	5.51	1.89
San Jose CA	1,096	6.9819	0.0899	0.7869	0.1102	-6.15	6.04	2.39
Seattle Ever WA	753	6.6181	0.0532	0.7309	0.0623	-7.30	5.82	1.40
Tampa St Pe FL	568	6.3379	0.0241	0.7052	0.1179	-7.32	5.43	2.45
Washington DC	858	6.7479	0.0266	0.6683	0.0766	-7.84	4.00	0.97

other cities in the system, and similarly, match efficiency shocks in Atlanta, Seattle, Miami, and Detroit are less correlated with other cities on average.

The Housing Markets

Table (2.3) presents the first stage estimates of housing supply elasticity, sorted from the largest estimated price elasticity to the lowest. Because virtually every city (and the national population) is growing over the time period, larger estimates of the price elasticity indicate higher marginal costs to additional housing in the MSA. The results are intuitive: the older, more congested cities of the Northeast and much of California, make up the top of the list as “steep” cities; the Sunbelt makes up most of the bottom as “flat” cities where quantity can be added without much price increase. The two exceptions to the growth trend are Cleveland, St. Louis, and especially Detroit, cities likely on the vertical part of their supply curves.

2.6.2 Second Stage Results: The Dynamic Model

Results for the structural parameters are presented in Table (2.4).³² The costs to moving to or from the outside option are estimated separately, as this location has no definable distance from actual cities. Note that the moving costs are “psychic” in the sense that they are costs to utility, and are therefore identified relative to the average level of utility. To convert to dollars, I need an estimate of the marginal utility of consumption, which due to concavity of the utility function varies with the level of consumption. Still, it is instructive to view the moving costs in dollars, so I evaluate the derivative at the national mean level of consumption at a unemployment rate of zero. At the mean, an arbitrary move (of 1 mile) would cost \$92,491 dollars of annual consumption.³³ The second panel of Table (2.4) presents moving costs in

³² The standard errors reported are least-squares standard errors and do not account for first stage error.

³³ This is on the order of magnitude of Bishop (2008) and Kennan and Walker (2011), though direct comparisons to those studies are difficult for several reasons, including their partial equilibrium

Table 2.3: Housing Price Elasticities

MSA	τ_1	τ_1 se	τ_0	τ_0 se	Mean weekly rent (dol)
Los Angeles CA	2.40	0.19	-46.89	4.10	450
New York Nor NY	1.60	0.13	-30.34	2.99	432
Detroit MI	1.58	0.08	-28.09	1.77	193
Miami Hialea FL	1.40	0.11	-23.90	2.28	252
Baltimore MD	1.38	0.10	-22.56	1.95	272
San Francisc CA	1.35	0.11	-22.53	2.25	536
Boston MA	1.25	0.09	-21.03	1.84	310
St Louis MO	1.18	0.08	-19.18	1.54	140
San Diego CA	1.18	0.08	-18.06	1.70	412
Minneapolis MN	1.16	0.06	-18.92	1.23	215
Portland Van OR	1.16	0.05	-17.89	1.04	261
Philadelphia PA	1.11	0.11	-18.39	2.30	243
Tampa St Pe FL	1.07	0.08	-16.66	1.65	161
Chicago Gary IL	1.02	0.08	-16.76	1.72	228
Washington DC	1.00	0.08	-15.62	1.66	343
San Jose CA	1.00	0.09	-14.29	1.95	573
Denver Bould CO	0.97	0.03	-14.58	0.59	258
Cleveland OH	0.94	0.06	-13.88	1.32	203
Seattle Ever WA	0.92	0.06	-13.32	1.33	324
Sacramento CA	0.87	0.10	-11.84	1.93	315
Riverside Sa CA	0.84	0.09	-11.32	1.75	273
Rural/O.O.	0.84	0.05	-15.15	1.22	181
Orlando FL	0.77	0.06	-10.07	1.29	190
Phoenix AZ	0.76	0.05	-10.48	1.05	198
Kansas City MO	0.76	0.03	-10.47	0.69	135
Pittsburgh B PA	0.76	0.05	-10.57	0.93	132
Cincinnati OH	0.52	0.02	-5.38	0.49	158
Atlanta GA	0.50	0.03	-5.25	0.56	187
Houston Braz TX	0.48	0.02	-5.28	0.49	143
Dallas Fort TX	0.28	0.02	-0.89	0.40	147

Table 2.4: Structural Parameters

Parameter	Estimate	SE
γ_1	1.4547	0.0042
γ_2	2.1403	0.0061
mc_0 : intercept	4.6353	0.0306
mc_1 : from City to O.O.	-3.1328	0.0608
mc_2 : from O.O. to City	5.8814	0.0478
mc_3 : Log Distance	1.3043	0.0068

<i>Selected move costs</i>			
Boston to:	New York, 6.70	Chicago, 8.36	San Francisco, 10.18
	\$ 136,591	\$ 175,679	\$ 205,652
Los Angeles to:	Riverside/SB, 5.32	San Diego, 6.20	Miami, 10.12
	\$ 106,240	\$ 123,705	\$ 201,958
Washington D.C. to:	Philadelphia, 6.31	Atlanta, 8.22	Dallas, 9.23
	\$ 125,870	\$ 163,937	\$ 184,114

utility and dollar terms (at the mean consumption level) for selected city pairs.

A baseline simulation uses the above structural parameters along with the first stage parameters to project the population sizes of each city. Figure (2.12) (end of article) compares the aggregate statistics from the baseline simulation with the data. Table (2.7) (end of article) presents a city-by-city summary of the model's endogenous variables, the local population, unemployment rate, housing prices, and estimated local fixed effects, μ . Since the labor market states are taken directly from the data, the labor market statistics from the baseline fit well. Table (2.5) presents the mean statistics from the cross sectional unemployment distribution in the data and the baseline model. The quartiles are plotted in Figure (2.7). The model fits the data well, though is slightly more disperse on average than the actual data. The decomposition of unemployment variance into temporal and local effects is reported specification and that wages/consumption enter linearly in their models.

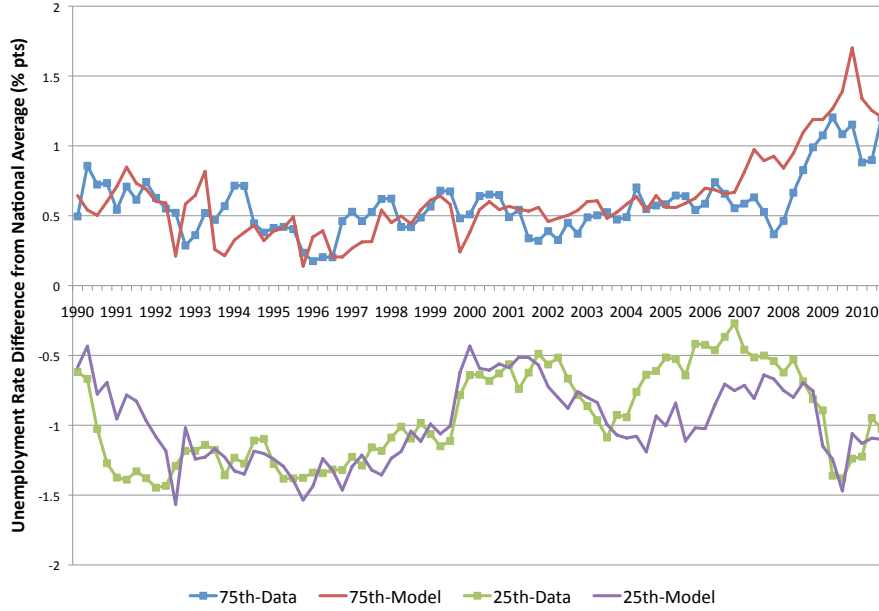


FIGURE 2.7: Cross Sectional Unemployment Distribution in the Data and Baseline Simulation

in Table (2.6).

Table 2.5: Cross Sectional Unemployment Distribution, Data and Baseline Model

	90 th pctile	75 th pctile	25 th pctile	10 th pctile	Std Dev
Data	1.37	0.58	-0.95	-1.49	1.14
Model	1.47	0.62	-0.99	-1.65	1.36

2.6.3 Simulations

I now turn to the counterfactual simulations. The simulations will be used to examine how the features of space-moving costs and the local housing market-affect the allocation of population to labor markets as the locations experience productivity shocks. In all simulations, the exogenous labor market states-wages and match efficiency-are taken unaltered from the data. Note that this preserves the local productivity shocks and the national component to the business cycle.

The first experiment removes moving costs, which reduces the dynamic equilibrium problem to a series of static equilibria. In a static, free mobility scenario, flow utilities across locations must be equilibrated; accordingly, I also refer to this case

Table 2.6: R^2 from Regressions of Unemployment Rates on Time and MSA Dummies, Data and Baseline Model

	Data	Model
Time dummies	.64	.51
MSA dummies	.16	.18
Time and MSA dummies	.82	.71

as the “Roback” scenario. Using the labor market state data and the estimated γ parameters, the equilibrium for each period is found by allocating population so that $u(w_j, \alpha_j, p_j) + \mu_j = u(w_k, \alpha_k, p_k) + \mu_k \forall j, k$.

The second experiment relaxes the link of the local housing market to the local labor market. I operationalize this by simulating a world in which residents receive a subsidy for half their rent in any location. In this scenario, which I refer to as “reduced rent,” housing is less important a determinant of the differences between locations, and wages and unemployment rates are more important. Both experiments are then compared to the baseline simulation of the model.

Before proceeding, a word about location fixed effects. In each of these counterfactual scenarios, I reset the fixed effects to match the average market shares of each location. Without this, the scenarios’ effects are primarily to the relative sizes of cities. Recall that the fixed effects are merely free parameters which soak up average differences in cities not captured by the model and are identified only in relation to the other parameters of the model. It is natural to think that, for example, San Francisco with its high wages might be much larger if people did not have to pay their rent in full. But these experiments are not about how cities got to their current sizes. Resetting the fixed effects allows me to focus on the business cycle dynamics of population adjustment without taking too literally the interpretation of the amenities.

Local Examples

Perhaps the easiest way to gain intuition from these experiments—and the working of the model itself—is to look closely at particular locations. Figure (2.8) shows the population and unemployment rates resulting from the three simulations for two example cities, Detroit and Seattle. Metropolitan Detroit had a well-known growth path from 1990-2010, with a significant and persistent decline in wages and employment after experiencing wage and employment growth in the 1990s. Looking at the time path of population in top left panel of the figure, the boom-bust cycle is evident. Notice how the counterfactual scenarios affect population volatility over the cycle. The Roback scenario admits more population during the growth period and allows more exit during the decline. The reduced rent scenario does the same to a lesser extent. The effect of greater population on the unemployment rate during the growth period is minimal, as the high employment efficiency (high α) exhibits a low rate of congestion at that time. When the labor market worsened (low α), however, and the rate of congestion increased, the flow out of Detroit had more impact on the unemployment rate. While the negative shock to α is still evident, the outflow of population in the counterfactual scenarios results in a lower unemployment rate relative to the baseline. The experience of Detroit is consistent with the notion that the fixity of housing contributes to urban decline, as suggested in Glaeser and Gyourko (2005).

Seattle is a location with a different path of labor market shocks, and while not as pronounced as Detroit, the effect of the counterfactuals is evident. Seattle saw growth in the late 1990s, then a rise in unemployment in the early 2000s, and then a rebound before another rise in the most recent recession. In the Roback scenario, population flowing in during the late 1990s expansion leads to an increase in unemployment over the baseline. Then when unemployment rose in the recession of 2008 and following, the Roback scenario permitted and outflow which reduces the unemployment rate relative to the baseline.

