Testing for Monocentricity

Daniel P. McMillen
Department of Economics (MC 144)
University of Illinois at Chicago
601 S. Morgan St.
Chicago, IL 60607

August 11, 2004

Prepared for *A Companion to Urban Economics*,
edited by Richard Arnott and Daniel McMillen

1. Introduction

The monocentric city model of Muth (1969) and Mills (1972) is still the dominant model of urban spatial structure. Its central predictions – that population density, land values, and house prices fall with distance from the city center – have been the subject of repeated empirical testing. Indeed, one objective of the model was to explain a set of stylized empirical facts, and extensions of the model were developed in response to empirical testing. This close cooperation between theory and empirical work is one of the hallmarks of the field of urban economics

A consensus appears to have developed that the monocentric city model is no longer an accurate depiction of urban spatial structure. This view is partly due to the unrealistic nature of the model's assumptions. Clearly not everyone works in the central city, and modern urban areas may be viewed more aptly as polycentric rather than monocentric. The central behavioral assumption of the model, that workers attempt to minimize their commuting cost, is called into question by the literature on "wasteful commuting" (Hamilton, 1982). O'Sullivan's (2002) popular textbook perpetuates the notion that the monocentric city model is designed to explain an old-fashioned city by listing as one of the assumptions "horse-drawn wagons," implying that the model does not apply to a modern city with cars.

In this chapter, I review some of the empirical evidence on the monocentric city model's predictions. I contend that the demise of the model is exaggerated. The central city still dominates urban spatial patterns, and the basic insights of the model apply to more complex polycentric cities. Much of the apparent decline in the explanatory power of the monocentric city model is actually a misunderstanding of the empirical evidence.

And importantly, many of the ways in which the model now fails are in fact explained by the comparative-statics predictions of the model itself. Although the model is over-simplified, it remains a useful analytical tool requiring only modest modifications to be remarkably accurate.

2. Empirical Predictions

2.1 Consumers

In the Muth-Mills version of the monocentric city model, consumers receive utility from housing and other goods. Housing is an abstract commodity in this model. It combines land, square footage, and all other housing characteristics into a single measure. The durability of housing is ignored because the static nature of the model is designed to focus on long-run equilibrium results. Each household has a worker who commutes each day to the central business district (CBD). The simplest version of the model includes neither congestion nor time costs of commuting. Instead, each round trip to the CBD costs \$*t* per mile. Since consumers have no direct preferences for one location over another, they would all try to live in the CBD in order to minimize their commuting costs unless house prices adjust to keep them indifferent between locations. In equilibrium, the price of housing must fall with distance from the CBD:

$$\frac{\partial P_h(d)}{\partial d} = \frac{-t}{H(d)} \tag{1}$$

where $P_h(d)$ is the price (or rent since the distinction is irrelevant in a static world) and $H(d)$ is the quantity of housing at a site d miles from the CBD.

Equation (1) is simply a formula for the slope of a function depicting the relationship between the price of housing and distance from the CBD. If the quantity of housing does not vary by location, equation (1) predicts that the price of housing is a linear function of distance. However, the model predicts that $H(d)$ is lower near the CBD than in more distant locations because consumers substitute away from housing and toward other goods when P_h is high. This substitution implies a particular shape for the house price function: the slope is steep when H is low, meaning that prices rise rapidly when approaching the CBD.

The first major implication of the monocentric model, then, is that, for a group of identical households, house prices decline with distance from the CBD according to a smooth, convex function. Figure 1 shows the general form of the function. In a world with different types of households, the general form of the relationship will continue to look much like the function shown in Figure 1 because the equilibrium house price function is the upper envelope of the functions for each household type. Since the quantity of housing is low where the price of housing is high, the function for the quantity of housing is upward sloping. Finally, since consumers substitute toward other goods as they consume less housing, the model predicts that Figure 1 also represents consumption of the non-housing good.

