The Composition of Metropolitan Employment and the Correlation of Housing Prices Across Metropolitan Areas

Christian L. Redfearn*
University of California, Berkeley

December 6, 1999

Abstract

"The variation in default rates by region is quite substantial. Default rates in the Northcentral states were about five times as large as default rates in the Southeastern states. These differences reflect the credit rate risk associated with the real estate markets in each of the regions, *the fortunes of the regional economies*, and the loan-to-value ratios and ages of the mortgages."

- Quigley and Van Order (1991), p. 358, italics added

"The pattern of covariances in these returns suggests that portfolio risk can be reduced by geographical diversification ..."

- Quigley and Van Order (1991), p. 361

Implicit in this argument is an assumption that the fundamentals that generate returns to housing are not perfectly correlated across space. However, if metropolitan areas are viewed as small open economies, they will share shocks to the prices of common inports and exports—shocks that may spill over to housing markets. This paper demonstrates that the correlation of returns to residential housing between two metropolitan areas is a function not only of their physical proximitiy but also the similarity of their industrial composition. This implies that as local economies evolve so will the covariance of housing returns—suggesting that the benefits derived from diversification are maximized by considering the industry risk inherent in the current metropolitan areas, not just the correlation of past returns.

*I would like to thank Michael Ash, Dirk Early, John Quigley, Chad Slawner, and the participants at the 1999 International Meetings of the American Real Estate and Urban Economics Association for helpful comments.

1 Introduction

In 1978 the composition of employment in Atlantic City, New Jersey shifted dramatically with the legalization of gambling. Thereafter, hotel and casino expansion further differentiated its economic base from the rest of the state. Through the first half of the next decade, housing price appreciation in Atlantic City followed a course that was less reflective of its location than that of its new industrial composition.

While the other metropolitan areas in New Jersey experienced similar housing price growth, the change in the value of Atlantic City's owner-occupied housing exhibited a pronounced cycle, around the rest of the state. Between 1978 and 1982, aggregate house price growth in Atlantic City grew by 55 percent more than the state average. By 1988 the index of aggregate housing prices in Atlantic City had depreciated relative to the state housing price level by 20 percent, a fall of almost half from its 1982 high. In contrast, average housing prices indexes in other major metropolitan areas in New Jersey deviated from the state average by more than 10 percent during only a handful of quarters during the entire 21-year sample period.

The introduction of gambling into Atlantic City induced both an increase in total employment and a shift in the fraction of employment dedicated to hotels and casinos. Over a short period of adjustment, Atlantic City's cross section of employment resembled that of the gambling cities of Nevada. Furthermore, for the years following the legalization of gambling, movements in house prices in Atlantic City more closely resembled those of Las Vegas and Reno than those of the metropolitan areas immediately surrounding it.

This paper examines the influence of industrial similarity on the correlation of aggregate house prices between metropolitan areas. The model parameterizes correlation as a function of national, regional, state, and local factors. By partitioning the fundamental determinants of house prices in this way it is possible to test the hypothesis that industrial similarity influences the correlation of housing returns independent of the effect of physical proximity. The analysis tests whether metropolitan areas that share more similar industrial composition share more similar movements in housing prices.

A reduced-form model of housing returns is presented. It relies on the spatial scope of the factors of supply and demand to identify the effect of the physical proximity and industry mix. That is, national interest rates, state taxes, and changes in the demand for land in a nearby city, all exert independent effects on the price of local housing. The empirical results suggest that housing returns across metropolitan areas are related by their industrial similarity. In each version of the model tested, the relative similarity between two metropolitan areas' industrial composition is a significant predictor of similarity in housing price movements between the two metropolitan areas. The relationship between two housing markets may change over time as their shared exposure to common shocks varies—historical time series may be less meaningful as the basis of diversification as metropolitan economies evolve.

Section 2 develops the concept of industrial similarity and provides a suggestive anecdote as to its influence by examining the experience of Atlantic City and the legalization of gambling in 1978. Section 3 reviews several papers that focus on the role of industry and space in housing markets. Section 4 describes the data. Section 5 discusses the role of industrial similarity and physical proximity in determining housing prices and outlines the research design. Estimation results are discussed in Section 6. Section 7 previews ongoing extensions and concludes.

2 Atlantic City and "Industrial Distance"

In 1978 gambling was legalized in Atlantic City. This exogenous shock greatly altered the structure of Atlantic City's employment, shifting it dramatically toward services, in particular hotel and casino employment growth was particularly large. Over the next five to seven years, housing prices in Atlantic City behaved remarkably unlike the other five metropolitan areas within New Jersey and more like three metropolitan areas well outside the state and region, New Orleans, Las Vegas, and Reno. I will refer to these cities as "destination cities."[1]

[1] More specifically Las Vegas and Reno are destinations for gambling. While New Orleans also offers some opportunity to gamble, employment is not dominated by this industry. The three share a high percentage of their workforce in services, especially hotels. Overall industrial similarity is discussed below.

