Do Cities and Suburbs Cluster?

William N. Goetzmann
Yale University

Matthew Spiegel
University of California, Berkeley

Susan M. Wachter
University of Pennsylvania

Abstract

This article addresses the issue of how closely the fortunes of suburbs are tied to the fortunes of the central city. We develop housing price indices for most of the zip codes in California and use them in a clustering procedure to determine whether city and suburban housing markets naturally aggregate or move separately. We find that central cities tend to group with their suburbs, suggesting that the housing markets of cities and suburbs are closely linked.

The interrelationship of cities and their suburbs is a central issue to the future of urban areas. In this article, we examine the degree to which suburban housing provides an economic substitute for housing in the core city. Our analysis addresses two fundamental questions: How is a city defined in terms of its inhabitants and its work force? Can the suburbs of a major metropolitan area effectively be considered extensions of their central city, or are they economically distinct regions whose fortunes rise and fall independently of the metropolitan core? The answers to these questions have profound implications for the optimal governance and financial structure of cities and metropolitan areas.

The most basic feature of residential real estate is location. Although two properties at the same location may not have the same price, due to differences in their size and amenities, changes in their value will be closely correlated. When two properties are separated in space but perceived by the market as substitutes for each other, their prices will also fluctuate together. Our study focuses on markets defined as in the latter.

A number of models of urban behavior explore the effect of distance and travel costs on residential housing decisions. In classic monocentric urban models, economically substitutable residential neighborhoods are arrayed in concentric rings around the central business district, with the highest population density in the center and with uniformly falling density with distance from the center. Clearly, such models no longer adequately explain residential spatial patterns.

Many major cities are losing population, while suburban areas are growing rapidly. Most U.S. cities that have declined in population size have done so despite strong regional growth. The growth of suburban population, coupled with urban decline, has given rise to public and scholarly debate over whether cities and suburbs are, in fact, economically linked.

In our study, we examine the degree to which suburban housing markets move together with urban housing markets by applying a clustering algorithm to housing price indices. If price fluctuations in these markets move in lockstep, the implication is that consumers consider them to be substitutes for each other. If not, the implication is that suburbanization has created independent markets. This implication goes beyond housing markets and may provide evidence about whether the fortunes of suburbs are tied to the fortunes of core metropolitan cities.

In this article, we present evidence from a newly developed database to address:

- Whether housing price indices for suburban zip codes in a major metropolitan area group more with one another than with housing price indices of the zip codes of their core cities.
- Whether housing price indices for core city zip codes group more frequently with other core city zip codes than with zip codes from their suburbs.
- Whether suburban zip codes group more frequently with the zip codes of suburbs of other cores than with the zip codes from their own core city.

Because clustering procedures result in statistics whose distribution is not well defined or understood, we use a nonparametric methodology to explore these hypotheses. Specifically, we generate samples of clustering results and test whether suburbs cluster more frequently with the inner city or among themselves.[1] We examine four major metropolitan regions in California: Los Angeles, San Diego, the San Francisco Bay Area, and Santa Barbara. Our results show that, generally, even properties in distant suburbs effectively belong to the same housing market as properties in the central city in these regions.

In the next section, we describe the methodology for constructing local housing price indices. Then we explain the method used to aggregate the indices and present the results of the aggregation. Finally, we report the results of these tests and provide conclusions.

Local Housing Price Indices

The ability to construct reliable housing price indices at the local level is usually hampered by a lack of data. The most commonly used procedure for constructing indices is the repeat-sales methodology, which relies on repeatedly observed transactions for properties in the geographical unit of observation. The smaller the unit, the less precise the index becomes. To address this problem, William Goetzmann and Matthew Spiegel (1996) develop a distance-weighted repeat-sales (DWRS) estimator, which draws upon information in the entire database to estimate returns for the specific unit.

