

Local Public Goods and the Market for Neighborhoods

By

Bryan Ellickson

**Discussion Paper Number 100
September 1977**

**Research for this paper was supported by
National Science Foundation Grant SOC 74-23109.
I would like to thank Joseph Ostroy for his helpful comments.**

1. Introduction

The notion of a market for neighborhoods can mean different things to different people. The phrase may be used, for example, simply to indicate that consumers shop for neighborhoods as well as houses. When an economist says that a market exists, however, he usually means that the market is competitive. To claim that a market is competitive implies in turn that it is efficient, that a prima-facie case exists against intervention in the market and that the proper role of government is to confine its actions to the redistribution of income.

To assert that the market for neighborhoods is competitive represents, therefore, a very strong claim. The main purpose of this paper is to establish conditions under which such a competitive market could exist and, assuming that these conditions are satisfied, to investigate the properties that such markets would exhibit.

Establishing conditions under which a competitive market for neighborhoods could exist is not a trivial task. The existence of neighborhoods implies the existence of externalities, commonly thought to be a source of market failure. Residence in a neighborhood can also be viewed as consumption of a public good, and public goods are another potential source of market failure. There is little basis in conventional economic theory to justify the belief that competitive behavior is possible under such conditions.

In this paper we view the development of a competitive theory of the market for neighborhoods as essentially equivalent to the formulation of a theory of competitive equilibrium in an economy with local public goods. Over the past two decades, economists have made considerable progress in justifying the notion of a competitive market for local public goods. But this theory has not been given a solid theoretical foundation. In a recent paper on this subject (Ellickson [1977c]) I have provided the necessary

foundations, an effort that has required a substantial rethinking of the entire theory of public goods.

This paper represents an application of my theory of local public goods to the market for neighborhoods. Application of the theory is not, however, immediate. In developing my theory of local public goods, I abstracted from any explicit consideration of the housing market. But to discuss the market for neighborhoods, it is impossible to separate the choice of a house from the choice of a neighborhood. Therefore, I begin this paper by presenting in Section 2 a theory of residential choice, a treatment that breaks new ground by developing a model capable of handling housing characterized by an arbitrary number of attributes. The resulting theory is shown to be empirically testable, and I end this section with a summary of some empirical evidence that, from the point of view of consumers, neighborhoods matter.

In Section 3 I present an overview of my theory of local public goods and its application to the market for neighborhoods. Sections 4 and 5 treat two specific examples of the theory designed to highlight certain aspects of particular relevance to neighborhoods.

At this juncture I want to comment on the style in which this paper is written. I have made a strenuous effort to avoid introducing mathematical technicalities wherever possible. Thus, in particular my discussion of residential choice in Section 2 and local public goods in Section 3 represents a paraphrase of the formal theory. I refer the reader interested in technical details to my other papers on the subject (Ellickson [1977b, 1977c]). In those instances where a mathematical statement seemed essential, I have followed the mathematics with a verbal summary of the main results.

2. Residential Choice

My main purpose in this paper is to consider the workings of a market for neighborhoods. A necessary first step is the development of a model of residential choice: How does a consumer choose a place to live?

Over the last decade and a half, economists have made considerable progress in answering this question. Against this background, it is easy to overlook the intrinsic difficulty that the problem poses for economic analysis. Economics works best when dealing with commodities that are homogeneous and perfectly divisible. The housing options open to a consumer in a particular metropolitan area, however, are (i) differentiated (because of location if nothing else) and (ii) indivisible (a consumer either chooses to reside in a particular house or he does not). The housing market should, therefore, be very difficult to model. Nevertheless, we now have quite successful models of such markets. How has this been accomplished?

a) The New Urban Economics

Consider the following stylized version of the standard model of urban residential location. Suppose that when a consumer chooses a house, all that he cares about is lot size, z_1 , and accessibility to the central business district, z_2 . Letting x denote the vector of commodities consumed other than housing, we assume that the consumer's utility function is given by:

$$U(x, z_1, z_2) \tag{1}$$

Assume that the consumer faces a budget constraint

$$p_x x + r(z_2)z_1 = y \tag{2}$$

where p_x is the vector of prices for the commodities included in x , $r(z_2)$ is the price of a unit of land as a function of accessibility and y is the consumer's income. Maximizing (1) subject to (2) yields the demand side of the standard model of Alonso [1964] and Muth [1969], where $r(z_2)$ is assumed

given. By introducing more explicit assumptions into the analysis regarding the form of consumer utility functions, the supply of land and the distribution of income, it is possible to derive explicitly the price function $r(z_2)$ which will sustain equilibrium^{1/}.

In this way we obtain a coherent model of equilibrium for an urban housing market. There can be no reasonable quarrel with the value of such efforts: they have been instrumental in giving direction and substance to urban economics as a field. But in some respects these models are not very satisfactory. The main thrust of the new urban economics has been toward the specific rather than the general, toward explicit computation of solutions for particular models rather than analytic characterization of solutions for general models. As a result we have no general proof that competitive equilibria will exist for urban residential housing markets and no proof that, if such equilibria exist, they will be Pareto optimal. Of greater concern for our present purpose, we have no practical way of characterizing housing market equilibrium if housing involves more than two characteristics: explicit computation is simply not feasible.

b) A General Model of Residential Choice

How then are we to develop a model of housing market equilibrium flexible enough to admit multiple characteristics so that, in particular, neighborhood characteristics can be allowed to influence housing choice? Suppose we return to our initial characterization of housing markets as involving a collection of differentiated indivisible commodities. When Alonso and Muth developed their model, economic theory provided little guidance on how to proceed: markets with indivisible commodities were uncharted terrain while product differentiation seemed to lead, if anywhere, in the direction of monopolistic rather than perfect competition. The situation is far different today. We now have

^{1/} See Mills and MacKinnon [1973] for a useful review of the "new urban economics" as this class of model is called.

general existence theorems for competitive equilibrium with indivisible commodities (Broome [1972], Mas-Colell [1975b]), and we know that product differentiation (with or without indivisibilities) can under appropriate circumstances be consistent with perfect competition (Mas-Colell [1975a], Hart [1977]). It seems reasonable to suppose that urban economics could profit from these new developments and, as we shall see, that is indeed the case.

What do these recent developments in economic theory imply about how to model the urban housing market? The introduction of housing characteristics does have a role to play: the assumption that houses with similar attributes will be treated as close substitutes by consumers provides an ingredient essential to the justification of the price-taking hypothesis underlying competitive analysis^{2/}. However, it is not necessary to assume that one of the characteristics functions as a divisible commodity^{3/}. This assumption is adopted in the new urban economics because it facilitates explicit computation, the only way to establish that equilibrium exists in such models. But in the more general setting, existence can be demonstrated non-constructively, freeing us to be more flexible in the models that we employ.

Thus, we are now free to assume that consumer utility depends on an arbitrary number, s , of housing characteristics so that the utility function for the n^{th} consumer is:

$$U_n(x_n, z_n) \quad (3)$$

^{2/} To be more precise, Mas-Colell [1975a] uses this condition on preferences to obtain equivalence between the set of competitive allocations and the core of the economy, justifying the competitive hypothesis in the manner of Edgeworth [1881].

