NBER WORKING PAPER SERIES

THE ECONOMICS APPROACH TO CITIES

Edward L. Glaeser

Working Paper 13696
http://www.nber.org/papers/w13696

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2007

I am grateful to the Taubman Center for State and Local Government for financial support. Kristina Tobio provided excellent research assistance. This essay was written for an edited multi-disciplinary volume on how different disciplines approach cities and it is meant to give non-economists a sense of the economic approach to cities. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

The Economics Approach to Cities
Edward L. Glaeser
NBER Working Paper No. 13696
December 2007
JEL No. R0

ABSTRACT

The economic approach to cities relies on a spatial equilibrium for workers, employers and builders. The worker's equilibrium implies that positive attributes in one location, like access to downtown or high wages, are offset by negative attributes, like high housing prices. The employer's equilibrium requires that high wages be offset by a high level of productivity, perhaps due to easy access to customers or suppliers. The search for the sources of productivity differences that can justify high wages is the basis for the study of agglomeration economies which has been a significant branch of urban economics in the past 20 years. The builder's equilibrium condition pushes us to understand the causes of supply differences across space that can explain why some places have abundant construction and low prices while others have little construction and high prices. Since the economic theory of cities emphasizes a search for exogenous causes of endogenous outcomes like local wages, housing prices and city growth, it is unsurprising that the economic empirics on cities have increasingly focused on the quest for exogenous sources of variation. The economic approach to urban policy emphasizes the need to focus on people, rather than places, as the ultimate objects of policy concern and the need for policy to anticipate the mobility of people and firms.

Edward L. Glaeser
Department of Economics
315A Littauer Center
Harvard University
Cambridge, MA 02138
and NBER
eglaeser@harvard.edu

I. Introduction

Why are some cities so much more productive than others? What are the environmental and social costs of density? Why are there ghettos? How does living close to others change us? Why do cities rise and fall? Why is housing so expensive in some places? Urban economics addresses all of these disparate questions and all of them can be seen as components of urban economics' great puzzle: why do so many people cluster next to each other in cities? That question is itself one part of the even grander quest of economic geography to understand all of the location decisions of people and firms.

The economic approach to understanding location choices, like living in cities, focuses on understanding the motives might underlie those choices. Are places attracting people by offering high wages or cheap housing or good weather? Why do firms stay in places where they must pay high wages? Since urban development reflects millions of individual choices to live in cities, understanding that development requires us both to understand the relative importance of the different urban attributes and to understand why cities have those attributes. For example, high wages certainly help attract people to New York City. However, for us to understand the eight million people who choose to live in that city, we would also need to understand why its wages were so high. The sub-discipline of agglomeration economics has developed to understand the productivity differences that presumably lie behind the observed income differences across space.

This essay explores the key elements of the economic approach to cities and how they reflect the core elements of my discipline. Economics has three great pillars, two of which help us to understand the world and one of which helps us to offer policy advice. The first pillar of economics is that people respond to incentives. This assumption is caricatured by some who suggest that economists think that people only respond to financial incentives, which is surely false. Still, it is true that the incentive principle leads economists to look at the financial incentives that might explain location choices.

The second pillar of economics is our concept of a no arbitrage equilibrium. Adam Smith used an early version of the no arbitrage equilibrium to make sense of wages; Milton Friedman popularized the concept with the phrase "there is no such thing as a free lunch."[1] This pillar enables us to not only examine individual decision, but also to make predictions about how an entire system will look.

In urban economics, there are three key no arbitrage relationships. First, individuals must be indifferent across space, which has been taken to mean that the flow of wages plus amenities minus housing costs is roughly equal in every location. Second, firms must be indifferent over space and over hiring new workers. This condition implies that differences in wages must be offset by differences in productivity. Third, builders must be indifferent about building or not building new units. This condition implies that housing prices cannot rise too far above the total costs of construction, as long as those costs are understood to include physical building costs, the price of land and the difficulties involved in dealing with land use regulations.

Economics' third pillar is the assumption that good policies increase the range of choices that an individual can make. Economists' enthusiasm for income is driven by the view that more wealth gives people more choices. Our enthusiasm for political freedom has the same source. Economists talk about good policies increasing "utility levels" which is often misunderstood as suggesting that these policies will make people happier. Happiness is an important emotion, but there is no sense in which it is particularly related to economists' definition of utility. Formally, higher level of utility is equivalent to having more options, not wearing a smile.

These three pillars have shaped the economic approach to cities. In Section II of this essay, I discuss the central theoretical construct of economic geography and urban economics: the spatial equilibrium. The power of the spatial equilibrium assumption is that it predicts that if something is particularly good in one location, then we should

[1] Robert Heinlein is usually given as the original source of the phrase, although there do appear to be earlier antecedents.

expect to see something bad offsetting it. In the intra-urban Alonso-Muth-Mills model, high prices close to the city center are offset by short commutes. In the inter-urban Rosen-Roback model, high incomes are offset by either high prices or disamenities. The spatial equilibrium assumption has been particularly effective in making sense of urban housing markets.

In Section III, I turn to the equilibrium condition for employers and builders. The firms' equilibrium condition leads us to explain the differences in incomes across space by understanding why productivity levels would differ across space. Productivity levels might be higher because of access to natural resources like productive land or rivers or because of increased ease of transportation to suppliers or customers. The builders' equilibrium condition means that to understand differences in housing costs across growing areas, we must understand why it costs more to build in some areas than in others.

Section IV turns to the empirical approaches favored by urban economists. Economics' theoretical definition of a city differs significantly from the empirical implementation of that definition. Conceptually, cities are the absence of physical space between people and firms. Cities are density or proximity, perhaps combined with sufficient scale. Empirically, cities are either the formal and somewhat arbitrary political units that bear that name or the name "metropolitan areas," which are themselves somewhat arbitrary combinations of counties, which are also arbitrary political units. While one might wonder about the mismatch between concept and data, economics is a pragmatic discipline that has generally happily used the imperfect available data.

The empirical methods used by urban economics are driven by the attachment of urban economists toeconomic theory. This attachment has produced two different styles of empirical research. One style of structural empirical research focuses on using data to estimate formal models. A second style of research emphasizes exogenous sources of variation, or instruments, such as rivers or sharp political boundaries. The importance that economists place on exogenous sources of variation comes from a core disciplinary

view that our theoretical models are meant to map connections between exogenous variables and outcomes.

Finally, in Section V, I turn to the economic approach to urban policy-making. The core insight of the field is the primacy of person over place. Economics judges policies by whether they increase the choices available to people, not on whether they help rebuild a particular locale. Economics does not preclude place-based policies, such as urban redevelopment, if they are the best way to help people, but economists do insist that these policies be judged on whether they improve individual's lives, not on whether they make a place more pleasant.

Beyond putting people first, urban economics has two other themes that run through its policy prescriptions. First, urban economics has often assumed that governments only imperfectly represent their constituencies. As a result, individual economists have offered favored institutions that might increase competition across governments and mitigate this problem. Second, since urban economics starts with the mobility decisions of people and firms, urban economists tend to argue that policies need to be designed not just on the basis of current location patterns but also with an understanding of how new policies will alter individual location choices.

II. The Spatial Equilibrium Approach

The theoretical centerpiece of urban economics is the concept of a spatial equilibrium which assumes that there are no free lunches to be gained by changing location. While this assumption is often treated with more general utility functions, economists often assume a linear utility function, which then implies that the elements of utility that are related to location choice are captured by:

(1) Income + Amenities – Housing Costs – Transportation Costs.

The spatial equilibrium assumption is that this flow is constant over space. While this assumption is obviously a simplification, it has had a remarkable ability to generate hypotheses that have been in accord with the evidence. Within metropolitan areas, the Alonso-Muth Mills model assumes that income is constant and looks at whether high housing costs are offset by low amenities or low transport costs. Across metropolitan areas, the Rosen-Roback model looks at the tradeoff between income, amenities and housing costs.

The spatial equilibrium approach is often augmented in different ways. For example, information about transportation costs and technologies can yield sharp predictions about how housing prices will change with distance from the city center. When there are different types of people, these models can predict where different people will live. The incorporation of housing supply into the model enables these equations to predict density levels within a city and population patterns across cities.

The Alonso-Muth-Mills Model

The most significant piece of urban economic theory remains the spatial equilibrium model of William Alonso (1964), which was extended by Mills (1967) and Muth (1969). Alonso's model looks within a metropolitan area and assumes that both income and amenities are constant. These assumptions then imply that housing costs plus transport costs are constant across space, which means that housing costs will decline as transport costs rise with distance to the city center. In the simplest case, where everyone works at the center of the city and transport costs rise linearly with distance to that center, i.e. if costs equal "t" times distance, then housing costs must equal costs at the center minus "t" times distance.