We apply the DWRS procedure to a database of repeated housing sales for the State of California during 1985–94. The data were filtered to eliminate non-arm's-length transactions, erroneous data, transactions with values below $50,000, and high-turnover properties (defined as those sold more than 6 times within 14 years). The estimation procedure results in 906 quarterly indices. Each index is a time-series of estimated arithmetic capital appreciation returns, measured as a percentage, equivalent to holding an equal-weighted portfolio of homes in that zip code area.[2] By equal-weighting zip code indices within a metropolitan area, we are able to construct citywide measures that do not suffer from a representativeness bias. By aggregating up from the zip code level, we can hold the relative importance of each neighborhood constant in the index.

The return indices from DWRS capture the dynamics of home prices within zip code regions. However, the procedure does not account for transaction costs, depreciation, upkeep costs, forgone rents, taxes, and other factors that would reasonably be of concern when calculating an investor's return. For our purposes, temporal price fluctuations capture the substitution effects with interregional absolute price level variations differenced out. Returns vary considerably across California zip codes. For the period 1980–94, the mean return was 4.3 percent, with a cross-sectional standard deviation of 1.3 percent.

Although not completely covering California, the data set includes the majority of the population of the State of California and covers all of its major cities (with the exception of Sacramento), including the four major metropolitan zones in this study: Los Angeles, San Diego, the San Francisco Bay Area, and Santa Barbara. We define a metropolitan core zip code as any code within one of the four major cities in the sample. For San Francisco, we define the core as both Oakland and San Francisco. For Los Angeles, we define the core as both Los Angeles and Long Beach. A suburban zip code is defined as any zip code that lies within 60 miles of the metropolitan core. For ties (for example, between San Diego and Los Angeles), we group the suburb with the larger city. Due to the distance rules, some zip codes have no associated metropolitan core. We find that all four cities reflect widespread common changes in housing prices over the 1980–94 period. Prices rose exponentially during the 1980s and then dropped or remained flat in the 1990s. While Los Angeles city and suburban price levels remained close throughout the study period, other cities manifested significant divergences between city and suburb.

Exhibit 1 reports summary statistics for the eight series and the entire sample, as well as t-tests for differences in means between city cores and suburbs. The t-tests on annualized average returns (using a modification for unequal variances) reject equality of suburb versus core mean returns for two cities: Santa Barbara and San Diego. Los Angeles and San Francisco show no evidence of differences in mean annualized returns between core and suburb for the entire period. Another approach to testing for differences in means is to aggregate the zip codes into time-series indices by whether they belong to a city core or a city suburb; t-tests on these time-series indices show no evidence of differential performance between city and suburb. Exhibit 2 reports the correlations across the city and suburb quarterly return series. Despite that the general trend of all markets was up in the 1980s and flat in the 1990s, the correlations are typically lower than 50 percent.

Exhibit 1

Summary Statistics for Housing Returns in Core and Suburban Zip Codes, 1980–94

Metropolitan Core/ Suburban Index	Number of Zip Codes	Mean	Standard Deviation	Skewness	Kurtosis	t-Value & p-Value on Annualized Returns	t- Value & p-Value on Time-Series
Entire sample*	909	0.043	0.013	-0.668	4.163	NA	NA
San Francisco core	42	0.054	0.005	-0.719	3.348	0.0547	-0.381
San Francisco suburbs	144	0.053	0.008	0.101	2.620	[0.585]	[0.704]
Santa Barbara core	8	0.042	0.008	0.240	1.487	2.25	-0.847
Santa Barbara suburbs	102	0.035	0.017	-0.621	2.810	[0.043]	[0.399]
Los Angeles core	72	0.041	0.009	-0.444	2.607	0.605	0.059
Los Angeles suburbs	278	0.040	0.009	-0.019	2.268	[0.546]	[0.953]
San Diego core	9	0.043	0.003	-0.715	2.613	5.209	-0.551
San Diego suburbs	80	0.035	0.003	-0.637	2.302	[0.000]	[0.583]

Note: Summary statistics are calculated for annualized geometric average returns to the capital appreciation of housing in zip codes making up either metropolitan core or suburban zones surrounding each city. Return indices are estimated via distance-weighted repeat-sales methods, as described in the text; t-values and p-values on annualized returns are based upon a Welch Modified Two-Sample t-test applied to the annualized geometric average returns for core versus suburban zones; t-values and p-values on time-series are the results of the same t-test applied to equal-weighted indices created for each city core and suburb.