Figure 1: Housing Prices-New Jersey MSAs Relative to New Jersey State

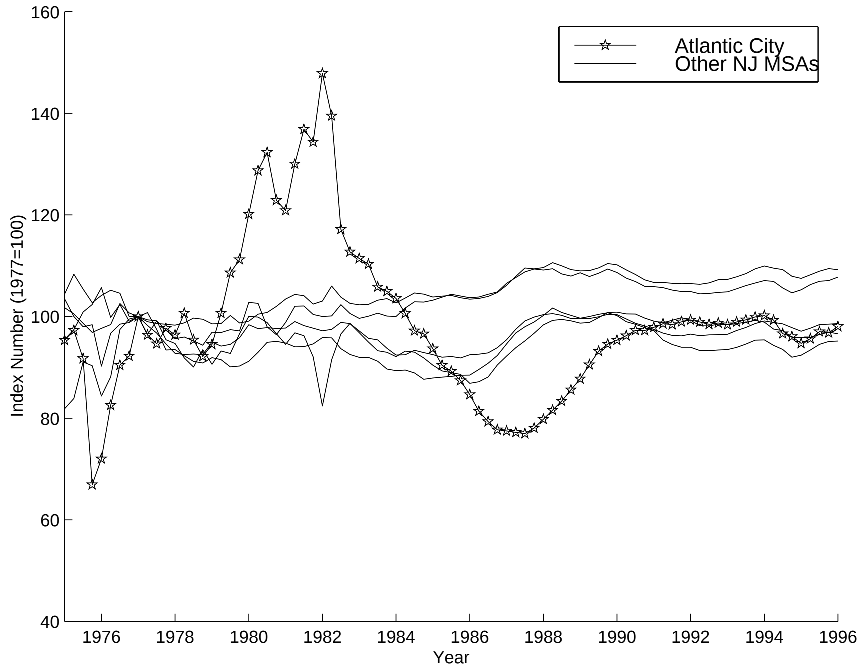

Figure 1 shows the course of aggregate housing prices in New Jersey's metropolitan areas normalized by aggregate New Jersey state housing price levels. Figure 2 shows the same for Atlantic City, Las Vegas, Reno, and New Orleans. Compared to the other of New Jersey's major metropolitan areas, the idiosyncratic movement of Atlantic City's housing prices is striking. Beginning in 1979 the evolution of prices in Atlantic City diverged and did not return to a pattern typical of the other New Jersey cities until more than a decade later. In the interim, housing prices followed a path similar to those of other "destination cities."

Another way to view the information shown in Figures 1 and 2 is to compare the correlation coefficients for the two groups of cities. For the entire sample period, from 1975 to the third quarter of 1996, the average of the correlations between Atlantic City and the five other New Jersey metropolitan area is 0.37; between Atlantic City and the "destination cities" it is 0.21. However, from 1978-1985 the average association with the Las Vegas, Reno, and New Orleans rose to 0.47, while the average correlation with Bergen-Passaic, Middlesex-Hunterdon-Somerset, Monmouth-Ocean City, Newark, and Trenton weakened slightly to 0.34.

Figure 2: Housing Prices-Other MSAs Relative to New Jersey State

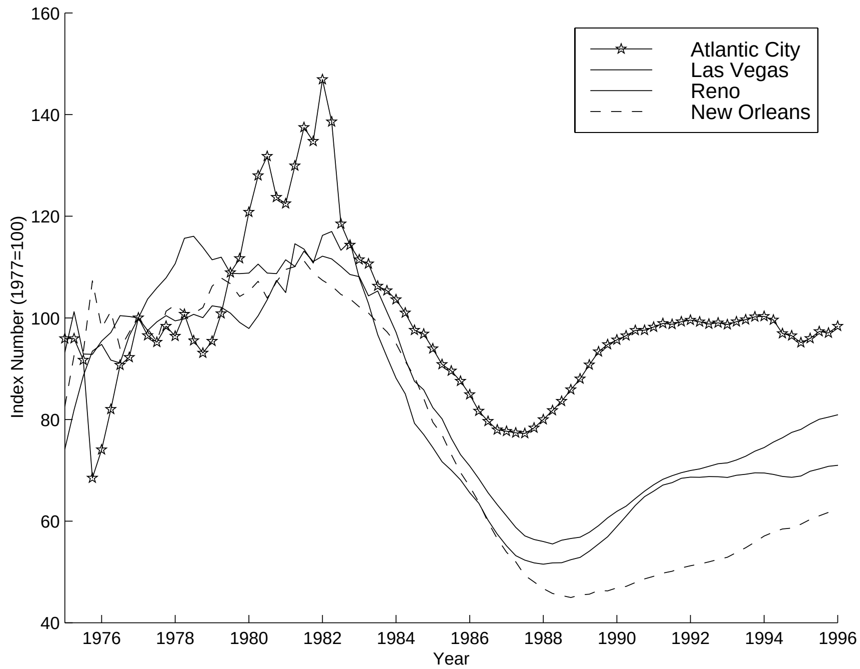

Physical proximity appears to be more influential than industrial similarity when, during the late 1980s, housing prices rose steeply in all of the metropolitan areas surrounding Atlantic City. At this point the correlation of housing prices with the "destination" cities weakens as the growth in Atlantic City's housing prices accelerates with the rest of New Jersey.

Figures 3 and 4 preview a key variable developed below, industrial distance. The figures show this measure of industrial similarity between Atlantic City and the other metropolitan areas over the sample period. Figure 3 shows how Atlantic City's cross section of employment diverged from the rest of the state—how much greater the "industrial distance" became—as the employment in Atlantic City evolved rapidly after the legalization of gambling. Conversely, Figure 4 demonstrates how this measure became smaller—the "industrial distance" narrowed—as the cross section of employment in Atlantic City grew more similar to those of Reno and Las Vegas.