The model is simple, elegant and contains a far amount of truth. Figure 1 shows the relationship between median housing prices across 187 cities in the greater Boston region and distance between those cities and downtown Boston.[2] On average, an extra mile

[2] Housing prices are based on the 2000 Census. Distance to Boston is defined in Glaeser and Ward (2006).

from the city center is associated with housing prices dropping by $1100. If an $1,100 drop in price is roughly associated with a $110 increase in annual cost, then the Alsonso-Muth-Mills model suggests that an extra mile of distance is associated with added travel costs of two dollars per week, which seems low but not implausible since so many people are not commuting into the city center.

The relationship is far from perfect. Distance explains only 15 percent of the heterogeneity in prices among those towns. Other factors, including housing quality and both exogenous and man-made amenities, differ across towns. Such amenities, like school quality, are often far more important in determining housing prices than proximity to downtown. Still, the model has made a prediction that is certainly not rejected by the data. Thousands of variants of this regression have been run by economists since Alonso's model was first published, and almost all of them have found this general pattern.

In some cases, economists have found a convex, rather than a linear relationship between distance and price, where housing prices drop steeply with distance over some initial distance and the relationship between prices and distance flattens out. A slight variant on the basic model can explain this pattern. Assume that individuals have access to two different transportation technologies, such as walking and driving. One technology involves no fixed costs, but imposes a cost of $\bar{t}$ times distance. The other technology involves a fixed cost, like having to buy a car, but imposes a lower cost of $\underline{t}$ times distance.

Optimal use of these technologies implies that the technology without fixed costs (perhaps walking) should be used until a point where $\bar{t} - \underline{t}$ times distance equals the fixed cost. People should be willing to pay the fixed cost only when the time savings from driving is big enough to pay for that fixed cost. Even if this calculation is a simplification, there is little doubt that car usage gets much higher further away from city centers (Glaeser, Kahn and Rappaport, 2007). This simple change to the model predicts

that housing costs will decline by $\bar{t}$ times distance in areas that are close to the city and decline by $\underline{t}$ times distance in areas that are further away from the city. The slope of prices on distance should therefore flatten out as distance from the city increases, which delivers the convex relationship seen in the data.

Alternatively, the cost of commuting "t" per unit distance might be a function of investment in technologies, like owning a bike or car, that reduces the price of commuting per unit distance. If we let k(t) denote cost of investing in these technologies where k(t) is continuously differentiable, k'(t)<0 and k"(t)>0, then cost minimization requires minimizing k(t)+t*distance. The first order condition is that -k'(t)=distance and the cost of commuting per unit distance falls with commute distance. People who live further away invest in better technology and this also implies the convex relationship that often appears in the data. Sensible use of transportation technologies implies that slower technologies will be used closer to the city center, which in turn implies that the relationship between prices and distance will be higher close to the center.

A second important permutation to the model allows densities to respond to demand. This change allows the model to speak not only to prices, but also to the degree of development at different locations. The model can be adapted either by making land consumption flexible so that people consume less land where it is expensive and close to the city. Alternatively, we can assume that living space is manufactured with land and physical capital. As land gets more expensive, builders use more capital and build up. High densities close to the city center can be seen as an application of the incentive principle, because a greater willingness to pay for living area that is valuable because of its proximity then induces a greater supply of living area per unit of land. Either version of the model predicts that densities will be higher closer to the city center.

Figure 2 shows the relationship between the logarithm of people per square mile and distance to Boston across the same 187 cities and towns. Again, there is a robust positive relationship and distance explains 45 percent of the variation in density levels across those cities and towns. For each extra mile of distance from Boston, the predicted

density level drops by .08 log points or about eight percent. Again, the predictions of the simple model are supported by the data.

The attraction of the Alonso-Muth-Mills model to economists illustrates a central point about economics that differentiates it from other fields. The model's strength lies in its ability to make predictions that hold generally, not in its ability to explain the exact peculiarities of particular places. While some disciplines place great value in adding complexity and nuance, economists like the ability to produce general rules that hold most of the time. Economists are usually more interested in common patterns than in particular idiosyncrasies.

One objection to the Alonso-Muth-Mills model is that it is increasingly at odds with a world that is no longer monocentric. In most cities, employment is located far away from the old city center (Glaeser and Kahn, 2001). But even this can be incorporated into the Alonso-Muth-Mills, framework. For example, Henderson and Mitra (1994) extend this model to include multiple employment centers. One general prediction of these extensions is that in cities where employment is more decentralized, the relationship between distance from the city and housing prices will be flatter. Glaeser and Kahn (2001) find exactly that. The tendency of prices to fall with distance to the center is much higher in those areas with centralized employment, which can be seen as support for using this type of model to explain differences in price patterns across cities.

The application of the spatial equilibrium concept in the Alonso-Muth-Mills model is useful not only in predicting housing prices and density levels, but also in predicting the locations of different population groups. In particular, the model has been particularly effective at looking at the location decisions of the rich and the poor. One of the most striking facts about American cities is that poverty rates are much higher in city centers than on the outskirts of town. This poverty does not mean that central cities make people poor, but rather that city centers attract the poor (Glaeser, Kahn and Rappaport, 2007).

This fact may seem initially to be a puzzle for the Alonso-Muth-Mills framework. Usually, we think that travel costs are higher for the rich than they are for the poor, because the rich have a higher opportunity cost of time since their wages are higher. After all, lost time is often the most important element in commuting costs. Economics suggest that someone who earns 20 dollars an hour should be willing to pay more to avoid an hour of commuting than someone who earns 10 dollars an hour. This logic suggests that the rich should pay more for low commuting costs at the city center, yet they generally are not the high bidders for land close to the urban core.

There are two explanations for this phenomenon. One is given by Becker (1965) who argues that the greater value of time for the rich can be offset if the rich also want to own more land. The greater desire for land pushes the rich to live where land is cheap, on the edge of the city, just as a greater desire for warmth pushes sun-lovers to move south. While this hypothesis is theoretically elegant, the desire of the rich to own more land does not seem to be great enough to justify their decisions to live on the urban edge.

An alternative view is that the rich and the poor use different transportation technologies (LeRoy and Sonstelie, 1983). If the rich drive and the poor take public transportation, then the rich can have lower costs of commuting per mile even if their time is worth more. Glaeser, Kahn and Rappaport (2007) find that access to public transportation does seem to explain the decisions of the poor to live in urban centers. Indeed, poverty rates even seem to rise in areas that are close to new subway stops. The connection between public transportation and poverty is not a problem, but rather a reflection of the valuable role that public transportation plays in serving and attracting the poor.

The Rosen-Roback Model

While the Alonso-Muth-Mills model is the core tool for understanding prices and density levels within metropolitan areas, the Rosen-Roback model is economics' core tool for understanding prices across metropolitan areas. The Rosen-Roback model is more complex than the Alonso-Muth-Mills model in that it allows income and amenities to

differ across space. The Rosen-Roback model compensates for this added complexity by treating the metropolitan area as a single homogeneous entity, so that everyone in the area is assumed to have the same housing costs, transport costs and amenity levels.

The core prediction of a spatial equilibrium across metropolitan areas is that high housing prices must reflect either high income or high amenities or both. People wouldn't be willing to pay for nice places if they weren't getting something for their money. Rosen (1979) presents the core theory that emphasizes that the price of amenities across space requires us to look at both wages and prices. Roback (1982) took this model to the data and found that people did take home less money, net of housing costs, in places with more amenities. Gyourko and Tracy (1991) expanded the set of studied amenities and focused on the willingness to pay for different types of government.

Figure 3 shows the basic empirical value of this approach by graphing housing costs on incomes across metropolitan areas. Forty percent of the variation in metropolitan area prices is associated with differences in income. On average, a one thousand dollar increase in income is associated with a $3,700 increase in housing values. While the strong association supports the model, the coefficient seems low. After all, a $3,700 increase in housing values suggests an increase in annual costs of no more than $400, given reasonable assumptions on interest rates, maintenance levels and local taxes. If it really only cost $400 dollars more to move into an area with $1,000 higher income levels, then people should flock to those high income, high cost areas.

There are at least two good explanations for why this relationship is still compatible with the Rosen-Roback model, but these explanations make it clear that the model's predictions are a little fuzzy. First, differences in income are not fully correcting for differences in human capital. It isn't clear that someone who moves to an area that we think has $1000 more income per year will actually earn that extra income, because the people in that high income area may also have higher skills. Second, high income places may also have unattractive amenities, like long commutes, that offset the higher income levels.

A second way to use the spatial equilibrium assumption is to work directly with incomes that correct for local prices. The American Chamber of Commerce Research Association (ACCRA) has created a set of these indices for a subset of metropolitan areas. Using these indices, we can look at real income levels which allegedly correct for housing price differences across the U.S. Perhaps the most natural amenity to investigate is warmth, and Figure 4 shows the relationship between real incomes and median January temperature across metropolitan areas.

Across the sample of approximately 200 metropolitan areas, median January temperature explores 23 percent of the variation. As January temperature rises by 10 degrees, real income drops by 720 dollars. People do seem willing to require higher real income to live in places that are colder, which certainly does support the idea that places with higher real income levels are worse in other dimensions, just as the spatial equilibrium assumption predicts.

