

Neighborhood Effects

Steven N. Durlauf¹

Department of Economics, University of Wisconsin

August 18, 2003

¹ Prepared for the Handbook of Regional and Urban Economics, volume 4, Economics, J. V. Henderson and J.-F. Thisse, eds. Please send all correspondence to Department of Economics, 1180 Observatory Drive, Madison WI, 53706-1393, sdurlauf@ssc.wisc.edu. The John D. and Catherine T. MacArthur Foundation, Institute for Research on Poverty, and University of Wisconsin Graduate School have provided financial support. I thank Ritesh Banerjee, Ethan Cohen-Cole, Artur Minkin, Giacomo Rondina, and Noam Yuchtman for excellent research assistance.

Abstract

This paper surveys the modern economics literature on the role of neighborhoods in influencing socioeconomic outcomes. Neighborhood effects have been analyzed in a range of theoretical and applied contexts and have proven to be of interest in understanding questions ranging from the asymptotic properties of various evolutionary games to explaining the persistence of poverty in inner cities. As such, the survey covers a range of theoretical, econometric and empirical topics. One conclusion from the survey is that there is a need to better integrate findings from theory and econometrics into empirical studies; until this is done, empirical studies of the nature and magnitude of neighborhood effects are unlikely to persuade those skeptical about their importance.

JEL Codes: C25, I30, R00

Keywords: identification, neighborhood effects, stratification, segregation, multiple equilibria, networks

“...in a neighborhood with a paucity of regularly employed families and with the overwhelming majority of families having spells of long term joblessness, people experience a social isolation that excludes them from the job network system that permeates other neighborhoods and that is so important in learning about or being recommended for jobs...And as the prospects of employment diminish, other alternatives such as welfare and the underground economy are not only increasingly relied on, they come to be seen as a way of life...Thus in such neighborhoods the chances are overwhelming that children will seldom interact on a sustained basis with people who are employed or with families that have a sustained breadwinner. The net effect is that joblessness, as a way of life, takes on a different social meaning: the relationship between schooling and post-school employment takes on a different meaning. The development of cognitive, linguistic and other education and job related skills necessary for the world of work in the mainstream economy is thereby relatively adversely affected. In such neighborhoods, therefore, teachers become frustrated and do not teach and children do not learn. A vicious cycle is perpetuated through the family, through the community and through the schools.”

William Julius Wilson, *The Truly Disadvantaged* (1987, pg. 57)

This chapter provides an overview of research on neighborhood effects in economics and other social sciences. In the last 15 years, there has been a renaissance of interest among economists in the social determinants of individual behavior and aggregate outcomes. Within the areas of urban, labor and family economics, much of this focus has centered on neighborhood effects. One reason for this is that a body of studies has argued that neighborhood influences, especially those defined by geographic residence, are important in understanding the persistence of inner city poverty; William Julius Wilson's *The Truly Disadvantaged* is probably the best known exposition of this argument. Another reason, stressed in Manski (2000), is that methodological advances in economic theory have provided ways to incorporate a range of spillover effects into rigorous microeconomic models. One source of these developments that Manski emphasizes is endogenous growth theory; many neighborhoods models incorporate spillover effects that, at least in their formal structure, are very similar to those found in growth contexts. In addition, the analysis of neighborhood effects in game theory (often called local interactions in this literature) has become an active area of research. This work does not require that neighborhoods be defined by geographic proximity; but does

rely on some notion of proximity versus distance in “social space,” a notion that is given content in Akerlof (1997). While the primary focus of this survey is residential neighborhood effects, a number of other forms of neighborhoods and associated effects will also be discussed.

The great bulk of the new neighborhood effects literature has focused on the consequences of neighborhood effects for individual behavior and for neighborhood composition; typically, the forms of these effects are taken as given. This is true even for the study of neighborhood effects in residential communities and schools, which constitute by far the main application of neighborhood effects to the study of substantive empirical phenomena. On the other hand, the microeconomic justifications given for residential neighborhood effects cover many of the motivations for the modern neighborhood effects literature.

What sorts of residential neighborhood effects are posited in the current literature?² One reason why neighborhoods affect their members, and the one whose microeconomic foundations are best justified (since they are merely a consequence of the rules for educational finance that exist in the US) is the local public finance of education. Hussar and Sonnenberg (2001) and Murray, Evans, and Schwab (1998) document the large differences in per pupil expenditures that persist across districts in the United States; this persistence is remarkable since state and federal programs at least implicitly designed to equalize expenditure are responsible for about 50% of public educational expenditure for kindergarten through grade 12 in the US.³ To be clear, there is considerable controversy concerning the relationship between school expenditures and inequality, well summarized in the survey by Hanushek (1986) and by the papers in Burtless (1996). My own reading of this literature suggests that while a general

²Neighborhood effects have generally not included agglomeration externalities, although these presumably play a role in cases such as the provision of education. See Fujita and Thisse (2002) for an outstanding survey of agglomeration models.

³Murray, Evans and Schwab (1998, pg. 799), find that the ratio of per pupil expenditures for the 95th percentile to the 5th percentile across US schools was 2.72 in 1972, 2.22 in 1982 and 2.40 in 1992. These authors also find that court mandated reduction of educational disparities within states have been efficacious. Hussar and Sonnenberg (2001) find some overall reduction in expenditure disparities at the district

relationship between school expenditures per capita and educational quality has proven hard to establish, schools in poor communities probably do suffer because of lack of resources. Hence very poor neighborhoods probably affect children along this dimension.

Outside of school finance, neighborhood effects have explicitly been justified along sociological and/or psychological lines. As such, they constitute forms of social interactions, which have themselves become a growing area of research; see Blume and Durlauf (2001), Brock and Durlauf (2001b) and Manski (2000) for surveys/overviews. One example of social interactions that is germane to neighborhoods is role model effects, in which the behavior of one individual in a neighborhood is influenced by the characteristics of and earlier behaviors of older members of his social group. Another form of social interactions is peer group influences; these differ from role model effects because they refer to contemporaneous behavioral influences and so may be reciprocal. Role model and peer group influences are both usually understood to produce some sort of imitative behavior, either contemporaneous or across age cohorts. This imitative behavior may be due to 1) psychological factors, an intrinsic desire to behave like certain others, 2) interdependences in the constraints that individuals face, so that the costs of a given behavior depend on whether others do the same, or 3) interdependences in information transmission, so that the behavior of others alters the information on the effects of such behaviors available to a given individual.⁴ Each of these types of imitative behavior implies that an individual, when assessing alternative behavioral choices, will find a given behavior relatively more desirable if others have previously behaved or are currently behaving in the same way. Hence the relative desirability of

level between 1980 and 1994, but caution that the decreases are not uniform across states and that large differences persist.

⁴Streufert (2000) shows how these interdependences can affect human capital formation by modeling how observations of adults in a community can affect inferences about the economic benefits to education. In his analysis, self-selection into neighborhoods means that inferences by the young about this trade off are biased. Streufert shows how under plausible assumptions, children in disadvantaged neighborhoods can be led to underestimate the returns to education, although this is not a logical implication of neighborhood self-selection. Related results are obtained by Roemer and Wets (1995).

staying in school is higher when adults in a community are college graduates or when one's peers are also staying in school.

One important reason why neighborhood effects have received much recent attention is that they provide a way of understanding why poverty traps might exist. To see the logic behind this claim, suppose that educational investment decisions exhibit strong role model influences, so that the decision to attend college for each high school graduate in a community is strongly (and positively) related to the percentage of college graduates among adults in a community. Such interdependence in behaviors creates the possibility that if one has two communities, one where the adults are all college graduates, and a second where none are, that these communities will converge to different levels of college attendance in a steady state. High and low college attendance rates are each reinforced across time as high (low) attendance rates among the current pool of adults lead to high (low) attendance rates among high school graduates, who in the future will influence the high school graduates to collectively exhibit high (low) rates as well. One way to think about a poverty trap is that a community, if initially comprised of poor members, will remain poor across long time periods, even generations. Intertemporal social interactions (i.e. social interactions in which choices made at one time affect others in the future) provide precisely this sort of dependence. An important antecedent to this idea is due to Loury (1977) in his construction of a theory of persistent racial inequality; ethnicity may be regarded as sort of neighborhood in social space.