*Not all California zip codes can be categorized as related to these four cities; 174 of them were beyond the range of these MSAs.

Exhibit 2

Correlations of Housing Return Indices, 1980–94

San Francisco metropolitan core	1						
Santa Barbara metropolitan core	0.046	1					
Los Angeles metropolitan core	0.523	0.118	1				
San Diego metropolitan core	0.431	0.211	0.315	1			
San Francisco suburban zones	0.457	-0.049	0.559	0.135	1		
Santa Barbara suburban zones	0.285	0.052	0.484	0.128	0.406	1	
Los Angeles suburban zones	0.467	0.053	0.735	0.565	0.534	0.396	1
San Diego suburban zones	0.422	-0.097	0.493	0.651	0.36	0.279	0.615

Note: Correlations are calculated for time-series of capital appreciation returns of housing in zip codes making up either metropolitan core or suburban zones surrounding each city. Return indices are estimated via distance-weighted repeat-sales methods, as described in the text.

Clustering Methodology

Another approach to determining which locations move together is the clustering methodology developed by Jesse Abraham, William Goetzmann, and Susan Wachter (1994). In this procedure, we define a distance metric in the space of returns. Unlike correlation studies, this study does not require long time-series for accuracy. It uses a *bootstrapping* method to test hypotheses such as whether housing-price-appreciation patterns in any suburban neighborhood of a given city are more similar to the patterns in other suburban neighborhoods or to the patterns in core neighborhoods of the same city. Bootstrapping allows researchers to perform statistical significance tests and construct confidence intervals around a test statistic even when the distribution of the test statistic is unknown.

In this case, the test statistic is a matrix of observations on the frequency with which two zip code areas cluster together, suggesting that the housing-price-appreciation patterns in the two areas are more similar to each other than to the patterns in other areas. We use a *k-means* clustering technique to construct clusters of zip codes whose housing-price-appreciation patterns are most similar to one another and most different from the price-appreciation patterns in other zip codes. K-means cluster analysis is a classification methodology that identifies natural groupings of objects (in this case, neighborhood housing price returns), so that objects in the same group are more like one another than they are like objects in other groups. The measure used in cluster analysis is Euclidean distance, which is the square root of the sum of the squared differences between values for local housing price returns and their respective group center.

Thus, for 10 years of annual data, the distance between 2 zip codes is the Euclidean distance in 10-dimensional space. Two zip codes with exactly the same series of returns will have a metric distance of zero. Two zip codes whose returns move in opposite directions will have a relatively large distance measure. We compute the distance measure for all zip codes and then use the k-means clustering algorithm to aggregate these into a prespecified number of groups; the details of this clustering procedure are explained in Abraham, Goetzmann, and Wachter (1994) and Goetzmann and Wachter (1995).

The question that we answer using bootstrapping is: Do two zip code areas cluster together because they truly have systematically similar housing-price-appreciation patterns or simply because their price-appreciation patterns seem similar in our particular observations? Furthermore, are these similarities due to a random sampling error or some other nonsystematic oddity in the data?

In more detail, we can regard the estimated quarterly housing-price-appreciation rates that we compute for each zip code area as having been drawn from a population of quarterly changes for that zip code. If two zip codes truly have systematically similar housing-price-appreciation patterns, then the underlying distributions from which each zip code's quarterly changes are drawn will be similar and correlated. Unfortunately, we cannot observe these underlying distributions. Instead, the only observations that we have are estimated quarterly changes. Bootstrapping, however, enables us to simulate the underlying distributions only on the basis of the estimated quarterly changes that we observe. We can then test whether the underlying distributions are actually similar and correlated or whether the observed price-appreciation patterns appear similar and correlated and are due to random sampling errors.