Another negative amenity is long commutes. Figure 5 shows the relationship between income and commute times across a sample of approximately 230 cities with populations of 100,000 or more. In this case, I am looking at purely nominal income levels and median commutes. As income increases by 10,000 dollars, median commute time increases by 1.5 minutes. This relationship confirms the view that high income places also have other negative amenities which offset high financial returns.

Hedonic Pricing and Housing Supply

The spatial equilibrium assumption holds out the possibility that housing prices can be used to accurately assess people's willingness to pay for location-specific amenities. If a particular amenity is always associated with a 500 dollar increase in housing costs, holding income and everything else constant, then it might be reasonable to infer that people place a value of 500 dollars on that amenity. In many cases—like crime and school quality—there may not be an independent way of assessing the value placed on an

amenity, so that hedonic pricing models offer the only hope of figuring out the value that people place on a particular public service. Since economists would like to weigh costs and benefits when they determine the optimal amount of public investment in these amenities, hedonic pricing models offer the hope of delivering an assessment of the public benefits associated with some forms of publicly provided goods.

A healthy body of research on housing price hedonics has emphasized at least two major issues with this type of analysis. First, on a purely conceptual level, housing prices can only tell us about the willingness to pay of the marginal resident in a particular area. This marginal resident's willingness to pay may not represent the average willingness to pay across the population. For example, the premium paid by the residents of tony Fifth Avenue to look out on Central Park, represents the valuation placed on park views by an extremely rich swath of the population. This valuation might be considerably higher than the valuation that a poor person might place on park views. There are no perfect solutions to this problem and it means that all hedonic estimates need to be thoughtfully interpreted.

The second problem with hedonic estimates is that they require us to hold everything else constant. Even the simplest models emphasize the difficulty of doing that. An attractive amenity will attract the people who are willing to pay the most for it. If those people are rich, and if people like living around rich people, then the first natural amenity will be correlated with the second amenity of living around richer people. A second example is that when comparing across jurisdictions, we have scores of differences in tax levels, public service provision and regulation that can all potentially impact the desire to live in an area. In principle, we can try to control for a rich array of area characteristics, but we may often doubt our ability to measure such things perfectly. As a result, hedonic estimates are almost always compromised of correlations between the observed neighborhood attributes and the error term in the regression.

This potential bias has generally led researchers away from metropolitan area level analysis, such as the correlation between weather and prices at the area level, to a lower

level of analysis where we can be more confident that neighborhoods are comparable, except for the variable in question. For example, older studies looking at the impact of school quality on housing prices had compared across school districts. These districts are usually coterminous with other governmental boundaries so it is hard to ensure that any observed effects reflect only schools. Moreover, high human capital people sort into area with better schools, so it is hard to know whether a correlation between schools and property values reflects the schools themselves or the people attracted by the schools.

In response to these problems, Black (1999) turned to using attendance districts. Attendance districts operate within jurisdictions and determine which of the lower schools that the children in a particular house will attend. People within the same school district who are in different attendance districts pay the same taxes and receive all of the same other government services. Black addressed the omitted neighborhood characteristic problem by comparing smaller and smaller geographic units, until she was literally comparing houses on opposite sides of the same street, which just happen to be associated with different lower schools. Black's estimates still find a significant willingness to pay for better schools, but the estimates are also much lower than those found in the earlier area level studies.

III. **Agglomeration Economies and Housing Supply**

In the previous section, I discussed the equalizing difference implications of the need for individuals to be in a spatial equilibrium. The spatial equilibrium framework focuses on individual location choice and the implication that bad things in a location, like high housing prices, are offset by good things in a location, like high wages. This indifference condition delivers only one side of the labor market and one side of the housing market. In the labor market, these individual choices deliver labor supply but do not give us labor demand. In the housing market, these individual choices deliver housing demand but not housing supply. To understand fully the distribution of people, prices and wages across space, we must also turn to the decisions of employers and builders or the determinants of labor demand and housing supply.

The usual purpose of economic models is to explain differences in endogenous variables, like prices and quantities, with exogenous variables, like geography or long-standing historical conditions. Since the spatial indifference condition of workers gives us little idea about where differences in wages or housing prices might come from, we need to bring in labor demand and housing supply to give us a chance of linking prices, wages and population with exogenous factors.

The Location of Firms and Labor Demand

In the previous section, we argued that a spatial equilibrium required that people be indifferent across space. A similar condition applies to firms. If we see firms operating in a particular locale, economists infer that the firm could not earn greater profits by moving somewhere else. As usual, economists think of this condition holding only approximately. Obviously, there are many factors, from moving costs to the spatial preferences of the CEO, that might mean that firms don't perfectly profit maximize when they choose locations. While economists' accept these caveats, the no-arbitrage equilibrium assumption still seems like the best available tool for understanding spatial decisions. This assumption pushes us towards the view that where firms face higher costs, then there must be something else that is good for those firms. Otherwise, they would leave.

High wages are perhaps the most obvious locational attribute which is bad from a firm's point of view. To workers, high wages are an attraction. To firms, high wages (holding worker quality constant) are a negative since high wages mean high costs. The logic of the spatial equilibrium suggests that there must be something good about high wage areas, from a firm's perspective, which offsets the high cost of doing business. The most natural explanation is that high wage regions are also areas that are more economically productive.

Labor economists would come to the same conclusion—high wages imply high productivity—as urban economists, but they might arrive through a slightly different

route. In a competitive labor market, economic theory suggests that firms should keep on hiring workers until the cost of an extra worker, i.e. the wage, is equal to the marginal benefit of that worker, i.e. the increase in total profits generated by that work. If final goods prices are constant across space, then this condition suggests that wages are, at least approximately, equal to the marginal product of labor. This reasoning suggests that wages are almost a direct measure of the marginal revenue product of labor in a given place.

Why are some places productive enough that firms are willing to pay substantially higher wages to locate there? Perhaps the simplest explanation is that a given place has some sort of innate geographic advantage, such as access to a valuable natural resource like a coal mine or a deep sea port. We should not be surprised that wages seem to have been high in nineteenth century Chicago, since that city had an unparalleled position as the hub of America's inland water network. Today, workers in Alaska earn high wages. From the worker's angle, these high wages compensate individuals for the discomfort of the cold. From the firm's angle, these high wages are compensated for by the natural-resource based Alaskan productivity.

While some locational advantages are innate, many more are man-made. A long tradition in urban economics, going back to von Thunen, emphasizes the advantages of proximity to customers and suppliers. If a firm is closer to its suppliers, then it is more productive because it can save on the shipping costs of its inputs. If a firm is closer to its customers, then it is more productive because it can save on the shipping costs of its final goods. The desire to save on these transport costs explains why great industrial cities grew up around America's early ports: an initial concentration of activity attracted firms that wanted to sell to the first residents and take advantage of the area's transport network. The ability to save on transport costs makes firms more productive and can offset the costs associated with higher wages.

Of course, throughout much of the 20th century, transport costs for goods have declined. Glaeser and Kohlhase (2004) document a roughly 90 percent reduction in the real cost of

moving a ton a mile by rail. This remarkable decrease becomes even more striking when we consider the general growth of the economy and the introduction of newer transport technologies like trucks and airplanes. The result of this decline is that erstwhile locational advantages associated with reduction in the cost of shipping goods have become far less important. The decline of cities that were built around the waterways of the Midwest was a natural consequence of the decreased importance of the costs of shipping goods. A second consequence of the decline in transport costs is that wages in areas that had a transportation advantage have fallen relative to the rest of the country.

While locational advantages associated with reductions in the cost of shipping goods are less important today than in the past, wage and productivity differences across space are as large as ever. The economic interpretation of that fact is that other forces must continue to make some places more productive than others. While transport costs for goods have dropped dramatically, the costs of moving people have not declined so steadily. Since time is the major input in the cost of moving people, and the value of time rises roughly with the wage rate, the cost of moving people remains high. As a result, many cities increasingly specialize in services that require face-to-face interactions.

The highest wage metropolitan areas tend today to specialize not in manufacturing, but in business services. These industries—including law, finance, consulting, and accounting— are more productive in some places than others because they are close to customers and each other. The complementarity between different types of services means that the productivity of a lawyer in Manhattan is generally higher than the productivity of the same lawyer in rural Montana, because in Manhattan the lawyer is physically proximate to potential clients in the financial services industry.

Proximity to people doesn't just reduce standard transport costs, but it also increases the access to the ideas of those people. A body of research following Alfred Marshall has emphasized a connection between local productivity and access to new ideas in particular locales. The great nineteenth century economist Alfred Marshall famously declared that in some concentrated locales "the mysteries of the trade become no mystery but are, as it

were, in the air." Jane Jacobs followed Marshall's lead in *The Economy of Cities* and emphasized the role that cities have historically played in generating new innovations. Jacobs argued that new ideas are formed by combining old ideas and the wealth of inspiration within dense urban areas made intellectual cross-fertilization easier.