A related notion of a poverty trap may be identified when one thinks about peer group effects. When the behavior of one member of a group is sufficiently positively dependent on the behaviors of others, this creates a degree of freedom in behavior of the group as a whole. Contemporaneous dependences in behavior mean that the members of a group will behave similarly. At the same time, these effects, when sufficiently strong, mean that the characteristics of the individuals involved will not uniquely determine what the group actually does. Dependence on history, reactions to common influences, etc. may determine which sort of average behavior actually transpires. The key idea, however, is that strong contemporaneous dependences in behavior lead to multiple possibilities for self-reinforcing behavior in groups. Within a given behavioral configuration, each individual is acting "rationally" in the usual sense. That does not

mean that each self-consistent configuration is equally desirable from the perspective of the members of the group. Another definition of a poverty trap is a socially undesirable (in the sense of producing poverty across a community) collection of behaviors in which the behaviors are mutually reinforcing and so individually rational.

Neighborhood effects are also important as they may reinforce the effects of changes in private incentives. Suppose one is considering whether to provide college scholarships to randomly chosen students across a set of high schools versus concentrating the scholarships among students within a given school. If the objective of the program is to alter high school graduation rates, then the presence of social interactions can, other things equal, mean that the concentration of the scholarships will be more efficacious. Assuming the direct incentive effect of the scholarships is the same for students across schools, the advantage of concentrating the scholarships in one school is that they will induce neighborhood effects that affect all students in the school, including those who have not been offered scholarships. More generally, neighborhood effects can amplify the effects of altering private incentives; this amplification is sometimes known as a “social multiplier” following Cooper and John (1988) and Manski (1993). The presence of social multipliers has important implications for the design of policies that have yet to be explored. In addition, perspectives on inequality that are driven by the effects of neighborhoods on individuals have important implications for how one conceives of notions such as equality of opportunity, and so has direct bearing on the assessment of alternative redistributive policies, see Durlauf (1996c,1999,2001) for discussion.⁵

The basic structure and implications of residential neighborhoods models suggest that the notion of neighborhoods may have general application in social science contexts. In fact, there are now a number of disparate research directions each of which focuses on

⁵In this work, I argue that neighborhood effects and related social interactions constitute a “memberships theory of inequality and poverty” in the sense that group memberships play a primary role in determining both cross-section inequality and social mobility over time. The role of groups in generating inequality is argued to matter both in assessing the ethical justifications for redistributive policies as well as having implications for the form of the policies. In the current context, neighborhood effects have obvious implications for the efficacy of government policies to redistribute neighborhood memberships as an antipoverty policy, for example.

populations of agents who are organized into groupings in which interactions occur; at least abstractly, these different analyses employ the idea of neighborhood effects. Examples of phenomena that have been studied using such ideas include 1) economic growth and industrial development in which spillovers occur between technologically similar industries (Durlauf (1993)), 2) economic development and market growth in which direct trading links between individual agents determine the extent of markets (Kelly (1997) and 3) business cycles in which local demand spillovers between industries arrayed on a lattice produce persistent aggregate fluctuations (Bak, Chen, Scheinkman, and Woodford (1993)). Hence, one part of this review will focus on contexts other than residence where neighborhood effects have proven to be a valuable research direction.

The neighborhood effects literature is quite disparate and no survey could cover the full range of theoretical models and empirical studies contained in the existing literature. Instead, the goal of this survey is to identify some of the major research questions that unify the many analyses as well as to highlight some of the weaknesses in the literature. To do this, the survey will work sequentially from theory to econometrics to empirical work. Section 2 provides an overview of theoretical work on neighborhood effects. The discussion focuses on both the implications of neighborhood effects on aggregate behavior and on the role of these effects in determining neighborhood composition. Section 3 provides some analysis of econometric issues that arise in the study of neighborhood effects. Section 4 reviews the existing empirical literature that directly addresses neighborhood effects. Section 5 reviews a body of studies that, while not directly focusing on neighborhood effects, nevertheless speaks to the importance of such effects. Section 6 contains summary and conclusions.

2. Theory

Theoretical models incorporating neighborhood effects involve two types of basic questions. First, how do the characteristics of a neighborhood affect the decisionmaking of its members, and in the aggregate, the behaviors in the neighborhood as a whole? Second, how does the presence of neighborhood effects influence memberships in

neighborhoods and in the aggregate, determine the configuration of the population across neighborhoods? While these questions of course interact throughout the theoretical literature, they are best understood as distinct.

i. Choice within neighborhoods

I first consider the abstract problem of how neighborhood effects influence individual choices and thereby produce interesting neighborhood behaviors in the aggregate. Consider n individuals who are members of a common neighborhood denoted as N . Each individual i , makes a choice c_i (a choice that is taken from the elements of some set of possible behaviors C). This individual-level decision will produce a probabilistic description of the choice given certain features of the individual and his neighborhood. The goal of the analysis is to construct a probability measure π^i for the vector of choices of all members of the group, c , that is consistent with these individual-level probability measures and understand how neighborhood effects determine its properties. Since individual agents are influenced by others, it is useful to define c_{-i} as the vector of all choices other than that of agent i .

From the perspective of individual decisionmaking, one may distinguish between four sorts of influences on observed behaviors. These influences have different implications for how one models the choice problem. These components are

α_i , a vector of deterministic (to the modeler) individual-specific characteristics associated with individual i ,

ϵ_i , a vector of random individual-specific characteristics associated with i ,

β , a vector of predetermined neighborhood-specific characteristics, and

π_i^j , the subjective beliefs individual i possesses about behaviors of others in his neighborhood, described as a probability measure over those behaviors.

Each of these components will be treated as a distinct argument in the payoff function that determines individual choices. The terms β_i and π_i^j constitute neighborhood effects. Following Manski (1993), who draws from the sociology literature in this regard (cf. Blalock (1984)), I will refer to β_i as a contextual effect and π_i^j as an endogenous effect, in order to distinguish those neighborhood effects that are determined by characteristics of neighborhood members as opposed to those effects that are determined by the contemporaneous behaviors of neighborhood members. The assumption that individuals are affected by their beliefs about the behaviors of others rather than by the actual behaviors is made for analytical convenience.

Even though the neighborhood effects terms may be “nonstandard” in the context of standard decision problems, individual choices are still defined via the maximization of some individual payoff function u_i ; given the notation we have introduced, individual choices are thus assumed to follow

$$c_i^j = \arg \max_{c_i^j} u_i(c_i^j, \beta_i, \pi_i^j) \tag{1}$$

In order to close this model, it is necessary to describe how beliefs about the behaviors of others are determined. The benchmark assumption in this literature is that beliefs are rational in the sense that

$$\pi_i^j = \Pr(c_j^k = c_j^k | c_i^j) \tag{2}$$

Notice that all uncertainty about the behaviors of others depends upon c_i^j by any agent other than i . The existence of an equilibrium set of behaviors within a neighborhood is thus a fixed point problem, i.e. determining what subjective conditional probabilities

concerning the behavior of others correspond to the conditional probabilities produced by the model when behaviors are based on those subjective beliefs.