Insert Figure 1 Here

One critical point to bear in mind is that the monocentric city model makes no direct predictions for the *value* of housing. The value of housing is the product of price times quantity, $V_h(d) = P_h(d)H(d)$. Since price falls with distance and quantity rises, the value of housing can go either way. Once we allow for differences in income among consumers, the model predicts that the value of housing is high where higher-income households choose to live. Again, this relationship between house values and distance from the CBD is ambiguous. Whether house values rise or fall with distance from the CBD has no direct empirical relevance for the monocentric city model. The trick is to isolate the price of a unit of housing from the quantity – a nearly hopeless task since housing is a complex, multi-dimensional good that cannot be measured simply.

2.2 Producers

Housing producers combine land and capital to produce housing. Producers will pay more for land near the CBD because consumers will pay more for housing there. Figure 1 thus can depict the equilibrium relationship between land values and distance from the CBD: land values decline at a decreasing rate with distance. Just as consumers substitute away from housing and toward other goods near the CBD, producers substitute away from land and toward capital where the price of land is high. This result implies that the ratio of capital to land declines with distance from the CBD. Thus, Figure 1 also represents the capital-land ratio. Indirectly, we also have a prediction that population density declines with distance because density must be high where the ratio of capital to land is high.

Lot sizes are easy to measure. But like housing itself, housing capital is a theoretical concept and is not easily measured. Producers substitute capital for land in various ways: building taller buildings, using more floor space, or simply by improving the quality of the non-land inputs. Empirically, building heights and floor areas are easy to observe. The most readily available measure of the capital-land ratio is the "floor-area ratio", which is simply building area divided by lot size. The model predicts that floor-area ratios fall with distance from the CBD as shown in Figure 1.

2.3 A Summary of Empirical Predictions for Distance from the CBD

The simple version of the monocentric city model produces an impressive number of predictions. The most important of these predictions are that the price of a unit of housing, land values, the capital-land ratio, and population density all decline smoothly with distance from the CBD, as shown in Figure 1. The only major alterations to the standard two-good consumer maximization problem are the assumptions that no two households can occupy the same space and workers must commute to their jobs. A full urban spatial structure follows from these assumptions.

The model has empirical content. Unlike many economic models, we have full functional form predictions. Figure 1 implies that distance from the CBD is the primary determinant of urban spatial relationships. For example, we should find that land values decline smoothly at a decreasing rate with distance from the CBD, the function should have no discontinuities, and this basic relationship should hold for different cities at different times.

2.4 Comparative-Statics Results

Although the simple functional form predictions are a powerful test of the model, many different models could produce the same functions. Another commonly used empirical approach is to test the monocentric city model's comparative-statics predictions. The predictions discussed here are based on the "closed-city" version of the model, in which the overall population of the city is an exogenous variable while the utility level of the representative household is endogenous. The model predicts that the function in Figure 1 shifts up as population or agricultural land values increase because either change increases the cost of land throughout an urban area. The increase in land values and house prices leads producers to build homes using higher capital-land ratios, which lead to higher population densities. Decreases in commuting costs make sites farther from the CBD relatively more valuable than closer sites. Thus, a decrease in commuting costs leads to a flatter slope for the functions depicted in Figure 1.

The results are ambiguous for the remaining important variable, income. An increase in income increases the demand for housing, which leads consumers to prefer sites farther from the CBD where the price of housing is lower. But it also increases the aversion to time spent commuting, which has an offsetting effect making sites closer to the CBD more valuable. Empirically, it appears that the former effect dominates as increases in income have generally led to declines in the slopes of the functions in Figure 1. However, the empirical tradeoff between the housing demand and commuting time cost elasticities has been the subject of very little empirical investigation.

Since the 1800s, most urban areas have enjoyed steadily rising incomes, lower commuting costs, and steady population growth. The path of agricultural land values is less clear; although they may well have declined in real terms, their effect is overwhelmed by the large increase in urban populations. Together, these changes should lead the functions depicted in Figure 1 to shift up and have flatter slopes. Thus, one way to test the comparative-statics predictions is to compare estimates for a single city over time. Alternatively, we might compare estimates across cities at a given time if measures are available for income, commuting cost, population, and agricultural land values.