Economists have become increasingly sympathetic to the view that cities serve as forges of human capital and incubators of innovation (Duranton and Puga, 2001). This view has been fused with the idea of human capital spillovers: location-specific productivity depends on the density of well educated smart people, since those people produce more ideas. One line of evidence in this area examines the determinants of wages within an urban area. Rauch (1993) showed that holding individual human capital constant, wages rise in high human capital areas. Figure 6 shows the relationship between the average income residual, i.e. income holding individual schooling and experience constant, and the share of the population in a metropolitan area with college degrees. People who live in more educated areas earn more, which may reflect higher productivity in those areas due to a faster exchange of ideas.

A second line of evidence on knowledge-based sources of local productivity has looked at patterns of urban growth. One fact that makes it seem that local knowledge is increasingly important for local productivity is the robust connection between initial human capital in an area and later growth. Local skills are among the best predictors of which metropolitan areas will succeed, especially in the older regions of this country. Figure 7 shows the correlation between the share of the population with bachelor's degrees in 1980 and the growth of the metropolitan area since then. This type of indirect evidence has pushed economists towards the view that locational productivity differences may owe something to the benefits that smart people gain from interacting with and learning from each other.

Housing Supply

The last major ingredient in the economic model of cities is housing supply. Economists are not so oblivious to the real world that they ignore the role of the built environment in the development of cities. Indeed, the percentage change in the number of people in a city is almost exactly the same thing as the percentage change in the number of homes in that city. Figure 8 shows this correlation across metropolitan areas in the U.S. between 1980 and 2000. If the number of homes is so tightly correlated with the overall growth of the area, then it is surely critical to understand the factors determining the growth in the number of homes in an area, which itself reflect the supply of housing.

Again, economists understand housing supply with the same no arbitrage tools that are used to understand the determination of wages and employment in the labor market. The key no arbitrage condition in the housing market is that housing prices must not be higher than the total cost of supplying new housing. If a builder can sell a home for more than it costs to build a home, then the logic of economics strongly suggests that a new home will be built. This logic implies that developers will keep on building to the point where the total costs of building a new unit are equal to housing costs.

Again, this condition is only meant to hold approximately. There are some market imperfections in the construction industry, although monopoly is likely to be rare since the number of builders in most metropolitan areas is quite large. Construction takes time, especially in the permitting process, and this means that builders are guessing the eventual sales price when they are initiating the permitting process. With these caveats, economists have still relied on the equilibrium condition that housing prices should be close to the total costs of building.

If every area had the same basic supply of housing, so that costs of production were more or less identical, then it would be perfectly reasonable to focus on factors other than housing supply. However, there is abundant evidence suggesting that housing supply differs from place to place. If there was one single housing supply for every place, then

we would expect to see high price in high construction areas and low prices in low construction areas.

Figure 9 shows the correlation between prices in 2005 and the number of new permits in a metropolitan area between 2000 and 2005. The figure shows that the expensive places in the U.S. have low development and the places in the U.S. with abundant development have low prices. This fact is not compatible with the view that places differ only in their level of housing demand. In that case, high price places would also have more development. Only differences in the supply of housing can explain why San Francisco has high prices and low permitting while Houston has high permitting and low prices.

Some of the differences in supply conditions reflect the actual physical costs of construction. Gyourko and Saiz (2006) estimate that one-fifth of the variation in housing prices across space can be attributed to differences in the cost of building homes. Materials, and especially labor, do cost more in some places than in others. Some part of the difference in housing supply also reflects physical conditions, like lack of land and natural barriers to development such as too much or too little water .

In my work in this area with Joseph Gyourko and others, I have argued that land use regulations, not construction costs or land density, explain the bulk of the differences in housing supply across space. This claim is based on both direct evidence linking land use controls with less construction and higher prices (Glaeser and Ward, 2006, Katz and Rosen, 1987) but also indirect evidence of many forms. For example, new construction per acre is more common in areas with higher density, not in areas with less land. Housing prices in Manhattan are far higher than the cost of supplying a new unit by building an extra floor on a new skyscraper. The discrepancy between price and cost seems to imply regulatory barriers on new construction. If this hypothesis is correct and differences in housing supply reflect differences in regulatory regimes, then these regulatory regimes are having a major impact on both the cost of living and urban growth.

Economics predicts that housing prices should not be higher than the costs of new building, but it does not predict any floor on housing prices. Housing prices can certainly fall below construction costs, but if they do, economics predicts that no new housing will be built. This scenario describes many of the inner cities of the rustbelt that have housing prices that are far below construction costs and almost no new construction (Glaeser and Gyourko, 2005). The durability of housing means that these places remain, despite the fact that their productivity and amenity levels are not high enough to justify housing costs that would cover the costs of construction. In these places, we expect to see continuing urban decline, because current low prices suggest that there is not nearly enough demand to prod the market into building new homes.

One implication of the abundant durable housing in declining cities is that attempts to revitalize these cities with new construction seem particularly odd. These declining cities have abundant housing and physical infrastructure relative to demand. That is why prices are so low. How could it make sense to respond to those conditions by building more infrastructure?

IV. The Empirical Approach of Economics to Cities

Disciplines are divided not only by core theoretical assumptions, but by their empirical methods. Over the past 50 years, economics has particularly distinguished itself by a focus on statistical work with large data sets, as opposed to case studies. Over the past 15 years, empirical economists have become particularly focused on causal inference. The strength of empirical economics is in quantitative empirical methods, especially sophisticated methods that are focused on exogenous sources of variation. By contrast, only a few economists have real expertise in non-quantitative forms of research, like ethnography.

The early classics of urban social science, none of which were written by economists, generally combine statistical work and descriptive observation. For example, the greatness of DuBois' *Philadelphia Negro* comes from its fusion of statistics and first

hand knowledge of the community. Over the next century, urban sociologists such as those in the "Chicago School" captured the rich details of urban life by developing tools of ethnographic research. Ethnographic classics, like Gans' *The Urban Villagers* or Liebow's *Talley's Corner*, presents rich depictions of particular neighborhoods.

Even the more statistical exercises in urban sociology and urban political science are often distinguished by their ability to capture more detailed information on particular areas. For example, over the course of the 20th century, sociological survey research became increasingly sophisticated about asking detailed questions about urban life. More recently, the work of Sampson, Raudenbusch and Earl (1997) use movie cameras and other technologies to acquire visual images of particular places that were then transformed into statistical measures of neighborhood activity that went far beyond the usual administrative data.

Economists have rarely been so willing to invest in getting richer depictions of any particular neighborhood. Since economists tend to be interested more in common patterns than in unique features of particular locales, few economists have tried to comprehensively measure any particular place. Moreover, since economists came late to the study of cities, there was little incentive to try to compete with sociologists in an area where that discipline is particularly strong. Finally, economists have a particular interest in financial variables like income and housing prices that do tend to be administratively available.

While economists have not contributed much to either ethnography or rich environmental measurement, economists have focused particularly on causal inference and exogenous variation. This focus comes from the strong attachment to formal economic theory. Economic models are generally geared towards predicting the relationship between an exogenous variable and endogenous variables, such as the link between an exogenous cost shock to an industry and both prices and quantities. Since our models are geared towards links between exogenous variables and outcomes, our theoretical work has also come to focus on exogeneity.

This theoretical background does not preclude interest in the correlation between two endogenous variables. For example, the relationship between income and housing prices shown in Figure 3 is certainly relevant to urban economics. However, the lack of exogenous variation in that figure makes it inherently unsatisfying to many economists. The link shown in the figure might mean that some places are intrinsically more productive than others and therefore have high wages which then cause high prices. Alternatively, the picture might mean that some places are intrinsically more expensive than others and those high prices then ensure the need to pay high wages. As discussed in the previous section a whole theoretical apparatus that embeds endogenous housing supply is needed to truly make sense of this correlation.

For an economist, making sense of Figure 3 requires a source of exogenous variation in the productivity of different regions. That exogenous variation would then enable us to trace out the impact that higher productivity has on wages and prices and housing supply. Unfortunately, such exogenous sources of variation are often difficult to come by. One approach has been to look at the changes over time in the international price of commodities produced by a particular locale and then to look at how wages, prices and construction change with respect to that price. A particularly clear example is the price of oil which strongly influences prices, construction and wages in Texas. Since the price of oil is determined by factors largely outside of the Texas economy, this provides us with one source of exogenous variation that can be used to examine the link between income and prices.

Over the past 15 years, economists have made much progress on using the tools of casual inference to examine urban issues. For example, Hoxby (2000) used the number of rivers across metropolitan areas to provide exogenous variation in the number of natural barriers which then predicts the number of governments in an area. Using this natural source of variation, Hoxby finds that more inter-governmental competition increases school performance.

No one has been more successful in identifying exogenous sources of variation than Steve Levitt, whose popular book *Freakonomics*, made his work world famous. Levitt's work often relies on exogenous events, like ACLU lawsuits against prisons, to identify important economic relationships, such as the impact of incarcerating criminals on crime and urban growth (Levitt, 1996, Berry-Cullen and Levitt, 1999). Without such sources of exogenous variation, economists would find it difficult to say anything meaningful about the relationship between incarceration and crime, since incarceration rates are themselves a function of the amount of crime in an area.