^{3/} Lot size, z_2 , played this role in the stylized version of the new urban economics presented above; in some variations, "housing services" assumes the same role.

where x_n is an r -dimensional vector of private divisible goods other than housing and z_n is an s -dimensional vector of housing characteristics. We assume that z_n belongs to some set K of potential housing characteristics where K is a compact metric space. Prices of the non-housing commodities are given by an r -dimensional price vector p_x . Housing prices are described by a function $h:K \rightarrow R$, traditionally called a hedonic price function. If K is infinite, then the assumptions imposed on preferences (that houses with similar attributes are close substitutes) imply that the function h is continuous; if K is finite, then this "function" simply associates to each house of type z a price $h(z)$. In competitive equilibrium the consumer then maximizes (3) subject to the budget constraint:

$$p_x x_n + h(z_n) = p_x \bar{x}_n \equiv y_n \quad (4)$$

where \bar{x}_n is the consumer's initial endowment^{4/}, and y_n is the consumer's income.

The formulation of residential choice we have presented is sufficiently general to meet our needs, but as it stands it is not very easy to use. In particular, there is no way to use this constrained maximum to derive a demand function for houses, the technique that economists normally use to convert the results of utility maximization to usable form. True the maximization of (3) subject to (4) yields a solution (x_n^*, z_n^*) , but how can we describe the dependence of this solution on the underlying parameters in a useful way? If we can't estimate demand functions, then what can we estimate to gain some insight into the nature of consumer behavior?

^{4/} Houses are produced, so none are initially owned. To keep matters as simple as possible, we will assume that, apart from housing production, we are dealing with an exchange economy.

One approach, common to several recent empirical studies of the housing market, is to estimate the hedonic price function h . Such estimates do provide a useful check on whether relevant characteristics have been left out of the model. However, it is clear within the present setting that hedonic price functions convey little information about consumer behavior: they are simply vehicles for describing equilibrium housing prices, and equilibrium prices result from the interplay of demand and supply^{5/}. Because the shape of the hedonic function depends on supply as well as demand, characteristics can be relevant to residential choice even if their introduction into the hedonic price function yields a coefficient not significantly different from zero^{6/}. For the same reason, there is little justification for assuming a particular functional form for h that persists over time, an assumption which -- were it justifiable -- would enable us to derive a demand function for characteristics. Clearly some other approach is needed if our model of residential choice is to yield testable implications about consumer behavior.

The problem we face is a reflection of the indivisibility of the housing commodities: what we want to explain is not the quantity but the type of housing that will be chosen. We need a way of viewing the decision process that focuses on this essential aspect of residential choice. The solution we propose involves the revival of one of the earliest devices employed in the theory of residential location, Alonso's bid price function.

c) Reformulation of the Model in Terms of Bid Price

Suppose we consider again the maximization of (3) subject to (4) where we now hold constant the choice of a house. z_n is fixed and the consumer is required to spend $h(z_n)$ on housing. His income net of housing expenses is

^{5/} For a forceful statement of this position, see Rosen [1974].

^{6/} This observation will be justified in Section 5.

$y_n - h(z_n)$, and he maximizes $U_n(x_n, z_n)$ subject to the constraint $p_x x_n = y_n - h(z_n)$. Assume that the utility function is strictly quasi-concave in x_n . We may then describe the solution to this constrained maximization problem in terms of an indirect utility function with all of the usual properties:^{7/}

$$\phi_n(p_x, z_n, y_n - h(z_n)) \quad (5)$$

The function ϕ_n gives the maximum utility the consumer can achieve given the price vector p_x , his income y_n and the type of house z_n he has been assigned. The consumer then chooses the type of house that maximizes utility: i.e., the vector of characteristics z_n^* which solves the problem

$$\max_{z_n \in K} \phi_n(p_x, z_n, y_n - h(z_n)). \quad (6)$$

To translate this solution into the language of bid price functions, we then consider a level curve for the indirect utility function, substituting bid price V_n for the hedonic price $h(z_n)$:

$$\phi_n(p_x, z_n, y_n - V_n) = u_n \quad (7)$$

where u_n represents some particular level of utility. For a given price vector p_x , income y_n and the utility level u_n , equation (7) defines an implicit relation between housing characteristics z_n and housing price V_n . Assuming consumers are not satiated, the indirect utility function will be a (strictly) monotonic increasing function of income net of housing cost, and therefore strictly decreasing as a function of housing cost. Hence, we can solve equation (7) to obtain the bid price function:

$$V_n = \psi_n(p_x, z_n, y_n, u_n) \quad (8)$$


Constrained maximization of (3) subject to (4) is equivalent to the maximization given by (6). The maximization represented by (6) is in turn

^{7/} Indirect utility functions have been employed in related contexts by Ellickson [1971], Solow [1973] and Polinsky and Shavell [1976].

equivalent to selecting a house that places the consumer on the bid price function corresponding to the highest achievable level of utility, illustrated in Figures 1a and 1b as the point of tangency of the hedonic price function h and the bid price function ψ . Figure 1a portrays the tangency condition for a characteristic consumers regard as desirable (e.g., lot or structure size, neighborhood quality) while 1b presents the corresponding tangency for an undesirable characteristic (e.g., level of pollution, age of the structure).


For the reader impatient with mathematical detail, we summarize where we have come so far. Houses are indivisible commodities which means that

Figure 1a.



Equilibrium Condition for a desirable characteristic

Figure 1b.



Equilibrium Condition for an undesirable characteristic

the conventional tools of consumer demand analysis are of limited relevance. Consumers in a competitive market maximize utility subject to a budget constraint, but the solution to this problem of constrained maximization cannot usefully be described in terms of demand functions. However, the solution can be described through the use of bid price analysis, where the hedonic function (representing the options offered to the consumer by the market) replaces the budget constraint and the family of bid price functions (one for each level of utility) replace the indifference curves of conventional theory. A consumer chooses the point on the hedonic price function which pleases him on his lowest bid price function (corresponding to the highest level of utility) just as, in the standard theory, a consumer chooses the point on the budget plane which puts him on his highest indifference contour.

At this juncture it is important to introduce a caveat. Figures 1a and 1b implicitly assume the existence of a continuous variety of housing types, but even if the set K of characteristics is finite all of the analysis presented so far goes through unscathed. The consumer's choice of a house no longer involves the tangency of two curves, but the consumer still ends up selecting the point on the hedonic function which places him on the lowest bid price function.

A primary virtue of the bid price analysis is the ease with which it accommodates the traditional hypotheses about housing markets that have been advanced in the urban economics literature. Much of urban economics consists of a variety of propositions asserting that the housing market sorts households of various types into different regions of the housing characteristics space. For example, high income households are assumed

to choose newer housing, larger lots and bigger houses, higher quality neighborhoods, better schools and less polluted locations than their low income counterparts. Since in equilibrium the hedonic price function will be an envelope of individual consumers' bid price functions, these hypotheses are equivalent to the assertion that the slope of a bid price function with respect to any of these characteristics is an increasing function of income.