Equation (1) is often simplified in two ways. First, the role of endogenous neighborhood effects is reduced to the expected value of the average choice of others, \bar{x}^i . This assumption reduces the possible endogenous effects to a single moment of the distribution of behaviors. Second, it is often the case that β is set equal to zero for all neighborhood members, eliminating uncertainty about the behavior of others. When these assumptions are made, individual decisions will solve

$$x_i = \frac{1}{1 + \beta} \left(\alpha + \beta \bar{x}^i \right) \quad (3)$$

This simplified version of the neighborhoods effects model makes clear that the interesting implications of the effects depends on the feedbacks between the choices of each individual. Neighborhoods models typically assume that these feedbacks represent a form of complementarity between choices. For the payoff function described in (2), complementarity⁶ means that if x_j increases, and x_k decreases, then

$$\frac{\partial^2 u_i}{\partial x_j \partial x_k} > 0 \quad (4)$$

(It is straightforward to modify the definition of complementarity in (4) to account for how beliefs about the average choice of others, rather than realized choices, affect individuals.) Complementarity is a fundamental property for interdependent decisionmaking because it implies that higher choice levels of others make higher choice levels relatively more attractive, thereby creating incentives for members of a common

⁶When u_i is twice differentiable, then the definition I employ is equivalent to requiring positive cross-partial derivatives between x_j and x_k . I prefer the definition of complementarity given by eq. (4) both because it is more general and because it better communicates the implications of complementarity for individual decisionmaking. I will use the cross-partial derivative definition when discussing papers that employ it.

neighborhood to behave similarly. In fact, the simplified neighborhood model (3) is an example of the class of coordination models studied in Cooper and John (1988) and generalized in Milgrom and Roberts (1990). Cooper and John establish, when θ is an interval of the real line and \hat{c} is twice differentiable, that complementarity is necessary for the presence of multiple equilibria.⁷ Intuitively, when complementarities are strong enough they imply that there exists a degree of freedom in the determination of individual and aggregate outcomes in the sense that complementarities dictate that individuals act similarly in equilibrium but do not specify their actual behaviors. In the neighborhoods context, this is important because multiple equilibria create the possibility that two neighborhoods with similar observables (i.e. distributions of \bar{c} within each neighborhood \bar{c} and levels of \bar{c}) can exhibit different aggregate behaviors.

The exact relationship between complementarities and multiple equilibria in a neighborhood will be quite complicated and will depend on the relationship between the joint distribution of the different determinants of individual choice within a neighborhood. The interplay of different choice determinants may be seen in a model due to Brock and Durlauf (2001b). In this model, members of a neighborhood make binary choices that are coded \bar{c} and 1. Agent payoffs are

$$u_i = \alpha \bar{c} + \beta c_i \quad (5)$$

where α and β are behavioral parameters and \bar{c} is the (rational) expected average choice in the neighborhood; this expectation is calculated under the assumption that each agent knows the values of the contextual effects and the individual characteristics of the neighborhood members.⁸ Notice that $\bar{c} = \frac{1}{n} \sum_{i \in N} c_i$, hence

⁷Cooper and John prove this under the additional condition that \hat{c} is concave, which allows them to model all choice decisions as identical.

⁸For ease of exposition, I ignore the fact that each individual knows his own realization of \hat{c} when forming beliefs about the average choice in the

an (expectations-based) complementarity is present and measured by this parameter. The random utility terms \hat{u}_i are assumed to be independent across individuals and to be negative exponentially distributed, i.e. $\hat{u}_i \sim \text{Exp}(-\lambda)$. Under these assumptions, it is straightforward to show that the expected average choice solves

$$\bar{x} = \frac{\sum_{i \in N} \exp(\lambda \bar{x})}{\sum_{i \in N} \exp(\lambda \bar{x})} \quad (6)$$

In this expression, \bar{x} and \bar{u} is the empirical distribution of \hat{u}_i within N . Characterizing the conditions for multiple equilibria is thus quite difficult even in this simple context, once one introduces heterogeneity across agents. For the special case where $\bar{u} = 0$, (so that all individual heterogeneity is generated by the unobserved \hat{u}_i) one can give a precise characterization of the number of equilibria.

Proposition 1. Number of equilibria in the Brock and Durlauf (2001a) model of binary choice with neighborhood effects

Under the model (5) and (6), if $\lambda > 0$, then there exists for each value of β a threshold $\bar{\beta}$, monotonically increasing in λ such that

- i. if $\beta < \bar{\beta}$, then there is a unique solution to (6)
- ii. if $\beta > \bar{\beta}$, there exist three solutions to (6)

One can show that the extremal equilibria (i.e. the equilibria associated with the values of \bar{x} that are largest and smallest in absolute value) are stable.

neighborhood. Brock and Durlauf (2001a) explicitly develop this model when agents are affected by the average choice of others, thereby dealing with this issue.

This proposition is useful as it illustrates how the predetermined incentives to make one choice or another, the strength of complementarity in choices, and the degree of heterogeneity across agents collectively determine the number of equilibria for a neighborhood. When predetermined incentives (β) are sufficiently strong relative to τ , then the equilibrium is unique. Intuitively, τ is smaller either when complementarities are relatively weak (β is small) or unobserved heterogeneity is large (σ is small). It is easy to see why weak complementarities militate against multiple equilibria. Why does τ play a role? Intuitively, high degrees of heterogeneity mean (since the density of ϵ is symmetric) that a large percentage of the agents in the neighborhood make choices having experienced relatively large draws of ϵ . When this is so, the percentage of agents whose decisions may be changed by the number of agents behaving similarly is relatively small, which can eliminate the possibility of multiple reinforcing configurations of behaviors.

Analogous results may be obtained when one moves to 3 or more choices (Bayer and Timmins (2002) and Brock and Durlauf (2002,2003)). To describe neighborhood effects in this context, it is necessary to generalize (5). When agents face K possible choices, this may be done by assuming that i 's utility for choice k is described by

$$u_{ik} = \beta v_k + \epsilon_{ik} \quad (7)$$

where τ_k is the expected percentage of agents in the neighborhood making that choice under when rational expectations are imposed. Relative to the binary case one important difference is that one needs to account for differential effects of individual, contextual, and endogenous effects on choice-specific payoffs; this is done by allowing the coefficients in (7) to be choice-dependent. When the random terms are double exponentially distributed, i.e. $\epsilon_{ik} \sim \text{Exp}(\tau_k)$, one has a multinomial

⁹ e is Euler's constant.

logit model of choice with neighborhood effects. Brock and Durlauf (2002,2003) provide results on the unique versus multiplicity of equilibria for this model and show how these findings generalize to other error densities.

Once one allows for general forms of heterogeneity, for example, by allowing the individual characteristics to vary across individuals, models of choice within neighborhoods rapidly become analytically intractable. This suggests that computer simulations will play an increasingly important role in studying these models, particularly as their microeconomic structures are enriched. Examples of local interactions structures that have been analyzed using a combination of analytics and computer simulations include Bell (2002) and Oomes (2003).

Complementarities within neighborhoods have implications for aggregate data beyond the possibility of multiple equilibria. This possibility was recognized in a paper by Glaeser, Sacerdote and Scheinkman (1996). The paper is also interesting as an example of a neighborhood interaction structure that is local in the sense that an individual is affected by immediate neighbors as opposed to the average behavior of the neighborhood as a whole.

In the Glaeser, Sacerdote, and Scheinkman (1996) model, agents face binary choices (in their context, on whether or not to engage in criminal activity). Agents are assumed to come in three types indexed by ω ; ω^i denotes the type of agent i . Agent heterogeneity is thus modeled via the distribution of types in a population. Mapping their model into the binary choice framework I have already described, each is associated with a payoff function

$$u_i = \beta \sum_{j \in N_i} \omega^j + \alpha \omega^i \tag{8}$$

The substantive difference between this payoff function and payoffs as described by (5) are twofold. First, there is no unobserved heterogeneity in individual payoffs due to differences in β or α ; rather heterogeneity appears because different types are associated with different values of the parameter ω in the payoff function. Second, neighborhood effects are restricted to pairs of nearest neighborhoods within a

neighborhood, since the decision of each agent is assumed to only depend on one other agent. Metaphorically, agents live on a 1-dimensional block and only care about their neighbor to the left.

The recursive structure of interactions makes the analysis of individual decisions relatively easy since there are no reciprocal feedbacks between individual decisions; while the choice of c_i affects c_{i+1} , the choice of c_{i+1} does not affect c_i . As a result, no issues arise of individuals making choices in order to alter the behavior of others. Nevertheless structures of this type can produce complex behaviors because the bilateral interactions across individuals indirectly connect all members of a neighborhood.