One approach to causal inference has been to use spatial discontinuities in public policy. Since policies change discontinuously at a border, it is possible to identify the impact of policy separately from other forces if those other forces are assumed to change more continuously. The work of Sandra Black on housing prices and school districts discussed above is one use of spatial discontinuities in economic policy. Holmes (1998) used these discontinuities to look at the impact of employment policies, like right-to-work laws on firm locations. While this approach produces something of a natural experiment associated with a somewhat random change in public policy at a border, it faces difficulties in identifying the impact of a particular policy if many policies simultaneously change at a border.

A final natural extension of the economic focus on causal inference has been an increasing interest in designing new policy experiments. Perhaps the most famous of these experiments is the Moving-to-Opportunity program funded by the Department of Housing and Urban Development. In this experiment, subjects that lived in high poverty areas were randomly allocated into three groups. One group was the control and received no assistance. A second group received a standard Section 8 rental voucher. The third group received a special voucher that committed the program to pay for housing if the recipient moved to a low poverty neighborhood. The recipients in the two treated groups both used their vouchers to move to areas with less poverty and less crime (Katz, Kling and Liebman, 2001).

This experiment was particularly exciting to economists because it seemed to offer the hope of identifying the impact of neighborhood choice on outcomes for children and adults. The basic correlation between poverty and growing up in poor neighborhoods seemed to suggest the presence of neighborhood effects, where children in those neighborhoods have unfortunate outcomes because of a lack of economic opportunity, role models or good schools. However, it was always possible that these correlations reflected the fact that parents in these neighborhoods were intrinsically different in some hard-to-measure way from parents in more affluent neighborhoods. The Moving-To-Opportunity experiment gave us the opportunity to compare parents who differed only on the basis of their receiving this voucher.

The basic findings of this experiment have suggested that neighborhood effects were probably over-rated. Parents who received the vouchers do seem to have become happier, but their economic success did not increase. Female children of voucher recipients did become somewhat more successful, but male children fared worse. The studies that came out of the Moving to Opportunity program have used exogenous variation to challenge prevailing wisdom that came from ecological correlations. These findings show that economists' use of exogenous variation can produce very different results than results using raw correlations.

The attention to causal inference is one empirical by-product of economists' focus on formal theory. Another by-product is the use of data to estimate formal models. For example, Ellison and Glaeser (1997) provide a formal model of firm location choice that then creates a natural index of geographic concentration of industry. Rosenthal and Strange (2001) use this index to understand the roots of agglomeration economies. Indeed, a common view in economics is that researchers should be able to justify any regression as an attempt to estimate the parameters of a formal model.

V. Economics and Urban Policy

The economic approach to urban policy combines the use of cost-benefit analysis and the assumption that the goal of policy is to increase the choices available to people. The most important part of this assumption is that people, not places, are the important outcomes. A policy that yields a beautiful place, but does little to increase the welfare of individuals has little appeal to most economists. Policies make sense to economists if their benefits to people outweigh their costs. This may not distinguish economists from sociologists, but it does distinguish economists from some architecturally oriented urban planners and from place-based politicians.

The economists' desire to put people first might seem obvious, but it is often in conflict with much place-based urban policy. For example, regional policies in Europe, such as the European Union's spending on infrastructure for poorer areas like the Mezzogiorno and in the U.S., such as the Appalachian Regional Commission, are classic place-based policies that aim to make particular regions wealthier. In the wake of Hurricane Katrina, many advocates for New Orleans called for hundreds of billions of dollars to be spent so that the city would come back.

The economics approach to public policy pushes us to ask whether this money would be better spent on people, rather than place based policies. For example, would 500,000 residents of New Orleans be better off with 200,000 dollars apiece or 100 billion dollars worth of government infrastructure. While it seems hard to argue that infrastructure spending was the best thing for the people of New Orleans, economics does not give us a universal rule against place-based spending. People's lives are certainly enriched when they live in a successful place and there surely are times when the best way to help people is to improve a place.

For example, late nineteenth century investments in water systems were placed-based policies. Massive public works projects, like the Croton Aquaduct, delivered clean water to residents of a particular locale. Clean water was critical for avoiding public health

disasters like Cholera epidemics and it is hard to argue that these place-based investments didn't deliver benefits that exceeded costs. Crime prevention is also generally handled place-by-place and investments in place-based policing also often seem to deliver benefits that outweigh costs.

However, there are other cases where place based systems do not seem so sensible. First of all, the beneficiaries of much place-based spending are often local landowners rather than current residents. For example, building a fancy contemporary art museum in a poor neighborhood that is short on contemporary art aficionados may well increase housing prices because of rich professionals who are willing to pay more for proximity to the museum. However, if the original residents don't care about the art then they will be made worse off because they need to pay more for their housing and haven't gained something that they particularly value.

A second issue with place-based policies is that they artificially distort migration decisions. One can easily argue that the best thing that can happen to the residents of a declining region is that they leave that region and move to areas with a brighter economic future. If place-based policies reduce the incentive to migrate, then they may reduce the beneficial process of moving from less productive areas to more productive areas. This argument is akin to arguments that economists make against propping up declining industries. Economic efficiency requires workers to move to more productive industries or places. Economic policy that tries to stop that process is reducing efficiency.

The third issue with place based policies is that these policies have often provided excuses for vast expenditures which benefit contractors more than target populations. Since amorphous appeals to the magic of a particular place are not amenable to conventional cost-benefit analysis, this means that all sorts of projects can be justified that do not provide people with benefits that are large enough to justify their costs. For example, much of the urban renewal spending in the 1960s was seen as being a way of rebuilding cities, but few people actually bothered to ask whether this spending would deliver people benefits large enough to offset costs.

Economics and Laissez-Faire

Some of the most famous results in formal economic theory suggest that competitive markets will deliver a socially desirable outcome in the absence of market failures. Of course, even these results do not claim that free markets equitably redistribute income. Moreover, market failures often appear endemic in urban settings. One example of market failures are externalities, which are defined as settings where one person's action impacts his neighbors in ways that are not mediated by the price system. A classic example is the production of pollution. A second example of a market failure is the existence of public goods, like the legal system, that are non-rival and non-excludable. The close connection of people in cities generally increases the level of externalities and the scope for government intervention.

The problems of externalities like pollution, the spread of disease and congestion mean that pure laissez-faire is rarely an option in urban economics. Few economists question the wisdom of disease-reducing expenditures on clean water or the need for local police. Moreover, the existence of agglomeration economies in cities makes it at least possible that economic externalities exist that could, in principle, justify significant government interventions in local economies.

Yet despite abundant externalities, economists remain more skeptical of many forms of government intervention than representatives of many other disciplines. This skepticism is based not on an uncritical faith in the free market but rather in doubts about the competence and benevolence of government. Adam Smith's *The Wealth of Nations*, which is the closest thing economics has to a founding document, was the product of the Scottish Enlightenment. Like many thinkers of the Scottish Enlightenment, Smith questioned the tendency of kings to perfectly care for their subject's interests. While there is little doubt that democracy has generally made governments more benign, economists have retained doubts about the perfection of government. As such, those economists who have a predilection for limited government base that view not on an

unquestioning faith in the market, but rather on doubts about the perfect benevolence of the state.

In the context of cities, doubts about governmental goodness and competence certainly have some basis in fact. Corruption was endemic in many American cities throughout much of the 19th century. Urban leadership has often been beholden to special interests on both the left and the right. Moreover, the limited resources available to local governments often ensures that even well intentioned city governments have difficulties figuring out the right course of action.

Skepticism about the limits of government is particularly important in debates about activist economic policy. On one side of this debate, pro-intervention advisors urge using government intervention to secure the growth of industries that will generate positive externalities for other people and firms in the city. Anti-intervention advisors doubt the ability of the government to get this right. Beason and Weinstein (1996) show that Japan's famous Ministry of International Trade and Industry was not good at directing funds to firms that were either ex post successful or ex ante likely to have a high return. The track records of city governments, which have often targeted elderly industries with little chance for growth, are hardly encouraging.

The opponents of firm-level targeting argue that the best economic development plan is to attract smart people and get out of their way. Within this group, there is some division on what attracts smart people. Some authors, like Richard Florida, emphasize hip downtowns. Others emphasize good schools, safe streets and fast commutes. While all of these debates are currently active, there is no sense that economists have come to any discipline-wide agreement on these issues.

City Limits and Migration

Urban economists are less divided on the view that good urban policy must reflect the fact that people and firms are mobile. Location choice lies at the heart of urban

economics, so it is unsurprising that urban economists have been particularly attuned to the connection between government policy and location choice. The seminal paper in local public finance is Tiebout (1956), which argued that many different localities would provide consumers with a variety of choices about the kind of public services that they would like to consume. Just as economists think that competition among firms ensures that consumers will have a rich set of choices about product choices, competition across governments seems to allow consumers to vote with their feet and choose the government that best fits their needs.