There is a danger, however, in treating characteristics one at a time as is done in much of the literature. It is easy to imagine, for example, that a high income household may choose an old house if it is located in a high quality neighborhood. The usual assertions about the slope of bid price functions make sense only ceteris paribus: e.g., holding other housing characteristics fixed, it does seem reasonable to suppose that high income households are willing to pay more for a reduction in the age of a house than are low income households. Thus, it is important to view this bid price analysis in a multi-dimensional setting.

d) Empirical Implementation of the Model

I have made the case that bid price analysis provides a useful way to describe consumer behavior in a competitive housing market. However, I have not yet resolved the question of how this approach can be given empirical content. One approach that has been attempted on a number of occasions is to estimate directly the bid price functions for various types of consumer. A major difficulty with this procedure, which I do not think can be eliminated in a satisfactory manner, is that market data provide information only on actual and not on bid price. There is another way to proceed, however, that seems to give quite satisfactory results. Since I have discussed this technique and the empirical results elsewhere (Ellickson [1977b]), I will present here only a few of the main conclusions.

Recall that the bid price function for consumer h is given by:

$$V_n = \psi_n(p_x, z, y_n, u_n) = \tilde{\psi}_n(z) \quad (9)$$

whereby the second equality we have suppressed the price vector p_x , household income y_n and the utility level u_n since all are held constant at their equilibrium values. Suppose now that we have classified the set of household into T groups indexed by t (for concreteness, suppose households are classified by income, race and family size). If all members of group t have the same income and preferences, if the characteristics represented by z capture all of the aspects of a house relevant to consumers and if there are no information costs to search in the housing market, then all of the consumers of type t will bid the same price for a house with characteristics z . In any empirical application, of course, none of these conditions will be met. Therefore, we replace the bid price function $\tilde{\psi}_t(z)$ for consumers of type t by the stochastic bid price function

$$V_t = \tilde{\psi}_t(z) + \epsilon_t \quad (10)$$

where ϵ_t is a random disturbance term. The deterministic proposition that a house with characteristics z will be occupied with probability one by a particular household type is then replaced by the probabilistic statement that

$$p(t|z) = \text{prob}\{\tilde{\psi}_t(z) + \epsilon_t > \tilde{\psi}_{t'} + \epsilon_{t'}, t' \neq t\} \quad (11)$$

giving the conditional probability that a house with characteristics z will be occupied by a consumer of type t .

Readers familiar with McFadden's [1974] approach to qualitative choice will recognize this formulation. If, following McFadden, we assume that the disturbance terms are independently and identically distributed Weibull, then equation (11) takes the form:

$$p(t|z) = \frac{\exp[\psi_t(z)]}{\sum_{t \in T} \exp[\psi_t(z)]} \quad (12)$$

Assuming that the bid price functions are linear in the parameters, we obtain

$$p(t|z) = \frac{\exp(\alpha_t z)}{\sum_{t \in T} \exp(\alpha_t z)} \quad (13)$$

a conditional logit model identical in form with McFadden's except that bid price functions replace the utility functions for the representative consumer. The parameters of this model can be estimated through maximum likelihood in exactly the same way that McFadden estimates his model where the parameters are now interpreted as the coefficients of the nonstochastic part of the bid price function for each type of household.

In the paper cited above (Ellickson [1977b]), I estimated this model using data drawn from a sample survey of 28,000 households in the San Francisco Bay Area conducted by the Bay Area Transportation Study Commission (BATSC) in 1965. Lack of space precludes a detailed recapitulation of the results obtained. Suffice it to say that the model performs extremely well. Higher income households exhibit a (statistically significant) stronger preference relative to low income households for more accessible locations⁷, newer housing, larger lots, more rooms, a better neighborhood (measured by median census tract income in 1960) and higher housing quality (measured by the residual from a hedonic regression). Furthermore, whites exhibit a stronger preference than blacks for housing located in census tracts that are predominantly white and in attendance areas of elementary schools whose students are predominantly white.

⁷ Precisely the result one expects since the model separates out the effect of savings in commuting time obtained by residing near the CBD and the savings in housing costs obtained by living in less accessible locations.

To summarize what has been accomplished thus far, we have a model of residential choice sufficiently flexible to accommodate a variety of housing characteristics including those that pertain to the neighborhood in which a house is situated. We have demonstrated that the model is empirically testable, and we have found that it performs extremely well. The traditional hypotheses regarding the effect of accessibility, age of the house, number of rooms, lot size and housing quality are supported by the data. What is more significant for our present purpose, neighborhood characteristics (median tract income, percent black in the census tract and in the elementary schools) have a strong impact on consumer behavior.

Thus, we now have one necessary ingredient for a theory of competitive equilibrium in a housing market with neighborhoods: a model of residential choice by consumers who act as price-takers. But this alone does not justify the conclusion that such an equilibrium can exist. We have not demonstrated the existence of a price vector p_x and a hedonic price function h that will support a competitive equilibrium. To answer that question we must venture outside the demand side to explore the more fundamental question: how are the various types of houses supplied in the housing market and, in particular, is it possible to have a competitive supply of neighborhoods? It is to this question that we will now turn.

3. The Competitive Supply of Local Public Goods

In the analysis presented in Section 2 we have skirted around the issue of existence of competitive equilibrium. We noted the recent progress that has been made in establishing existence of competitive equilibrium for economics with indivisible (Broome [1972], Mas-Colell [1975b]) and differentiated (Mas-Colell [1975a]) commodities. But while this work holds the key to solving our problem, the connection is not immediate. All of these models are confined to the case of a pure exchange economy.

a) Local Public Goods and Indivisible Private Goods

My basic claim is that if we introduce production (subject to initial increasing returns to scale) into the models of competitive equilibrium with indivisible commodities developed by Broome [1972] and Mas-Colell [1975b] we obtain as well a model of competitive equilibrium with local public goods. To put the matter differently, the notion of a local public good is not a logically distinct concept in economics -- it is simply a special case of an indivisible private commodity. A formal justification for this claim is presented elsewhere (Ellickson [1977b]). My aim here is to give a heuristic argument in favor of this point of view which I hope will convince the reader that the approach makes sense.

To keep matters simple, ignore for now the interpretation of neighborhood characteristics as local public goods. Consider an economy in which consumers must choose among several alternative public goods, each provided by a different "firm": elementary schools can serve as a concrete illustration. The conjecture that a competitive equilibrium could exist for such an economy is due to Tiebout [1956], and over the years the Tiebout model has attracted considerable attention. However, despite its intuitive appeal, there is no satisfactory existence proof for Tiebout equilibrium.

The key to developing a formal theory of competitive equilibrium with local public goods is to recognize that the choice of a particular public good by a consumer involves an indivisibility: the consumer either chooses to consume the public good or he does not. This is a simple and intuitively obvious observation, but its implications are profound. What it means is that if we refer to the quality z of a public good, we must be careful to recognize that this quality is just a label for a type of public good. Consumers don't buy public goods by the pound, they choose among alternative public goods. To use the terminology we employed in the Section 2, the quality of a public good is simply a characteristic and, as in the model of residential choice, we want to avoid treating characteristics as though they were divisible commodities.

Once we have adopted this point of view, there is no longer any particular advantage to assuming that alternative public good types can be described by a scalar quality index z . There is no difficulty, for example, in allowing z to be some finite dimensional vector of characteristics as in Section 2. In fact, if the set K of alternative types of public good that could be produced is finite, it is not necessary to introduce characteristics into the analysis at all. If the set K is infinite the introduction of characteristics may be decisive in establishing existence of a competitive equilibrium, but this I have not proved^{8/}. However, if K is finite, then the use of characteristics is unnecessary in proving existence, and in the paper cited above (Ellickson [1977c]) I make no mention of them at all.