To simulate the underlying distributions, we create a *pseudo-history* by sampling with replacement from the entire set of estimated quarterly price-appreciation rates. That is, we randomly select 1 quarterly price change from the set of 40 quarterly changes (4 quarters for each of 10 years) estimated for a particular zip code and consider that observation—along with the observations for the same year and quarter in all other zip codes—as the first quarter of the pseudo-history. We then replace the selected observation, make a second random selection from the same set of 40 quarterly changes for the same zip code, and consider that observation—and the corresponding observations for all other zip codes—as the second quarter of the pseudo-history. By repeating this method, we create one complete pseudo-history of 40 quarters. This pseudo-history has the same expected mean, variance, correlation across zip codes, and other properties as the true but unobservable underlying distribution, so it simulates one observation for the true underlying distribution. To simulate the entire distribution, we repeat the process many times to create a large number of pseudo-histories (1,000 in this application). Then we can derive reliable estimates of the underlying distribution by estimating sample statistics, including clusters, from this sample of pseudo-histories.

After generating our sample of 1,000 pseudo-histories, we use the k-means clustering method to group together zip codes with similar price-appreciation patterns. In this way, one zip code can cluster with another in as many as 1,000 of the pseudo-histories—strongly suggesting that the underlying price-appreciation patterns are truly similar—or as few as none of them—suggesting that the underlying price-appreciation patterns are not at all similar. The relative frequency with which each pair of zip codes clusters together measures the similarity of the price-appreciation patterns.

We use this methodology to establish whether:

- Suburban zip codes group with their core city zip codes.
- One core city's zip codes group with other core city zip codes, rather than with its suburban zip codes.
- Suburban zip codes from one core city group with suburban zip codes from other cities, rather than with their core city zip codes.

To perform a statistical test of these hypotheses, we use a statistical procedure called the Wilcoxon rank-sum test, which gives a test statistic with the familiar standard normal distribution that enables us to perform a straightforward hypothesis test.

Results of Testing for Suburb-City Relationships: Which Locations Move Together?

In this section, we report results of tests for suburb-city relationships. We use the bootstrapping methodology of Abraham, Goetzmann, and Wachter (1994) and Goetzmann and Wachter (1995) to generate distributions about the k-means clusters. For each sample, we rerun the k-means algorithm, specifying 20 clusters. As described above, this is done 1,000 times, and the results are saved. We then count how often any two zip codes cluster with each other. If two zip codes are close substitutes, then their returns will move together. Thus they will tend to cluster with each other, even when some years are left out and others appear multiple times in the sample.

With this frequency-of-association matrix, we then determine whether associations between suburbs of a single metropolitan area are higher, on average, than associations between the suburbs and the core metropolitan area. The Wilcoxon rank-sum test for unequal samples is used to test the null hypothesis that suburbs cluster more frequently with their inner city than with one another. In particular, we compare the vector of association frequencies between suburb-suburb zip code pairs with the vector of association frequencies between city-suburb zip code pairs. For suburb-suburb pairs, we do not allow zip codes to be paired with themselves. This comparison tests whether city and suburban zip codes are linked. A link suggests that housing markets in suburbs of a specific metropolitan area have more in common with their core city housing markets than with one another.

The results of this test are reported in exhibit 3. We find that for the largest of the four metropolitan areas, Los Angeles, housing markets in suburbs do group more with Los Angeles central city zip codes than with other suburban zip codes. For San Francisco and San Diego, we find the contrary. For Santa Barbara, the evidence is weak and no conclusion is possible.

Exhibit 3

Tests of Equality of Means for the Frequency of Associations of Suburb-City and Suburb-Suburb k-Means Clusters

Metropolitan Area	**Wilcoxon Z-Statistic**	**Probability Value of Rejecting the Null**
San Francisco	-2.77	0.0064
Santa Barbara	-5.78	0.5633
Los Angeles	28.66	0
San Diego	-3.85	0

Note: A positive Z-statistic indicates that the bootstrapped association frequency between suburb and core city zip codes is higher than the frequency between suburb and suburb. For suburb-suburb associations, the own-zip-code association frequency of 100 percent is omitted.