Figure 3: Industrial Distance-New Jersey MSAs Relative to Atlantic City

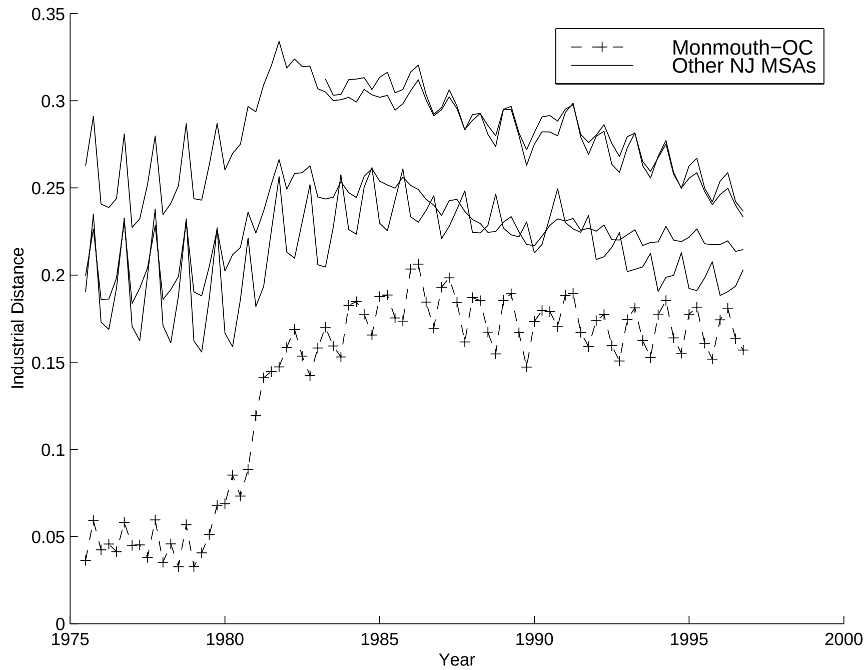

Figure 4: Industrial Distance-Other MSAs Relative to Atlantic City

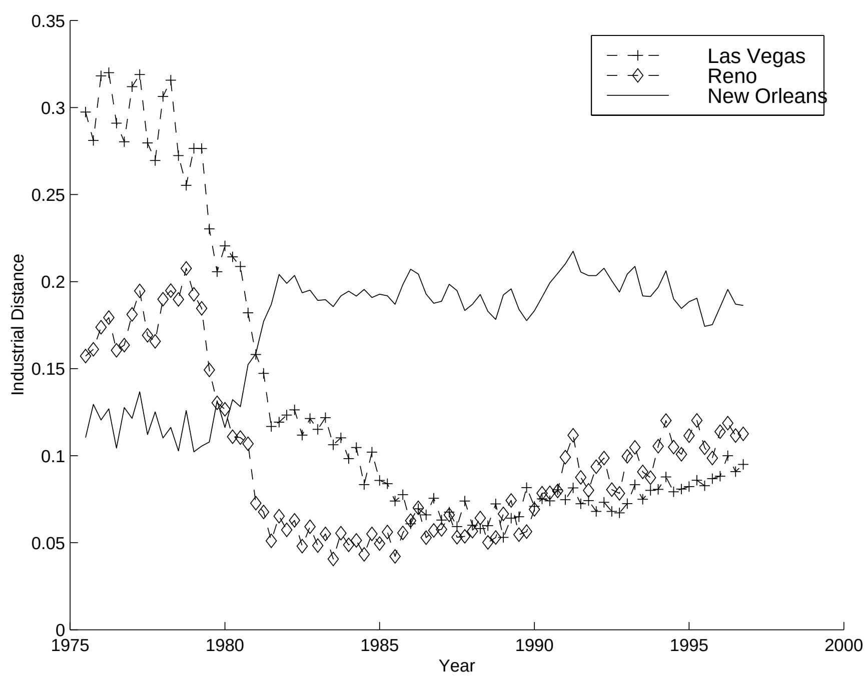

3 Related Research

There is little existing literature on spatial correlation in housing prices across metropolitan areas and even less on spatial correlation across regions. Space and geography have received more research attention recently, but the majority of this work has concentrated on spatial correlation in housing prices within housing markets. [2]

Clapp, Dolde, and Tirtiroglu (1995) find a significant spatial diffusion process in housing prices across neighboring municipalities within larger metropolitan areas in their study of San Francisco and Connecticut. Pollakowski and Ray (1997) obtain similar results, finding a significant lead/lag structure in intrametropolitan housing prices within a large urban area but do not reach the same conclusion for Census divisions.

The economic rationales that support these types of results include informational decay as distance increases, spatial spillovers of shocks, or “ripples” (Meen and Andrew 1998, Cromwell 1992), and broadly similar economic fundamentals within regions.

Others have examined the impact of industry mix on economic outcomes. Terkla and Doeringer (1991) use a modified shift-share analysis to examine the relative importance of industry mix and local cost factors on employment growth. They find “that industry mix interacting with national trends dominates the economic performance of regions over the short-run periods.” Clark (1998) finds that industry-specific shocks are important, but concludes that they are dominated by region-specific shocks. Browne (1992) compares the industrial structures of New England and Texas in an attempt to explain the boom and bust cycles in each. Case and Mayer (1996) find that the share of local employment in manufacturing was a significant influence on the course of house prices within the Boston consolidated metropolitan statistical area.

The mechanism by which housing markets may be influenced by changes in employment is developed in Blanchard and Katz (1992) and, more recently, in Johnes and Hyclak (1999). In both, shocks to demand for locally produced goods leads to changes in local employment, with migration restoring equilibrium. That is, the level of local employment, not the level

[2] Volume 14(3), 1997, of the Journal of Real Estate Finance and Economics is dedicated to spatial correlation at the micro level.

of wages, adjusts to aggregate shocks. The link to local housing markets is through the aggregate demand for housing, falling with net emigration and higher short-run employment. While not specifically addressing industry mix, these papers illustrate how the composition of local economic activity might influence housing markets and suggest that aggregate industry shocks might systematically influence geographically independent housing markets. Clearly industry shocks occur, recent shocks to industries such as motor vehicles and the energy sector and their asymmetric impacts on Detroit and Texas, respectively, are now familiar. The impacts of changes in defense spending influenced New England and Los Angeles disproportionately from the rest of the country as their economies were relatively heavily invested in defense activities.