The latter approach – comparing estimates across cities – is far less common because it is more difficult to acquire data for a cross section of cities than for a single city over time. Excellent examples of the approach include Mills (1972) and Brueckner and Fansler (1983). Mills compares population density estimates for Baltimore, Milwaukee, Philadelphia, and Rochester for 1880-1963. He finds some evidence that intercepts are higher and slopes are flatter when cities have higher populations and incomes and lower commuting costs. However, by far the most important explanatory variable is the lagged dependent variable, indicating that inertia is a critical determinant of the density function coefficients. Brueckner and Fansler compare total land areas across 40 American cities in 1970. As predicted by the model, they find that land areas are lower when population and incomes are lower and when agricultural land values are higher. Evidence on the effect of commuting costs on land area is less clear: their attempts to measure this variable produce the right signs but the coefficients are statistically insignificant. In general, this approach is hampered by the difficulty inherent in measuring variables such as income and commuting costs.

3. Empirical Modeling Approaches

3.1 Regression-Based Approaches

The functions shown in Figure 1 are estimated easily by ordinary least squares regression procedures. The most commonly used functional form is the simple negative exponential function:

$$\ln y_i = \alpha - \beta x_i + u_i \qquad (2)$$

where x_i is distance from the CBD at location i, u_i is an error term, and α and β are parameters. The dependent variable, y_i, may be the price of a unit of housing, land value, the capital-land ratio, or population density. The negative exponential function generally fits urban spatial relationships well. In this formulation, β is the "gradient" because each additional mile from the CBD causes y to fall by 100×β percent. Additional terms can easily be added to the estimating equation. Equation (2) imposes the structure implied by the monocentric city model: $\partial y_i / \partial x_i = -\beta y_i < 0$, and $\partial^2 y_i / \partial x_i^2 = \beta^2 y_i > 0$.

Although equation (2) is the most commonly used estimating equation, it may not be flexible enough for many urban spatial relationships. The land value gradient, for example, tends to be higher near the city center than in more distant locations. Adding higher-order terms – x^2 or x^2 and x^3 – may improve the fit significantly. One particularly attractive formulation is the cubic spline. In this approach, the distance variable, x, is split into equal intervals and a separate cubic function is applied to each region. The function is constrained to be smooth at the boundaries between regions (which are known as "knots"). For example, in the empirical section of this chapter, I divide distance from the CBD into four intervals. The minimum value is x_0, the boundaries between regions are x_1, x_2, and x_3, and the maximum value is x_4. The distance between each knot is (x_4-x_0)/4. The estimating equation is

$$\ln y_i = \alpha + \beta_1(x_i - x_0) + \beta_2(x_i - x_0)^2 + \beta_3(x_i - x_0)^3 + \gamma_1(x_i - x_1)^3 \times D_1 + \gamma_2(x_i - x_2)^3 \times D_2 + \gamma_3(x_i - x_3)^3 \times D_3 + u_i \quad (3)$$

The D_k terms are dummy variables that equal one when $x_i \geq x_k$. Additional flexibility can be added by defining more regions. Each additional region simply involves a new definition of the knots and an additional interaction term between a dummy variable and a cubic term of the form $(x_i - x_k)^3$.

A good example of the use of spline functions is Anderson (1982). Alternative flexible estimators such as Fourier expansions (Gallant, 1982) are also useful. Nonparametric and semiparametric estimators such as that used by McMillen (1996) are popular alternatives. However, nonparametric estimators are far more difficult to use and have few advantages when nonlinearity is confined to a single variable. Interactions between variables may be easier to model using nonparametric techniques, though.

3.2 The Two-Point Method for Population Density Functions

The most widely studied urban spatial relationship is population density. The only necessary data to estimate a population density function are population, land areas, and distance from the CBD for small geographic areas. Such data now are readily available for zip codes, census tracts, and even smaller areas. However, historical data are more apt to be reported only for larger geographic areas, such as municipalities and counties, making it difficult to compare density gradients over a long time.

Mills (1972) proposed a clever method for estimating historical population density gradients using extremely limited data sources. The starting point for his procedure is the simple negative exponential function, equation (2), where *y* is defined as population density. Population density at location *d* can be written as *P*(*d*)/*A*(*d*), where *P*(*d*) is the number of residents and *A*(*d*) is the land area in an infinitesimal ring *d* miles from the city center. In a circular city, *A*(*d*) = 2π*d*. Mills generalizes this specification somewhat by assuming that $A(d) = \phi d$. For example, ϕ is approximately equal to π in Chicago or Milwaukee and to the full 2π in Indianapolis. The approach does not work in cities such as San Francisco.