Tiebout mentions, but does not emphasize, a second potential benefit from competition among government: improving the incentives facing governments. Since economists tend to doubt the innate benevolence of the public sector, there may be benefits from forcing governments to compete to attract people and firms. Just as private sector monopolies don't serve consumers well, public sector monopolies can be just as problematic. Hoxby (2000) is a classic paper on the connection between inter-governmental competition and better school provision.

Economists' emphasis on the mobility of people and firms guides policy recommendations at both the local and national level. At the local level, the mobility of factors means that policy makers face strong limits in the extent that they can pursue redistribution policies. Raising taxes on the rich to fund the poor may be socially attractive, but at the urban level such redistribution seems to lead to a quick flight of the rich. The classic statement of this point is Paul Peterson's *City Limits*. Peterson is a political scientist, although one who was strongly influenced by economics and especially Tiebout. One hypothesis is that an increasing recognition of the mobility of wealth explains why socially visionary mayors of the 1960s, like John Lindsay, have been replaced by mayors who look more like city managers, like Richard M..Daley (Glaeser and Kahn, 1999).

The mobility of the rich suggests both that mayors should avoid repelling the well-to-do, but also that they might want to pursue policies that will attract high human capital

residents. The robust correlation between urban success and a skilled population suggests that attracting skilled people might be a particularly sensible policy for a mayor who is trying to generate long run economic success for his area. Of course, there is less consensus about what things will attract a skilled population to a city.

At the national level, the emphasis on the mobility of people and firms often leads economists to favor spatial neutrality. Many economists think that the government should not interfere in the competition between firms, each of which is trying to maximize their own profits. Analogously, many economists think that it makes sense for the national government to restrain from interfering in the competition between cities, each of which is trying to attract people and firms. Policies that favor particular cities will dilute the incentives to compete and create spatial distortions that artificially push people to one place or another. There is no conflict between economists advising mayors on how to build up their cities while simultaneously advising the national government to stay neutral, any more than there is a conflict between economists advising individual firms and simultaneously opposing government intervention in the marketplace.

VI. Conclusion

Urban economics is only one of the many disciplines that contribute to our understanding of cities. Among the social sciences, the economic approach to cities is distinguished by its strong ties to formal economic theory based on the concept of a no-arbitrage equilibrium. Economists seek to understand cities with a framework that requires people to be indifferent over space, employers to be indifferent over where to locate and how many people to hire and builders to be indifferent about whether or not to build more or taller buildings. This attachment to theory provides urban economics with discipline and with a clear structure. Almost everything urban economists do can be understood as part of the large question of understanding why people choose to locate in urban areas.

While the economic approach to cities has many strengths, the economic approach also has profound limitations. Economists have never acquired the skills to study the built

environment or the intricacies of individual neighborhoods. Our overwhelming focus on quantitative methods has left us poorly suited to treat historical narratives in a scientific fashion. While I believe that no one can make sense of cities without the tools of economics, I also believe that no economist can make sense of cities without borrowing heavily from other disciplines.

References

Alonso, William (1964) *Location and Land Use*. Cambridge: Harvard University Press.

Beason, Richard and David Weinstein (1996) "Growth, Economies of Scale and Targeting in Japan (1955-1990)" *Review of Economics and* Statistics 78(2): 286-295.

Becker, Gary S. (1965) "A Theory of the Allocation of time," *Economic Journal* 75(299): 493-517.

Black, Sandra E. (1999) "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics* 114(2): 577-599.

Duranton, Gilles and Diego Puga (2001) "Nursery Cities: Urban Diversity, Process Innovation and the Life Cycle of Products," *American Economic Review* 91(5): 1454-1477

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

Glaeser, Edward L. and Matthew Kahn (2001) "Decentralized Employment and the Transformation of the American City," *Brookings-Wharton Papers on Urban Affairs* 2 (2001).

Glaeser, Edward L. and Matthew Kahn (1999) "From Lindsay to Giuliani: The Decline of the Local Safety Net?", *Federal Reserve Bank of New York Economic Policy Review* 5(2): 117-132.

Glaeser, Edward L., Kahn, Matthew and Jordan Rappaport (2007) "Why do the poor live in cities?" *Journal of Urban Economics*, forthcoming.

Glaeser, Edward L. and Janet E. Kohlhase (2004) "Cities, Regions and the Decline of Transport Costs," *Papers in Regional Science* 83(1): 197-228.

Glaeser, Edward L., Kolko, Jed and Albert Saiz (2001) "Consumer City," (joint with J. Kolko and A. Saiz), *Journal of Economic Geography* 1: 27-50.

Glaeser, Edward L. and Bryce Ward (2006) "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston," NBER Working Paper # 12601.

Gyourko, Joseph and Albert Saiz (2006). "Construction Costs and the Supply of Housing Structure", *Journal of Regional Science*, Vol. 46, no. 4 (November 2006): forthcoming.

Gyourko, Joseph and Joseph Tracy (1991) "The Structure of Local Public Finance nad the Quality of Life," *Journal of Political Economy* 99(4): 774-806.

Holmes, Thomas J. (1998) “The effects of state policies on the location of industry: Evidence from state borders,” *Journal of Political Economy* 106 (4): 667 – 705

Hoxby, Caroline M. (2000) "Does Competition Among Public Schools Benefit Students and Taxpayers?" *American Economic Review*, 90 (5).

Katz, Lawrence, Kling, Jeffrey and Jeffrey Liebman (2001) “Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment,” *Quarterly Journal of Economics* 116(2): 607-654.

Katz, L. and K.T. Rosen (1987) “The Interjursdictional Effects of Growth Controls on Housing Prices,” *Journal of Law and Economics* 30(1): 149-160.

LeRoy, Stephen and Jon Sonstelie (1983) “Paradise Lost and Regained: Transportation Innovation, Income and Residential Location,” *Journal of Urban Economics* 13: 67-89.

Mills, Edwin S. (1967) “An Aggregative Model of Resource Allocation in a Metropolitna Area,” *American Economic Review 57*(2): 197-210.

Muth, Richard (1969) *Cities and Housing.* Chicago: University of Chicago Press.

Rauch, James E. (1993) “Productivity gains from geographic concentration of human capital: Evidence from the cities,” *Journal of Urban Economics* 34 (3): 380 – 400.

Roback, Jennifer (1982). “Wages, Rents, and the Quality of Life”, *Journal of Political Economy*, Vol. 90, no. 4 (December 1982): 1257-78.

Rosen, Sherwin (1979). “Wage-Based Indexes of Urban Quality of Life”. In *Current Issues in Urban Economics*, edited by Peter Mieszkowski and Mahlon Straszheim. Baltimore: Johns Hopkins Univerity Press, 1979.

Rosenthal, Stuart and Will Strange (2001) “The Determinants of Agglomeration,” *Journal of Urban Economics* 50(2): 191-229.

Sampson, R.J.., S.W. Raudenbush, and F. Earl. 1997. “Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy.” *Science* 277(5328): 918-24.

Tiebout, Charles (1956) “A Pure Theory of Local Expenditures” *Journal of Political Economy* 64(5): 416-424.