In this paper I have introduced characteristics for two reasons: (i) to clarify the connection between the traditional theory of public goods and that presented here; and (ii) to establish a bridge between the treatment

^{8/} The conjecture is based upon the work of Mas-Colell on differential commodities.

of public goods and the theory of residential location outlined in the preceding Section. However, in this paper I assume that K is finite, and in that case the introduction of characteristics is inessential. They should be regarded as simply a mnemonic device, a convenient way to label the alternative types of public good that could be provided by the economy.

If we consider schools, for example, the components of z could represent pupil-teacher ratios, racial composition, dummy variable indexes for teaching style (Montessori, military, etc.) and so on. The preferences of the n^{th} consumer can then, just as in Section 2, be described by a utility function of the form $U_n(x_n, z_n)$. Note that for consumers choosing the same public good, the characteristic vector z enters all utility functions just as in the traditional theory of public goods.

b) The Production of Local Public Goods

Up to this point I have said nothing about how public goods are produced. It is here that the alternative theory of public goods I have been describing begins to exhibit its decisive advantage over the traditional theory.

In the standard literature there are two notions of scale that one has to contend with, scale with respect to "output" z and scale with respect to the number of consumers provided the public good. For pure public goods only the former type of scale is relevant since the public good is available to everyone, but for local public goods it seems necessary to consider scale of the latter type as well. However, when viewed from our perspective, it makes little sense to talk about scale with respect to z (which we call the quality or the type, not the output, of the public good). Even if z is a scalar, what does it mean to double quality? Once we realize that z may be a collection of dummy variables, the whole notion of scale with respect to z becomes nonsensical. Fortunately, it turns out that scale with respect to quality is irrelevant to

the issue of whether a competitive equilibrium exists. All that matters is scale with respect to the number of consumers provided with the public good.

Consider now the production of a public good of type z , say a school described by a particular set of characteristics. We assume that no one owns the public good initially, but that it is produced using inputs of private commodities. To facilitate the use of a diagram, suppose that only one type of input (x) is used to produce the public good. Figure 2 illustrates a typical production set where, bowing to the usual conventions in general equilibrium theory, we represent inputs as negative numbers and outputs as positive. The horizontal lines represent the combinations of inputs and outputs that are feasible for the firm producing the public good of type z . Note in particular that output is measured in terms of the number n_z of consumers provided the public good. Contrary to the usual practice in public good theory, we do not refer to z as output but rather as the type (or occasionally the quality) of the public good.

Figure 2 illustrates an assumption that is basic to the proof of existence of a competitive equilibrium: after a certain scale is reached, production of the public good exhibits constant returns to scale (when restricted to integer outputs)^{9/}. We will refer to the minimum scale \tilde{n}_z after which one has constant returns as the optimal size of the jurisdiction; in Figure 2, $\tilde{n}_z=3$.




Figure 2: Production Set for Public Good of Type z .

^{9/} Diminishing returns to scale can also be incorporated into the model without much difficulty.

c) Competitive Equilibrium with Local Public Goods

We are now ready to state

Tiebout's conjecture:

if the optimal size of jurisdictions is small relative to the size of the economy for each type of local public good and the usual assumptions are imposed on consumer preferences, then a competitive equilibrium exists for the economy.

Unfortunately, Tiebout's conjecture is false. It is not difficult to exhibit economies satisfying the conditions of his theorem for which there exists no competitive equilibrium. But his result is nearly true in a sense that can be made quite precise.

The reason why Tiebout's conjecture fails is the presence of non-convexity, holes in the production sets of firms and the consumption sets^{10/} of consumers that preclude the use of the standard procedures for proving existence of equilibrium. In the present instance, these non-convexities arise as a consequence of the indivisibilities and the increasing returns that have been introduced into the model. In the last few years, techniques have been developed to handle this type of situation^{11/}. The basic idea is to construct an artificial version of the original economy by filling in the gaps using as little filler as possible. The artificial economy so constructed satisfies the convexity assumptions needed to establish existence of a competitive equilibrium. The final step is to demonstrate that the resulting equilibrium can

^{10/} For the reader unfamiliar with this bit of terminology, a consumption set is the set of all commodity bundles that a consumer is physically able to consume in the absence of a budget constraint; in the standard model, it is often assumed to be the entire non-negative orthant.

^{11/} The basic idea was first conceived by Farrell [1959] and Rothenberg [1960] and given a precise formulation by Starr [1969]. Arrow and Hahn [1971] extended Starr's results to economies with production. A number of papers have elaborated on the idea in recent years.

nearly be attained by the original economy provided that the non-convexities are "small" relative to the size of the economy.

While the mathematic tools needed to prove the result are rather sophisticated, the conclusion should be intuitively obvious to most economists. The key idea is the treatment of local public goods as indivisible private goods. Once that step is accepted, the rest follows directly. The approximation result is not much different from that needed to justify almost any application of competitive equilibrium analysis: most commodities come in discrete units, and u-shaped cost curves are accepted by most economists as the norm; approximation theorems of the sort I have been discussing are, therefore, basic to essentially all of economics.

Returning to the context of local public goods, we may loosely summarize the main result as follows:

Tiebout's Theorem: If the optimal size of firms producing public goods is small relative to the size of the economy and the usual assumptions are imposed on consumer preferences, then an approximate competitive equilibrium exists for the economy. If the firms are producing public goods in the region where economies to scale are exhausted, the equilibrium can be regarded as exact.

d) A Competitive Market for Neighborhoods

With this theory of competitive equilibrium with local public goods as background, we are now ready to discuss the notion of a competitive market for neighborhoods. In treating neighborhood attributes as local public goods, it is convenient to distinguish between two types of neighborhood characteristics:

Type A designates a neighborhood characteristic that is independent of the characteristics of the houses or the consumers that are located in the neighborhood;

Type B denotes neighborhood characteristics defined in terms of the people or the types of houses in the neighborhood.

I believe that most examples of neighborhood characteristics are of type B. It is useful to begin with type A, however, because this sort of characteristic fits more directly into our theoretical framework. An example would be neighborhood schools if consumers cared only about characteristics such as pupil/teacher ratios or availability of a football field and not about the racial or socioeconomic composition of the student body.

What would a competitive market for neighborhood characteristics of type A look like if it existed? The outcome would essentially be that described by McGuire [1974] in his refinement of the Tiebout model: consumers in a particular jurisdiction would all pay the same amount to receive the public good (equal tuition per child in the case of schools since it is children who "use" the schools); and the jurisdictions would be stratified, grouping together households who have the same tastes regarding public goods.

How likely is it that a competitive equilibrium of this sort could exist? In the case of schools, I suspect that scale economies are exhausted quite quickly, probably upon reaching the size of a typical elementary public or private school. Were school services provided privately it seems quite probable that the market for schools would be quite competitive. However, most elementary and secondary schooling in the United States is not provided privately but by relatively large public school districts, and the taxes used to pay for these schools presumably exhibit some degree of progressivity. Consumers without children who desire to contribute nothing to the public schools seldom if ever have that option. It is not difficult to see the reasons for these deviations from a competitive solution: existing political jurisdictions can prevent entry into the business of providing alternative schools (including no schools), and such barriers to entry can frustrate the operation of a competitive process. It

is not hard to construct examples where such behavior on the part of existing jurisdictions is perfectly rational^{12/}. I will return briefly to this issue in the concluding section.