Mills' trick is to integrate the population density function so that only observable variables remain. Given the expression for *A*(*d*), we can re-write equation (2) as *P*(*d*) = $\phi d e^{\alpha - \beta d}$. Integrating this function by parts from *d*=0 to d=d_c, we have P_c $= (\phi e^{\alpha} / \beta^2)[1 - (1 + \beta d_c)e^{-\beta d_c}]$. If d_c represents the central-city radius, then P_c is the population of the central city. Letting d_c go to infinity, the population of the entire metropolitan area is approximately $P = \phi e^{\alpha} / \beta^2$. The assumption that the metropolitan extends forever simplifies the calculations and causes little bias. Finally, the ratio P_c/P leads to a tractable equation:

$$\frac{P_c}{P} = \left[1 - (1 + \beta d_c)e^{-\beta d_c}\right] \tag{4}$$

Given the central city population, the population of the entire metropolitan area, and the radius of the central city, equation (4) has only one remaining unknown, β. Thus, this equation can be solved iteratively by choosing a value of β that provides the closest match between the two sides of equation (4).

Population data are readily available from the Census Bureau. However, d_c can be difficult to calculate. Although ϕ does not enter equation (4) directly, it may be easier to estimate than d_c. The two parameters are related by the identity $A_c = (\phi/2)d_c^2$, where A_c is the land area of the central city. Since data on central city land areas are available from the Census and ϕ is usually easy to approximate, the calculations can be simplified by calculating an implicit value for the central-city radius: $d_c = \sqrt{2A_c/\phi}$. Thus, Mills' procedure makes it possible to estimate theoretically appropriate population density gradients with readily available historical data – population of the central city and metropolitan area. Mills (1972) and Macauley (1985) are good examples of the technique.

4. Empirical Results – Land Values and Population Density

Figure 2 shows the trends in population density and land value gradients since Chicago's founding in the 1830s. McDonald (1997) calculated the density gradients for 1870-1990 using Mills' two-point method, and I used data from 2000 to update the estimates. The land value gradients come from McMillen (1996). For 1836-1928, the estimates are based on data from Hoyt (1933). Hoyt presents maps of average land values for tracts of land with the City of Chicago that typically are about a square mile in area. Mills (1969) was the first to use this classic data source to estimate land value functions. I updated the estimates for 1960-1990 using data from *Olcott's Land Values Blue Book of Chicago*. Until recently, *Olcott's* presented land value estimates annually for every block in the city.

Insert Figure 2 Here

The important trend to note in Figure 2 is the long-running decline in both land value and population density gradients. Land values decline by more than 60% per mile with distance from the CBD in 1836 and 1857. Over the rest of the nineteenth century, the gradient declines to about 50%, and it falls to about 20% in 1928. By 1960, the land value gradient hovers near zero, but the gradient rises back to 14% in 1990. The population density gradient falls from over 60% in 1870 and 1880 to 20-30% for 1920-1950 to between 10% and 20% thereafter.

During this long time period, commuting costs fell dramatically, incomes increased, and Chicago grew from a small village of a few thousand people to a city with about 3 million residents and a metropolitan area with a population of about 8 million. The monocentric city model predicts that both the land value and population density functions should shift up and the gradients should decline over time. The data strongly support both predictions. Although the two-point method does not produce standard error estimates, the land-value regressions fit the Hoyt data well. The R^2s indicate that the single explanatory variable, distance from the CBD, explains more than 80% of the variation in the natural logarithm of land values in 1836 and 1857, and the R^2 remains as high as 0.61 in 1910. However, the R^2 falls to 0.24 in 1928, and is nearly zero in 1960. The R^2 is 0.10 in 1990. Thus, the evidence is mixed: the model appears to fit well in the 1800s, and the decline in the gradients matches theoretical predictions. But low R^2s indicate that the model may have little predictive power now.