The same methodology allows us to test for intercity relationships. We can test whether housing price indices in core zip codes in different cities are more closely associated with core zip codes in other cities than they are with zip codes in their own suburbs. The result for each city is reported in exhibit 4. The exhibit shows that the Z-statistic for equality of association frequencies between the core zip codes of Los Angeles and San Francisco and the core zip codes of Los Angeles and suburban zip codes of Los Angeles is -4.45. The negative number indicates that the core-suburb association frequencies are higher than the core-core association frequencies, and the number's magnitude is greater than 2, indicating a high probability level. In other words, the Los Angeles core is significantly more closely related to its suburbs than to the San Francisco core. Most of the test statistics in exhibit 4 are significantly negative, indicating that this is a general result. The exception is Santa Barbara, which has a positive association with the San Francisco core.[3]

Exhibit 4

Tests of Equality of Means for the Frequency of Associations of Core-to-Core Versus Suburb-Core k-Means Clusters

Metropolitan Area	**SF Core/ Suburb**	**SB Core/ Suburb**	**LA Core/ Suburb**	**SD Core/ Suburb**
San Francisco (SF)		0.74	-4.45*	-5.56*
Santa Barbara (SB)	9.95*		-0.22	-0.03
Los Angeles (LA)	-6.58*	-14.61*		-11.06*
San Diego (SD)	-2.08	-3.75*	-4.10*	

Note: A negative Z-statistic indicates that the association frequency between core-to-core is lower than the frequency between suburb-to-core city zip codes. For example, the value of -6.58 in the third row and first column indicates that the San Francisco core-suburb association frequencies are significantly higher than the association frequencies between the San Francisco core zip codes and the Los Angeles core zip codes. *Indicates significance levels exceed 99 percent.

Exhibit 5 reports the suburb-suburb relationships in the same manner as exhibit 4. In almost all cases, the connection between suburb and city is tighter than the relation across cities, whether core to core or suburb to suburb. On the basis of these findings for major cities in California, on average, we find that suburbs group with their own core cities.

Exhibit 5

Tests of Equality of Means for the Frequency of Associations of Suburb-to-Suburb Versus Suburb-Core k-Means Clusters

Metropolitan Area	SF Core/ Suburb	SB Core/ Suburb	LA Core/ Suburb	SD Core/ Suburb
San Francisco (SF)		-34.56*	-29.90*	-26.12*
Santa Barbara (SB)	2.61*		-0.07	0.68
Los Angeles (LA)	-54.90*	-72.10*		-47.93*
San Diego (SD)	-7.46*	-11.33*	-6.50*	

Note: A negative Z-statistic indicates that the association frequency between suburb-to-suburb is lower than the frequency between suburb-to-core city zip codes. For example, the value of -54.90 in the third row and first column indicates that the San Francisco core-suburb association frequencies are significantly higher than the association frequencies between the San Francisco suburb zip codes and the Los Angeles suburb zip codes.
*Indicates significance levels exceed 99 percent.

Conclusions

In this article, we report results that are part of an ongoing quantitative study of the relationship between housing markets in suburbs and their core cities. We use estimates of residential housing returns to develop measures of the substitutability of spatial markets. Price fluctuations in housing markets can indicate which communities are close substitutes for one another. While politics and geography define cities, the market defines their boundaries. These market boundaries may be contiguous, and they may group suburbs with their core cities.

To define the boundaries of housing markets and to analyze whether these markets in cities and suburbs aggregate or move separately, we develop zip-code-based housing price indices for most of California. The database provides local housing price indices within four major metropolitan zones: Los Angeles, San Diego, the San Francisco Bay Area, and Santa Barbara. Each index is a time-series of estimated arithmetic capital appreciation returns, measured as a percentage, equivalent to holding an equal-weighted portfolio of homes in that zip code, for the period 1980–94.