Neighborhood characteristics of type B are more interesting from a theoretical point of view, and I think that there is ample reason to assume they constitute an empirically significant phenomenon. Returning to the case of schools, the heated resistance to bussing provides evidence of a sort, though of course we could accept the testimony of parents who say they simply don't like to have their children ride buses. Large lot zoning, segregation of land uses and the like provide evidence of a different sort. I will admit to a strong preference for two-acre zoning, a woodsy environment and an absence of motor homes. But I am not going to delve into empirical issues here. The main question I want to address is whether such neighborhood characteristics could be supplied through a competitive process.

In a sense the answer is quite straightforward: in a competitive equilibrium, firms produce entire neighborhoods. It seems plausible to assume that economies to scale are exhausted relatively quickly, perhaps at about the size of a typical housing tract designed by a developer. If consumers are sensitive to the average housing quality, minimum lot size or ethnic composition of neighborhoods, then the competitive model implies that developers will cater to these tastes. It is tempting to point to the behavior of developers (and, more dramatically, the designers of "new towns") as evidence of this competitive behavior. But what makes type B characteristics interesting and worth a separate discussion is that usually we do not observe neighborhoods being produced by single firms. More typically the characteristics of a neighborhood are determined by the actions of many different landlords and homeowners, each presumably acting in his own best interests.

^{12/} See Ellickson [1977a].

The absence of single firms producing neighborhoods does not necessarily imply that a competitive market for neighborhoods will fail to exist. The landlords and homeowners in a neighborhood may behave as though their actions were guided by a single firm. The widespread reliance on housing codes, zoning ordinances and restrictive covenants as devices to restrict the behavior of individual economic agents can be taken as evidence of the coordination of decision-making needed to sustain competitive allocations. It is even possible to argue that competition among neighborhoods will tend to encourage the development of such institutions: neighborhoods that fail to engage in this cooperative behavior will fail to "survive"; consumers will move out of neighborhoods in which non-cooperative behavior leads to under-maintenance. Following Alchian's [1950] lead, we can argue that this Darwinian struggle results in the competitive solution.

Nevertheless it seems clear that in raising this issue, the question whether the market for neighborhoods will induce economic agents to act cooperatively, we have reached the core of the problem of justifying the competitive theory of neighborhoods. I believe that it is this issue, rather than that of whether scale economies are exhausted, that is responsible for most of the divergence in opinion among economists on what would constitute an optimal housing policy. I suspect that most economists would be willing to grant that increasing returns to scale is not a significant barrier to the realization of a competitive equilibrium for neighborhoods, but they would differ on the question of whether the requisite coordination of actions would be forthcoming in the presence of fragmented ownership of housing parcels. At the level of generality of the theory presented in this section it is difficult to get a handle on this problem. Therefore, in the following two sections we turn to much more specific models intended to bring the issue of coordination into sharper focus.

4. Cooperation, Competition and the Supply of Neighborhoods

In the preceding Section I argued that neighborhood characteristics of type B raise the most interesting questions regarding the existence of a competitive market for neighborhoods. Neighborhood characteristics of type B are characteristics defined either in terms of the type of housing or in terms of the types of consumer located in the neighborhood. In this section we will be concerned with characteristics of the first kind, characteristics which depend on the type of housing in the neighborhood. In Section 5, we will treat characteristics of the second kind.

a) Competitive Equilibrium and Neighborhoods: An Example

The model I will present is the simplest model I have been able to construct capable of illustrating the phenomenon we are interested in, cooperation among distinct economic agents to produce a neighborhood. We will assume there is only one housing characteristic consumers care about, the average "quality" of housing in the neighborhood in which they reside. Let N index the set of consumers in the urban area and $J \subset N$ the subset who live in the j^{th} neighborhood. Let z_n denote the amount of housing purchased by the n^{th} household where z_n is produced subject to constant returns to scale using az_n units of the divisible private good x where a is a positive constant. The average housing quality in neighborhood J is then defined as $\bar{z}_J = \frac{1}{|J|} \sum_{n \in J} z_n$ where $|J|$ equals the number of consumers residing in neighborhood J .

All consumers have the same utility function,

$$U_n = x_n \bar{z}$$

where x_n is the amount of the private good consumed by the n^{th} consumer and \bar{z} is the average quality of housing in the neighborhood where consumer n lives. The n^{th} consumer is assumed to have an initial endowment \bar{x}_n of the divisible commodity.

We have assumed that the production of housing exhibits constant returns to scale. However, the production of neighborhood quality is subject to increasing returns at least over an initial range where the neighborhood is small: a single house certainly does not constitute a neighborhood. Recall that for a given type (quality) of neighborhood \bar{z} we measure output in terms of the number of consumers residing in the neighborhood. We will assume that there exists some minimum size \bar{n} of neighborhood after which scale economies are exhausted. After that point the neighborhood can be expanded at constant returns to scale.

If we follow the procedure described in Section 3 to construct an artificial ("convexified") version of this economy, and let the price of the divisible commodity x serve as numeraire with $p_x = 1$, then the marginal (= average) cost of supplying an addition unit of the neighborhood of type \bar{z} is equal to $a\bar{z}$; i.e., we are adding one consumer to the neighborhood, and in order to maintain neighborhood quality at \bar{z} the consumer added must be supplied a house of quality $z_n = \bar{z}$. Using the terminology of Section 2, this result implies that the hedonic price function faced by consumers is given by $h(\bar{z}) = a\bar{z}$. The n^{th} consumer then maximizes $U_n = x_n \bar{z}$ subject to the budget constraint $p_x x_n + h(\bar{z}) = p_x \bar{x}_n$, or making the substitutions $p_x = 1$ and $h(\bar{z}) = a\bar{z}$, $x_n + a\bar{z} = \bar{x}_n$. It is an easy exercise to show that the consumer will choose

$$(x_n^*, z_n^*) = (\bar{x}_n/2, \bar{x}_n/2a).$$

Thus, in the artificially constructed economy a consumer with initial endowment \bar{x}_n will choose to live in a neighborhood of quality $\bar{x}_n/2a$, paying a price $h(\bar{x}_n/2a) = (\bar{x}_n/2)$. In the original economy before convexification it may not be possible to achieve this allocation for all consumers. We must have enough consumers of each type (i.e., with the same initial endowment and hence choosing the same neighborhood type) to enable consumers of the same type to form at

least one neighborhood of size greater than or equal to \tilde{n} . If this condition is satisfied, we will then have a competitive equilibrium for this economy and, in particular, a competitive market for neighborhoods. Neighborhoods will be stratified with all consumers in the neighborhood having the same initial endowment^{13/}. All houses in a neighborhood will be identical in quality with neighborhood quality equal, by definition, to the quality of a typical house.

The example presented above is not intended to be realistic, but rather to provide a vehicle for studying the properties of a competitive market for neighborhoods. At the end of Section 3 I remarked that coordination of the actions of individual homeowners and landlords is perhaps the central issue in justifying the claim that a competitive market for neighborhoods can exist. It is this issue that we now wish to explore.

b) Non-Cooperative Behavior and the Quality of Neighborhoods

Considering a competitive equilibrium for the economy described above, we fix our attention on some particular neighborhood. Using the properties we have established for the competitive equilibrium, we know that this neighborhood will be populated by a group of consumers with identical endowments $\bar{x}_n = b$ and that each consumer has purchased a house of quality $b/2a$. Neighborhood quality is then also equal to $b/2a$.