Abraham, Goetzmann, and Wachter (1994), henceforth referred to as AG&W, establish metropolitan area groupings based on the "closeness" of their housing price movements. Similarity in this regard is agnostic as to any dimension of comparison save quarterly returns to owner-occupied housing. They find "geography dominates economics when it comes to differentiating housing markets." This inference is based on the city groupings they observe by clustering them according to similar movements in house prices. Membership in a typical cluster appears to be primarily a function of geographical proximity and secondarily as a function of gross concentrations of economic activity, i.e. the Rust Belt, the Oil Patch states, etc. Their results and supporting arguments are intuitively appealing, but unsatisfactory on several points.

The first of these is the lack of systematic analysis of their clusters. Visual inspection of the results supports the notion that geography is relatively more important than economics in determining the relationship between outcomes in housing markets across metropolitan areas, but no formal test is undertaken. Asecond weakness is that the research method employed in AG&W is essentially static. While the location of metropolitan areas is fixed over time, their economies are not—as local employment evolves so will the similarity in housing returns. Finally, the AG&W results are based on only 30 metropolitan areas, limiting the extent to which industry concentration effects might be found.[3]

[3] Henderson (1997) finds that "medium-sized" cities are much more likely to specialize in a particular

This paper examines the extent to which industrial similarity influences the correlation of returns to housing across metropolitan areas. The use of historical returns in any analysis is valid only if the data generating process in time-invariant. The results of AG&W suggest that industrial concentration has guided return similarity. This implies that as the underlying economies evlove so may the similarity of returns across metropolitan areas. Testing this proposition is the major goal of this paper.

4 The Data

This paper utilizes three sources of data: the Fannie Mae/Freddie Mac Conventional Mortgage Home Price Indexes and the state and area employment time series from the Bureau of Labor Statistics (BLS).

The Conventional Mortgage Home Price Indexes are based on the combined history of mortgages purchased by Fannie Mae and Freddie Mac. The indexes are constructed using the weighted repeat sales method described in Case and Shiller (1989), and so hold quality constant.[4] The combined mortgage pool includes both refinances as well as house sales, so not all house prices are derived from market transactions, using instead the appraised value when refinancing occurs.[5]

The strength of the Fannie Mae/Freddie Mac indexes is the breadth of their coverage. The data include indexes for all 50 states and the District of Columbia, as well as 151 metropolitan areas. The indexes are published quarterly, beginning in 1975. The indexes employed in this paper extend though the third quarter of 1996.

Employment information comes from the Bureau of Labor Statistics (BLS). These time

industry. It is more likely to find industry effects where concentration is high, expanding the analysis to the largest 150 cities is likely to help identify industry effects.

[4]This method is not without its detractors. Meese and Wallace (1991) find evidence that it substantially overstates price increased in rising markets. Englund, Quigley, and Redfearn (1999a, 1999b) and Gatzlaff and Haurin (1997, 1998) find significant evidence that the use of only those dwellings that sell twice imparts selection bias and results in a biased price index.

[5]For the period 1975-1994, Pollakowski and Ray (1997) report a total pool of 17.5 million mortgages yielding 4.6 million matching transactions. The presence of refinances in the pool does not seem to cause any consistent bias. Where possible the Freddie/Fannie indexes have been compared with commercially available repeat-sales indexes, which use only transactions data. While there are short-run deviations between the indexes, there is a high correlation between the indexes over the entire sample period.

series include aggregate employment data by state and metropolitan area, as well as by 1-, 2-, 3-, and 4-digit SIC codes. The length of the time series is, in general, inversely related to its specificity, with very little comprehensive data existing for most 3- and 4-digit industry classifications. For the broad employment categories the data are excellent; the 1-digit time series are available monthly beginning in 1939.

The final source of data is the location coordinates of each metropolitan area—these are metropolitan "centers" as defined by the U.S. Census. The latitude and longitude of each city are used to establish the physical distance between each pair of cities using an adjustment to Pythagoras' Theorem to account for the curvature of the earth. They are great circle distances that do not consider natural features, such as lakes, mountains, etc. that would influence actual travel distances.

Table 1 describes the data. Panel A summarizes the movement of housing returns for four consecutive five-year periods. It is immediately clear that aggregate house price movements have varied substantially, both across time and Census division. The divisional returns are calculated using the unweighted quarterly housing returns from the metropolitan areas within each of the nine Census divisions. The returns are nominal, which is readily apparent from the consistently high returns during the high-inflation period in the late 1970s. Also noteworthy is the variation in returns within each cross-section, suggestive of idiosyncratic regional movement in housing prices.

Panel B reports the summary statistics for the correlations in housing returns between metropolitan areas, the dependent variable in the first model presented below. Not surprisingly, there is considerable variation in each of the three measures of correlation. Panel B shows that the price indexes vary considerably from quarter to quarter, but generally rise together with the overall price level.