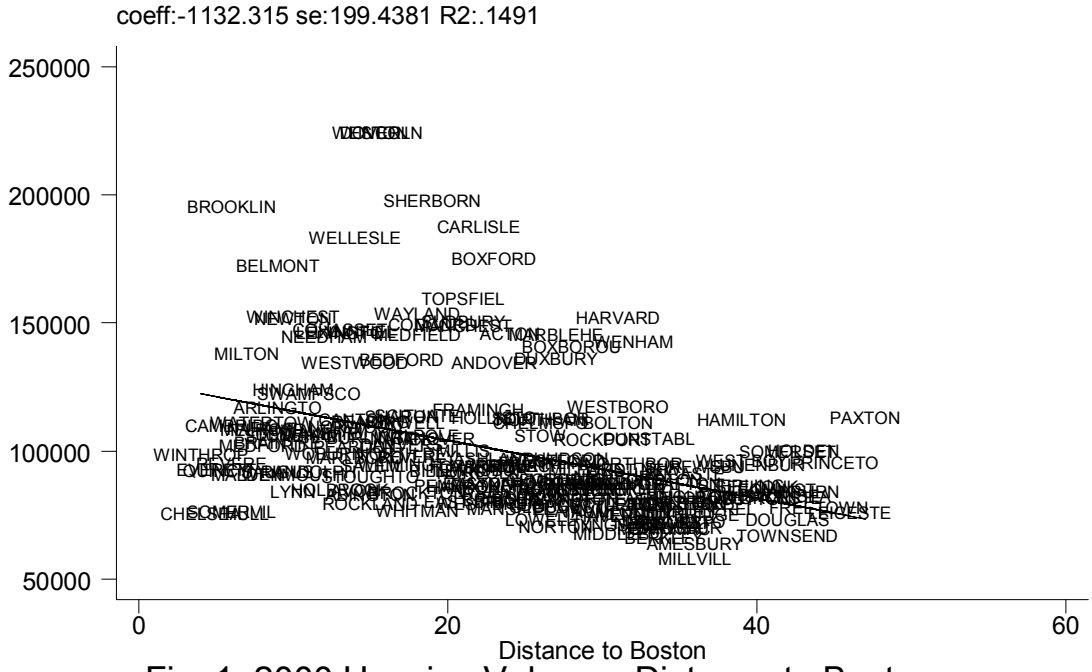

Fig. 1: 2000 Housing Value on Distance to Boston

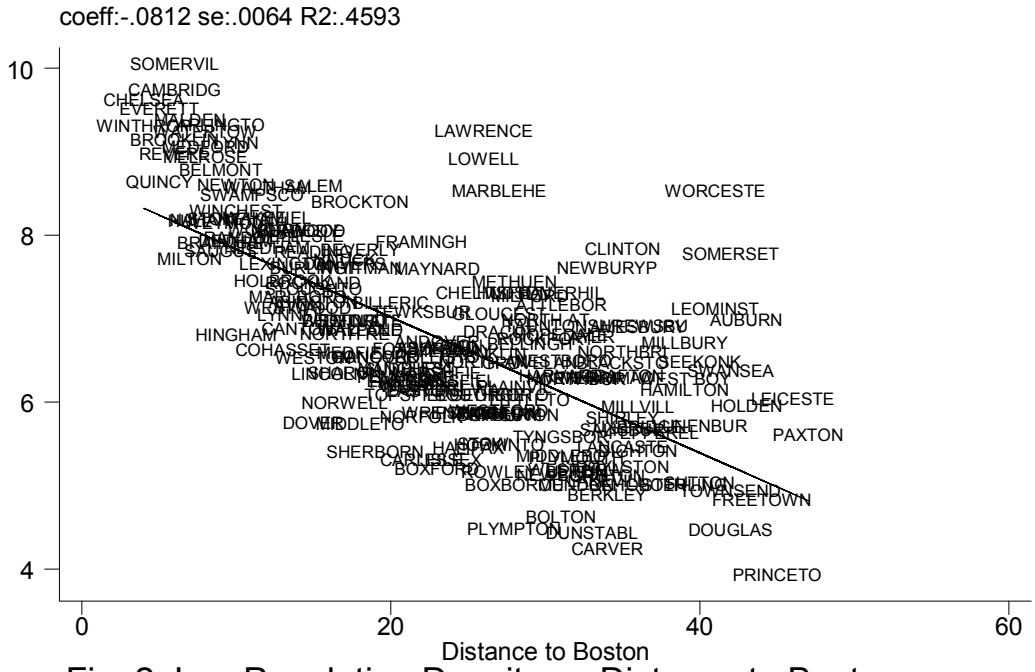

Fig. 2: Log Population Density on Distance to Boston

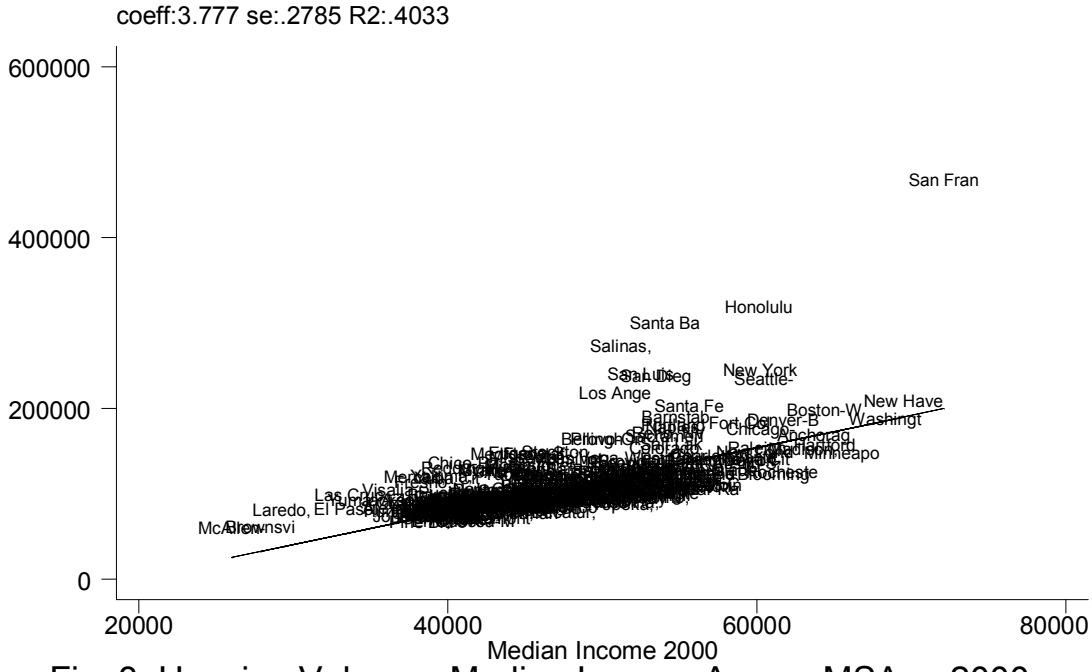

Fig. 3: Housing Value on Median Income Across MSAs - 2000

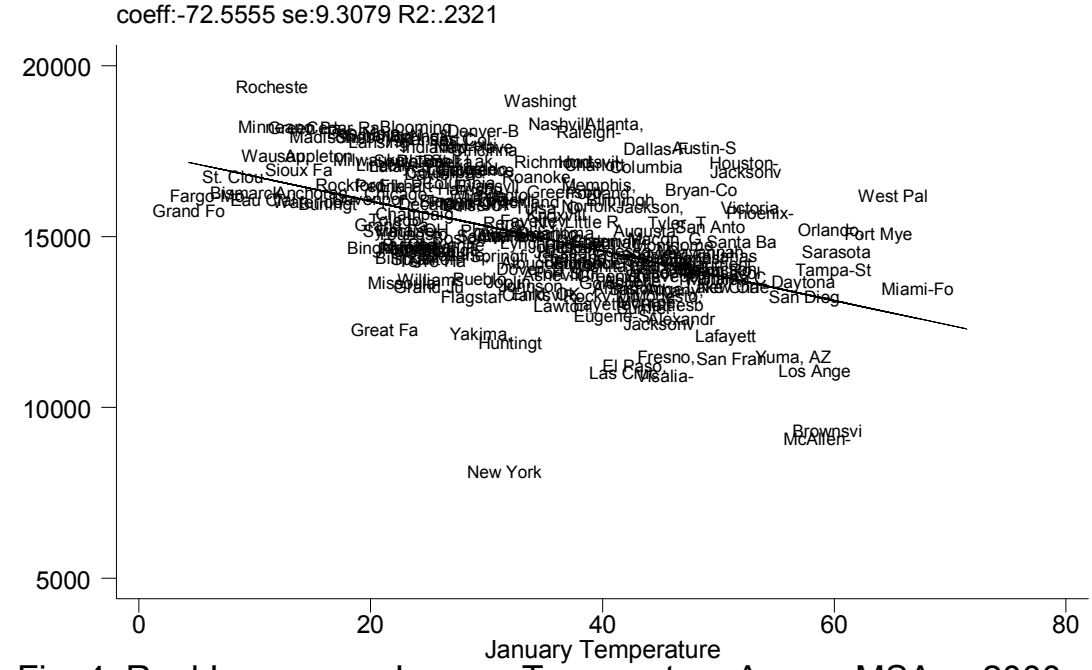

Fig. 4: Real Income on January Temperature Across MSAs - 2000

Fig. 5: Transportation Time on Median Income Across Cities - 2000