The more recent evidence against the model's predictive ability is not as strong as appears at first glance. In a model with a single explanatory variable, the formula for the R^2 is $b^2 s_x^2 / s_u^2$ where b is the estimated gradient and s_x^2 and s_u^2 are the variances of the explanatory variable and the regression residuals. If the locations for the observations do not change over time and the residual variance stays the same, the R^2 will decline whenever the gradient declines. Theory has nothing to say about trends in the model error variance. Thus, the model *predicts* that R^2s should fall over time as commuting costs decline and income increases.

Another commonly overlooked feature of land-value and population density gradients is the way data sets have changed over time. Data collection has improved greatly in recent years. In the past, data were much more likely to be reported as simple averages for large areas. For example, Hoyt's square-mile tracts are unusually small for historical land-value data. Olcott's presents much more detailed data: the numbers sometimes vary within a single city block. For population density, Clark's (1951) classic study uses data for square mile rings around the CBD. A city with a 10-mile radius will only have 10 observations. Modern data sets are much larger. For example, McMillen and Lester (2003) estimate population density gradients using more than 10,000 observations in the Chicago metropolitan area. R^2s are usually much higher for aggregated data.

Further, in the past researchers usually confined their attention to the central city. Suburban functions are often nearly flat; the R^2 formula implies that the estimated function's predictive power will be low for these observations. Thus, adding suburban data usually drives down the R^2. Finally, the monocentric city model only predicts that the functions should look like those in Figure 1; it does not predict that the negative exponential function is the correct one. In McMillen (1996), I find that a more general functional form and the addition of a few explanatory variables – distance from Lake

Michigan and distance from O'Hare and Midway Airports – raises the R^2 for the 1990 Olcott's land value function from 0.10 to 0.87. Whereas a theory than only accounts for 10% of the variation in land values is unimpressive, 87% is impressive indeed. No one ever claimed that distance from the CBD was the only variable that matters. If flexible functional forms and a few variables account for nearly 90% of the spatial variation in land values across small geographic areas, then the theory is alive and well.

Other evidence does point to important deficiencies of the monocentric city model. The primary problem is the static nature of the model. In fact, age matters. McDonald and Bowman (1979) find that land values increase with distance from the CBD in older areas of Chicago, and McDonald (1979) finds little relationship between land values and population density in these areas. Densities reflect the past whereas land values reflect expectations about the future. Brueckner (1986) finds evidence in favor of a vintage model of urban spatial structure: population densities increase discretely across distances where building ages are different. To be realistic, the simple model must be altered to take into account the fact that buildings last a long time and are not destroyed whenever economic conditions change. Other changes such as suburban subcenters and the effect of rivers and lakes on urban spatial form are more marginal and can be handled by introducing a few additional explanatory variables. Vintage effects are a more fundamental challenge to the monocentric model requiring significant changes to the theory (e.g, Anas, 1978; Wheaton, 1982).

5. Empirical Results – Floor-Area Ratio

Although floor-area ratios have not been analyzed as extensively as population densities or land values, they are a critical part of the monocentric city model. The model predicts that low commuting costs for sites near the city center lead to high land values, which in turn lead to high floor-area ratios. In this section, I present estimates of the relationship between floor-area ratios and distance from the CBD in Cook County Illinois, which includes Chicago. The data set includes single-family residential homes selling between 1983 and 1999. Data on lot size, floor area, and addresses come from the Cook County Assessor's Office, and the list of sales comes from the Illinois Department of Revenue.

The first two columns of results in Table 1 present the results of regressions using the natural logarithm of the floor-area ratio as the dependent variable and distance from the CBD as the explanatory variable. Only data from the City of Chicago are used for these regressions. The first column shows that floor-area ratios decline by 5.5% with each additional mile from the Chicago CBD. The R^2 indicates that only about 15% of the variation in this variable is explained by distance. The second column of results is based on a smooth cubic spline with four regions. The R^2 rises to 0.207, and the coefficients for the additional explanatory variables are all statistically significant. Figure 3 shows the estimated functions. The spline function's additional explanatory power comes from the sharp rise in estimated floor-area ratios near the CBD. Both functions are consistent with the predictions of the monocentric city model.