Will consumers living in this neighborhood, each acting in his own best interests, agree to maintain this pattern of behavior? Our competitive analysis suggests that they will, but there is an alternative model of consumer behavior that implies they will not: the Prisoner's dilemma analysis of neighborhood blight first formulated by Davis and Winston [1961] and elaborated by Schall [1976].

^{13/} Recall that in this example utility functions are the same for all consumers.

To illustrate the Davis-Whinston approach to this problem, suppose that there are n_o consumers in the neighborhood we have selected. Assume that no consumer has the option to leave and no additional consumer can enter. Each consumer has an endowment of b units of the divisible good x , and the problem is to determine how much of this endowment each consumer will spend on housing. The average amount of housing produced in the neighborhood now takes the form of a pure public good in the sense of Samuelson [1954] (because we have assumed consumers cannot enter or leave the neighborhood). Therefore, we can determine the amount of housing that must be produced in the neighborhood to achieve a Pareto optimum by equating the sum of marginal rates of substitution to the marginal rate of transformation:

$$\sum_{n=1}^{n_o} MRS_n = MRT \quad (14)$$

Because of the assumed form of the utility functions,

$$MRS_n = \left(\frac{\partial U_n}{\partial \bar{z}} \right) \left(\frac{\partial U_n}{\partial x_n} \right) = x_n / \bar{z} \quad (n=1, \dots, n_o)$$

We have assumed that to produce z_n units of housing services, consumer n must use az_n units of x as input. So the neighborhood faces the production constraint

$$\sum_{n=1}^{n_o} x_n + a \sum_{n=1}^{n_o} z_n = n_o b$$

or letting $x = \sum_{n=1}^{n_o} x_n$ and $\bar{z} = \frac{1}{n_o} \sum_{n=1}^{n_o} z_n$,

$$x + an_o \bar{z} = n_o b$$

The production possibility frontier for this neighborhood is then given by:

$$F(x, \bar{z}) = x + a n_0 \bar{z} - n_0 b = 0 \quad (15)$$

and therefore $MRT = \left(\frac{\partial F}{\partial z} \right) / \left(\frac{\partial F}{\partial x} \right) = a n_0$; substituting into equation (14), we

obtain $\frac{1}{z} \sum_{n=1}^{n_0} x_n = a n_0$ or $x = a n_0 \bar{z}$. Substituting this into equation (15) yields:

$$\bar{z} = b/2a \quad (16)$$

Thus, a cooperative solution leads to precisely the neighborhood quality produced by the competitive process. Given the symmetry among consumers, it is natural to assume that each will contribute equally to the provision of the aggregate quantity of housing services, and that is precisely the competitive allocation.

The basic claim of Davis and Whinston is that consumers will not act in this cooperative manner. Consider some particular consumer living in this neighborhood. Given the housing consumption of the other consumers in the neighborhood, consumer n will maximize $x_n \bar{z}$ subject to the constraint $x_n + a z_n = b$

where as before $\bar{z} = \frac{1}{n_0} \sum_{n=1}^{n_0} z_n$. The constrained maximum is then determined by

the first-order condition $x_n = a n_0 \bar{z}$ and the budget constraint. Substituting the former into the latter yields the reaction functions for each of the n consumers (in implicit form):

$$a n_0 \bar{z} + a z_n = b \quad (n = 1, \dots, n_0) \quad (17)$$

Because of the symmetry among consumers, it is clear that the solution to this system of equations will have all consumers choosing the same level of z_n , so we have $a z_n (1 + n_0) = b$ or

$$z_n = b/a(1 + n_0) \quad (18)$$

as the non-cooperative solution. The quality of the neighborhood is then also equal to $b/a(1 + n_0)$.

Comparing the cooperative (competitive) solutions given by (16) and the non-cooperative solution given by (18), we see that if $n_0 > 1$ (i.e., more than one consumer resides in the neighborhood) then non-cooperative behavior leads to neighborhood quality that is less than optimal, precisely the "blight" phenomenon of the Davis-Whinston model.

How is blight avoided in the competitive theory of neighborhoods? In his perceptive critique of the Davis-Whinston analysis, Rothenberg [1967] observes that their model reaches too far: it implies that all neighborhoods and not just slum neighborhoods will be subject to blight, an implication which Rothenberg finds unacceptable on intuitive grounds. Thus, he concludes that homeowners in high quality neighborhoods must be finding some way to cooperate, either through explicit mechanisms such as zoning ordinances and housing codes or through less visible forms of social pressure. Our competitive model suggests that another mechanism may be at work, ignored by Davis-Whinston and Rothenberg alike. If a neighborhood is subject to blight, consumers may simply move to a better neighborhood, one in which the optimal quality is maintained either through the actions of a single firm or because zoning, housing codes and social pressure are effective in that neighborhood. Because the non-cooperative solution is not Pareto optimal, such moves will make consumers better off.

I do not think that on a priori grounds we can conclude that the competitive model is correct and non-competitive models are wrong. It is true that the failure to allow consumers to move is a severe weakness in the Davis-Whinston analysis. But Hurwicz [1974] has demonstrated that essentially all

competitive models are vulnerable to strategic maneuvers of the sort that Davis and Whinston emphasize. Given the current state of the art, economists' belief in the relevance of the price-taking hypothesis to actual market behavior has to be regarded as primarily a matter of faith, and this paper is no exception.

I conclude this section with a brief comment on matters of terminology. It is common in the literature to see arguments of the Davis-Whinston sort referred to as "supply models" and competitive models as "demand models". This represents an abuse of language. All competitive equilibrium models have a supply side. What sets competitive analysis apart is its reliance on the price-taking hypothesis. The proper distinction is between models that are competitive and those that are not. In this light, Schall's [1976] designation of non-cooperative solutions of the sort we have been discussing as "competitive" seems particularly inappropriate.

5. The Competitive Supply of Neighbors

In the preceding section we presented an example of one form that neighborhood characteristics can take, attributes described in terms of the housing in the neighborhood. In this section we will provide an example of the other form, characteristics defined in terms of the people who live in the neighborhood. Income, occupation and ethnicity could each serve as an example. I will focus on race.

We will assume that there exists two types of consumer, blacks and whites. Suppose that all consumers have utility functions of the form:

$$U_n = \begin{cases} x_n z_n (1 + w), & 0 \leq w \leq t_i \\ x_n z_n (1 + t_i), & t_i \leq w \leq 1 \end{cases}$$

where t_i is a "tolerance" parameter equal to t_B for blacks and t_w for whites. We will assume $0 \leq t_B \leq t_w \leq 1$. As in Section 4, z_n represents housing consumption and x_n the consumption of a non-housing divisible commodity by consumer n . We will assume again that consumer n can be provided z_n using az_n units of the non-housing commodity as input. All consumers have an identical initial endowment $\bar{x}_n = b$.

While housing is produced under conditions of constant returns, we assume that the production of neighborhoods exhibits increasing returns when neighborhoods are small (e.g., a single consumer does not constitute a neighborhood). But we also assume that after neighborhoods reach some critical minimum size, they can be expanded subject to constant returns to scale while maintaining the same racial composition (neighborhood type).