The reduced rent scenario actually exhibits an increase in unemployment in the the most recent recession. While this appears to be opposite of what the simulation

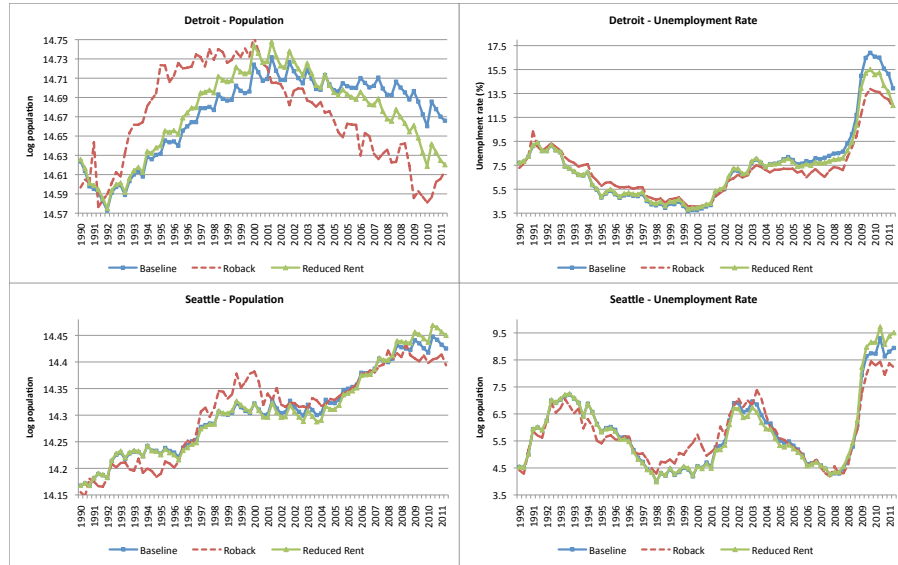


FIGURE 2.8: Time Paths of Population and Unemployment in Detroit and Seattle

did to Detroit, the behavior of the model in the counterfactual is in some sense the same. Recall that the model is working in a closed system. For population to move out of Detroit or other significantly depressed cities, it has to go somewhere, so while Seattle also had an increase in unemployment in 2008-2010, relative to Detroit, the shock was smaller. (Note the difference between the cities in the actual values of unemployment on the vertical axis.) With a reduced impact of housing price, Seattle adds population flowing from the worst cities. Detroit was the worst city in terms of unemployment, so population wants to flow out in both scenarios. Unemployment in Seattle, on the other hand, was not as extreme. For utility to be equalized in a static equilibrium, workers still leave Seattle, but workers' in the reduced rent scenario—still solving a dynamic problem—value Seattle's relative long term prospects more highly and flow into the city, arbitraging the dynamic differences in labor markets. This case is illustrative of how the two scenarios need not have the same effects on all cities.

Comparing Scenarios

The purpose of these counterfactuals is to analyze the impact of spatial frictions on the spatial arbitrage of labor, as evidenced by the cross sectional distribution of

Table 2.7: Mean Cross-Sectional Unemployment Dispersion Statistics from Simulation Scenarios

Statistic	Baseline	Free Mobility	Reduced Rent
Standard Dev	1.364	1.359	1.221
Interquartile Range	1.611	1.813	1.441
Interdecile Range	3.128	3.595	2.994
Range	6.487	5.350	5.578

unemployment rates. Table (2.7) presents several measures of dispersion in the unemployment rate. The numbers reported are the means over the 21 years of simulated data of the cross-sectional statistics. For example, the first row, first column indicates that in the baseline simulation, the average cross-sectional standard deviation in unemployment was 1.36. From this table we can see that reducing the importance of housing prices had a greater effect than the migration cost on the dispersion in unemployment; each statistic is smaller than in the baseline simulation, indicating, on average, a compression of the unemployment distribution. However, the Free Mobility scenario shows the largest effect on the range (maximum-minimum), indicating a greater effect on outliers in the distribution, while having less effect on the center of the distribution as measured by the other statistics. These effects are also illustrated by the kernel density plot of the cross sectional unemployment rate distribution, displayed in Figure (2.9).

Figure (2.10) illustrates one of the unemployment dispersion measures, the standard deviation, over the simulated time series. The figure presents the standard deviation in the data in blue, and the simulations points as differences to baseline simulation relative to the data; i.e., for plotted point y_t , $y_t = data_t + (counterfactual_t - baseline_t)$. The reduced rent scenario, in green, is more consistently below the data plot, indicating a compression of the dispersion in virtually all states of the economy in this time period. Note that the effect of the reduced rent scenario on cross-sectional dispersion of the unemployment rate was larger in later years. There are two potential explanations here. First, the recession of 2008 and following created a

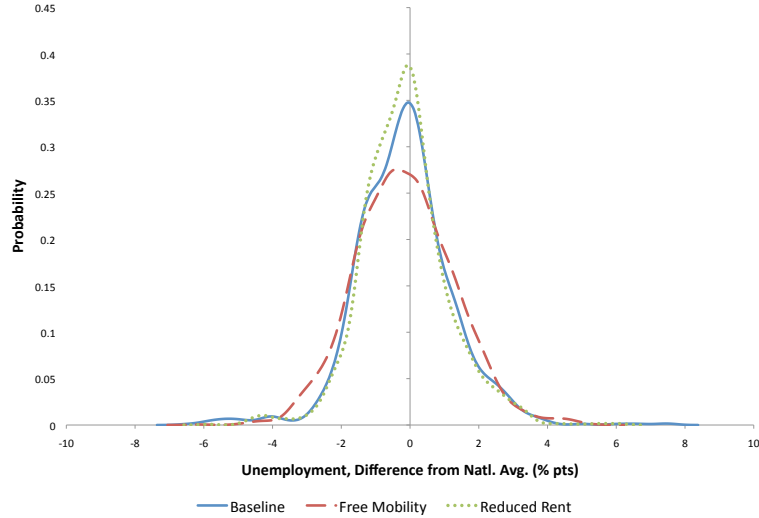


FIGURE 2.9: Cross Sectional Unemployment Distribution, Kernel-Density Estimates, by Simulation

large degree of cross sectional dispersion in the unemployment rate, so there is simply more geographic differentials to arbitrage. Also consider that incomes continued to grow throughout the time period, raising housing prices. Because prices grow at different rates in different markets, the cross sectional dispersion in housing prices also rose.³⁴ Therefore the experiment has more bite in the later years. The Free Mobility scenario, in red, displays less of an effect, but it does cause a reduction in the unemployment dispersion in the latter half of the data, especially in the most recent recession.

The primary question is, what do these scenarios reveal of the completeness of the spatial arbitrage of labor supply? A reduction in the cross-sectional dispersion in unemployment should mean less local variation, with local unemployment tracking more closely to the national average. To examine whether this is occurring in the counterfactuals, I repeat the exercise reported in Table (2.1): I take the simulated paths of unemployment rates, pool together all MSAs, and regress these data points on a vector of time and MSA dummies.³⁵ The R^2 from these regressions indicate how

³⁴ A fact documented on a larger scale by Van Nieuwerburgh and Weill(2010).

³⁵ A similar exercise was conducted by Blanchard and Katz (1992) on employment growth instead of unemployment rates.




FIGURE 2.10: Cross Sectional Standard Deviation in Unemployment Rate by Scenario, 1990-2010

much variation is due to national business cycles or, in the time and MSA dummies case, the combination of national business cycles and persistent unemployment differences. Table (2.8) presents the adjusted R^2 statistics from each regression.³⁶ That the R^2 is higher in the data generated by the reduced rent scenario suggests more local/temporal unemployment variation is arbitrated away in this scenario than in the other simulations. This is the key result of the paper.

Table 2.8: Adjusted R^2 Statistics from Pooled Regressions of Unemployment on Time and Location Dummies

Regression	Baseline	Free Mobility	Reduced Rent
Time dummies	0.64	0.63	0.70
Time, MSA dummies	0.82	0.78	0.93

Other measures of the local/temporal unemployment variation are reported in Table (2.9). This table reports the mean squared error (MSE) of, in the first row,

³⁶ Reported as $y_t = data_t + (counterfactual_t - baseline_t)$.

local unemployment rates from the national average, $\frac{1}{T} \frac{1}{J} \sum_t \sum_j (ur_{tj} - \bar{ur}_t)^2$, and in the second row, of the time-volatility of local unemployment to the volatility of the national unemployment, $\frac{1}{J} \sum_j (\sigma_j^{ur} - \bar{\sigma}^{ur})^2$. Like Table (2.8), these demonstrate that the reduced rent scenario results in less local/temporal unemployment variation and more spatial arbitrage of labor. The effect is also evident in Figure (2.9), with the Reduced Rent scenario displaying more density around the center of the distribution (i.e. the national average) than the baseline or Free Mobility scenarios.

Table 2.9: Mean Squared Error from National Unemployment Rate and National Unemployment Volatility

MSE	Baseline	Free Mobility	Reduced Rent
Unemployment Rate	1.96	2.06	1.57
Unemployment Volatility	0.324	0.334	0.228

In a similar vein, the reduced rent scenario causes the local unemployment and income volatilities to be more similar to that of the national time series, further indication that local variation is diminished in that scenario. Table (2.10) presents evidence to this effect. For each city, I calculate the volatility (time series standard deviation) in unemployment and in per-capita income³⁷ and then find the cross-sectional dispersion from the national average in each statistic; the cross-sectional dispersion is reported in Table (2.10). The lower amount of dispersion in the reduced rent scenario indicates that the cities' time series movements in these variables track the national cycles more closely. The Free Mobility scenario does not exhibit this effect.

How much is the average city changing in size in response to labor market shocks? Table (2.11) presents measures of a location's population volatility. The coefficient of variation is the (detrended) standard deviation of population divided by the mean; the statistic reported in table (2.11) is the average CV taken over all MSAs. Both

³⁷ The per-capita income is $pci = er \times w$, a statistic closely related to the expected income in the utility specification.

Table 2.10: Cross-sectional Dispersion in Unemployment and Income Volatility

Cross-sectional standard deviation in:	Baseline	Free Mobility	Reduced Rent
Employment Volatility	8.16	8.76	5.73
Per-capita Income Volatility	189.34	193.02	180.57

Table 2.11: Mean Population Volatility by Simulation

Statistic	Baseline	Free Mobility	Reduced Rent
Average City Population			
Coefficient of Variation	0.0265	0.0323	0.0279
Out-Migration Rate	0.009	(na)	0.009

counterfactuals increase the population volatility, with the Free Mobility scenario having the greatest effect. The second row of the table shows that the reduced rent scenario does not change the average mobility of workers.³⁸

How is the Free Mobility scenario changing local population volatility so much, while having less effect on the unemployment rate dispersion? The answer requires some unpacking of the simulation results. To begin, Table (2.12) reports correlation of population differences with local wage premia; a positive correlation indicates that population from the counterfactual is higher than the baseline in periods when wages are above the city’s trend, and population is lower than the baseline when below trend. Essentially, a positive correlation indicates more responsiveness of the labor supply to labor market shocks.³⁹ The reduced rent scenario exhibits a much higher correlation between population differences and wage premia than does the Free Mobility scenario. The correlation with α premia is smaller, but at least positive under the reduced rent scenario. There is also a positive correlation between migration differences in the reduced rent scenario and both labor market states.

³⁸ The flow condition of the Free Mobility static population equilibrium pins down only the allocation of population, and not its origin-destination flows. Hence, there is no migration statistic for the Free Mobility scenario.

³⁹ Recall that average population of each city is preserved in the counterfactuals, so that a population premium in one period is balanced by a deficit in another.

Table 2.12: Correlation of Population Differences to Baseline Simulation with Labor Market States

Statistic	Free Mobility	Reduced Rent
Correlation of Population Difference:		
with wage premia	0.077	0.413
with α premia	-0.010	0.028
Correlation of Out-Migration Difference:		
with wage premia		-0.376
with α premia		-0.150

That is, though average migration is not changing in the reduced rent scenario, the migration is more “targeted” in the sense that it is going to positively shocked labor markets. A comparable migration statistic does not exist for the Free Mobility scenario.

The picture of housing prices as a pro-cyclical moving cost is starting to emerge. When the housing market’s tie to the labor market is relaxed, population is more able to respond to productivity shocks since the labor market differential is less absorbed by a coincident move in the price of housing. Housing prices move in the direction of the labor market shock, mitigating the incentives to relocation; there is both “lock-in” in the downside of labor market shocks and “lock-out” on the upside. Reducing the effect of housing, as in the reduced rent scenario, allows the labor force more easily to follow local labor shocks, hence the positive correlation of the labor market states with population and the negative correlation with out-migration.

Why does the free mobility scenario yield a weaker correlation? Part of the answer is that even with free mobility, housing prices still congest good labor markets and mitigate poor ones. Thus, even if there is no moving cost threshold to overcome, the consumption differentials between markets will be reduced, and the incentive for reallocation is smaller. The pro-cyclical cost is still there even when the permanent cost is removed.

Secondly, there are other sources of utility outside the labor market. Workers in high amenity locations are more willing to accept poor labor market conditions, as

Table 2.13: Correlations of Wage Premium/Population Difference Related to Locational Attributes

Wage Prem/Pop Diff Correlation's Correlation to:	Free Mobility	Reduced Rent
Baseline FE, μ	-0.48	0.30
Mean Rent	-0.51	0.47
Estimated Rent Elasticity	-0.37	0.47

their utility can still be relatively high even when the location is hit with a negative productivity shock. To the extent that amenities are incorporated into housing prices, these two effects can be related.

Table (2.13) presents evidence for these arguments. For each location, I found the correlation of population differences to wage premia, as in Table (2.12). I then found the correlation of these city-level correlation statistics with city-level measures of housing prices and amenities. Table (2.13) reports the “correlation of correlations.” The city-level statistics are reported individually in Table (2.7) at the end of the article.⁴⁰ For the Free Mobility scenario, cities with higher amenities (as measured by μ) show less population responsiveness to labor market shocks. Cities with higher and more elastic housing prices also show weaker responsiveness. Under the reduced rent scenario, the opposite relationship holds: these are the more relatively more responsive locations since they are, in a sense, most affected by the treatment.