Insert Table 1 Here

Insert Figure 3 Here

The last two columns of results present comparable estimates for all of Cook County. The data for these regressions are averages by census tract. The point of these regressions is to show how dramatically the explanatory power increases when aggregate data are used rather than data for individual homes. The R^2 for the simple negative exponential function rises to 0.556 even though the estimated gradient does not change at all. The R^2 for the spline rises to 0.692, indicating that an impressive 69.2% of the variation in the logarithm of the floor area ratio is explained by a single variable, distance from the Chicago CBD. A graph of the estimated function (not shown here) reveals that distance from the CBD no longer has much explanatory power beyond about 18 miles. In fact, the R^2 is only 0.021 for a spline function with four equally-spaced intervals from 15 miles from the CBD to the maximum value of 34 miles. The model works well in the area of metropolitan area that is still dominated by the CBD; it does not work well in newer suburban areas

6. House "Prices"

Other than population density, the most common test of the monocentric city model is the relationship between house "prices" and distance from the CBD. The word "price" is in quotes because the unit price of housing is not observed in practice. The actual dependent variable for these regressions is the value of housing, which is the product of price and quantity. As we have seen, the model makes no predictions for this variable. The apparent failure of the model is illusory.

How can researchers claim that a regression of house values on distance from the CBD is a test of the monocentric city model? Suppose that the unit price of housing declines with distance from the CBD according to the function $P_h(d) = e^{\mu-\delta d}$. Then the relationship between house values and distance is $\ln V_h(d) = \ln H(d) + \mu - \delta d$. If it were possible to control for all house characteristics (Z), then we would have $\ln V_h(d) = Z\theta + \mu - \delta d$, which is a function that can be estimated easily. A good, careful example of this approach is Coulson (1991), but bad examples are much more abundant.

The problem with this approach is that the full set of housing characteristics is never observed. Suppose we have two houses that have three bedrooms and a garage, are on 1/3-acre lots, and appear to be identical in all measurable ways. House A is 5 miles from CBD and costs \$300,000 and House B is 10 miles from the CBD and costs \$500,000. There appears to be an upward-sloping house price gradient. Unfortunately, missing variables include school quality, other local public services, and vague concepts like house quality. If we could control for these variables, we might find a downward-sloping house-price function. But the family living in the higher-priced home undoubtedly has a higher income than the family in Home A. Thus, what the regression really does is to trace out the places where higher-income families live. As we have seen, the model predicts that the sign of this gradient can go either way.

7. Conclusion

By definition, a model is a simplification of reality. The monocentric city model uses a very simple idea – that people will pay a premium for sites that lead to lower commuting costs – to generate a complete model of the spatial structure of an entire city. Although the assumptions are not literally accurate, they produce a mathematically tractable model with remarkable predictive power. Even modern urban areas tend to be dominated by the traditional CBD. Land values, population density, and floor-area ratios decline markedly with distance from the CBD. Over time, rates of decline have fallen to the point that the CBD no longer dominates the entire urban area to the same extent as before. But the theory actually predicts the decline in the importance of the CBD: declining commuting costs and increasing incomes lead to significantly lower gradients.

Much criticism of the monocentric city model comes from a misunderstanding of the empirical results. Lower gradients produce lower R^2s. More disaggregated data sets also lead to lower R^2s. Suburban data have never been explained well by the monocentric city model, and the low gradients in these areas tend to further reduce R^2s. The model makes no predictions regarding the spatial pattern of house values, and empirically it is nearly impossible to measure the theoretically relevant unit price of housing.

Within central cities, the most important deficiency of the monocentric city model is its failure to take into account the longevity of the capital stock. Land values will be low in areas where floor-area ratios are high if existing buildings, which reflect past economic conditions, are costly to demolish. The model's failure is more serious in the suburbs. Modern cities often have large suburban employment centers with marked effects on the spatial structure of the suburbs. Suburban areas have a fragmented system of government. Local variations in zoning practices, tax rates, and the provision of local public goods have significant effects on the spatial patterns of variables such as house prices and population densities. Although distance from the CBD has little predictive power in the suburbs, spatial relationships can often be modeled accurately with a few additional variables such as distance from employment subcenters and access to the transportation system. With a few modifications, the monocentric city model remains a useful tool for understanding urban spatial relationships.

References

Anas, A. (1978). Dynamics of Urban Residential Growth. *Journal of Urban Economics* 5, 66-87.