Glaeser, Sacerdote and Scheinkman close this model by specifying the three types of agents. Type 0 agents always choose $c_i = 0$; this may be generated by $c_i = 0$. Type 1 agents always choose $c_i = 1$; this behavior may be generated by $c_i = 1$. Type 2 agents always make the same choice as his neighbor, which will occur if $c_i = c_{i-1}$. The distribution of agent types is assumed to be random; this randomness makes the behavior of each individual unpredictable to the modeler. These assumptions on heterogeneity may be shown to imply that there is a unique equilibrium set of choices in the neighborhood. Glaeser, Sacerdote and Scheinkman assume agent types are independent across individuals with distributed with fixed probabilities θ .

These assumptions on heterogeneity may be shown to imply that there is a unique equilibrium set of choices in the neighborhood. This uniqueness is consistent with the goal of the analysis, which is to show how endogenous neighborhood effects influence the volatility of average crime rates across neighborhoods. Defining $\bar{c}_n = \frac{1}{n} \sum_{i=1}^n c_i$ and $\bar{c}_n^2 = \frac{1}{n} \sum_{i=1}^n c_i^2$, Glaeser, Sacerdote and Scheinkman prove that the normalized sample average $\sqrt{n}(\bar{c}_n - \theta)$ obeys the following central limit theorem:

Proposition 2. Limiting behavior of average choice in the Glaeser, Sacerdote and Scheinkman (1996) model

What makes this proposition important is that it provides a description of how endogenous neighborhood effects influence the dispersion of average behaviors across neighborhoods. To see this, observe that if there were no neighborhood effects, this would mean that all agents are of type 0 or 1. In this case, the variance term in the proposition would reduce to σ^2 . As the percentage of type 2 agents increases, the variance term also increases, and becomes unbounded when $\theta = 1$. This model thus illustrates that even when neighborhoods exhibit unique equilibria, neighborhood effects multiply the effects of idiosyncratic shocks and thereby increase the variance in the average behaviors that are observed.

The statistical mechanics of neighborhood interactions

The Brock and Durlauf (2001a) and Glaeser, Sacerdote, and Scheinkman (1996) models are mathematically equivalent to models that have appeared in the statistical mechanics literature in physics and mathematics. (Models of this type are also known as interacting particle system models or random field models). Statistical mechanics models specify a particular topology of interactions for agents within a population and explore the implications of this topology for the distributions of behaviors of the objects. These interaction topologies are often local in the sense that agents are usually arrayed in some space with the strength of direct interactions determined by distance; one may think of collections of elements that directly interact as neighborhoods. Mathematical structures that describe populations that are comprised of neighborhoods in this sense were in fact originally developed in the context of statistical mechanics, where many fundamental phenomena such as magnetization may be described as aggregate properties of a large number of interacting elements. For the purposes of socioeconomic phenomena, statistical mechanics models allow one to analyze alternative forms of neighborhood interactions. For example, one can analyze environments where the direct payoff interdependences between members of a neighborhood are uniform as is done in Brock

and Durlauf or analyze cases where direct interdependences occur between nearest neighbors as in Glaeser, Sacerdote, and Scheinkman. One of the key ideas in this literature is that overlapping neighborhoods create population-wide interdependences in behaviors, which is the basis of Proposition 2 above. Durlauf (1997), Ioannides (1997) and Kirman (1997) provide different perspectives on the use of these models in social science contexts.

Ideas from the statistical mechanics literature have begun to appear in the economics literature as these models possess a number of interesting properties. Among these properties I would identify

1. Multiple equilibria and nonergodicity. One way to understand the structure of statistical mechanics models is to think of them as probability descriptions. Following the language employed in the description of the basic model of choice with neighborhood effects, the preferences, constraints and beliefs of the agents produce a set of conditional probability measures that describe each individual's behavior given the behaviors and/or characteristics of others. An equilibrium in the model is a joint probability measure that is consistent with these conditional probability measures. When more than one such measure exists, the model exhibits multiple equilibria. The mathematical term for such multiplicity is nonergodicity.

2. Phase transition. A model exhibits phase transition if its properties qualitatively change for a small change in a parameter value. Phase transitions are thus a way of describing when threshold effects occur in an environment. The Brock and Durlauf model exhibits a phase transition around the threshold $\bar{\theta}$; the number of equilibria in the model changes as the compound parameter θ is moved above the threshold.

3. Universality. Statistical mechanics models often have aggregate properties that do not depend on the specifics of the interaction structure; so that different interaction structures will produce identical aggregate behaviors. For an example of universality, one can show that equilibrium average choice levels in the Brock and Durlauf model are unchanged if one replaces the payoff function with

(10)

That being said, universality is largely an unexplored property in economics contexts. From the perspectives of neighborhood effects, these models are interesting as they allow for the exploration of a range of alternative interaction structures. I regard the study of universality to be one of the most important future directions in this literature since there are generally no good empirical reasons to favor any particular neighborhood interaction structure over another.

The first example of a statistical mechanics model to appear in economics appears to be due to Föllmer (1974). Föllmer considered a population of individuals arrayed on a torus formed by the 2-dimensional integer lattice \mathbb{Z}^2 . (To visualize this torus, think of a \mathbb{Z}^2 lattice where the edges are matched up, producing a doughnut shape, to ensure each element has 4 neighbors within 1 unit distance, and let \mathbb{T}^2 denote the torus.) In this model, each individual is associated with one of two utility functions; following the earlier discussion, the utility function associated with an agent is indexed by $i \in \mathbb{T}^2$. Föllmer's objective was to determine whether, if each individual has an equal unconditional probability of being associated with either utility function, the aggregate population would converge (as population size grew) to one where half of the population is associated with each utility function. Föllmer assumed that each individual's probability of having a particular utility function was a function of his nearest neighbors, i.e

$$p_i = \frac{1}{2} + \frac{1}{2} \sum_{j \in N(i)} \sigma_j \quad (11)$$

with $\sigma_j \in \{-1, 1\}$. Föllmer (1974) shows that the average preference structure in the population, $\bar{\sigma}$ obeys the following Proposition.

Proposition 3. Average preferences in the Föllmer (1974) model of interdependent utility functions

There exists a $\bar{\beta}$ such that under the conditional probability structure (11)

- i. if $\beta > \bar{\beta}$, then $\bar{u} = u^*$ for every realization of the process
- ii. if $\beta < \bar{\beta}$, then there exist values $\epsilon > 0$ and $\delta > 0$, $\left| \frac{\beta}{\delta} \right| > \left| \frac{\epsilon}{\delta} \right|$ such that

$$\lim_{\delta \rightarrow 0} \frac{\beta}{\delta} > \frac{\epsilon}{\delta} \tag{12}$$

In the Föllmer model, β measures the strength of interdependence in preferences between nearest neighbors on the lattice; notice how the analogous parameter appears in the Brock and Durlauf (2001a) and Glaeser, Sacerdote, and Sheinkman (1996) models. The theorem states that when this interdependence exceeds a certain threshold, then each sample path realization will tip toward one of the two utility functions predominating in the aggregate economy, even though there are no ex ante differences in the probabilities of either outcome. As such, the model is nonergodic.

Part of the growth of statistical mechanics models has occurred in the context of game theory; this work is an important advance over Föllmer (1974) in that this original work did not have any behavioral foundations. In a seminal paper, Blume (1993) has provided a rigorous game-theoretic formulation of lattice interaction structures of this type. This paper is also important in that it illustrates how the technical assumptions that underlie these models have substantive implications for individual behavior. One important advance Blume (1993) makes relative to Föllmer (1974) is the construction of a continuous time version of the model which allows agents to update their choices asynchronously. This allows Blume to use a Nash equilibrium concept to analyze the model. Ellison (1993) provides a related analysis of local interactions of this type, with an emphasis on the speed at which the population converges to its limiting behavior. Ellison shows how local interactions, placed in a learning context, can slow down convergence relative to a model where agents react to all other choices, as occurs in

Brock and Durlauf. Anderlini and Ianni (1996) extend work of this type to a context where individuals play a sequence of games with randomly selected neighbors, showing how the steady state behavior of such a model differs from the case where interactions are not local. Put differently, Anderlini and Ianni show that local interactions can allow a diversity of strategies to coexist in equilibrium even when all agents are playing the same (in this case, coordination) game. Recent analyses of equilibrium properties of statistical mechanics models include Blume and Durlauf (2002) and Scheinkman and Horst (2003); the former paper studies the dynamics of statistical mechanics models where interactions are common across all members whereas the latter is useful in its analysis of the relationship between the description of local neighborhoods for complementarities and the description of conditional probabilities for individual choices given the choices of others in the population.