We use clustering techniques to test whether housing markets in central cities have more in common with those of other central cities than with housing markets in their surrounding suburbs. We perform similar tests to determine whether housing markets in the suburbs of a metropolitan area have more in common with those of suburbs of other metropolitan areas than those in their core city. We address this by implementing a procedure, based on housing price fluctuations, to test whether central core zip codes group more with central cities of other metropolitan areas than with their suburbs, and whether suburbs group more with suburbs of other central cities than their central city. We found, with the exception of Santa Barbara, that housing markets within the central city are significantly more closely related to those in their suburbs than to those in other central cities. We also found that the connection between suburb and city is tighter than the relation of suburb to suburb across metropolitan areas, again with the exception of Santa Barbara, on the basis of housing price fluctuations. For most areas, whether core-to-core or suburb-to-suburb, zip codes group with other zip codes from their metropolitan region and suburbs generally group with their metropolitan core, providing evidence that is consistent with housing markets that are spatially linked.

Authors

William N. Goetzmann is professor of finance and real estate at the Yale School of Management. He is an expert in a diverse range of investments, including art as investment, the performance of the single-family home in the investment portfolio, long-term patterns in the stock market, and the history of the world's financial markets. His co-authored work with Matthew Spiegel and Susan M. Wachter has focused on the development of econometric methods for constructing indices of property returns and on the analysis of the geographical aggregation of real estate markets.

Matthew Spiegel, an associate professor of finance at the University of California at Berkeley's Haas School of Business, has written a number of influential papers on finance, real estate, and economics. His editorial assignments include membership on program committees for the Winter Finance Conference and the Western Finance Meetings. In addition, he co-edits the Journal of Financial Markets *and is an associate editor of the* Review of Financial Studies.

Susan M. Wachter is chair of the real estate department and professor of real estate and finance at The Wharton School of the University of Pennsylvania. Dr. Wachter is co-editor of Real Estate Economics and past president of the American Real Estate and Urban Economics Association (AREUEA).

Notes

1. The difficulty with this approach is that we must rely on the existing definitions of the city (specified by the zip code directory). Thus we cannot distinguish between a neighboring urban community and a suburb, and the test is actually a test of whether residential indices for zip codes surrounding the city move more closely with the city or with one another.

2. Because of well-known bias in the repeat-sales procedure due to Jensen's inequality, we use an approximation to the arithmetic return developed in Goetzmann (1992). See also Case, Pollakowski, and Wachter (1996).

3. See Goetzmann, Spiegel, and Wachter (1997) for a formal discussion of the methodology and for maps of clustering results.

References

Abraham, Jesse, William Goetzmann, and Susan M. Wachter. 1994. "Homogeneous Groupings of Metropolitan Housing Markets," *Journal of Housing Economics* 3:186–206.

Case, Bradford, Henry O. Pollakowski, and Susan M. Wachter. 1996. "Frequency of Transaction and Housing Price Modeling," *Journal of Real Estate Finance and Economics*.

Goetzmann, William. 1992. "The Accuracy of Real Estate Indices: Repeat Sales Estimators." *Journal of Real Estate Finance and Economics* (Spring).

Goetzmann, William, and Matthew Spiegel. 1996. "A Spatial Analysis of Residential Housing Returns," *Journal of Real Estate Finance and Economics* (forthcoming).

———. 1995. "Non-Temporal Components of Residential Real Estate Appreciation," *Review of Economics and Statistics* (September):199–206.

Goetzmann, William, and Susan M. Wachter. 1995. "Clustering Methods for Real Estate Portfolios," *Real Estate Economics* 23(3):271–310.

Goetzmann, William, Matthew Spiegel, and Susan M. Wachter. 1997. "Cities and Suburbs." Wharton Real Estate Center Working Paper, Presented at the AREUEA Meeting, New Orleans, LA, January.