A measure of similarity in metropolitan employment structure is constructed. The measure, referred to in this paper as industrial distance, is the Euclidean distance between the industrial composition of two cities. That is, the industrial distance between city i and city

Table 1: Descriptive Statistics

A. Average Quarterly Returns (nominal, in percent)		Year 1976-1980	1981-1985	1986-1990	1991-1995
Census Division	East North Central	2.23%	0.29%	1.28%	1.26%
	East South Central	2.63	0.87	0.83	1.07
	Mountain	3.20	0.64	0.38	1.84
	Middle Atlantic	2.16	1.90	1.81	0.38
	New England	2.47	3.16	1.46	-0.07
	Pacific	3.90	0.62	2.45	0.09
	South Atlantic	2.17	1.13	1.06	0.74
	West North Central	2.36	0.49	0.69	1.15
	West South Central	3.01	1.01	-0.51	1.03

B. Average Intermetropolitan Correlations	Mean	Std. Dev.	Minimum	Maximum
Index Levels	0.89	0.10	0.08	1.00
Quarterly Returns	0.06	0.21	-0.72	0.93
Year-over-Year Returns	0.23	0.24	-0.71	0.96

C. Share of Metropolitan Employment by Industry (percent)	Mean	Std. Dev.	Minimum	Maximum
Services	24.27%	5.81%	8.96%	50.17%
Wholesale & Retail Trade	23.50	2.57	13.40	37.02
Manufacturing	18.77	8.84	2.87	47.64
Government	17.23	6.07	6.17	42.98
FIRE	5.94	2.25	2.21	19.11
Trade	5.18	1.61	1.58	11.29
Construction	5.09	1.97	1.36	16.65

D. Industrial Distance (times 100)	Mean	Std. Dev.	Minimum	Maximum
Industrial Distance	15.35	7.59	0.53	55.65

E. Physical Distance (in miles)	Mean	Std. Dev.	Minimum	Maximum
Physical Distance	1332.27	897.65	11.67	3899.99

j is defined as

$$(1) \quad IndustrialDistance \equiv ID_{ij} = \left(\sum_k (share_{ikt} - share_{jkt})^2 \right)^{\frac{1}{2}},$$

where k and t index industry and time, respectively. $share_{ikt}$ is the proportion of total metropolitan area i employment involved in industry k at time t.

Industrial distance is calculated using one-digit SIC codes. This is a coarse level of aggregation, which hides much of the specific industrial composition of a metropolitan economy. However, it does identify broad industrial structure—cities identified by concentrations in government, finance, and manufacturing are differentiated.[6]

The components of industry mix are given in Panel C. The share of local employment has been calculated for the one-digit industries with the exception of mining, which employs an insignificant proportion of the labor force in the metropolitan areas included in this research. Panel C emphasizes the substantial variation in metropolitan employment, with the larger variances occurring in the proportion of employment in manufacturing and government.

From the industry shares, industrial distance is calculated, the summary statistics for industrial distance are reported in Panel D. The minimum of .005 describes the close similarity in the shares of employment between Nashua, New Hampshire and Rockford, Illinois in 1975. Nashua is also part of the pair of cities that is "furthest" as measured by industrial distance. At .556, Nashua and Las Vegas, Nevada in 1980 are the most dissimilar metropolitan areas during the sample period.

Panel E reports the summary statistics for the physical distance measure. The maximum distance of 3899 miles may be surprising—it is the distance from Honolulu, Hawaii to Portland, Maine.

5 Spatial Factors, Industry Mix, and Housing Returns

The anecdotal evidence presented in Section 2 suggests that at least two local factors play a significant role in generating correlated outcomes across housing markets. These factors

[6] The appropriate industrial classification level is the topic of ongoing research. The key issue for this research is extent to which similarity arises from the specific bundle of goods a city produces or whether broader types of activities, such as manufacturing or finance, insurance, and real estate (FIRE), are sufficient.

are similarity in industrial composition and physical proximity. In order to estimate the magnitudes of their influence, it is necessary to control for other sources of correlation in housing returns. Previous research provides guidance in identifying them (see especially Reichert (1990)).

In general, changes in inflation, interest rates, and the federal tax code have a nation-wide impact on housing prices. These factors set the parameters of user cost (see Poterba (1984)) and, therefore, influence the price of owner-occupied housing systematically across the country. Gross migratory trends such as the exodus from the Midwest to the Sunbelt or to the Rocky Mountain states are region-wide phenomena, and should have distinct effects on housing prices in the affected areas. State tax and expenditure policies directly, and indirectly, influence house prices within a state's borders. Exposure to common policies that affect the value of housing will cause house prices in different metropolitan areas to be correlated as a result of their colocation within different goverment jurisdications. The focus of this paper is local metropolitan characteristics as an additional source of common movements in housing prices.

If industries concentrate to the point that are entirely located within one geographical area, then it will not be possible to differentiate between their respective influences on housing prices. Identification requires variation in metropolitan industry mix, within and across regions. Ellison and Glaeser (1997) find that spatial concentration varies substantially across industries. Henderson (1997) finds that medium-sized cities are more likely to have a portion of their employment concentrated in one industry than are larger cities. Diminishing economies of scale in combination with increasing transportation costs to suppliers and markets ensure that the production of goods and services is located throughout the country.[7] It is through this variation in local employment that aggregate industry shocks link local economies, even those separated by great physical distances. These shocks, in turn, affect housing markets.

Changes in the health of local economies affect house prices through shifts in the demand for residential real estate. Housing cannot migrate in response to changing demand, so the

[7] See (Kim 1995) for a discussion of long-run trends in deconcentration in manufacturing.

transmission of shocks to the local economy should be apparent in local housing prices as a result of their inelastic supply. It should be noted that supply responses to these shocks are likely to differ across metropolitan areas. To the extent that supply responses differ, this will dampen the measured correlation in housing returns. This will make finding a statistically significant relationship among the factors causing correlated movements in housing prices more difficult.

Neighboring metropolitan areas may experience similar comovements in housing prices due, not to similar industrial compositions, but rather to competition for land. In this sense all houses within commuting distance can be viewed as imperfect substitutes—the degree of substitutability depending in part on the relative proximity of the metropolitan areas. For example, the success of Silicon Valley has been felt in every part of the San Francisco Bay Area, including the Central Valley city of Tracy, a city very far from San Jose in many dimensions, but not distance. In this case, housing price correlation is driven by physical proximity and not the underlying structure of the local economy.