So while free mobility can reduce some dispersion in unemployment, it depends on where the shocks are hitting. This is the intuition behind why the Free Mobility scenario reduces the range of the unemployment distribution, but not the inner percentiles. For locations out in the tail of the unemployment distribution and with low amenities, such as Detroit in the 2000s as shown in Figure (2.8), free mobility

⁴⁰ Table (2.7) also shows that house prices and amenities affect the change to per-capita income volatility induced by the counterfactual. In the reduced rent scenario, income volatility is reduced in higher amenity, higher house price, and higher rent elasticity cities, but raised in the opposite types of cities. This effect is not present in the Free Mobility scenario. As in Tables (2.8),(2.9), and (2.10), this shows that it is the reduced rent scenario that makes cities’ local cycles more like the national cycle.

can cause significant reallocation of population.

In contrast, the reduced rent scenario always has a margin for reallocation so long as housing prices are related to labor market shocks. It affects all locations to some degree and in any state of the world. To illustrate the more consistent effect of housing prices, Figure (2.11) displays a scatterplot of each simulation’s effect on the unemployment distribution versus the degree of underlying dispersion; each point represents one quarter in the simulation. While the Free Mobility scenario sometimes reduces dispersion, it sometimes does not. The regression line for the Free Mobility points is nearly flat and has no explanatory power as measured by the R^2 . On the other hand, the reduced rent scenario consistently compresses the dispersion in unemployment and exhibits an effect that grows in magnitude with the degree of underlying dispersion.

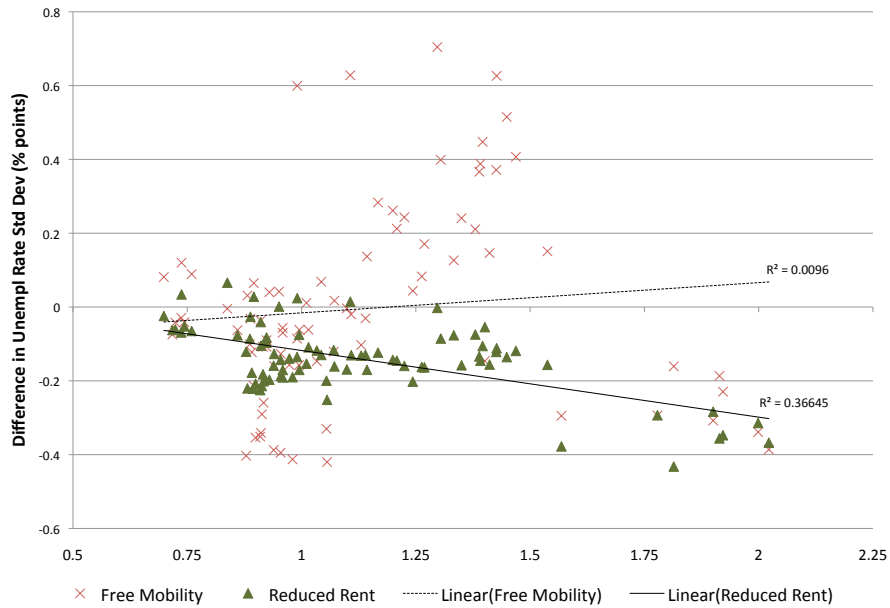


FIGURE 2.11: Experiments’ Differences in Dispersion Plotted Against Dispersion

In conclusion, the marked cross-sectional differences in labor markets owe more to the pro-cyclical nature of locally priced goods than to generic moving costs. While both are spatial frictions in some sense, they are not of the same type—one is a direct cost of distance, while the other is an indirect result of distance, that of local market clearing conditions for spatially delineated goods. This difference translates

to dissimilar effects on the dynamics of individual labor markets and on the resulting dispersion between them. It is the pro-cyclical friction which impacts the spatial arbitrage of labor.

2.7 Conclusions and Directions for Future Work

This paper considered the impact of spatial frictions on the national economy. Cities are heterogeneous in their evolution of productivity, but spatial arbitrage by labor suppliers is incomplete as evidenced by substantial geographic dispersion in unemployment rates. Migration costs, location preferences, and locally traded goods are non-labor market features that are salient to workers location decisions. To the extent that these impede labor allocation across space, these features are frictions to the labor market.

The goal of this line of research is to better understand how space affects the allocation of labor supply, and ultimately, what geographic segmentation means for models of the national labor market. This paper quantified the relative importance of migrations costs and housing prices as frictions to the labor market, finding that unemployment dispersion—and generally, the spatial arbitrage of labor—owes more to reactions in housing prices than to generic migration costs. More generally, housing prices and moving costs may be thought of as pro-cyclical and neutral frictions, respectively. Understanding the nature of spatial frictions—particularly the correlation of spatial frictions with local labor demand shocks—may be an important insight for studying the allocation of labor over space.

The model presented herein is well suited for extensions studying the role of space on labor market and population dynamics at the local and national level. The current paper has not even leveraged all the attributes of location heterogeneity available within the model. Future work will study the effect of spatial frictions at a more local level, such as the effect of geographic isolation or housing supply elasticity on the differential responses of cities to productivity shocks. Extensions introducing heterogeneity of workers and/or local industries is another topic left for future work.

Additional Exhibits

Table 2.14: State Transition Processes

MSA	ψ_0 (w cons)	ψ_1 (lag w)	ψ_2 (w trend)	a_0 (a cons)	a_1 (lag α)	a_2 (a trend)
Rural/O.O.	3.285 **	0.470 **	0.001 **	0.000	1.000 **	0.000
Atlanta GA	0.000	1.000 **	0.000	0.025	0.612 **	-0.004
Baltimore MD	3.270 **	0.486 **	0.002 **	0.000	1.000 **	0.000
Boston MA	2.062 **	0.685 **	0.001 **	0.000	1.000 **	0.000
Chicago Gary IL	2.841 **	0.564 **	0.001 **	0.000	1.000 **	0.000
Cincinnati OH	2.084 **	0.671 **	0.001 **	0.000	1.000 **	0.000
Cleveland OH	2.947 **	0.540 **	0.001 **	0.000	1.000 **	0.000
Dallas Fort TX	0.000	1.000 **	0.000	0.000	1.000 **	0.000
Denver Bould CO	0.000	1.000 **	0.000	0.000	1.000 **	0.000
Detroit MI	2.285 **	0.653 **	0.000 *	0.000	1.000 **	0.000
Houston Braz TX	4.055 **	0.372 **	0.002 **	0.000	1.000 **	0.000
Kansas City MO	3.099 **	0.510 **	0.001 **	0.000	1.000 **	0.000
Los Angeles CA	3.466 **	0.467 **	0.001 **	0.000	1.000 **	0.000
Miami Hialeah FL	3.464 **	0.451 **	0.002 **	0.000	1.000 **	0.000
Minneapolis MN	2.114 **	0.671 **	0.001 **	0.000	1.000 **	0.000
New York Nor NY	5.857 **	0.131	0.003 **	-0.469 **	0.469 **	0.002
Orlando FL	1.990 **	0.680 **	0.001 **	0.000	1.000 **	0.000
Philadelphia PA	4.532 **	0.301 **	0.002 **	0.000	1.000 **	0.000
Phoenix AZ	2.150 **	0.658 **	0.001 **	0.000	1.000 **	0.000
Pittsburgh B PA	5.147 **	0.189 *	0.002 **	0.000	1.000 **	0.000
Portland Van OR	0.000	1.000 **	0.000	0.000	1.000 **	0.000
Riverside Sa CA	1.881 **	0.699 **	0.001 **	0.000	1.000 **	0.000
Sacramento CA	2.572 **	0.596 **	0.001 **	0.000	1.000 **	0.000
St Louis MO	4.937 **	0.226 **	0.002 **	0.000	1.000 **	0.000
San Diego CA	1.685 **	0.734 **	0.001 **	0.000	1.000 **	0.000
San Francisc CA	0.000	1.000 **	0.000	0.000	1.000 **	0.000
San Jose CA	0.000	1.000 **	0.000	0.000	1.000 **	0.000
Seattle Ever WA	0.000	1.000 **	0.000	0.000	1.000 **	0.000
Tampa St Pe FL	3.619 **	0.416 **	0.002 **	0.000	1.000 **	0.000
Washington DC	3.500 **	0.466 **	0.002 **	0.000	1.000 **	0.000

⁴¹Asterisks (*) indicate significance at 0.01(**) and 0.05(*) levels.