Anderson, J. E. (1982). Cubic-Spline Urban-Density Functions. *Journal of Urban Economics* 12, 155-167.

Brueckner, J. K. (1986). A Switching Regression Analysis of Urban Population Densities. *Journal of Urban Economics* 19, 174-189.

Brueckner, J. K. and Fansler, D. A. (1983). The Economics of Urban Sprawl: Theory and Evidence on the Spatial Sizes of Cities. *Review of Economics and Statistics* 65, 479-482.

Clark, C. (1951). Urban Population Densities. *Journal of the Royal Statistical Association Series A* 114, 490-496.

Coulson, N. E. (1991). Really Useful Tests of the Monocentric City Model. *Land Economics* 67, 299-307.

Gallant, A. R. (1982). On the Bias in Flexible Functional Forms and an Essentially Unbiased Form: The Fourier Flexible Form. *Journal of Econometrics* 15, 211-245.

Hamilton, B. W. (1982). Wasteful Commuting. *Journal of Political Economy* 90, 1035-1051.

Hoyt, H. (1933). *One Hundred Years of Land Values in Chicago*. University of Chicago Press, Chicago.

Macauley, M. K. (1985). Estimation and Recent Behavior of Urban Population and Employment Density Gradients. *Journal of Urban Economics* 18, 251-260.

McDonald, J. F. (1979). An Empirical Study of a Theory of the Urban Housing Market. *Urban Studies* 16, 297-297.

McDonald, J. F. (1997). *Fundamentals of Urban Economics*. Prentice-Hall, Upper Saddle River NJ.

McDonald, J. F. and H. Woods Bowman (1979). Land Value Functions: A Reevaluation. *Journal of Urban Economics* 6, 25-41.

McMillen, D. P. (1996). One Hundred Fifty Years of Land Values in Chicago: A Nonparametric Approach. *Journal of Urban Economics* 40, 100-124.

McMillen, D. P. and T. W. Lester (2003). Evolving Subcenters: Employment and Population Densities in Chicago, 1970-2020. *Journal of Housing Economics* 12, 60-81.

Mills, E. S. (1969). The Value of Urban Land. In Perloff, H. (ed.) *The Quality of the Urban Environment*. Resources for the Future, Washington, DC.

Mills, E. S. (1972). *Studies in the Structure of the Urban Economy*. Resources for the Future, Baltimore.

Muth, R. F. (1969). *Cities and Housing*. University of Chicago Press, Chicago.

O'Sullivan, A. (2002). *Urban Economics*. Irwin/McGraw-Hill., New York.

Wheaton, W. C. (1982). Urban Residential Growth under Perfect Foresight. *Journal of Urban Economics* 12, 1-21.

Table 1

Floor-Area Ratio Regressions

	Chicago Homes	Chicago Homes	Cook County Census Tracts	Cook County Census Tracts
Constant	-0.700	0.782	-0.626	0.268
	(268.518)	(34.604)	(35.829)	(4.465)
x	-0.055		-0.055	
	(205.472)		(39.540)	
x-x_0		-0.903		-0.490
		(41.215)		(14.099)
$(x$-$x_0)^2$		0.151		0.061
		(23.229)		(10.427)
$(x$-$x_0)^3$		-0.009		-0.003
		(15.332)		(9.534)
$(x$-$x_1)^3\times(x\geq x_1)$		0.005		0.004
		(6.814)		(8.776)
$(x$-$x_2)^3\times(x\geq x_2)$		0.007		-0.001
		(33.797)		(4.367)
$(x$-$x_3)^3\times(x\geq x_3)$		0.007		0.001
		(14.510)		(0.773)
R^2	0.151	0.207	0.556	0.692
Number of Observations	237,420	237,420	1,251	1,251

Notes. The dependent variable is the natural logarithm of the floor-area ratio. Absolute *t*-values are in parentheses. The evenly spaced knots for the Chicago spline function begin at x_0 = 0.780 with an increment of 4.007 between knots. Comparable values for the Cook County spline function are 0.782 and 8.312.