The research design follows from the proposition that existing industry structure is itself a filter through which aggregate industry shocks are passed to local housing markets. This implies that the correlation between the returns to housing between two cities will then be a function of the similarity of their local industrial composition.

AG&W employ the k-means algorithm[8] to allocate metropolitan areas to one of a predetermined number of clusters.[9] In the context of housing returns, the statistical relationship that drives the results of the k-means algorithm is the correlation between housing returns between metropolitan areas. That is, metropolitan areas are grouped to maximize the correlation between the return series within clusters and minimize the correlation across clusters. In order to compare the relative importance of geography and economics, the first model presented in this paper also exploits the correlation in returns to owner-occupied housing.

[8]Simply put, the k-means algorithm clusters data so that the within cluster variation in minimized while the across cluster variation is maximized. The k refers to the specified number of groups into which the data is clustered

[9]AG&W attempt to endogenize the number of clusters, but find little evidence that strongly favored one partition over another. They argue that "meaningful divisions can be identified at several levels of aggregation."

The correlation between the time series of quarterly house price changes in two cities is regressed on measures of physical proximity and industrial similarity. That is,

$$(2) \quad \rho_{ij} = \alpha + \gamma ID_{ij} + \delta PD_{ij} + \beta_C CMSA_{ij} + \beta_S State_{ij} + \beta_R Region_{ij}$$

where PD_{ij} and ID_{ij} are physical distance and the industrial distance between metropolitan areas i and j, respectively. α is an intercept, δ and γ are the effects of the two distances. $CMSA_{ij}$ and $State_{ij}$ are dummy variables that the value one if metropolitan areas i and j are in the same consolidated metropolitan statistical areas or state, respectively. $Region_{ij}$ is also a dummy variable, indicating whether the two cities are in the same region. Region is defined by Census region, division, or Bureau of Economic Analysis region, depending on the specification.

One potential challenge to the preceding analysis is the static treatment of industrial composition. The economies of major U.S. metropolitan areas have changed substantially over the twenty-year sample period. If this transformation has occurred unevenly across metropolitan areas the shared exposure to aggregate industry shocks may have changed substantially over time.

In order to capture the evolution of industrial composition and, therefore, of industrial distance between two metropolitan areas, an analogous model is estimated using quarterly data. The model is dynamic, using quarter-by-quarter returns and industrial distances. The static elements, physical distance, state and region dummies, are defined as above.

Three examples of regressions will help demonstrate the mechanics of the second model. Consider the case that correlation is purely a national effect—possibly through inflation, national economic growth, etc. In this instance, regressing the period t return of city i on the period t return of city j for every combination of cities will obtain the correct estimate of the average intratemporal correlation. That is,

$$(3) \quad r_{it} = \rho \cdot r_{jt}.$$

If, on the other hand, correlation in housing returns in purely an intrastate phenomenon—returns to aggregate metropolitan housing prices move together within state borders—then equation (3) is modified slightly:

$$(4) \quad r_{it} = \rho \cdot I_s \cdot r_{jt},$$

where I_s is an indicator taking the value of one when cities i and j are in the same state and zero if not.

Finally, consider the case that the correlation in housing returns declines linearly with the physical distance between the cities, that is

$$(5) \quad \rho_{ij} = \frac{\rho}{PD_{ij}}.$$

In this case, the regression described by

$$(6) \quad r_{it} = \rho_{ij} \cdot r_{jt} = \left(\frac{\rho}{PD_{ij}}\right) \cdot r_{jt} = \rho \cdot \frac{r_{jt}}{PD_{ij}}$$

will yield an estimate of the distance-based intermetropolitan correlation, ρ.

In addition to testing for national, state, and distance effects, the influences of colocation within a region or consolidated metropolitan area and similarity of industrial mix can also be examined. These tests are undertaken simultaneously by executing the following regression,

$$(7) \quad r_{it} =)\alpha_t + \gamma ID_{ij} + \delta PD_{ij} + \beta_C CMSA_{ij} + \beta_S State_{ij} + \beta_R Region_{ij}) \cdot r_{jt} + \omega_{ijt}.$$

In equation (7), r_{it} and r_{jt} are the returns to housing in metropolitan areas i and j, respectively, at time t.[10] α, $CMSA_{ij}$, $State_{ij}$, and $Region_{ij}$, are defined as above with an added time index, t, on the intercept α. These terms should capture systematic national, regional, and state effects. $\delta \cdot r_{jt}$ is the interaction of the return in city j at time t with the physical distance between two cities—the parameter δ captures the effect of distance-weighted house price changes outside of city i on house price changes in city i. $\gamma \cdot r_{jt}$ and γ are defined analogously, with the important difference that the industrial distance varies with time. ω_{ijt} is white noise.

The model is slightly unusual in that all of the coefficients in this model are the effects due to interaction terms. However, interpretation is straight-forward. For example, if there exists

[10]The return to housing is a measured as a function of the change in the aggregate housing price index. Three series were utilized, quarterly percent change, year-over-year percent change, and kernel estimator of quarterly percent change. The results were robust to choice of return measure.

a common return element to the correlation between metropolitan house prices within the same region, the interaction of the state dummy variable with return should be significant and positive. Similarly, a significantly negative coefficient on the industrial distance interaction term can be interpreted as evidence that correlation in housing returns declines with the dissimilarity between the industrial structure of two cities.

6 Estimation and Results

For each possible pairing of metropolitan areas, the correlation of yearly change in house prices was calculated, as were the physical and industrial distances between the two cities. The industrial distance used in this static analysis is the average of the quarterly industrial distances. A variety of other measures of industrial distance were examined without any significant effect on the results described below.