A number of papers have employed statistical mechanics models to understand substantive phenomena. The notion of neighborhood differs across contexts; what makes the abstract notion of neighborhood useful is the implicit assumption in all of these models that agents are arrayed in some space in which distances between agents may be defined, and where these distances influence interactions. Some examples of these applications include:

Business cycles. Bak, Chen, Scheinkman and Woodford (1993) develop a model of sectoral interactions based on order flows between intermediate and final good producers. When production exhibits nonconvexities, idiosyncratic producer-level shocks are able to produce aggregate business cycles. Oomes (2003) studies the role of local demand interactions in producing spatial persistence in unemployment. Interestingly, she finds that when labor may be hired in continuous quantities, the long run distribution of unemployment is uniform across locations whereas if labor is a binary variable, spatial differences may persist indefinitely; this role for nonconvexities echoes Bak, Chen, Scheinkman and Woodford. Recently, Bell (2002) provides an interesting mix of local and global interactions in a model in which preferences are influenced by neighborhood effects whereas global effects are introduced via prices. Bell shows that these two effects can combine to produce rich cross-section diversity and aggregate cycles.

Economic growth. Durlauf (1993) studies a model of interacting industries in which individual industries are arrayed on a 1-dimensional lattice. Distances between industries on the lattice are interpreted as reflecting technological similarity. In the model, the choice between a high productivity (with high fixed cost) versus a low productivity technology by industry i at t depends on the technology choices of industries j to $i + \Delta$ at $t - 1$, thereby allowing each industry to be affected by an arbitrarily large (although finite) span of neighboring industries. Durlauf (1993) shows how this model can produce development traps and explain industrial development using the idea of backwards and forwards linkages due to Albert Hirschman. Kelly (1997) develops a model of local transportation linkages and economic development that captures Adam Smith's insights on market size and specialization; this model is discussed in the context of networks and neighborhoods below.

Land use. Irwin and Bockstaed (2002) show how neighborhood interactions can affect long term patterns of land use. In this model, the decision to develop a given parcel of land is affected by the development state of neighboring parcels. This model is shown to produce fragmented patterns of development on urban/rural fringes that are consistent with observed patterns. One interesting feature of this analysis is that it finds that the neighborhood effects for land development are negative, which may reflect congestion effects.

Technology adoption. Allen (1982) uses statistical mechanics models to analyze technology adoption when network externalities are restricted to local neighborhoods. Anand and Kiefer (1995) also consider the question of technological adoption and provide conditions under which more efficient technologies are adopted and when they are not.

Financial markets. Brock (1993) develops a general statistical mechanics approach to financial markets and shows how the interactions that underlie these models may be used to interpret a range of stylized facts concerning time varying volatility and market volume in financial markets. These ideas have been pursued in recent work such as Cont

and Bouchard (2000) and Focardi, Cincotti, and Marchesi (2002) who consider neighborhoods defined by random connections to show how stock market crashes may arise. One problem with efforts such as these (unlike Brock's work), which are part of the growing "econophysics" literature (cf. Mantegna and Stanley (1999) is a well written introduction) is that they have weak microeconomic foundations, as discussed in Durlauf (2003).

Finally, it should be noted that within game theory there has emerged a rich literature on local interactions that does not use statistical mechanics properties explicitly, although the mathematical structures are similar. Eshel, Samuelson, and Shaked (1998) explore the evolution of different behavioral types when the payoffs from different strategy choices are determined by play with neighbors and show how altruistic behavior can survive in such a context; Nowak, Bonhoeffer, and May (1994) and Nowak and May (1993) study similar environments. Perhaps the most general and deep treatment of local interactions models in game theory is Morris (2000) who analyzes these structures with great generality and establishes conditions under which behaviors that originate among a finite set of players can spread across a population given arbitrary local interaction structures.

ii. Equilibrium neighborhood configurations

The presence of contextual effects and endogenous effects in individual payoff functions implies that a complete theory of neighborhood effects must account for how neighborhoods are formed in the presence of these effects. From the perspective of an abstract choice problem, individual neighborhood choices may be thought of as

$$\dots \tag{13}$$

where u_n is the expected utility associated with neighborhood n , and \mathbf{r} denotes any variables that affect this payoff; typically, this will be the rental price to residence. Notice that this payoff calculation requires that agents form beliefs about which equilibrium will emerge in a neighborhood when multiple equilibria exist. An equilibrium in an endogenous neighborhoods model is a configuration of agents across neighborhoods such that each agent solves (13) and the resulting values of \mathbf{r} are those that are generated by this configuration. Put differently, the neighborhood choices of individual agents help determine the effects in a neighborhood, so location decisions need to exhibit self-consistency.

The key theoretical feature of interest in these types of models concerns how individuals with different attributes are allocated across neighborhoods; this allocation will determine the extent to which different neighborhoods will exhibit different contextual effects; without such differences, the only way that neighborhoods may affect inequality is via multiple equilibria. The most common attribute that is studied in this literature is income, although other attributes have been considered, as is discussed below. Much of the interest in neighborhood configurations, in turn, focuses on the extent to which neighborhoods are stratified by income or other attributes. Neighborhoods are said to be stratified with respect to an attribute \mathbf{r} if it is the case that the supports of the of intra-neighborhood distributions of \mathbf{r} do not overlap except at endpoints. Stratification by income, for example, provides a basis for understanding persistence in economic status across generations: poor families are consigned to poor neighborhoods, whose effects make it more likely their children are poor, etc.

As is well understood in the urban economics literature, modeling equilibrium allocations of individuals across neighborhoods involves a number of subtleties. In order to understand the essential features of these types of models, I describe a simple framework developed in Bénabou (1996a), which is adapted to the notation I have employed and makes clear the issues that arise with respect to the question of whether equilibrium neighborhood configurations are stratified.

I work with a slightly simplified version of the Bénabou model that is modified to correspond to previous notation. The population is assumed to consist of a continuum of

agents with associated measure μ_i who may live in one of two neighborhoods, each of size μ_i and denoted as A and B . Residence in a neighborhood entails the payment of rent r_i ; these rents go to an absentee landlord. There are two types of agents, denoted by whether they are associated with individual characteristics θ_i or θ_j ; the characteristic is a scalar. Let α_i denote the percentage of θ_i agents. All contextual effects derive from the average value of θ_i within the neighborhood, which is denoted as $\bar{\theta}_i$. The payoff to a given neighborhood can therefore be written as $\mu_i (r_i - \bar{\theta}_i)$.

What conditions determine whether the equilibrium allocation of families across neighborhoods is stratified by θ_i ? Bénabou answers this question by considering how individual agents are willing to trade off the rental price r_i against the neighborhood characteristic $\bar{\theta}_i$. Specifically, for each individual, one may define the function $\mu_i \frac{\partial r_i}{\partial \bar{\theta}_i}$ which characterizes the marginal tradeoff between r_i and $\bar{\theta}_i$ that leaves utility unchanged at a given $\bar{\theta}_i$. This tradeoff implicitly describes how different agents will react to changes in relative neighborhood rental prices. This function may be derived by differentiating the payoff function with respect to r_i and $\bar{\theta}_i$ when utility is held constant, i.e.

$$\mu_i \frac{\partial r_i}{\partial \bar{\theta}_i} = \frac{\partial \mu_i (r_i - \bar{\theta}_i)}{\partial \bar{\theta}_i} = \mu_i \frac{\partial r_i}{\partial \bar{\theta}_i} - \mu_i \tag{14}$$

The properties of $\mu_i \frac{\partial r_i}{\partial \bar{\theta}_i}$ determine whether neighborhoods are stratified in equilibrium. The following Proposition is proven in Bénabou (1996a).

Proposition 4. Stratification in the Bénabou (1996a) model of endogenous neighborhoods

If \bar{w}_i is increasing in x_i , then the only stable equilibrium configuration of families is one that is stratified.