Table 2.15: Correlation Across Locations in Shocks to Wages

MSA	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)	
(1) Rural/O.O.	1.00	0.68	0.89	0.86	0.90	0.84	0.87	0.78	0.65	0.58	0.81	0.72	0.91	0.86	0.86	0.71	0.84	0.93	0.89	0.90	0.74	0.85	0.75	0.64	0.89	0.69	0.44	0.47	0.92	0.83	
(2) Atlanta GA	0.68	1.00	0.61	0.73	0.68	0.60	0.61	0.84	0.71	0.39	0.37	0.39	0.65	0.62	0.69	0.51	0.56	0.64	0.67	0.39	0.74	0.64	0.53	0.40	0.71	0.73	0.53	0.55	0.63	0.55	
(3) Baltimore MD	0.89	0.61	1.00	0.76	0.74	0.64	0.71	0.66	0.52	0.50	0.70	0.54	0.86	0.76	0.67	0.55	0.83	0.83	0.74	0.76	0.62	0.79	0.77	0.58	0.81	0.62	0.39	0.32	0.85	0.79	
(4) Boston MA	0.86	0.73	0.76	1.00	0.88	0.80	0.79	0.82	0.66	0.49	0.81	0.64	0.78	0.80	0.81	0.73	0.67	0.85	0.81	0.80	0.73	0.67	0.53	0.83	0.82	0.56	0.38	0.80	0.71		
(5) Chicago Gary, IL	0.90	0.68	0.74	0.88	1.00	0.89	0.91	0.82	0.72	0.50	0.84	0.71	0.85	0.82	0.92	0.87	0.71	0.92	0.86	0.80	0.75	0.69	0.61	0.62	0.80	0.74	0.43	0.47	0.84	0.77	
(6) Cincinnati OH	0.84	0.60	0.64	0.80	0.89	1.00	0.87	0.81	0.67	0.62	0.72	0.70	0.79	0.77	0.85	0.72	0.64	0.82	0.83	0.79	0.72	0.68	0.56	0.55	0.75	0.71	0.47	0.56	0.79	0.72	
(7) Cleveland OH	0.87	0.61	0.71	0.79	0.91	0.87	1.00	0.75	0.59	0.64	0.76	0.70	0.83	0.80	0.86	0.80	0.73	0.88	0.86	0.89	0.69	0.67	0.60	0.62	0.74	0.62	0.46	0.42	0.86	0.70	
(8) Dallas Fort TX	0.78	0.84	0.66	0.82	0.82	0.81	0.75	1.00	0.79	0.50	0.70	0.54	0.75	0.74	0.80	0.66	0.62	0.73	0.77	0.71	0.84	0.67	0.54	0.49	0.74	0.83	0.55	0.55	0.72	0.64	
(9) Denver Bould CO	0.65	0.71	0.52	0.66	0.72	0.67	0.59	0.79	1.00	0.23	0.58	0.50	0.63	0.59	0.72	0.52	0.45	0.57	0.65	0.62	0.78	0.58	0.42	0.42	0.67	0.77	0.56	0.52	0.57	0.60	
(10) Detroit MI	0.58	0.39	0.50	0.49	0.50	0.62	0.64	0.50	0.23	1.00	0.34	0.52	0.34	0.49	0.35	0.56	0.51	0.68	0.51	0.47	0.58	0.51	0.43	0.54	0.32	0.36	0.42	0.60	0.44		
(11) Houston Braz TX	0.81	0.57	0.70	0.81	0.84	0.72	0.76	0.70	0.58	0.34	1.00	0.60	0.81	0.78	0.77	0.85	0.61	0.84	0.73	0.84	0.63	0.58	0.53	0.53	0.68	0.67	0.34	0.32	0.74	0.70	
(12) Kansas City MO	0.72	0.39	0.54	0.64	0.71	0.70	0.70	0.54	0.50	0.52	0.60	1.00	0.62	0.39	0.71	0.60	0.56	0.68	0.70	0.69	0.56	0.52	0.46	0.51	0.60	0.46	0.39	0.29	0.63	0.60	
(13) Los Angeles CA	0.91	0.65	0.86	0.78	0.85	0.79	0.83	0.75	0.63	0.53	0.81	0.62	1.00	0.88	0.79	0.69	0.79	0.87	0.83	0.84	0.70	0.81	0.73	0.58	0.84	0.67	0.46	0.43	0.90	0.81	
(14) Miami Hialeah FL	0.86	0.62	0.76	0.80	0.82	0.77	0.80	0.74	0.59	0.54	0.78	0.59	0.88	1.00	0.76	0.70	0.77	0.82	0.87	0.82	0.68	0.74	0.66	0.55	0.80	0.63	0.50	0.46	0.87	0.78	
(15) Minneapolis MN	0.86	0.69	0.67	0.81	0.92	0.85	0.86	0.80	0.72	0.49	0.77	0.71	0.79	0.76	1.00	0.80	0.68	0.87	0.83	0.84	0.69	0.72	0.55	0.56	0.77	0.70	0.39	0.49	0.80	0.75	
(16) New York Nor NY	0.71	0.51	0.55	0.73	0.87	0.72	0.80	0.66	0.52	0.35	0.85	0.60	0.69	0.70	0.80	1.00	0.53	0.83	0.68	0.79	0.50	0.49	0.36	0.48	0.55	0.52	0.21	0.30	0.66	0.63	
(17) Orlando FL	0.84	0.56	0.83	0.67	0.71	0.64	0.73	0.62	0.45	0.56	0.61	0.56	0.79	0.77	0.68	0.53	1.00	0.78	0.80	0.74	0.60	0.78	0.69	0.57	0.75	0.52	0.41	0.43	0.83	0.74	
(18) Philadelphia PA	0.93	0.64	0.83	0.85	0.92	0.82	0.88	0.73	0.57	0.51	0.84	0.68	0.87	0.82	0.87	0.83	0.78	1.00	0.81	0.90	0.62	0.75	0.67	0.62	0.79	0.66	0.36	0.38	0.90	0.77	
(19) Phoenix AZ	0.89	0.67	0.74	0.81	0.86	0.83	0.86	0.77	0.65	0.68	0.73	0.70	0.83	0.87	0.83	0.68	0.80	1.00	0.83	0.77	0.77	0.77	0.64	0.57	0.84	0.64	0.54	0.55	0.87	0.79	
(20) Pittsburgh B PA	0.90	0.59	0.76	0.80	0.89	0.79	0.89	0.71	0.62	0.51	0.84	0.69	0.84	0.82	0.84	0.79	0.90	0.83	1.00	0.65	1.00	0.65	0.70	0.63	0.61	0.79	0.62	0.41	0.38	0.88	0.74
(21) Portland Van OR	0.74	0.74	0.62	0.75	0.73	0.72	0.69	0.84	0.78	0.47	0.63	0.56	0.70	0.68	0.69	0.50	0.60	0.62	0.77	0.65	1.00	0.62	0.57	0.40	0.74	0.77	0.69	0.54	0.69	0.60	
(22) Riverside Sa CA	0.85	0.64	0.79	0.64	0.69	0.68	0.67	0.67	0.58	0.58	0.58	0.52	0.81	0.74	0.72	0.49	0.78	0.75	0.77	0.70	0.62	1.00	0.69	0.55	0.79	0.54	0.34	0.50	0.81	0.74	
(23) Sacramento CA	0.75	0.53	0.77	0.67	0.61	0.56	0.60	0.54	0.42	0.51	0.53	0.46	0.73	0.66	0.55	0.36	0.69	0.67	0.64	0.63	0.57	0.69	1.00	0.59	0.79	0.56	0.43	0.26	0.76	0.58	
(24) St Louis MO	0.64	0.40	0.58	0.53	0.62	0.55	0.62	0.49	0.42	0.43	0.53	0.51	0.58	0.55	0.56	0.48	0.57	0.62	0.57	0.61	0.40	0.55	0.59	1.00	0.56	0.36	0.18	0.19	0.59	0.57	
(25) San Diego CA	0.89	0.71	0.81	0.83	0.80	0.75	0.74	0.74	0.67	0.54	0.68	0.60	0.84	0.80	0.77	0.55	0.75	0.79	0.84	0.79	0.74	0.79	0.79	0.56	1.00	0.73	0.54	0.53	0.85	0.76	
(26) San Francisco CA	0.69	0.73	0.62	0.82	0.74	0.71	0.62	0.83	0.77	0.32	0.67	0.46	0.67	0.63	0.70	0.52	0.52	0.66	0.64	0.62	0.77	0.54	0.56	0.36	0.73	1.00	0.59	0.42	0.65	0.52	
(27) San Jose CA	0.44	0.53	0.39	0.56	0.43	0.47	0.46	0.55	0.56	0.36	0.34	0.39	0.46	0.50	0.39	0.21	0.41	0.36	0.54	0.41	0.69	0.34	0.43	0.18	0.54	0.59	1.00	0.41	0.46	0.41	
(28) Seattle Ever WA	0.47	0.55	0.32	0.38	0.47	0.47	0.56	0.42	0.52	0.42	0.32	0.29	0.43	0.46	0.49	0.30	0.43	0.38	0.55	0.38	0.54	0.50	0.26	0.19	0.53	0.42	0.41	1.00	0.40	0.46	
(29) Tampa St Pe FL	0.92	0.63	0.85	0.80	0.84	0.79	0.86	0.72	0.57	0.60	0.74	0.63	0.90	0.87	0.80	0.66	0.83	0.90	0.87	0.88	0.69	0.81	0.76	0.59	0.85	0.65	0.46	1.00	0.77		
(30) Washington DC	0.83	0.55	0.79	0.71	0.77	0.72	0.70	0.64	0.60	0.44	0.70	0.60	0.81	0.78	0.75	0.63	0.74	0.77	0.79	0.74	0.60	0.74	0.58	0.57	0.76	0.52	0.41	0.46	0.77	1.00	
Mean Corr	0.78	0.61	0.68	0.73	0.76	0.72	0.73	0.70	0.60	0.49	0.67	0.58	0.75	0.72	0.73	0.61	0.66	0.75	0.73	0.66	0.66	0.59	0.51	0.73	0.62	0.44	0.43	0.75	0.67		

Table 2.16: Correlation Across Locations in Shocks to Employment Efficiency

MSA	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)	(25)	(26)	(27)	(28)	(29)	(30)			
(1) Rural/O.O.	1.00	0.10	0.82	0.69	0.80	0.79	0.83	0.79	0.73	0.67	0.72	0.70	0.79	0.24	0.72	0.41	0.66	0.80	0.41	0.84	0.81	0.59	0.71	0.60	0.71	0.68	0.46	0.59	0.75	0.74			
(2) Atlanta GA	0.10	1.00	0.38	0.11	0.04	0.02	0.13	0.00	-0.13	0.25	0.06	0.12	0.22	-0.14	-0.02	0.13	-0.22	0.15	0.15	0.11	0.02	0.25	0.27	0.13	0.21	0.18	0.11	0.06	0.14	0.10			
(3) Baltimore MD	0.82	0.38	1.00	0.64	0.63	0.57	0.68	0.60	0.53	0.55	0.60	0.58	0.74	0.17	0.50	0.32	0.49	0.71	0.30	0.67	0.68	0.54	0.70	0.57	0.69	0.68	0.43	0.43	0.66	0.69			
(4) Boston MA	0.69	0.11	0.64	1.00	0.74	0.76	0.74	0.74	0.47	0.38	0.83	0.56	0.71	0.39	0.76	0.61	0.45	0.91	0.35	0.79	0.69	0.54	0.45	0.50	0.81	0.85	0.61	0.40	0.72	0.69			
(5) Chicago Gary IL	0.80	0.04	0.63	0.74	1.00	0.81	0.80	0.73	0.64	0.46	0.77	0.64	0.67	0.29	0.76	0.56	0.44	0.79	0.49	0.72	0.77	0.50	0.41	0.33	0.63	0.69	0.43	0.42	0.61	0.67			
(6) Cincinnati OH	0.79	0.02	0.57	0.76	0.81	1.00	0.81	0.82	0.54	0.55	0.79	0.60	0.77	0.30	0.77	0.60	0.53	0.81	0.43	0.78	0.70	0.58	0.48	0.47	0.74	0.71	0.53	0.54	0.71	0.73			
(7) Cincinnati OH	0.83	0.13	0.68	0.74	0.80	0.81	1.00	0.79	0.57	0.61	0.76	0.60	0.77	0.16	0.72	0.55	0.59	0.78	0.40	0.81	0.75	0.57	0.56	0.49	0.71	0.66	0.56	0.50	0.69	0.68			
(8) Dallas Fort. TX	0.79	0.00	0.60	0.74	0.73	0.82	0.79	1.00	0.65	0.47	0.84	0.58	0.74	0.28	0.79	0.55	0.69	0.75	0.49	0.73	0.74	0.61	0.50	0.52	0.71	0.65	0.51	0.51	0.71	0.68			
(9) Denver-Bould CO	0.73	-0.13	0.53	0.47	0.64	0.54	0.57	0.65	1.00	0.23	0.60	0.58	0.58	0.22	0.68	0.37	0.52	0.54	0.40	0.61	0.70	0.39	0.44	0.48	0.48	0.48	0.56	0.43	0.43	0.51	0.54		
(10) Detroit MI	0.67	0.25	0.55	0.38	0.46	0.55	0.61	0.47	0.23	1.00	0.33	0.50	0.32	0.05	0.32	0.09	0.49	0.43	0.24	0.50	0.51	0.54	0.56	0.46	0.45	0.34	0.33	0.32	0.57	0.49			
(11) Houston Braz TX	0.72	0.06	0.60	0.83	0.77	0.79	0.76	0.84	0.60	0.33	1.00	0.62	0.71	0.28	0.78	0.71	0.48	0.82	0.48	0.77	0.67	0.56	0.43	0.51	0.72	0.75	0.47	0.44	0.65	0.63			
(12) Kansas City MO	0.70	0.12	0.58	0.56	0.64	0.60	0.60	0.58	0.58	0.50	0.62	1.00	0.53	0.11	0.66	0.41	0.38	0.64	0.34	0.62	0.60	0.41	0.47	0.52	0.53	0.58	0.43	0.38	0.54	0.49			
(13) Los Angeles CA	0.79	0.22	0.74	0.71	0.67	0.77	0.77	0.74	0.58	0.52	0.71	0.53	1.00	0.27	0.62	0.54	0.59	0.79	0.42	0.77	0.72	0.67	0.66	0.50	0.82	0.75	0.60	0.57	0.70	0.76			
(14) Miami Hialeah FL	0.24	-0.14	0.17	0.39	0.29	0.30	0.16	0.28	0.22	0.05	0.28	0.11	0.27	1.00	0.23	0.16	0.27	0.33	0.26	0.29	0.25	0.15	0.03	0.16	0.32	0.33	0.23	0.07	0.36	0.28			
(15) Minneapolis MN	0.72	-0.02	0.50	0.76	0.76	0.77	0.72	0.79	0.68	0.32	0.78	0.66	0.62	0.23	1.00	0.65	0.48	0.80	0.41	0.72	0.65	0.52	0.36	0.45	0.66	0.68	0.43	0.44	0.60	0.67			
(16) New York Nor NY	0.41	0.13	0.32	0.61	0.56	0.60	0.55	0.55	0.37	0.09	0.71	0.41	0.54	0.16	0.65	1.00	0.18	0.66	0.26	0.56	0.35	0.35	0.17	0.26	0.48	0.55	0.28	0.25	0.32	0.47			
(17) Orlando FL	0.66	-0.22	0.49	0.45	0.44	0.53	0.59	0.69	0.52	0.49	0.48	0.38	0.59	0.27	0.48	1.00	0.48	1.00	0.48	0.17	0.54	0.60	0.60	0.51	0.37	0.56	0.47	0.47	0.50	0.71	0.55		
(18) Philadelphia PA	0.80	0.15	0.71	0.91	0.79	0.81	0.78	0.75	0.54	0.43	0.82	0.64	0.79	0.33	0.80	0.66	0.48	1.00	0.32	0.85	0.70	0.60	0.53	0.50	0.80	0.82	0.48	0.44	0.71	0.71			
(19) Phoenix AZ	0.41	0.15	0.30	0.35	0.49	0.43	0.40	0.49	0.40	0.24	0.48	0.34	0.42	0.26	0.41	0.26	0.17	0.32	1.00	0.41	0.42	0.30	0.30	0.25	0.38	0.31	0.31	0.21	0.42	0.35			
(20) Pittsburgh B PA	0.84	0.11	0.67	0.79	0.72	0.78	0.81	0.73	0.61	0.50	0.77	0.62	0.77	0.29	0.72	0.56	0.54	0.85	0.41	1.00	0.70	0.55	0.64	0.54	0.79	0.75	0.55	0.56	0.77	0.70			
(21) Portland Van OR	0.81	0.02	0.68	0.69	0.77	0.70	0.75	0.74	0.70	0.51	0.67	0.60	0.72	0.25	0.65	0.35	0.60	0.70	0.42	0.70	1.00	0.50	0.54	0.50	0.65	0.68	0.65	0.59	0.76	0.67			
(22) Riverside Sa CA	0.59	0.25	0.54	0.50	0.58	0.57	0.61	0.39	0.54	0.56	0.41	0.67	0.15	0.52	0.52	0.35	0.60	0.60	0.30	0.55	0.50	1.00	0.48	0.39	0.66	0.54	0.36	0.45	0.66	0.62			
(23) Sacramento CA	0.71	0.27	0.70	0.45	0.41	0.48	0.56	0.50	0.44	0.56	0.43	0.47	0.66	0.03	0.36	0.17	0.51	0.53	0.30	0.64	0.54	0.48	1.00	0.51	0.65	0.55	0.49	0.45	0.68	0.47			
(24) St Louis MO	0.60	0.13	0.57	0.50	0.53	0.47	0.49	0.52	0.48	0.46	0.51	0.52	0.50	0.16	0.45	0.26	0.37	0.50	0.25	0.54	0.50	0.39	0.51	1.00	0.49	0.46	0.28	0.33	0.48	0.46			
(25) San Diego CA	0.71	0.21	0.69	0.81	0.63	0.74	0.71	0.71	0.48	0.45	0.72	0.53	0.82	0.32	0.66	0.48	0.56	0.80	0.38	0.79	0.65	0.66	0.65	0.49	1.00	0.81	0.68	0.54	0.84	0.75			
(26) San Francisco CA	0.68	0.18	0.68	0.85	0.69	0.71	0.66	0.65	0.56	0.34	0.75	0.58	0.75	0.33	0.68	0.55	0.47	0.82	0.31	0.75	0.68	0.54	0.55	0.46	0.81	1.00	0.67	0.45	0.74	0.66			
(27) San Jose CA	0.46	0.11	0.43	0.61	0.43	0.53	0.56	0.51	0.43	0.33	0.47	0.43	0.60	0.23	0.44	0.28	0.47	0.48	0.31	0.55	0.65	0.36	0.49	0.28	0.68	0.67	1.00	0.47	0.68	0.53			
(28) Seattle Ever WA	0.59	0.06	0.43	0.40	0.42	0.54	0.50	0.51	0.43	0.52	0.44	0.38	0.57	0.07	0.44	0.25	0.50	0.44	0.21	0.56	0.59	0.45	0.45	0.33	0.54	0.45	0.47	1.00	0.60	0.56			
(29) Tampa St Pe FL	0.75	0.14	0.66	0.72	0.61	0.71	0.69	0.71	0.51	0.57	0.65	0.54	0.79	0.36	0.60	0.32	0.71	0.42	0.77	0.76	0.66	0.68	0.48	0.84	0.74	0.68	0.60	1.00	0.69				
(30) Washington DC	0.74	0.10	0.69	0.69	0.67	0.73	0.68	0.68	0.54	0.49	0.63	0.49	0.76	0.28	0.67	0.47	0.55	0.71	0.35	0.70	0.67	0.62	0.47	0.46	0.75	0.66	0.53	0.56	0.69	1.00			
Mean Corr	0.66	0.10	0.57	0.62	0.60	0.63	0.63	0.63	0.49	0.43	0.61	0.51	0.64	0.22	0.58	0.41	0.47	0.64	0.34	0.64	0.61	0.50	0.48	0.44	0.63	0.60	0.46	0.44	0.62	0.62	0.59		