With these assumptions we can again apply the technique described in Section 3 to find a competitive equilibrium. If we construct the artificial

(convexified) economy corresponding to the one described above, and let the non-housing commodity serve as numeraire with $p_x = 1$, then we obtain the hedonic price function $h(z,w) = az$. Since both blacks and whites prefer living in a white neighborhood, it appears that all consumers will reside in all-white neighborhoods. Clearly this is impossible for blacks, and it is here that we discover why neighborhood characteristics that depend on the attributes of ones neighbors are of independent analytic interest.

At this stage it is convenient to distinguish between two cases: Case I in which firms are unable to charge different prices to blacks and whites living in the same neighborhood (presumably because it is against the law), and Case II in which such "price discrimination" is possible.

We begin with Case I. If $t_w = 1$ (so that whites have no tolerance for living with blacks) and there are enough whites and blacks to form segregated neighborhoods that exhaust the economics of neighborhood formation, then competition results in complete segregation. It is easy to see why this must be so. To turn a non-negative profit, a firm producing an integrated neighborhood must charge each consumer a price $p_z \geq a$ for each unit of housing. But in that case another firm offering an all-white neighborhood and charging a price $p_z = a$ will attract all of the white consumers, and the integrated neighborhood will "tip" to all black.

Note that in this competitive equilibrium, the price of housing is the same ($= a$) in the black and in the white neighborhoods. Thus, contrary to the conclusion of the popular Bailey [1959] model, whites do not have to pay for their prejudice. And this is true despite the fact that in this equilibrium blacks would be willing to pay more than whites are paying in order to live in the white neighborhoods.

This indictment of the Bailey model is not confined to the special case we have been considering. Even if we relax the assumption that blacks and whites have identical endowments, it is easy to construct examples in which we get complete segregation even though blacks would be willing to pay more than whites actually pay to live in white neighborhoods and even though the price of housing in black and white neighborhoods is identical. The fatal flaw in the Bailey model is its neglect of the supply side of the market^{14/}.

An objection could be raised to the preceding analysis that we have relied heavily on the notion that single firms produce neighborhoods. The response is identical to that of Section 4: a "survival of the fittest" argument implies that consumers in white neighborhoods will act as though their actions were guided by a single firm. The restriction against black entry can be enforced through restrictive covenants, the actions of real estate brokers or simply by making life miserable for blacks who have the temerity to buy into the neighborhood. Thus, we reach the conclusion, implicit in the "cooperative behavior" of Section 4, that competition implies collusion.

The notion that competition and discrimination can go hand in hand should come as a rather unpleasant surprise. We tend to think of competition as good and discrimination as bad, so their conjunction is dissonant. To gain some perspective, it is worthwhile considering a related example involving discrimination that seems less objectionable than that based on race. Schelling [1969] tells a story about an ice cream parlor that tipped when it became a teen-age hangout, no longer frequented by its former clientele. Since the teen-agers evidently were unwilling to spend much money, the place eventually went out of

^{14/} Arrow [1971] has made the same point in much the same way in the context of segregated labor markets.

business. Obvious the firm was not a profit maximizer. If it were, the owner would have hung up a sign saying "No Loitering", perhaps selectively enforcing this discriminatory rule against the younger customers while allowing regular customers to remain. He then would be acting competitively, as do the firms producing racially segregated neighborhoods in our model, and he would be able to stay in business.

It is worth noting at this point the implications of the model developed in this section for the theory of residential choice presented in Section 2. I remarked in Section 2 that hedonic price functions do not necessarily convey much information because these functions result from the interplay of supply and demand. This remark is graphically illustrated by the results of this section: it is possible to get complete stratification with respect to racial composition even though whites pay no premium for living in all-white neighborhoods^{15/}.

If $t_w < 1$, then the implications for blacks are not quite so bleak. If, as seems reasonable, we assume that the number of blacks exceeds that which could be accommodated by integrated neighborhoods with percentage black less than $1 - t_w$ in each neighborhood, then it is easy to show that competitive equilibrium will involve the existence of two sorts of neighborhoods: one set that is integrated with precisely the fraction $1 - t_w$ black and a residual category of completely black neighborhoods. When blacks enter a metropolitan area they will then be channeled to black neighborhoods with entry to white neighborhoods governed by the tolerance of whites for living among blacks.

^{15/} Much the same objection can be raised to Oates' [1969] test of the Tiebout model. See Hamilton [1976].

Up to this point we have assumed that firms are unable to charge different prices to blacks and whites living in the same neighborhood. It is worth noting that such price discrimination (Case II) is quite consistent with a competitive model. Firms producing integrated neighborhoods can be regarded as producing a joint product, a neighborhood with a fraction w of white slots and $1 - w$ black slots. It is easy to show that an integrated neighborhood may be viable provided that (i) whites are offered housing at a discount relative to housing in all-white neighborhoods and (ii) blacks are charged a premium for the "privilege" of living in an integrated neighborhood. What accounts for this rather bizarre phenomenon is the presence of a hidden input: only whites are able to contribute to the whiteness of a neighborhood^{16/}.

In concluding this section I feel that it is necessary to issue a warning that would be unnecessary were the topic less ladden with emotional content: the demonstration that an allocation is Pareto optimal does not mean that it is socially acceptable. In the present instance I have indicated that a competitive allocation may involve complete segregation. A competitive allocation is Pareto optimal so that it is not possible to make some consumers better off without making others worse off. But that does not mean that we have to accept such an allocation as in any sense just, and there is nothing inconsistent about espousing open housing legislation to frustrate this competitive process.

^{16/} It is somewhat questionable whether one should call this behavior "price discrimination" since firms are simply producing joint products.

6. Conclusion

In this paper I have sketched the outlines of a competitive theory of neighborhoods. In some respects the theory seems to be relatively complete, but it is clear that important issues remain to be resolved. For example, some means needs to be developed to give the supply side of these models empirical content comparable to the methods available for the demand side (as presented in Section 2). Without a supply side we have no way to account for the options that actually are made available to consumers, options which the demand analysis takes as given.

If we are to apply the theory of local public goods to neighborhoods with any degree of confidence, we also need a better way to define neighborhoods. In the examples presented in Sections 4 and 5, I was necessarily rather hazy in justifying the assumptions about returns to scale in the production of neighborhoods required for the theory. At this level of abstraction, it does not seem possible to be more specific. The reason, I believe, for this inherent ambiguity is that a proper notion of neighborhood (more precisely, those aspects of neighborhood involving what I have called characteristics of type B) must be defined spatially.

The most promising approach to a better definition of neighborhood seems to be that of Schelling [1969] who essentially defines neighborhoods in terms of the characteristics of adjacent neighbors; more generally we would include characteristics of adjacent houses as well. If we consider his simplest model of consumers strung out along a line, then from the point of view of any consumer the relevant "neighborhood" is defined in terms of the consumers within, e.g., two places to the left or right of his position. Approaching the problem in this way builds in a natural notion of initial increasing returns (neighborhoods always include five people) and a justification for eventual constant returns (a particular sequence of blacks and whites can be

repeated indefinitely). Racial composition of these neighborhoods in the Schelling model functions much as it does in the model presented in this paper: it can be viewed as a form of local public good. But there the similarity ends. Schelling has no prices or competitive markets in his model, and it seems very difficult to capture his notion of neighborhood within the confines of more standard economic analysis. Nevertheless, I suspect that it is possible, and that the resulting theory would be far richer than the one I have presented.