Bénabou's result may be stated directly in terms of willingness to pay; here I follow the excellent exposition in Becker and Murphy (2000). For each agent, let w_i denote the willingness to pay for neighborhood A . Without loss of generality, assume that w_i is increasing in x_i . Under the assumption that this function is strictly increasing in x_i , stratification in the sense described in Proposition 4 will occur if (letting neighborhood A denote the "better neighborhood") there exists a pair of rental prices r_A and r_B such that

$$w_i(r_A, x_i) > w_i(r_B, x_i) \quad (15)$$

and

$$w_i(r_A, \bar{x}_A) > w_i(r_B, \bar{x}_A) \quad (16)$$

where \bar{x}_A denotes the average x value for a neighborhood that contains n_A agents. These conditions are equivalent to those in the Bénabou theorem, since a greater willingness to pay by affluent families is an implication of the assumption that w_i is increasing in x_i .

Why does stratification require that the willingness to pay for a better neighborhood is increasing in the individual attribute x ? As Bénabou explains, when this condition holds, any deviation of neighborhood composition away from symmetry will induce n_A families to move to the higher x neighborhood. Bénabou's analysis makes clear that the presence of neighborhood effects is not sufficient to produce equilibrium stratification; what is additionally required is that there is a positive correlation between the attribute and willingness to pay for the average of the attribute.

What factors contribute to the fulfillment of conditions necessary for stratification, as delineated in Proposition 4 or in eqs. (15) and (16)? Bénabou (1996a) provides a typology of different factors. With respect to neighborhood effects per se, one factor that contributes to the conditions necessary for stratification to occur is associated with complementarity between individual and neighborhood attributes, i.e. $\alpha_i \geq \alpha_j$; this condition is in fact sufficient in under the complete set of assumptions in the Bénabou (1996a) model (see his Proposition 2) to ensure stratification. This is unsurprising as complementarity between individual and group attributes means that the benefits to high α neighborhoods are greater for high α types. Bénabou also shows that other factors can contribute to stratification even when complementarity does not hold. One reason why stratification may still occur is via financial market imperfection, e.g. the ability of α_i agents to borrow at lower interest rates than others. This argument is closely related to the role of capital market imperfections in perpetuating inequality originally studied by Loury (1981). In Loury's model, families are unable to borrow to educate their children because they are unable to commit their children to repay these loans; thereby producing interfamily inequality in the marginal value of educational expenditures. A related phenomenon occurs in the context of neighborhoods models in that financial market imperfections can keep poorer families from being able to enter neighborhoods even when they would benefit more from the various neighborhood effects than the current residents.

The neighborhoods literature contains a range of alternative specifications of neighborhood models; recent examples include Bénabou (1993,1996a), de Bartolome (1990), Durlauf (1996a,b), Epple, Filimon, and Romer (1984), Epple and Platt (1998), and Fernandez and Rogerson (1996,1997), Hoff and Sen (2000), and Nechyba (1997). Table 1 is a summary of some of the features of these models. Earlier models that are driven by similar ideas include Miyao (1978) who did early work on the stability of community segregation, though for that model the main driving mechanism is a desire to be with others of one's type. There is still no consensus as to the most useful general framework. One reason for this is that multiple neighborhoods models, both in the specific context I address as well as the broader urban economics literature, pose a

number of theoretical complications. One complication arises from the need to ensure the existence of majority voting equilibria for neighborhood tax rates.¹⁰ More generally, the construction of a neighborhoods model in which an equilibrium exists require a number of modeling choices, including both the details of neighborhood structure (number, fixed or variable neighborhood size, nature of housing market) and the nature of heterogeneity across individuals.

One important set of differences in various neighborhood models concerns the specification of neighborhood structure. Models assumptions vary in terms of the number of neighborhoods, the size of neighborhoods, the degree of heterogeneity in agents, and the mechanisms that determine neighborhood membership. In Bénabou (1996a) and most of the analysis in Becker and Murphy; it is assumed that there are two types of agents and two equal sized communities. Associated with these two neighborhoods are rental prices; these prices determine why some families prefer to be in the lower quality neighborhood. Epple, Filimon, and Romer (1984) study an environment with a fixed number of communities but with no restrictions on community size. Neighborhood entry is free; differences in neighborhood choices occur on the basis of the tax/education combinations that they offer; there are no costs to residing in a particular neighborhood. Fernandez and Rogerson (1996) analyze a similar model with two communities and two (greater than two) income types. By eliminating the housing consumption decision and placing taxes on income rather than housing consumption, as is done in Epple, Filimon and Romer (1984), they achieve some valuable simplifications that render the environment useful for policy analysis, albeit at the cost of some realism. Durlauf (1996a) allows for the number of neighborhoods and their respective sizes to be endogenous, but does this at the cost of modeling neighborhood entry as a club in which existing members may be vetoed by current members using income requirements for neighborhood membership; Durlauf (1996b) shows how to support this allocation with house prices for particular choices of preferences and the human capital production function. Hoff and Sen (2000) makes the interesting extension to these types of models

¹⁰Typically, expenditures are determined by taxes (i.e. there is no interneighborhood borrowing) with an equilibrium tax rate defined as one such that no single alternative would defeat it in an election.

by enriching the housing options a neighborhood offers to allow for renters as well as homeowners. For this range of specifications, the equilibrium neighborhood allocations are, under intuitive assumptions on preferences and the various technologies that appear (e.g. production of human capital), stratified by income which is the single dimension along which families differ.

While the stability of stratified equilibria appears to be robust to a range of alternative specification of neighborhood structure and the rules that determine neighborhood membership, it is possible to identify two ways in which this finding is sensitive to substantive microeconomic assumptions. First, it is possible that the willingness to pay differential needed for equilibrium stratification will fail to hold: de Bartolome (1990), analyzing a model in which families differ according to the ability level of their children, provides conditions under which two communities can coexist, each of which is mixed. Becker and Murphy (2000) work with a version of the two-neighborhood Bénabou (1996a) model and illustrate how mixed communities can coexist; they do this by relaxing the ability to pay assumption that underlies the Proposition 4.

Perhaps more fundamental, stratification depends on the assumption that families are distinguished along a single dimension. As such, much of the richness of neighborhood decisions has been assumed away. This is understandable given the difficulties that exist in developing neighborhood models with multiple levels of heterogeneity. However, there have been important recent advances in this regard. While not focusing explicitly on neighborhoods (they study allocations across private and public schools) Epple and Romano (1998) consider a model in which families differ by both child ability and parental income. Their model produces an equilibrium in which schools are ordered according to quality but in which sorting occurs along the two dimensions of income and ability. While extrapolating their results to a neighborhood context is not straightforward (since schools are allowed to offer financial incentives to influence their composition), their findings suggest an important dimension along which current models of neighborhood formation should be generalized. In the context of neighborhood formation, Epple and Platt (1998) and Nechyba (1997) provide equilibrium models where agents differ with respect to preferences as well as income. Epple and

Platt (1998) allow for costless movement across neighborhoods whereas Nechyba (1997) restricts neighborhood sizes and introduces a fixed stock of heterogeneous houses that may be traded. Two attractive features of the Epple and Platt (1998) model are 1) voters account for the consequences for neighborhood composition when voting over tax rates, a possibility whose importance was earlier shown in Epple and Romer (1990) and 2) renters and homeowners may both be present in a neighborhood. While neither of these analyses includes spillover effects, they are, in my view, an important stepping stone for future work. Even though there exists substantial stratification of neighborhoods by income, there is also considerable overlap in the income supports of neighborhoods (see Jargowsky (1997) for a broad overview); an empirically successful theory needs to account for this mixing. And of course, the logic of the Tiebout approach to community formation places primary weight on preferences as the source of segregation, not income per se.¹¹