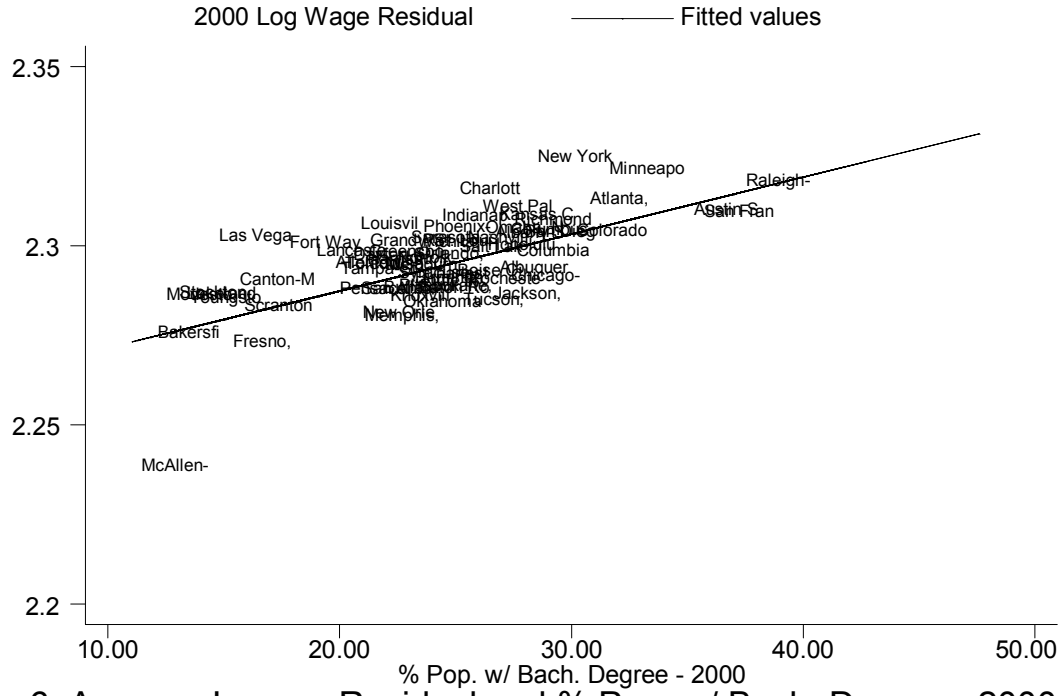

Fig. 6: Average Income Residual and % Pop. w/ Bach. Degree - 2000

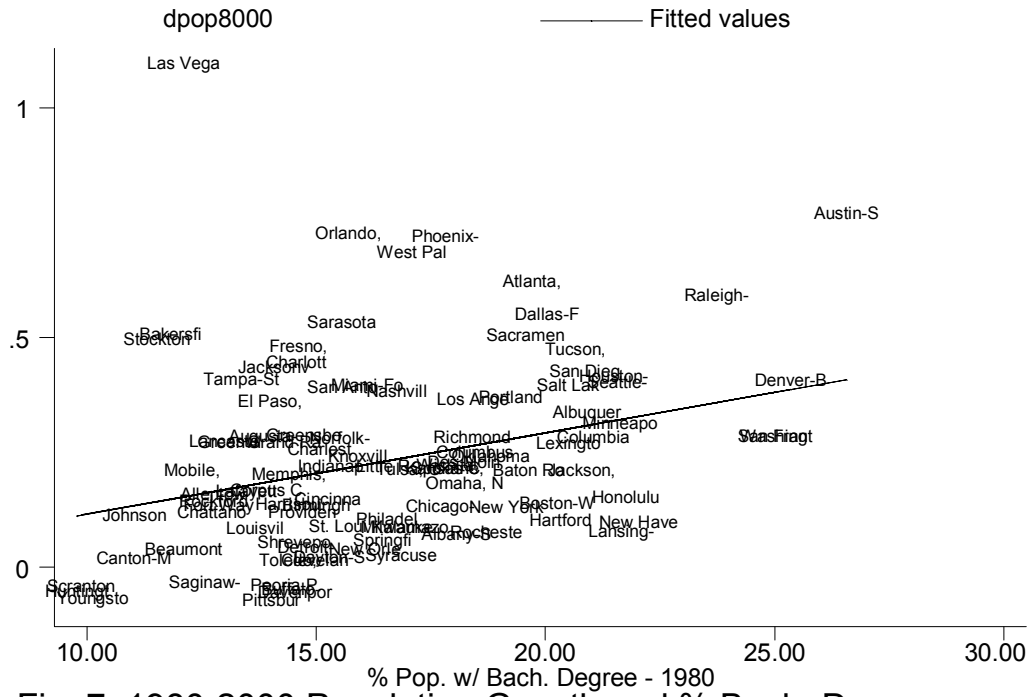

Fig. 7: 1980-2000 Population Growth and % Bach. Degree

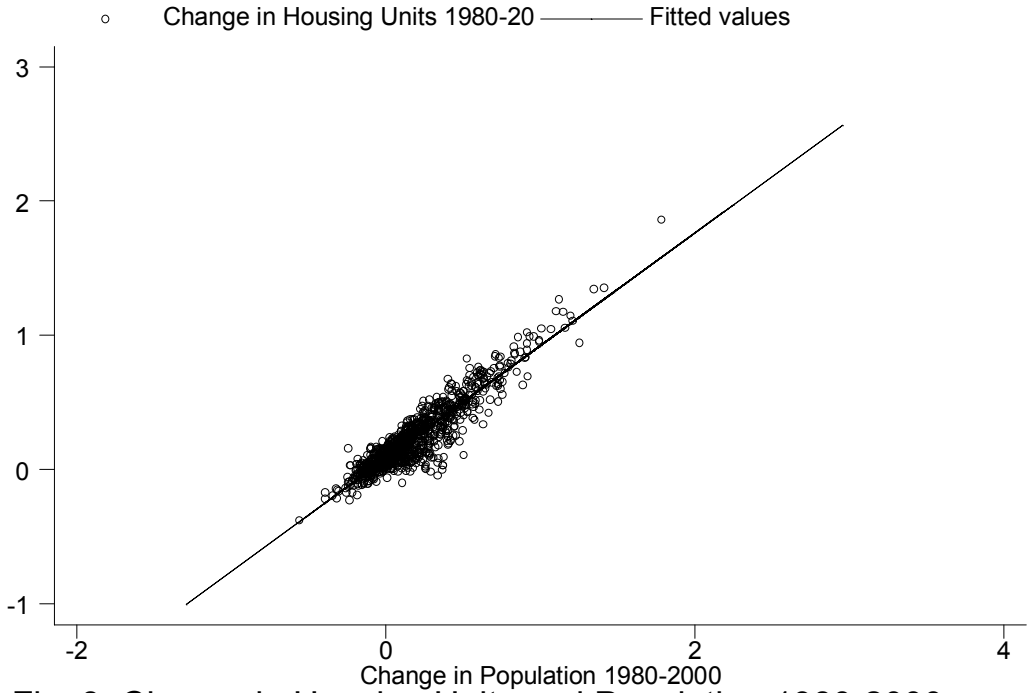

Fig. 8: Change in Housing Units and Population 1980-2000

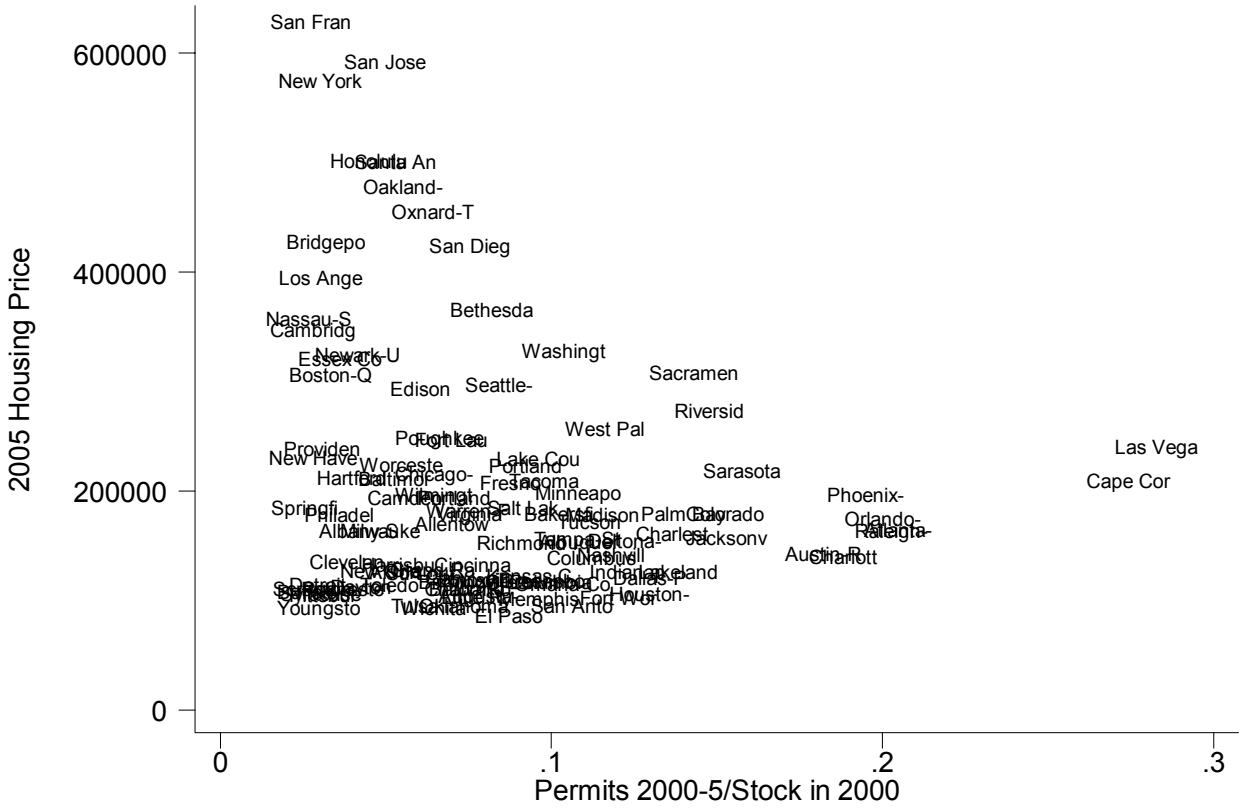

Fig. 9: Correlation Between 2005 Housing Prices and New Permits 2000-2005