Table 2.17: Local Statistics From Baseline Simulation

	Avg Pop		Std Dev Pop		Avg Unempl Rate		Std Dev Unempl Rate		Avg Rent		Std Dev Rent		FE	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model		
Rural/O.O.	76.05	76.18	0.59	0.35	6.44	6.46	1.44	1.54	180.86	203.35	27.19	22.84	0.00	
Atlanta GA	2.26	2.26	0.07	0.06	5.04	4.81	1.84	2.04	186.96	197.48	22.34	21.07	0.07	
Baltimore MD	1.32	1.32	0.01	0.02	5.04	5.56	1.29	1.28	271.77	306.62	69.21	58.08	0.04	
Boston MA	2.34	2.35	0.03	0.05	5.42	5.49	1.74	1.67	310.06	306.20	77.60	62.50	-0.12	
Chicago Gary IL	4.61	4.62	0.07	0.05	6.55	6.59	1.71	1.83	227.98	249.41	31.94	31.00	-0.08	
Cincinnati OH	1.03	1.03	0.01	0.02	5.57	5.57	1.66	1.52	158.26	164.20	12.82	10.97	-0.14	
Cleveland OH	1.07	1.07	0.02	0.02	6.48	6.57	1.52	1.78	203.42	200.10	15.80	17.00	-0.10	
Dallas Fort TX	2.74	2.74	0.04	0.05	5.38	5.36	1.43	1.38	147.40	155.31	10.31	8.88	-0.04	
Denver Bould CO	1.58	1.58	0.04	0.04	4.98	5.03	1.48	1.50	257.65	285.20	55.44	52.87	-0.07	
Detroit MI	2.35	2.36	0.06	0.07	7.40	7.53	2.89	3.25	192.95	221.48	32.62	35.66	-0.15	
Houston Braz TX	2.41	2.41	0.03	0.03	5.79	5.79	1.20	1.20	142.83	150.44	16.69	14.56	-0.03	
Heuston Braz TX	6.98	6.98	0.02	0.02	5.39	5.41	1.43	1.52	135.37	139.54	16.32	13.11	-0.13	
Kansas City MO	6.98	6.09	0.12	0.07	6.83	6.86	2.06	2.04	449.68	525.34	145.38	130.71	0.47	
Los Angeles CA	1.37	1.37	0.03	0.03	5.99	5.76	2.15	1.92	232.44	131.69	82.48	37.16	-0.10	
Miami Hialeca FL	1.79	1.79	0.03	0.04	4.30	4.42	1.38	1.19	215.26	225.57	47.84	38.20	-0.12	
Minneapolis MN	8.82	8.83	0.10	0.07	6.63	6.66	1.53	1.53	432.29	484.78	115.08	115.64	0.05	
New York Nor NY	0.89	0.89	0.02	0.03	5.31	5.30	2.39	1.99	189.69	191.55	50.03	28.81	-0.01	
Orlando FL	3.02	3.03	0.03	0.03	5.76	5.82	1.45	1.46	243.00	256.21	46.13	34.27	-0.06	
Philadelphia PA	1.64	1.64	0.03	0.04	4.74	4.52	1.62	1.81	197.51	205.99	56.82	36.14	-0.02	
Phoenix AZ	1.19	1.19	0.01	0.01	6.35	6.46	1.12	1.37	132.16	139.57	10.32	10.86	-0.11	
Pittsburgh B PA	1.04	1.03	0.03	0.18	6.24	5.41	1.99	2.92	261.28	286.48	62.50	56.25	0.13	
Portland Van OR	1.48	1.48	0.04	0.05	7.55	7.62	2.76	2.38	272.75	350.09	88.45	53.02	0.36	
Riverside Sa CA	0.96	0.96	0.02	0.03	7.35	7.38	2.24	2.00	315.37	345.09	88.39	55.22	0.17	
Sacramento CA	1.39	1.39	0.01	0.02	5.96	6.00	1.68	1.75	140.13	153.60	18.39	17.03	-0.12	
St Louis MO	1.37	1.37	0.03	0.02	5.57	5.51	1.96	1.81	412.30	470.42	134.31	92.74	0.38	
San Diego CA	2.39	2.38	0.06	0.10	5.51	4.79	1.89	1.85	535.95	575.65	159.50	75.42	0.54	
San Francisc CA	0.90	0.90	0.05	0.05	6.04	6.22	2.39	2.41	572.67	540.71	166.55	124.02	-0.09	
San Jose CA	1.64	1.64	0.03	0.03	5.82	5.79	1.40	1.33	324.02	335.93	76.28	54.75	-0.03	
Seattle Ever WA	1.18	1.18	0.03	0.04	5.43	5.34	2.45	2.06	160.55	168.05	44.56	27.27	0.00	
Tampa St Pe FL	2.67	2.67	0.04	0.03	4.00	3.96	0.97	0.86	343.34	345.16	93.50	62.81	-0.05	
Washington DC														

Table 2.18: Correlation of Population and Income Volatility Differences to Baseline Simulation with Rents and Amenities,
By MSA

MSA	Free Mobility		Reduced Rent		Free Mobility		Reduced Rent		Baseline FE μ	Mean Rent	Estimated Rent Elasticity
	Wage Prem/Pop Diff	Correlation	Wage Prem/Pop Diff	Correlation	Per-Capita Income	Volatility Ratio	Per-Capita Income	Volatility Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Rural/O.O.	0.00	0.00	0.99	1.01	0.00	180.86	0.84	0.00	180.86	0.84	
Atlanta GA	-0.41	-0.38	1.19	1.04	0.07	186.96	0.50	0.07	186.96	0.50	
Baltimore MD	-0.82	0.29	1.05	1.00	0.04	271.77	1.38	0.04	271.77	1.38	
Boston MA	0.45	-0.28	0.99	1.00	-0.12	310.06	1.25	-0.12	310.06	1.25	
Chicago Gary IL	0.55	-0.38	0.98	1.01	-0.08	227.98	1.02	-0.08	227.98	1.02	
Cincinnati OH	0.90	-0.21	0.85	1.02	-0.14	158.26	0.52	-0.14	158.26	0.52	
Cleveland OH	0.71	0.82	1.00	1.01	-0.10	203.42	0.94	-0.10	203.42	0.94	
Dallas Fort TX	0.90	-0.74	0.89	1.01	-0.04	147.40	0.28	-0.04	147.40	0.28	
Denver Bould CO	-0.02	0.82	1.01	0.98	-0.07	257.65	0.97	-0.07	257.65	0.97	
Detroit MI	0.59	0.78	0.97	0.97	-0.15	192.95	1.58	-0.15	192.95	1.58	
Houston Braz TX	0.97	-0.77	0.84	1.01	-0.03	142.83	0.48	-0.03	142.83	0.48	
Kansas City MO	0.90	-0.30	0.89	1.02	-0.13	135.37	0.76	-0.13	135.37	0.76	
Los Angeles CA	-0.72	0.67	1.15	0.98	0.47	449.68	2.40	0.47	449.68	2.40	
Miami Hialea FL	0.45	0.67	1.00	0.96	-0.10	252.44	1.40	-0.10	252.44	1.40	
Minneapolis MN	0.56	-0.11	0.99	1.01	-0.12	215.26	1.16	-0.12	215.26	1.16	
New York Nor NY	-0.82	0.25	1.07	1.00	0.05	432.29	1.60	0.05	432.29	1.60	
Orlando FL	0.71	-0.07	0.97	0.97	-0.01	189.69	0.77	-0.01	189.69	0.77	
Philadelphia PA	0.62	-0.89	0.98	1.02	-0.06	243.00	1.11	-0.06	243.00	1.11	
Phoenix AZ	-0.69	0.78	1.15	0.98	-0.02	197.51	0.76	-0.02	197.51	0.76	
Pittsburgh B PA	0.50	-0.01	0.91	1.03	-0.11	132.16	0.76	-0.11	132.16	0.76	
Portland Van OR	0.13	0.45	0.82	0.92	0.13	261.28	1.16	0.13	261.28	1.16	
Riverside Sa CA	-0.20	-0.47	0.88	0.93	0.36	272.75	0.84	0.36	272.75	0.84	
Sacramento CA	-0.51	0.68	1.04	0.98	0.17	315.37	0.87	0.17	315.37	0.87	
St Louis MO	0.79	0.11	0.98	1.02	-0.12	140.13	1.18	-0.12	140.13	1.18	
San Diego CA	-0.96	0.92	1.11	0.96	0.38	412.30	1.18	0.38	412.30	1.18	
San Francisc CA	0.67	0.86	0.93	0.89	0.54	535.95	1.35	0.54	535.95	1.35	
San Jose CA	-0.50	0.82	1.02	0.98	-0.09	572.67	1.00	-0.09	572.67	1.00	
Seattle Ever WA	0.56	0.17	0.99	1.00	-0.03	324.02	0.92	-0.03	324.02	0.92	
Tampa St Pe FL	0.70	-0.13	0.97	1.00	0.00	160.55	1.07	0.00	160.55	1.07	
Washington DC	0.93	-0.68	0.96	1.01	-0.05	343.34	1.00	-0.05	343.34	1.00	
Correlation with (1)					-0.48	-0.51	-0.37	-0.48	-0.51	-0.37	
Correlation with (2)					0.30	0.47	0.47	0.30	0.47	0.47	
Correlation with (3)					0.22	0.33	0.38	0.22	0.33	0.38	
Correlation with (4)					-0.67	-0.53	-0.32	-0.67	-0.53	-0.32	




FIGURE 2.12: Aggregate Statistics from Baseline Simulation

Chapter Appendix: More on the Equilibrium Approximation Method

The PPI Algorithm

The approximation is to the value of choosing a location, conditional on that location being the optimal choice. My application of PPI is similar to that of Sweeting (2010), with the major difference being that instead of firms, I have locations, and the identity of “players” is therefore not consistent. Indeed, the choice facing the worker is which player they want to be, i.e where they want to locate. The difficulty this imposes is twofold. First, the value function approximation parameters, λ , must be allowed to be unique for each location, due to the heterogeneity of locations in terms of local markets and position in space. But second, these parameters must all converge together. I will illustrate the system of equations with a simple two location example.