Table 2 reports the results of estimating several specifications of equation (2). The explanatory power of the models is low, accounting for only fourteen percent of the variation in the observed correlation in housing price changes between two cities. However, each of the parameters is of the predicted sign and all are highly significant.

The coefficients on industrial distance[11] strongly suggest that similarity in industrial structure indeed influences the comovement of return to housing across markets—that *ceteris parabis* greater industrial similarity leads to more highly correlated returns. The estimated parameter on industrial distance is highly stable across the different models, indicating that effect of industrial distance is robust to the parameterization of physical proximity.

The variables capturing spatial proximity—physical distance and the dummies for same CMSA, state, and region—are also highly significant and of the predicted sign. Specifically, the coefficient on physical distance is significant and negative in each model, varying with the combination of included spatial dummies. The difference between the regional partitions is somewhat surprising. The Census divisions are smaller, and allow for specific divisional

[11]The models presented use log transformation of both industrial distance and physical distance. The relationship between these variables and the correlation of housing returns may be highly nonlinear. A limited set of other nonlinear transformations of these distances distance yielded no improvement in explanatory power.

Table 2: Static Regression Results - equation (2)
(t-statistic in parentheses, 8911 observations)

Model	1	2	3	4	5	6	7	8	9	10
R-squared	0.008	0.107	0.110	0.112	0.118	0.128	0.122	0.119	0.137	0.127
Intercept	0.252	1.101	1.026	0.982	0.924	0.758	0.831	0.896	0.643	0.760
	(19.05)	(48.12)	(37.58)	(34.29)	(31.44)	(22.74)	(25.82)	(29.61)	(18.12)	(22.58)
ln(Industrial Dist.)	-0.054		-0.031	-0.032	-0.029	-0.027	-0.029	-0.030	-0.027	-0.029
	(8.35)		(4.99)	(5.20)	(4.74)	(4.48)	(4.73)	(4.92)	(4.49)	(4.82)
ln(Physical Dist.)		-0.108	-0.106	-0.100	-0.092	-0.070	-0.079	-0.088	-0.054	-0.070
		(32.73)	(31.94)	(28.47)	(24.94)	(16.76)	(19.53)	(23.18)	(12.00)	(16.24)
CMSA dummy				0.163				0.125	0.152	0.150
				(5.01)				(3.84)	(4.70)	(4.59)
State dummy					0.164			0.154	0.131	0.094
					(9.15)			(8.56)	(7.32)	(4.92)
Census Reg. dummy						0.121			0.118	
						(13.80)			(13.36)	
Census Div. dummy							0.124			0.106
							(11.31)			(9.08)

effects to be captured, but the models indicate the the broader Census regions are more appropriate.

All of the national, regional, and state effects discussed above are visible in the different specifications. The intercept, interpreted as the national contribution to intermetropolitan correlation, is positive and highly significant. Regional and state effects are also important. Taken together, the static model suggests that the strength of the correlation between housing returns in two metropolitan area dissipates with the dissimilarity of their underlying economies and the physical distance between them.

This static model is restrictive in the sense that the industrial distance between each pair of cities is the average industrial distance over two decades. Clearly, holding relative industrial similarity constant over this period is unrealistic. The second model allows for the underlying metropolitan economies and their relative industrial similarity to evolve over time.

While the form of the second model, defined in equation (7), differs from the first, the interpretation of the estimated coefficients is the same. If industrial similarity is a determinant of housing price correlation, then the interaction term relating the housing returns of two cities as a function of the industrial distance between the two should be negative and significant. The same logic applies for the physical distance-return interaction terms.

Table 3 reports the results of the estimation of the time-varying correlation model, and offers more support for the hypothesis that both linear and industrial distances influence the correlation of housing returns across metropolitan areas.

With the exception of the CMSA dummy, the results presented in Table 3 are similar to those presented in Table 2. The dynamic models explain about 24 percent of the variation in annual housing returns, approximately twice that of the static models. Again, the coefficient on industrial distance is highly significant and stable. The regional dummies are similar in magnitude and similar in relative magnitude, with the coefficient on the state dummy being the largest, with census division and census region following.

The coefficient on the CMSA dummy is much larger relative to the other dummies in the dynamic model than in the static model. This indicates that there is considerable comove-

Table 3: Dynamic Regression Results - equation (7)
(t-statistic in parentheses, 152267 observations)

Model	1	2	3	4	5	6	7	8	9	10
R-squared	0.229	0.240	0.240	0.241	0.241	0.242	0.241	0.242	0.244	0.243
Intercept	0.042	0.041	0.041	0.042	0.041	0.041	0.041	0.042	0.042	0.042
	(222.41)	(221.37)	(221.38)	(222.19)	(221.81)	(221.61)	(221.66)	(222.47)	(222.73)	(222.60)
ln(Industrial Dist.)	-0.025		-0.012	-0.014	-0.011	-0.011	-0.011	-0.013	-0.013	-0.012
	(7.87)		(3.84)	(4.46)	(3.56)	(3.59)	(3.38)	(4.16)	(3.99)	(3.90)
ln(Physical Dist.)		-0.080	-0.080	-0.067	-0.066	-0.049	-0.056	-0.057	-0.026	-0.041
		(46.76)	(46.26)	(35.81)	(33.76)	(21.17)	(25.86)	(27.66)	(10.25)	(17.46)
CMSA dummy				0.254				0.227	0.250	0.244
				(17.28)				(15.33)	(16.79)	(16.38)
State dummy					0.142			0.122	0.105	0.066
					(15.03)			(12.75)	(10.99)	(6.42)
Census Reg. dummy						0.100			0.102	
						(19.32)			(19.48)	
Census Div. dummy							0.112			0.098
							(17.85)			(14.44)

ment of housing returns within large conurbations that is measured only after controlling for the evolution of the underlying economies.