The vast majority of studies of neighborhood formation focus on rents or prices as the mechanism by which neighborhood memberships are restricted. As such, these models typically ignore the role of zoning restrictions in influencing neighborhood composition. Zoning, in this context, may be thought of as placing a requirement on the number of units of housing (perhaps measured in terms of quality) that a homeowner must possess in order to reside in a given neighborhood. As originally shown by Hamilton (1975), these effects can be first order. For example, Hamilton develops a model in which perfect sorting of neighborhoods by income occurs when the number of neighborhoods is variable. One exception to this lack of attention to zoning is Fernandez and Rogerson (1997), who study zoning in the context of a two community model with local finance of education. Working with a model in which equilibrium allocations of families across neighborhoods are stratified by income, they show that the introduction of zoning (both exogenous and endogenously chosen by neighborhood members) alters the allocation of families across neighborhoods by making the high income neighborhood

¹¹This discussion does not exhaust the reasons why neighborhoods are mixed by income. Frankel (1998) for example, shows that information imperfections can generate incentives for high income families to live near low income families because of the effects of low income families on neighborhood consumption goods prices. It is not clear how important this is empirically, although the argument is quite clever.

more exclusive. Interestingly, zoning does not necessarily increase the disparity in per pupil educational expenditures across districts. The reason for this is that the greater exclusivity of the richer neighborhood raises the average income of the poorer neighborhood and therefore can, in equilibrium, reduce the gap in interneighborhood expenditures. Fernandez and Rogerson (1997) do not analyze the role of zoning in models with neighborhood effects other than those induced by the public provision of education; this would be a useful extension of their work.

Segregation

Most of the modern neighborhood effects literature has assumed that preferences over neighborhood composition are associated with the effects of composition on “economic” outcomes such as human capital accumulation. Segregation models, in contrast, focus on how neighborhoods evolve when individuals have a preference with others of similar ethnicity; these preferences are taken as a primitive with no suggestion that ethnicity matters in any other sense.

Modern models of neighborhood effects and segregation originate in the classic work of Schelling (1969,1971). Schelling’s work was motivated by the desire to understand the phenomenon of neighborhood “tipping,” which refers to the claim that the black/white composition of neighborhoods would, when the percentage of blacks reached around 30%, precipitously shift to a very high percentage of blacks. Schelling’s goal was to see whether this type of behavior (as well as the fact that racial segregation levels in the US are very high) could be explained without recourse to the assumption that whites preferred all white or nearly all-white neighborhoods. In contrast, Schelling considered preferences in which individuals were indifferent to the racial composition of a neighborhood until the percentage of individuals of the same type dropped below a certain threshold. Preferences of this type are “nonracist” in the sense that segregation is not strongly preferred to alternative configurations. Schelling considered a model of a population of blacks and whites arrayed along points in a 2 dimensional lattice. In this model, agents sequentially consider whether to move to another site on the lattice (moves are only permitted to open sites.) Agents would choose to move if and only if 1) the

percentage of their nearest neighborhoods of different ethnicity was above a certain percentage and 2) a preferred location was open. Remarkably, Schelling found that for a range of moving thresholds, this population would essentially completely segregate itself over time.

Why does this occur? Notice that Schelling assumes that agents act myopically and sequentially. These assumptions ensure that when an agent moves, it is always towards a neighborhood with a higher percentage of his own group than otherwise. Hence, any initial configuration can only become more segregated. What happens when an initial configuration contains agents that are discontent? A more segregated outcome must emerge. This intuition has been formalized in theoretical work by Granovetter and Soong (1988) and Young (1998); the latter is noteworthy in recasting Schelling's work in the context of evolutionary game theory. In related work, Mobius (2000) shows how random fields methods may be applied to Schelling-type models in ways that allow for richer neighborhood configurations than have previously been studied; this latter work suggests methods in statistical mechanics (cf. Brock and Durlauf (2001a,b) and Durlauf (1997)) may prove to be generally useful in the modeling of segregation dynamics.

Recent work has shown how the emergence of segregation occurs in the presence of alternative preferences to those studied in the original Schelling work. Panes and Vriend (2003) show how segregation emerges even when agent preferences prefer integration to all other outcomes. This work is a valuable extension to Schelling in that it both shows that even a strict preference for integration can be overcome by dislike of being isolated. Bruch and Mare (2003) consider alternative preferences which smooth out the behavior of individuals in the sense that the decision to move from a neighborhood is no longer a zero/one function of a threshold; simulation evidence in that paper finds introducing smoothness in individual decisions strongly reduces long run levels of segregation.

While Schelling-type models have provided fundamental insights into the emergent nature of segregation, it is important to recognize their limitations. First, I am unaware of any version of the Schelling model that introduces a well developed housing market. (Mobius (2000) and Zhang (2003) introduce housing prices, but these prices are not generated as a consequence of market clearing; rather, they constitute an additional

cost associated with neighborhood composition that is related to demand conditions.) Second, these models have yet to incorporate forward-looking behavior on the part of agents. Hence, agents make neighborhood location decisions without regard to whether these decisions will be changed in the future. As a result, these models rely on agents making residential location decisions without consideration of the future characteristics of a neighborhood. These limitations have rendered the Schelling approach somewhat hard to interpret in light of the rationality postulates of standard economic theory.

One additional concern stems from the fact that the Schelling model has not been subjected to much empirical scrutiny. For this reason, Easterly (2003) is important. Easterly argues that while approximately 10% of American urban census tracts shifted from majority white to majority black between 1970 and 2000, there is little evidence that these changes are well explained by the initial shares of whites in the population, as the pure Schelling model predicts. Hence, it seems important to embody microeconomic foundations into the Schelling framework in a way which can incorporate other factors that influence neighborhood choice, factors that of course are precisely what the neighborhood effects literature studies. Sethi and Somanathan (2001) is a valuable contribution in this regard as it models neighborhood choice when individuals have preferences over both the ethnicity and the income levels of neighbors. An evaluation of the empirical implications of this model would facilitate an extension of Easterly's analysis.

iii. Neighborhoods and networks

A final area where neighborhood effects have been analyzed concerns their role in constituting social networks. To make this discussion concrete, it is useful to think about a population of agents in which some topology describes who directly communicates with whom. Information can only travel across direct communication links between individuals. One can then think about a neighborhood as a set of agents who are all capable of mutual communication via the network. Conceiving of neighborhoods in this way may be useful in understanding phenomena such as the transmission of information

about job opportunities, as is discussed below. Social networks as a whole constitute a sufficiently large area to be far beyond the scope of this chapter (see Jackson (2003) for a survey of recent theoretical developments in modeling networks). Rather, I focus on a few aspects of this literature that most closely relate to issues of how neighborhood effects produce inequality.

One of the important general questions concerning networks is what sorts of networks emerge from a given set of direct links between agents. Following work such as Kirman (1983) and Ioannides (1990) this question may be studied formally using a model in which there are n agents. Suppose each agent has a certain number of direct connections and that these connections are reciprocal, so that if i is connected to j , then j is connected to i . This environment is thus an example of a graph where the agents constitute vertices and the direct connections constitute edges. Neighborhoods are defined as groups of agents all of whom are directly or indirectly connected; i is indirectly connected to j if there is a set of agents $\{k_1, k_2, \dots, k_m\}$ such that i is directly connected to k_1 , k_1 is directly connected to k_2 and so on, with k_m directly connected to j . Different configurations of a fixed number of edges will thus induce different neighborhood structures.

From this perspective, it is clear that the structure of networks in a population will depend on both the number of direct connections between agents as well as their distribution. It is possible to form a common network for n agents from as few as $n-1$ direct links. Similarly, since for any subset of n agents, it is possible to identify n distinct links, it is possible for n links to leave most agents isolated. The natural question that arises from this structure concerns identifying which network structures are most “likely” in the space of possible configurations of direct links. This question may be answered using random graph theory, which originates in the work of Erdos and Renyi (1960). Formally, suppose that each pair i and j are directly connected with probability p and that direct connections are independently distributed. The dependence of the probability of a direct connection on population size is important. If the probability is not so dependent, then the number of direct connections

for each agent grows linearly in population size; in contexts such as word-of-mouth communication linear growth of direct connections would be unappealing.