The choice values can be written as

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

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

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 . γ is Euler’s constant, the mean of the TIEV 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) as well as Sweeting (2010).

Recall that the approximation of v_j uses a linear-in-parameters function of the states. Let Z denote the matrix representing the terms $g_k(x)$ comprising this function for location j ; let λ_j denote the parameters on this function for location j ; \hat{Z}' denotes the expected value of this function in the next period, with the expectation incorporating the exogenous process for the states and the current guess at the choice probabilities. Note that because choice probabilities of atomistic agents correspond to migration flows, the expectation of future population is degenerate; that is, choice probabilities predict one spatial distribution of population. The system of equations is represented in matrix form as:

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

The lefthand side matrix is block diagonal of the row vectors Z_j . This can be rearranged as

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

This matrix equation is evaluated for N points in the basis state space. Shortly I describe how the basis is drawn.

The terms λ are chosen to minimize the distance between the two sides of the approximated contraction. (This is step 5 of the algorithm.) Doing so represents the policy iteration step of PPI. Since the λ terms enter linearly, the distance can be minimized and the λ updated by a simple ordinary least squares formula:

$$\hat{\lambda} = (Q'Q)^{-1}(Q'y)$$

For step 6 of the algorithm, the choice probabilities (“policies” in PPI) can then be updated according to (2.8) using:

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

This update returns the algorithm to step 4, the guess of the policy. This routine is iterated until convergence in the policy function, which is the migration share matrix. Convergence in the parameters λ occurs concomitantly.

State Aggregation

The states used for approximating the conditional value function are the location's own wage, match efficiency α , and population, and the average wage and α of other locations. All states are expressed relative to the value of states in the outside option. I also experimented with use of other locations' population but found that in the closed system, it was too highly correlated with own population to be of use. To see this, consider two nearby locations, j and k at an iteration of PPI away from the fixed point. Conditional of all other states, higher population in k should be bad for j since it means those workers could flow into j and congest utility. The intermediate guess of the policy is 'incorrect' and, without loss of generality, puts too much population into location j at the expense of k . Utility in j is then too low, but at a time when population in k is low. The correlation then picked up by OLS in the step recovering λ is that low population in k is associated with low utility in j . Thus population in k is deemed good, in contradiction of the theory. Because of this bias in the OLS step, convergence was much faster and more reliable when population of other locations was left out of the aggregated state space.

Z and E(Z)

With the five state variables—'own' three and aggregated two—I construct a fully interacted polynomial of the states. For illustration, the terms contained in the polynomial were

$$Z = [1 \ z_1 \ z_2 \ z_3 \ z_4 \ z_5 \ z_1z_2 \ z_1z_3 \ \dots \ z_2z_3 \ z_2z_4 \ \dots \ z_1z_2z_3 \ \dots \ z_1z_2z_3z_4z_5]$$

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 fifth order interaction of all the states. Some columns of

the Z matrix were highly collinear which was troublesome in the inversion of the matrix Q . In these cases, the extraneous columns of Z were simply suppressed to facilitate matrix operations.

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.

The $E(Z)$ matrix is found through simulation using the residuals from the state transition process estimation as the empirical distribution of shocks. To allow for correlation across locations in the labor market states, simulation draws use the empirical shock covariance matrix.

Choosing the States

Just as any regression is determined by its data, 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 estimation, the more data points the better, but these basis points come at a computational cost. Some judgment on the part of the researcher is necessary to parsimoniously choose points that still yield a good approximation to the underlying value function. I draw basis points in three ways. First, I use states that actually occur in the data, 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. (In the first two draws, the states are drawn independently. In the data, the states are naturally correlated.) Which pairs of the three states are negatively correlated and the actual magnitude of the correlation is chosen by a preliminary random number draw for each vector of points. 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.

The same set of basis points is used for all evaluations of the approximating function, both within estimation and for counterfactuals. Otherwise, the simulation error inherent in the draw of basis points would affect comparison of candidate values for the structural parameters, impeding convergence. Similarly, the same basis was used for the counterfactual simulations.

Housing Regulation and Labor Market Growth

3.1 Introduction

The following chapter describes another type of experiment conducted using the modeling and empirical framework of the previous chapter. Unless otherwise noted, the setting is same as the last chapter, and I refer the reader to it for any further details.

The counterfactual experiments from the previous chapter were designed to understand how workers respond to high frequency productivity shocks when faced with different moving and housing cost regimes. The location fixed effects were recalibrated under each experiment so that the average size of the city over the span of the data was unchanged. In contrast, the following experiment seeks to understand how city size would evolve under a different housing market regime.

The intuition of the experiment is straightforward. The housing price elasticities were estimated for each city, and they enter the dynamic model as parameters calibrated to most nearly replicate the response of local housing prices to demand. Changing any of the housing price elasticity parameters creates an alternate sce-

nario: cities with counterfactual housing supply functions. The experiment is then simply to alter the housing price elasticities in a particular way, simulate the model, and see how city sizes evolve. The productivity states—the local wages and matching efficiencies—are left unchanged in order to isolate the impact of the housing supply functions.

Any alteration to a housing price parameter creates some type of experiment, and one of the key advantages of the equilibrium framework is that the simulations predict how all cities evolve, even if the experiment affects the the housing price elasticities of a subset only. While many experiments can yield interesting results and helpful intuition, I focus on a particular type of policy-relevant experiment, the effect of local housing supply regulation on city growth. This question has been examined in the reduced form by Saks (2008), but to my knowledge, never in a structural setting capable of policy simulations. The experiment asks, if cities altered their housing supply regulation, how would it affect their housing price elasticities and ultimately the constraint on city growth?

3.2 Developing The Experiment

Housing supply regulation is an important determinant of housing price volatility, as recognized by the housing literature.¹ However, natural constraints to supply—chiefly the amount of available land—are very important as well. The housing price estimates I reach are based on high frequency demand variation which did not distinguish between inputs in housing construction. To conduct the desired counterfactual, I must first ascertain how much of the price elasticity estimate is due to local regulation. And to do this, I lean heavily on Saiz (2010).

Saiz’s 2010 study “The Geographic Determinants of Housing Supply” addresses this question directly. Using detailed geographic information, he quantifies for many

¹ See Paciorek (2010) and numerous references therein.

U.S. metro areas how much land is available for building; that is, not ocean, inland lake, river or wetland, too steeply sloped for construction, etc.² Saiz also reports a measure of local housing supply regulation, the Wharton Regulation Index (WRI) developed in other studies with his colleagues.³ Saiz reports both of these measures in his Table I; I re-report the measure for the cities in my data in Table 3.1 below. Figure 3.1 plots each of these versus my estimates of housing price elasticity. There is clear positive correlation of both measures with my estimates.

Saiz then uses a cross section of metro areas to regress housing supply elasticity on the determinants of housing supply, including useable land and the WRI. This is exactly the decomposition of housing price elasticities that I need.

Saiz's estimation of housing price elasticities uses data very different from what I employ, which is both explanation and disclaimer. First, his demand side variable is quantity of housing, whereas mine is income. He uses a long horizon of three decades (1970-2000), whereas I use quarterly variation in income. He observes one data point per city and uses the spatial variation in the cross section to identify the average contribution of available land and supply regulation to price elasticity, whereas I use temporal variation in demand to estimate average price elasticity separately for each city. The correlation displayed in Figure 3.1 is reassuring that the techniques are not drastically different, but rather than try to rationalize all the differences and evaluate the strengths and weaknesses of each, I simply normalize his estimates in order to convert my own.

Specifically, my experiment changes all cities' level of regulation to that of the average city; Saiz (2010) reports the average value of the WRI to be -0.1. Table 3.1 displays the conversion of my estimates under this counterfactual regime. Table III of Saiz (2010), column (5), reports the average contribution of the percent unusable

² Note that "usable" land does not necessarily mean vacant.

³ See Gyourko et al. (2008)

land to be 0.516, while the impact of regulation is estimated to be 0.268. (Saiz uses the transformation of $\widehat{WRI}_j = \log(3 + WRI_j)$, where WRI_j is the index as reported in Table 3.1 in his estimating equation. I make the same transformation.) I use these point estimates and the local measures of unusable land and regulation to generate a baseline local housing price elasticity (pe) for each city. This calculation is $pe(baseline)_j = 0.516 \times UL_j + 0.268 \times \log(3 + WRI_j)$, where UL_j is the city j measure of unusable land and WRI_j is the city j regulation measure. The counterfactual is then $pe(counterfactual)_j = 0.516 \times UL_j + 0.268 \times \log(3 - 0.1)$. I use the ratio of the “Saiz-normalized” estimates, $\frac{pe(baseline)_j}{pe(counterfactual)_j}$, to find the percentage change to the price elasticity for each city, which I then use to scale my own price elasticity estimates. These latter two numbers are reported in the last two columns of Table 3.1.

Note that the experiment is twofold: 1) how regulated is the city, and 2) how much does the regulation affect its price elasticity vis-a-vis geographic constraints? As Table 3.1 shows, the size and direction of the experiment can vary between locations for either or both reasons.

To conduct the experiment, I feed in the counterfactual price elasticities⁴ to the PPI algorithm to approximate the value functions under the counterfactual regime. Then I feed in local productivities and the approximating parameters from PPI to simulate the model from 1990 to 2010, generating the endogenous variables: local population, employment, and housing rental price.

⁴ Recall that the price elasticity—the slope of the housing price function—was estimated jointly with the intercept of the housing price function. Tilting the slope without adjusting the intercept would affect the average housing price at any and any demand levels. Therefore I adjust the intercepts of each city’s housing price function so that the housing price in the initial period is consistent with the baseline simulation.

Table 3.1: Conversion of Price Elasticities Using Saiz’s Estimates

City	My Elasticity Estimate	Saiz: Unusable Land (%)	Saiz: Regulation Index	Pct Change to Elasticity When WRI=-0.1	Price Elasticity In Experiments
Atlanta GA	0.50	4.08	0.03	-3.69	0.48
Baltimore MD	1.38	21.87	1.60	-23.69	1.05
Boston MA	1.25	33.90	1.70	-21.95	0.98
Chicago Gary IL	1.02	40.01	0.02	-2.16	0.99
Cincinnati OH	0.52	10.30	-0.58	16.72	0.60
Cleveland OH	0.94	40.50	-0.16	1.15	0.96
Dallas FortW TX	0.28	9.16	-0.23	3.84	0.29
Denver CO	0.97	16.72	0.84	-16.84	0.81
Detroit MI	1.58	24.52	0.05	-3.18	1.53
Houston Braz TX	0.48	8.40	-0.40	9.77	0.53
Kansas City MO	0.76	5.82	-0.79	30.02	0.99
Los Angeles CA	2.40	52.47	0.49	-8.19	2.20
Miami Hialea FL	1.40	76.60	0.94	-10.77	1.25
Minneapolis MN	1.16	19.23	0.38	-9.64	1.05
New York Nor NY	1.60	40.42	0.65	-11.10	1.42
Orlando FL	0.77	36.13	0.32	-7.14	0.71
Philadelphia PA	1.11	10.16	1.13	-21.91	0.87
Phoenix AZ	0.76	13.95	0.61	-14.11	0.66
Pittsburgh B PA	0.76	30.02	0.10	-3.90	0.73
Portland Van OR	1.16	37.54	0.27	-6.29	1.09
Riverside Sa CA	0.84	37.90	0.53	-9.87	0.76
Sacramento CA	0.87	12.05	0.59	-14.13	0.75
St Louis MO	1.18	11.08	-0.73	23.71	1.46
San Diego CA	1.18	63.41	0.46	-7.17	1.09
San Francisco CA	1.35	73.14	0.72	-9.15	1.23
San Jose CA	1.00	63.80	0.21	-4.24	0.95
Seattle Ever WA	0.92	43.63	0.92	-13.66	0.79
Tampa St Pe FL	1.07	41.64	-0.22	2.32	1.10
Washington DC	1.00	13.95	0.31	-9.02	0.91

3.3 Results

This section reports the results from two experiments in this counterfactual regime. In the first experiment, the price elasticities of all cities are changed to reflect their projected values at a mean level of regulation; this results in lesser regulated cities having steeper housing price functions, and more regulated cities having flatter housing price functions. The second experiment only affects the more regulated cities and leaves the lesser regulated cities unchanged.