The coefficient on industrial distance is highly significant statistically, but less so economically. In order to understand what the magnitude of the estimate implies a simple counterfactual can be calculated. If two cities were to converge slightly in their industrial composition, that is if the industrial distance were to be reduced by 10 percent, what would the outcome be on aggregate correlation in housing returns? The coefficients from model 9 predict that a one standard deviation increase, from the mean, in industrial similarity would lead to a three percent increase in the correlation between their aggregate housing returns. An symmetric change the physical proximity of the two cities has a much larger effect, increasing the correlation of the returns by 19 percent.

7 Conclusion

The inference that should be drawn from this research is that industrial composition matters in a particular way to the comovement of aggregate housing returns across metropolitan areas. The relationship is consistent with the theory that aggregate industry shocks are transmitted to local economies as a function of the types of economic activity that the city undertakes.

This finding should inform an investor attempting to hedge residential housing risk in two ways. The first is that understanding portfolio risk in residential real estate requires understanding the industry risk metropolitan areas face. The second is that the use of backward-looking correlations may not be wise as the industrial composition of cities evolve. Optimal hedging will have to consider the underlying process that produces returns in housing markets.

Clearly more work is necessary to understand and control for the influence of spatial proximity and is ongoing. Additional research is also under way to better characterize regional business cycles and the way in which shocks are propagated across industries. It is likely that the marginal propensity to consume owner-occupied housing differs across industries, implying that implicit weighting scheme used in calculating the industrial similarity between cities could be improved upon. A more relevant measure would account for asymmetries along these lines.

Advances in the treatment of these variables should only increase the reliability in the results presented in this paper. However, the findings of this paper strongly suggest that the exposure to common industry shocks systematically influence outcomes in residential housing markets.

References

Abraham, J. M., W. N. Goetzmann, and S. M. Wachter (1994): "Homogeneous Groupings of Metropolitan Housing Markets," *Journal of Housing Economics*, 3(3), 186–206.

Blanchard, O. J., and L. F. Katz (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 0(1), 1–61.

Browne, L. E. (1992): "Why New England Went the Way of Texas Rather than California," *New England Economic Review*, 0(0), 23–41.

Case, K. E., and C. J. Mayer (1996): "Housing Price Dynamics within a Metropolitan Area," *Regional Science and Urban Economics*, 26(3-4), 387–407.

Case, K. E., and R. J. Shiller (1989): "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 79(1), 125–37.

Clapp, J. M., W. Dolde, and D. Tirtiroglu (1995): "Imperfect Information and Investor Inferences from Housing Price Dynamics," *Real Estate Economics*, 23(3), 239–69.

Clark, T. E. (1998): "Employment Fluctuations in U.S. Regions and Industries: The Roles of National, Region-Specific, and Industry-Specific Shocks," *Journal of Labor Economics*, 16(1), 202–29.

Cromwell, B. A. (1992): "Does California Drive the West? An Econometric Investigation of Regional Spillovers," *Federal Reserve Bank of San Francisco Economic Review*, 0(2), 13–23.

Ellison, G., and E. L. Glaeser (1997): "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach," *Journal of Political Economy*, 105(5), 889–927.

Englund, P., J. M. Quigley, and C. L. Redfearn (1999a): "Do Housing Transactions Provide Misleading Evidence About the Course of Housing Values?," *unpublished manuscript*, pp. 1–26.

——— (1999b): "The Choice of Methodology for Computing Price Indexes: Comparisons of Temporal Aggregation and Sample Definition," *Journal of Real Estate Finance and Economics*, 19(3), 91–112.

Gatzlaff, D. H., and D. R. Haurin (1997): "Sample Selection Bias and Repeat-Sales Index Estimates," *Journal of Real Estate Finance and Economics*, pp. 33–50.

——— (1998): "Sample Selection and Biases in Local House Value Indices," *Journal of Urban Economics*, 43(2), 199–222.

Henderson, V. (1997): "Medium Size Cities," *Regional Science and Urban Economics*, 27(6), 583–612.

Johnes, G., and T. Hyclak (1999): "House Prices and Regional Labor Markets," *Annals of Regional Science*, 33(1), 33–49.

Kim, S. (1995): "Expansion of Markets and the Geographic Distribution of Economic Activities: The Tends in U.S. Regional Manufacturing Structure, 1860-1987," *Quarterly Journal of Economics*, 110(4), 881–908.

Meen, G., and M. Andrew (1998): *Modelling Regional House Prices: A Review of the Literature. A Report Prepared for the Department of the Environment, Transport, and the Regions.*

Meese, R., and N. Wallace (1991): "Nonparametric Estimation of Dynamic Hedonic Price Models and the Construction of Residential Housing Price Indices," *American Real Estate and Urban Economics Association Journal*, 19(3), 308–32.

Pollakowski, H. O., and T. S. Ray (1997): "Housing Price Diffusion Patterns at Different Aggregation Levels: An Examination of Housing Market Efficiency," *Journal of Housing Research*, 8(1), 107–24.

Poterba, J. M. (1984): "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach," *Quarterly Journal of Economics*, 99(4), 729–52.

Quigley, J. M., and R. Van Order (1991): "Defaults on Mortgage Obligations and Capital Requirements for U.S. Savings INstitutions: A Policy Perspective," *Journal of Public Economics*, 44(3), 353–69.

Reichert, A. K. (1990): "The Impact of Interest Rates, Income, and Employment upon Regional Housing Prices," *Journal of Real Estate Finance and Economics*, 3(4), 373–91.

Terkla, D. G., and P. B. Doeringer (1991): "Explaining Variations in Employment Growth: Structural and Cyclical Change among States and Local Areas," *Journal of Urban Economics*, 29(3), 329–48.