Figure 1

Functional Form Prediction

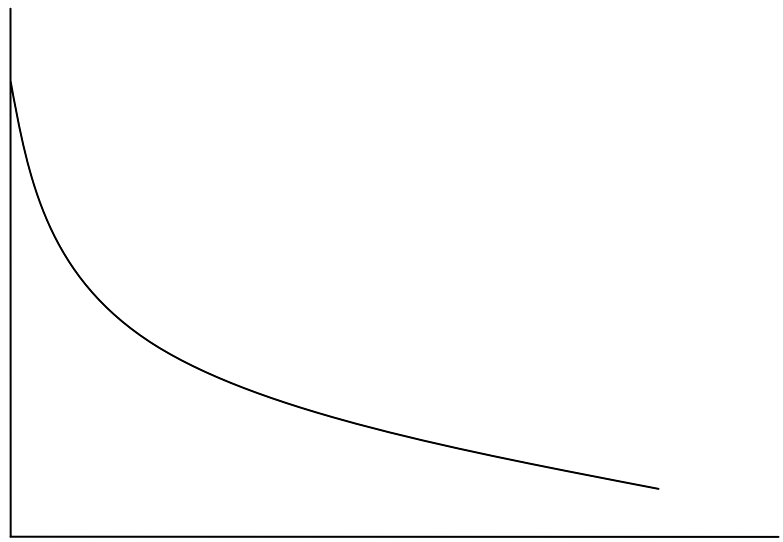

Figure 2

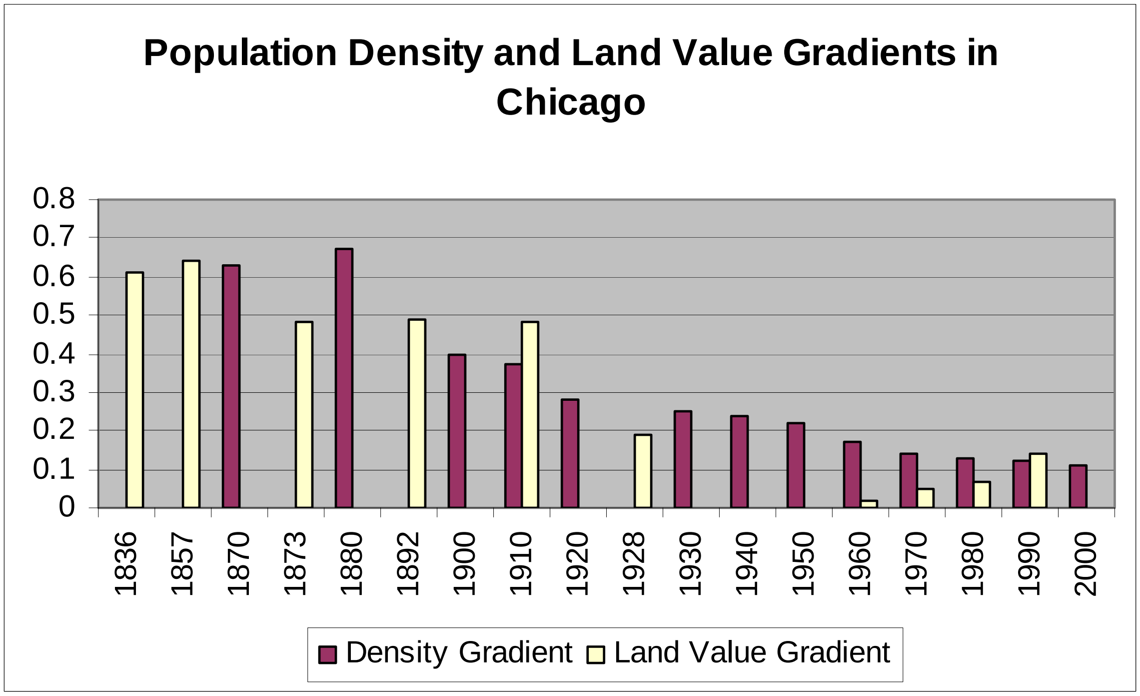
Population Density and Land Value Gradients in Chicago
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1836
1857
1870
1873
1880
1892
1900
1910
1920
1928
1930
1940
1950
1960
1970
1980
1990
2000
Density Gradient
Land Value Gradient

Figure 3
Floor-Area Ratios for Individual Homes in Chicago