3.3.1 Experiment 1: All Cities

Table 3.2 reports the results from the first experiment. The simulated endogenous variables are reported for each city. The values reported are the average percentage

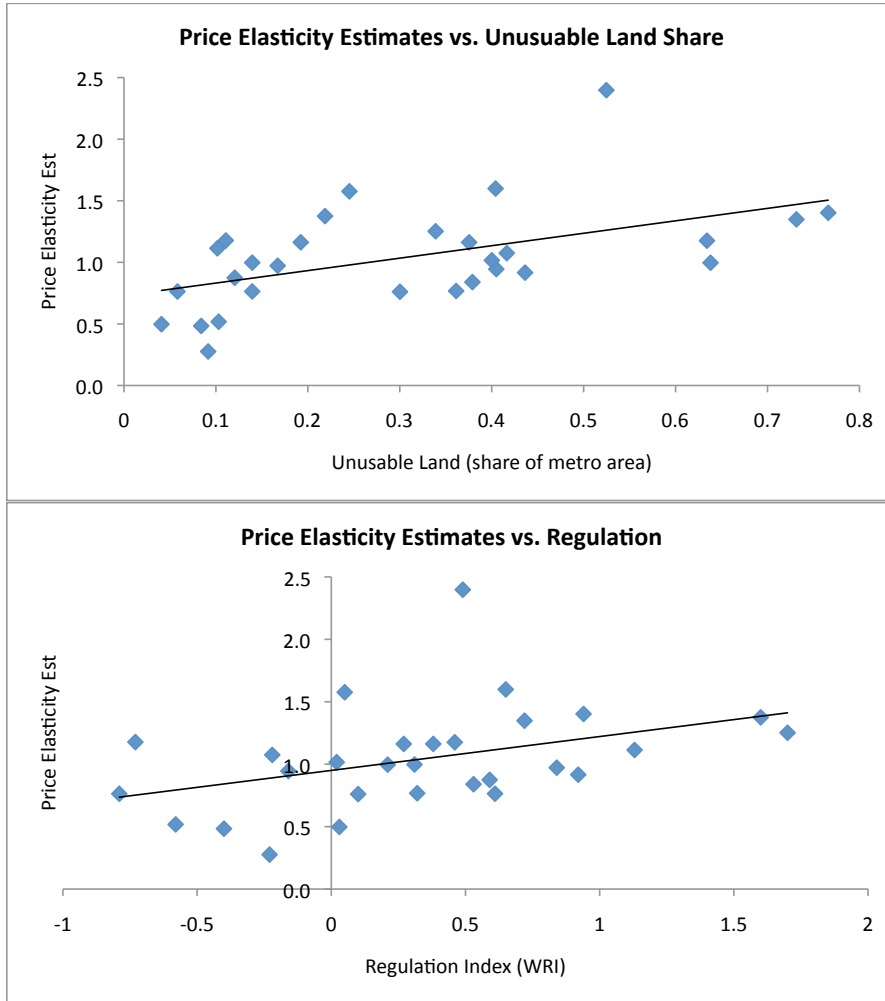


FIGURE 3.1: Price Elasticity Estimates Plotted Against Measures of Unusable Land and Regulation

difference from the baseline simulation of the last two years of the counterfactual simulation; the end of the data period was selected in order to allow the simulation to have taken effect.⁵

By and large, lowering the slope of the housing price function permitted more labor market growth in terms of population (Denver, San Francisco, Seattle), and increasing the slope restricted growth (Dallas, Houston, Tampa). However, it is not

⁵ First running the economy to a new steady state at the beginning of the data—a simulation “burn-in”—had little impact on the end of period values.

a perfect mechanical correlation, and there are several counterexamples of cities that saw no gains or even losses of population despite being deregulated (Boston, New York, Orlando, Miami). Figure 3.2 displays the scatterplot of population change to price elasticity change. The direction of the correlation is clear, but so is the heterogeneity in the size of the effect.

The reason for this is twofold, with both parts highlighting the subtleties of using an experiment in a dynamic equilibrium framework. First recall that the economy is a closed system—adding population anywhere means subtracting it from at least one other place—so some places may see declines despite flatter housing price functions by a simple adding-up constraint. But why so much heterogeneity in the magnitude of the effect? This reflects the dynamic decision process embedded in the model. The values functions which are directing location choices are capturing not just current housing rents, but expectations of future rents, future wages, future employment growth, the option value of a city’s point in space, and so on. Thus some places with minor deregulation (e.g. San Francisco) may be a lot more attractive than cities with larger deregulation (e.g. Philadelphia) because they have more valuable amenities and labor markets, both today and in the future.

The next column of Table 3.2 reports the corresponding city housing price difference. (Recall that the housing prices are expressed as one-period rents.) Here also we see two competing effects in play, something like a price and income effect, because of the equilibrium effects. On one hand, a flatter housing price function means lower housing rents (the price effect), but lower housing rents, especially when coupled with good current and future labor market prospects, means more demand for the city (the income effect), so the net effect is ambiguous a priori. Thus a city like San Francisco sees a rise in its rents despite lower regulation, Dallas sees lower rents despite the increase in regulation, and Cincinnati sees higher rents despite the fall in demand. The housing price effect is not ambiguous if a city is both deregulated *and*

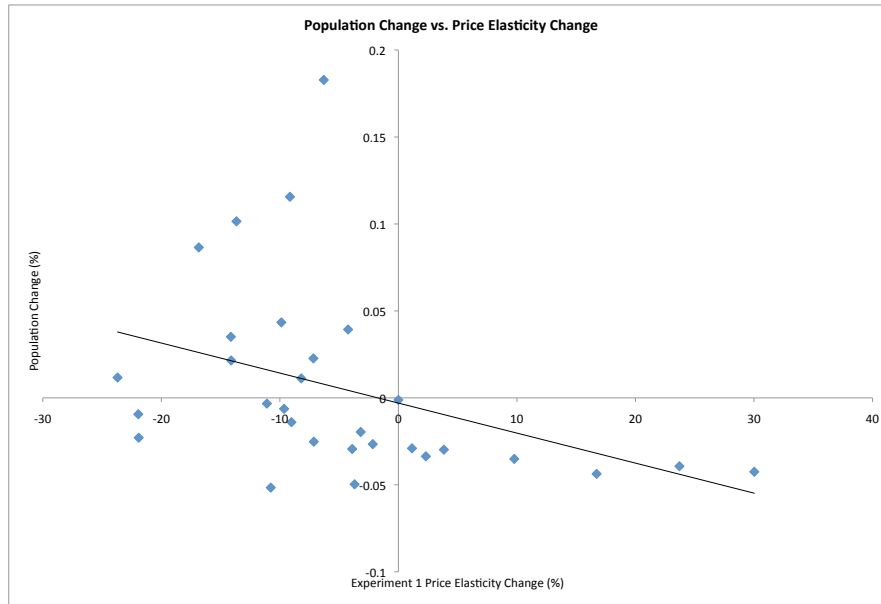


FIGURE 3.2: Experiment 1: Population Change Plotted Against Price Elasticity Change

sees a fall in demand due to the adding-up constraint, such as Atlanta.

The last two columns show the effect on the local labor markets. By the model's design, a rise in population necessarily means a rise in the unemployment rate. However, a rise in population does mean a rise in the employment level, which translates to a rise in the local total income.

The last two rows of Table 3.2 report the aggregate statistics for the nation (the whole system, including the rural/outside option), and the urban system (the named cities only). Compositional effects in the average lead to a slight increase in average housing prices. There are minor effects to the unemployment rates, though there is a modest but meaningful increase in the total income of the economy, resulting from workers on average being able to have more access to higher wage markets.

Table 3.2: Experiment 1: All Cities at Average Regulation Level–Mean Percentage Difference in Simulated Variables at End of Period

City	Experiment Size	Population	House Price	Unemployment Rate (pct point)	Local Total Income
Rural/O.O.	0.00	-0.120	-0.070	-0.030	-0.080
Atlanta GA	-3.69	-4.970	-2.220	-1.560	-3.260
Baltimore MD	-23.69	1.170	-10.330	0.400	0.730
Boston MA	-21.95	-0.940	-6.670	-0.250	-0.670
Chicago Gary IL	-2.16	-2.660	-2.690	-0.590	-1.990
Cincinnati OH	16.72	-4.370	1.460	-1.340	-2.900
Cleveland OH	1.15	-2.900	-1.870	-0.720	-2.100
Dallas FortW TX	3.84	-2.980	-0.480	-0.670	-2.250
Denver CO	-16.84	8.650	-6.060	2.070	6.380
Detroit MI	-3.18	-1.960	-2.360	-0.630	-1.210
Houston Braz TX	9.77	-3.510	0.350	-0.750	-2.710
Kansas City MO	30.02	-4.250	3.740	-1.070	-3.080
Los Angeles CA	-8.19	1.120	-0.870	0.280	0.810
Miami Hialea FL	-10.77	-5.160	-1.580	-1.670	-3.320
Minneapolis MN	-9.64	-0.630	-4.600	-0.230	-0.380
New York Nor NY	-11.10	-0.330	-5.040	-0.080	-0.240
Orlando FL	-7.14	-2.520	-3.170	-0.600	-1.860
Philadelphia PA	-21.91	-2.290	-5.940	-0.730	-1.490
Phoenix AZ	-14.11	2.150	-5.490	0.460	1.650
Pittsburgh B PA	-3.90	-2.940	-2.580	-0.680	-2.180
Portland Van OR	-6.29	18.280	11.040	3.980	13.830
Riverside Sa CA	-9.87	4.340	-1.070	1.200	2.970
Sacramento CA	-14.13	3.510	-2.390	1.070	2.300
St Louis MO	23.71	-3.930	1.670	-0.980	-2.850
San Diego CA	-7.17	2.270	-1.560	0.770	1.440
San Francisco CA	-9.15	11.560	7.360	3.720	7.630
San Jose CA	-4.24	3.930	1.650	1.050	2.730
Seattle Ever WA	-13.66	10.150	-0.120	2.140	7.780
Tampa St Pe FL	2.32	-3.360	-1.400	-1.220	-2.010
Washington DC	-9.02	-1.400	-3.860	-0.460	-0.910
National		0.000	1.730	-0.014	0.066
Urban		0.000	3.070	0.013	0.200

3.3.2 Experiment 2: Reducing Regulation in Highly Regulated Cities Only

The second experiment is the same as the first, except that the lightly regulated cities—those with below average measures of regulation—are left out of the treatment. The reason is to highlight “push” versus “pull” effects; that is, how much of the population in lightly regulated markets is there not because it was pulled by the low regulation and low housing prices, but because it was pushed out of more regulated markets. (Detroit was also left unaltered despite having above-average regulation because it was experiencing declining housing demand; that is, its housing supply curve was likely estimated on the “vertical” downside rather than the upside.)

The results are much the same as the first experiment, so the discussion will focus on the push versus pull effects. Note that several of the lightly regulated markets (Dallas, Houston, Tampa, Cleveland) see losses in population very similar to the first experiment. This suggests that the population losses in the first experiment were not so much due to increases in the housing price function as they were due to being the losers in the population adding-up constraint, because of poorer labor markets and/or amenities. Thus, some measure of their population is there for push reasons, not pull reasons.

Figure 3.3 plots the population change versus the experiment size; the points on the y-axis are the untreated markets. I included two trend lines, one for treated markets, and one for all markets. The line for all markets displays the negative correlation of the first experiment because the untreated markets are still losers of population. For the treated markets, the heterogeneity of the regulation effect is evident, and there is little correlation between size of the experiment and population change. That is, the housing is not a binding constraint all cities. This is because of the varying dynamic values of the cities, as discussed earlier.

3.4 Conclusion

This chapter presented a policy-related counterfactual exercise of my dynamic equilibrium model of local labor markets. The experiment focused on the impact of local housing supply regulation on population growth. Decreasing regulation lowers the constraint to population growth, but does not necessarily increase population growth, as workers are choosing some deregulated cities over others. The effect of changing regulation on housing rents is ambiguous a priori and depends on the size of the change in demand for housing. Finally, a second experiment shows how some population in lightly regulated markets is there from being pushed out of more regulated markets, and not because of the inherent demand for the low regulation cities.