To the theorist the construction of competitive models is an end in itself. For the policy-maker, on the other hand, abstract analysis may seem pretty useless. I think that conclusion is unwarranted. Perhaps the most important lesson to be drawn from work such as this is an appreciation of just how subtle the competitive process can be. Based on experience with a few models, it is often assumed that externalities wreck havoc with the invisible hand, that collusion and competition are antithetical and that a neighborhood going down hill is evidence of suboptimality. The models we have presented imply that all of these phenomena can be consistent with an efficient competitive process.

Of course, the fact that competitive equilibrium is consistent with such phenomena does not prove that the world is competitive. I am certain that the same facts could be explained by alternative non-competitive models. The discipline of constructing a competitive model can contribute to the evaluation of non-competitive models: we have seen, for example, that the Bailey model of segregation and the Davis-Whinston model of blight do not stand up under such scrutiny. But the development of more logically consistent non-competitive models may be the most useful by-product of competitive analysis.

Economists have not made much progress in going beyond competitive analysis. Recent work by Lancaster [1975], Spence [1976] and Dixit-Stiglitz [1977] represents one approach that can be applied to the market for neighborhoods. One

advantage of my interpretation of local public goods as indivisible private goods is that such efforts to develop a theory of product differentiation extend immediately to this new context.

I will admit to a vague sense that the approach taken in the recent literature on product differentiation is not very satisfactory. More germane to the subject of this paper, I do not believe that it points to the basic source of suboptimality in the market for neighborhoods. It seems evident that a much more serious problem arises because of the barriers to entry erected by existing political jurisdictions. In our competitive analysis we tacitly assumed that any group of consumers wishing to form a neighborhood could do so, subject to the constraints imposed by the market. But high income consumers cannot form a neighborhood in the central city without buying into the problems of the central city as well, subjecting themselves to taxes intended to help the poor, sharing schools of poor quality and incurring the risk of bussing. Low income consumers are unable to form lower quality neighborhoods in affluent suburbs because the residents of the suburb want to avoid the erosion of their tax base and the possibility of contributing more to income redistribution^{17/}. I find it hard to believe that this is a portrait of a fully competitive process.

^{17/} For a more explicit treatment of some of these issues, see Ellickson [1977a]

References

- Alchian, Armen A., 1950, "Uncertainty, Evolution and Economic Theory," Journal of Political Economy 58, 211-221.
- Alonso, William, 1964, Location and Land Use, Harvard University Press, Cambridge, Mass.
- Arrow, Kenneth J., 1971, Some Models of Racial Discrimination in the Labor Market, The Rand Corporation, RM-6253-RC.
- Arrow, Kenneth J., and Frank H. Hahn, 1971, General Competition Analysis, Holden-Day, San Francisco.
- Bailey, Martin J., 1959, "Note on the Economics of Residential Zoning and Urban Renewal," Land Economics 35, 288-92.
- Broome, John, 1972, "Approximate Equilibrium in Economics with Indivisible Commodities," Journal of Economic Theory 5, 224-249.
- Davis, Otto A., and Andrew B. Winston, 1961, "The Economics of Urban Renewal," Law and Contemporary Problems 26, 105-17
- Dixit, Avinash K. and Joseph E. Stiglitz, 1977, "Monopolistic Competition and Optimal Product Diversity," The American Economic Review 67, 297-308.
- Edgeworth, Francis Y., 1881, Mathematical Psychics, P. Kegan, London.
- Ellickson, Bryan, 1971, "Jurisdictional Fragmentation and Residential Choice," The American Economic Review 61, 334-339.
- Ellickson, Bryan, 1977a, "The Politics and Economics of Decentralization," Journal of Urban Economics 4, 135-149.
- Ellickson, Bryan (with Barry Fishman and Peter A. Morrison), 1977b, Economic Analysis of Urban Housing Markets: A New Approach, The Rand Corporation, R-2024-NSF.
- Ellickson, Bryan, 1977c, "Competitive Equilibrium with Local Public Goods," Discussion Paper Number 96, Department of Economics, University of California, Los Angeles.
- Farrell, M. J., 1959, "The Convexity Assumption in the Theory of Competitive Markets," Journal of Political Economy 67, 377-391.
- Hamilton, Bruce W., 1976, "The Effects of Property Taxes and Local Public Spending on Property Values: A Theoretical Comment," Journal of Political Economy 84, 647-650.
- Hart, Oliver D., 1977, "Monopolistic Competition in a Large Economy with Differentiated Commodities," working paper, Churchill College, Cambridge.

- Hurwicz, Leonid, 1974, "The Design of Mechanisms for Resource Allocation," in M.D. Intriligator and D.A. Kendrick (eds.), Frontiers of Quantitative Economics, Vol. II, American Elsevier, New York.
- Lancaster, Kelvin, 1975, "Socially Optimal Product Differentiation," The American Economic Review 65, 567-585.
- Mas-Colell, Andreu, 1975a, "A Model of Equilibrium with Differentiated Commodities," Journal of Mathematical Economics 2, 263-295.
- Mas-Colell, Andreu, 1975b, "Indivisible Commodities and General Equilibrium Theory," Working Paper IP-228, University of California, Berkeley.
- McFadden, Daniel, 1974, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka (ed.), Frontiers in Econometrics, Academic Press, New York, 105-143
- Mills, Edwin S., and James MacKinnon, 1973, "Notes on the New Urban Economics," Bell Journal of Economics and Management Science 4, 593-601.
- Muth, Richard F., 1969, Cities and Housing, University of Chicago Press, Chicago.
- Oates, Wallace E., 1969, "The Effects of Property Taxes and Local Public Spending on Property Values: An Empirical Study of Tax Capitalization and the Tiebout Hypothesis," Journal of Political Economy 77, 957-970.
- Polinsky, A Mitchell, and Steven Shavell, 1976, "Amenities and Property Values in a Model of an Urban Area," Journal of Public Economics 5, 119-129.
- Rosen, Sherwin, 1974, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," Journal of Political Economy 82, 34-55.
- Rothenberg, Jerome, 1960, "Non-convexity, Aggregation, and Pareto Optimality," Journal of Political Economy 68, 435-478.
- Rothenberg, Jerome, 1967, Economic Evaluation of Urban Renewal, The Brookings Institution, Washington, D.C.
- Samuelson, Paul A., 1954, "The Pure Theory of Public Expenditures," Review of Economics and Statistics 36, 387-89.
- Schall, Lawrence D., 1976, "Urban Renewal Policy and Economic Efficiency," The American Economic Review 66, 612-628.
- Schelling, Thomas C., 1969, Models of Segregation, The Rand Corporation, RM-6014-RC.
- Solow, Robert M., 1973, "On Equilibrium Models of Urban Location," in M. Parkin (ed.), Essays in Modern Economics, Longmans, Green & Co., London.
- Spense, A. Michael, 1976, "Product Selection, Fixed Costs, and Monopolistic Competition," Review of Economic Studies 43, 217-235.
- Starr, Ross, 1969, "Quasi-Equilibria in Markets with Non-Convex Preferences," Econometrica 37, 25-38.
- Tiebout, Charles, 1956, "A Pure Theory of Local Expenditures," Journal of Political Economy 64, 416-424.