Table 3.3: Experiment 2: Reducing Regulation in Highly Regulated Cities—Mean Percentage Difference in Simulated Variables at End of Period

City	Experiment Size	Population	House Price	Unemployment Rate (pct point)	Local Total Income
Rural/O.O.	0.00	-0.120	-0.070	-0.030	-0.080
Atlanta GA	-3.69	-5.000	-2.230	-1.570	-3.280
Baltimore MD	-23.69	1.110	-10.370	0.380	0.700
Boston MA	-21.95	-1.000	-6.710	-0.270	-0.710
Chicago Gary IL	-2.16	-2.710	-2.730	-0.600	-2.030
Cincinnati OH	0.00	-3.320	-1.140	-1.010	-2.210
Cleveland OH	0.00	-2.980	-2.030	-0.740	-2.150
Dallas FortW TX	0.00	-3.010	-0.670	-0.680	-2.280
Denver CO	-16.84	8.570	-6.100	2.050	6.330
Detroit MI	0.00	-2.420	-2.360	-0.780	-1.500
Houston Braz TX	0.00	-3.370	-1.310	-0.710	-2.590
Kansas City MO	0.00	-2.980	-1.650	-0.750	-2.160
Los Angeles CA	-8.19	1.100	-0.910	0.270	0.790
Miami Hialeah FL	-10.77	-5.210	-1.620	-1.690	-3.350
Minneapolis MN	-9.64	-0.660	-4.620	-0.240	-0.400
New York Nor NY	-11.10	-0.390	-5.100	-0.100	-0.280
Orlando FL	-7.14	-2.560	-3.180	-0.610	-1.890
Philadelphia PA	-21.91	-2.360	-5.980	-0.750	-1.540
Phoenix AZ	-14.11	2.070	-5.530	0.440	1.590
Pittsburgh B PA	-3.90	-3.010	-2.630	-0.700	-2.240
Portland Van OR	-6.29	18.170	10.940	3.950	13.740
Riverside Sa CA	-9.87	4.280	-1.100	1.190	2.930
Sacramento CA	-14.13	3.460	-2.420	1.060	2.270
St Louis MO	0.00	-2.920	-2.490	-0.720	-2.110
San Diego CA	-7.17	2.240	-1.580	0.760	1.420
San Francisco CA	-9.15	11.490	7.300	3.700	7.580
San Jose CA	-4.24	3.890	1.620	1.040	2.700
Seattle Ever WA	-13.66	10.050	-0.180	2.120	7.700
Tampa St Pe FL	0.00	-3.190	-2.050	-1.160	-1.900
Washington DC	-9.02	-1.460	-3.900	-0.490	-0.950
National		0.000	1.610	-0.015	0.640
Urban		0.000	2.850	0.012	0.190

These experiments highlighted the power of the dynamic equilibrium framework, illustrating how to implement locally-heterogenous policy experiments. The framework picks up subtleties of a dynamic equilibrium setting: the price and income effects, adding-up constraints, and option value. The dynamic equilibrium model can be a powerful tool for studying cities and labor market dynamics.

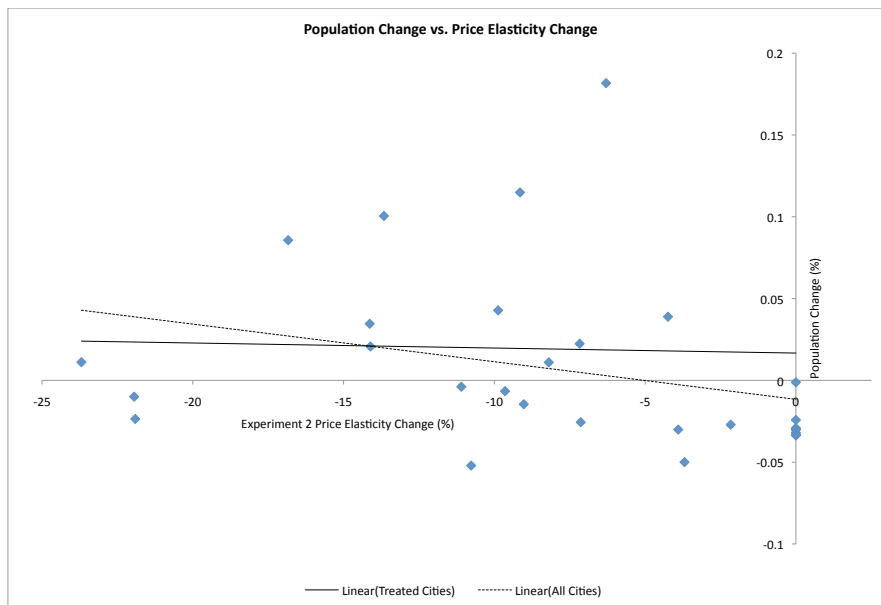


FIGURE 3.3: Experiment 2: Population Change Plotted Against Price Elasticity Change

References

- Aguirregabiria, V. and P. Mira. (2007). "Sequential Estimation of Dynamic Discrete Games." *Econometrica*, 75(1). 1-53.
- Ahn, T., P. Arcidiacono, W. Wessels. (2011). "The Distributional Impacts of Minimum Wage Increases When Both Labor Supply and Labor Demand Are Endogenous." *Journal of Business and Economic Statistics*. 29(1). 12-23.
- Bajari, P., C.L. Benkard and J. Levin. (2007). "Estimating Dynamic Models of Imperfect Competition." *Econometrica*, 75(5). 1331-1370.
- Benitez-Silva, H., G. Hall, G. Hitsch, G. Paultetto, and J. Rust. (2000). "A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems." mimeo, Yale University.
- Berry, S. T. (1994). "Estimating Discrete Choice Models of Product Differentiation." *RAND Journal of Economics*, 25(2). 242-262.
- Bishop, K. C. (2008). "A Dynamic Model of Location Choice and Hedonic Valuation." mimeo, Olin School of Business, Washington University in Saint Louis.
- Blanchard, O.J. and L. F. Katz. (1992) "Regional Evolutions." *Brookings Papers on Economic Activity*, 1. 1-37.
- Carrere, C. (2006). "Revisiting the effects of regional trade agreements on trade flows with proper specification of the gravity model." *European Economic Review*, 50(2) 223-247.
- Chan, S. (2001) "Spatial Lock-in: Do Falling House Prices Constrain Residential Mobility?" *Journal of Urban Economics*, 49. 567-586.
- Dahl, G. (2002). "Mobility and the Returns to Education: Testing a Roy Model With Multiple Markets." *Econometrica*, 70(6). 2367-2420
- Davis, M., J. Fisher, and M. Veracierto. (2010). "The Role of Housing in Labor Reallocation." mimeo, Federal Reserve Bank of Chicago.
- Ferreira, F., J. Gyourko, and J. Tracy. (2010). "Housing Busts and Household Mobility." *Journal of Urban Economics*, 68. 34-45.

- Gemici, A. (2011). "Family Migration and Labor Market Outcomes." mimeo, New York University.
- Glaeser E. L. and J. Gyourko. (2005). "Urban Decline and Durable Housing." *The Journal of Political Economy*. 113(2). 345-365.
- Glaeser, E.L., J. Gyourko, and R.E. Saks. (2006) "Urban Growth and Housing Supply." *Journal of Economic Geography*, 6. 71-89.
- Greenwood, M. J. (1975). "Research on Internal Migration in the United States: A Survey." *Journal of Economic Literature*, 13. 397-433.
- Greenwood, M. J. (1985). "Human Migration: Theory, Models, and Empirical Studies." *Journal of Regional Science*, 25(4). 521-544.
- Greenwood, M. J. and J. Sexton (2012). "On the Temporal Stability of Gravity Models of Internal Migration." mimeo, University of Colorado.
- Gyourko, J. C. Mayer, and T. Sinai. (2006). "Superstar Cities." NBER Working Paper.
- Harris, J.R. and M. P. Todaro. "Migration, Unemployment, and Development: A Two Sector Analysis." *American Economic Review*, 60(1). 126-142.
- Karahan, F. and S. Rhee (2011). "Housing and the Labor Market: The Role of Migration on Aggregate Unemployment." mimeo, University of Pennsylvania.
- Karahan, F. and S. Rhee (2011). "Housing and the Labor Market: The Role of Migration on Aggregate Unemployment." mimeo, University of Pennsylvania.
- Kennan, J. and J.R. Walker. (2011). "The Effect of Expected Income on Individual Migration Decisions." *Econometrica*, 79. 211251.
- Kennan, J. and J.R. Walker. (2011). "The Effect of Expected Income on Individual Migration Decisions." *Econometrica*, 79. 211251.
- Kolaczyk, E. D. (2009). *Statistical Analysis of Network Data: Methods and Models*. New York: Springer.
- Krusell, P. and A.A. Smith, Jr. (1998). "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106(5). 867-896.
- Lee, D. (2005). "An Estimable Dynamic General Equilibrium Model of Work, Schooling, and Occupational Choice." *International Economic Review*, 46. 1-34.
- Lee, D. and K. I. Wolpin. (2006). "Intersectoral Labor Mobility and the Growth of the Service Sector." *Econometrica*, 74. 1-46.

- Lucas, R. E. Jr. and E. C. Prescott. "Equilibrium Search and Unemployment." *Journal of Economic Theory*, 7. 188-209.
- Minnesota Population Center. National Historical Geographic Information System: Pre-release Version 0.1. Minneapolis, MN: University of Minnesota 2004. www.nhgis.org.
- Molloy, R. S., C. L. Smith, and A. Wozniak (2011). "Internal Migration in the United States." *Journal of Economic Perspectives*, 25(3). 173-196.
- Niedercorn, J. H. and B. V. Bechdolt Jr. (1969). "An Economic Derivation of the 'Gravity Law' of Spatial Interaction." *Journal of Regional Science*, 9(2). 273-282.
- Notowidigdo, M. (2010). "The Incidence of Local Labor Demand Shocks." mimeo, MIT.
- Paciorek, A. (2010). "Supply Constraints and Housing Market Dynamics." mimeo, Wharton School, University of Pennsylvania.
- Pakes, A. and P. McGuire. (1994) "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model." *RAND Journal of Economics*. 25(4). 555-589.
- Pesendorfer, M. and P. Schmidt-Dengler. (2008). "Asymptotic Least Squares Estimators for Dynamic Games." *Review of Economic Studies*, 75. 901-928.
- Petrongolo, B. and C. A. Pissarides. (2001). "Looking Into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature*, 39(2). 390-431.
- Pissarides, C. A. (1985) "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages." *American Economic Review*. 75(4). 676-690.
- Pissarides, C. A. and J. Wadsworth (1989). "Unemployment and the Inter-Regional Mobility of Labour." *The Economic Journal*, 99. 739-755.
- Roback, J. (1982) "Wages, Rents, and Quality of Life." *The Journal of Political Economy*. 90(6). 1257-1278.
- Saiz, A. (2010). "The Geographic Determinants of Housing Supply." *The Quarterly Journal of Economics*, 125 (3). 1253-1296.
- Saks, R. E. (2008) "Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth." *Journal of Urban Economics*. 64. 178-195.

- Shimer, R. (2005). "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*. 95(1). 25-49.
- Sinai, T. M. and N. S. Souleles. (2009). "Can Owning a Home Hedge the Risk of Moving?" NBER Working Paper.
- Sjaastad, L. A. (1962). "The Costs and Returns of Human Migration." *Journal of Political Economy*. 70 (5). 80-93.
- Sweeting, A. (2010). "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry." Mimeo. Duke University.
- Topel, R. (1986). "Local Labor Markets." *Journal of Political Economy*. 94(3), S111-S143.
- Van Nieuwerburgh, S. and P-O. Weill. (2010). "Why Has House Price Dispersion Gone Up?" *Review of Economic Studies*, 77(4). 1567-1606.
- Weintraub, G. Y., C. L. Benkard, and B. Van Roy. (2008) "Markov Perfect Industry Dynamics With Many Firms." *Econometrica*. 76(6). 1375-1411.
- Winkler, H. (2011). "The Effect of Homeownership on Geographic Mobility and Labor Market Outcomes." mimeo, Ziman Center for Real Estate, UCLA.
- Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.
- Zipf, G. K. (1946). "The P1 P2/D Hypothesis: On the Intercity Movement of Persons." *American Sociological Review*, 11(6). pp. 677-686

Biography

Kyle Douglas Mangum was born December 15, 1982 in Sellersville, Pennsylvania, and grew up in the Philadelphia area. He graduated summa cum laude from Taylor University, Upland, Indiana in 2005 with a bachelor of arts in economics and political science and a minor in mathematics. At the conclusion of his undergraduate studies, he received the Wall Street Journal Student Achievement Award for excellence in the study of business and economics. After a stint in private industry, he went on to study in the doctoral program in economics at Duke University, Durham, North Carolina, also earning a master of arts in economics in 2008. While a graduate student, he received a full tuition scholarship, Duke's conference travel fellowship, and, twice, Duke's summer research fellowship.