Under these assumptions on direct connection probabilities, the following proposition (taken from Palmer (1985, chapter 4)) holds:

Proposition 5. Neighborhood size and direct connection probabilities

- i. if $\frac{c}{N} > \frac{1}{2}$, then as N grows, the largest neighborhood in the population will be of order N^2
- ii. if $\frac{c}{N} < \frac{1}{2}$, then as N grows, the largest neighborhood in the population will be of order N
- iii. if $\frac{c}{N} = \frac{1}{2}$, then as N grows all agents will be members of a common neighborhood

This proposition illustrates how the distribution of neighborhoods exhibits two phase transitions. First, as the probability of a direct connection moves above $\frac{1}{2}$, so that the expected number of connections in the aggregate population exceeds $\frac{N}{2}$, the size of the largest neighborhood qualitatively changes. While the population is broken up into many relatively small (when compared to the overall population size) neighborhoods, once the direct connection probability exceed the first threshold, a large neighborhood emerges. (This is known as the giant component in the random graph literature). One can show that the order of this neighborhood is unique. Second, as the probability moves above $\frac{1}{2}$, multiple neighborhoods disappear as N becomes large. What sort of substantive interpretation may be given to this property? If neighborhood effects derive

from the role of social networks in pooling information, as might occur in labor contexts, then the “quality” of the information available to a given individual may depend on the size of his network. One interpretation of the emergence of an underclass, consistent with Wilson (1987) and formalized in work such as Finnerman and Kelly (2003) is equivalent to the emergence of a subset of the population that is isolated from the information flows available to others.

Random graph techniques have been applied in several contexts. Kirman (1983) appears to be the first to suggest this methodology to study how local interactions lead to aggregate interactions. This analysis focuses on the question of whether local trading networks will lead to price uniformity in a population; the assumption being that arbitrage ensures prices are common across all bilateral trades in a common neighborhood. Ioannides (1990) is the first to recognize the full power of the threshold theorems in random graph theory and uses a random graph framework to study risk sharing among traders, analyzing the relationship between trading probabilities and the number of agents whose risk is pooled via within neighborhood trading. Other work has employed a random graph structure to analyze dynamic questions. For example, Kelly (1997) develops a model of specialization and economic growth in which transportation linkages between individual sites increase the extent of the market and can produce accelerated growth. His analysis exploits a theorem due to Bollobas and Thomason (1987) that shows that threshold properties such as the Proposition 5 are in a certain sense generic to random graphs, a result that should prove to have wide applicability in economic contexts. Durlauf (1996d) uses random graph methods to build a model of business cycles out of neighborhood interactions; in this context neighborhoods define groups of traders who are connected directly and indirectly with larger neighborhoods producing demand externalities for individual agents. However, it seems fair to conclude that random graph methods for linking neighborhoods and aggregate outcomes have not proven to be a major tool in modeling neighborhood effects, at least yet. I believe this is for two reasons. First, random graph methods are difficult to work with and so relaxation of various assumptions that are unnatural from an economist’s perspective (e.g. the independence of bilateral connection assumption or the lack of heterogeneity in bilateral connection probabilities) has yet to be achieved and is certain to be extremely difficult.

Second, the random graph models that have appeared, with the exception of Kelly (1997) typically treat bilateral connections nonstrategically, whereas in many cases such as the development of communications relationships, one would expect these links to be choice variables.

Networks and neighborhoods have also been jointly studied in the context of labor markets and information transmission. From the perspective of neighborhood effects, this work is especially important as it represents the best articulated effort to develop microeconomic foundations for neighborhood effects in the sense that networks models develop neighborhood effects as a consequence of information imperfections; as I have argued above, the literature typically takes neighborhood effects as a primitive. The basic objective of the labor markets/social networks literature is to understand the role of local contacts in the transmission of information about jobs. It is generally accepted that approximately $\frac{1}{3}$ of all job vacancies are filled by individuals who knew someone working at the firm with the vacancy. To the extent job information is transmitted across neighbors, this creates a basis for neighborhood effects: higher employment rates in a neighborhood will imply more rapid transmission of information about job openings in a neighborhood. Montgomery (1990a) provides an early formal analysis of job acquisition and social networks. Montgomery (1990b) and Finnerman and Kelly (2003) explicitly consider the implications of differences in labor market conditions across neighborhoods for persistent inequality. Finnerman and Kelly (2003) is noteworthy in developing a model in which the relationship between job opportunities and the density of job referrals in a neighborhood exhibits threshold behavior so that neighborhoods with referral densities below a certain level produce a pool of never employed workers with probability 1. Krauth (2003) provides a model in which neighborhoods are defined by groups of directly and indirectly connected individuals; by embedding this information structure in a Mortensen-Pissarides job matching model, Krauth (2003) shows how small changes in neighborhood composition can induce large changes in unemployment. In very recent work, Calvo-Armengol and Jackson (2003) develop a sophisticated analysis of labor market inequality and social networks for general graph structures to describe the networks. This model provides ways of understanding wage inequality in the context of neighborhoods when understood as social networks. An interesting outstanding question

is whether the findings of authors such as Calvo-Armengol and Jackson (2003) can be enriched with random graph ideas in order to produce a theory of persistent differences in networks.

iv. Stratification and efficiency

A final area of theory concerns the efficiency of allocations of agents across neighborhoods in the presence of neighborhood effects.

One question concerns the efficiency of stratified allocations. To understand this, suppose that a planner faces the problem of allocating n agents in coalitions of size k . Agents are associated with exogenous characteristics x_i and each coalition has a payoff $\hat{v}(S)$ where S denotes the vector of characteristics of the k coalition members. $\hat{v}(S)$ is strictly increasing in S , so the individual characteristics may be thought of as measures of ability or education. Finally, assume that $\hat{v}(S) = \hat{v}(S')$, if S' is a permutation of the elements of S ; this simply means that the order of elements in S does not matter.

As stated, the allocation problem here is essentially the marriage problem studied in Becker (1973), in which the question was what allocation of men and women into couples will produce the highest quality children. In Becker's original formulation, the efficiency of mating the highest quality male with the highest quality female, etc. depended on whether the marriage production function exhibited complementarities between the qualities of the spouses. Extending the definition described in eq. (4), complementarity in $\hat{v}(S)$ means that if one fixes a subset of k elements of S at some value \tilde{x} and partitions $S \setminus \tilde{x} = S_1 \cup S_2$, then the function $\hat{v}(S) - \hat{v}(\tilde{x})$ has the property that

Recent developments in economic theory make it simple to study this problem of efficient stratification in the presence of increased differences. Under mild technical conditions (see discussion in Milgrom and Roberts (1990) and Topkis (1998)),

complementarity is equivalent (under weak technical conditions) to supermodularity of a function. Let \hat{x}^i denote a vector formed by the maximum of each of the corresponding elements of x^i and y^i and let \hat{x}^j denote the vector formed by the minimum of each of the corresponding elements of x^j and y^j . The function \hat{v}^i is supermodular if

$$\hat{v}^i(\hat{x}^i, \hat{x}^j) \geq \hat{v}^i(x^i, x^j) \quad (17)$$

As shown in Durlauf and Seshadri (2003), this condition, when combined with the assumption that \hat{v}^i is permutation invariant, immediately implies that stratification is efficient.

As argued in Durlauf and Seshadri (2003), however, the relationship between complementarity and efficiency of stratification relies on two assumptions beyond complementarity: 1) all coalitions are of equal size and 2) coalition memberships do not alter individual behaviors. Hence, there is no presumption that stratification is efficient in more general contexts even when complementarities are present within neighborhoods. A separate question concerns the efficiency of competitive equilibria in neighborhoods models. This question has been studied in a range of contexts, e.g. de Bartolome (1990), Becker and Murphy (2000) and Bénabou (1993,1996a). These different models produce the common result that there is no guarantee that equilibrium allocations of families across neighborhoods are efficient; further, it is possible to identify conditions under which the level of stratification is higher than is dictated by efficiency.

Bénabou's (1996a) argument is very insightful and follows from an analysis of the total human capital produced by an allocation of families. For a given family, the human capital of an offspring is determined by $\hat{x}^i - \bar{x}^i$, where as above, \bar{x}^i is the average of x in neighborhood n . Notice that $\hat{x}^i - \bar{x}^i = \frac{1}{2}(x^i + y^i) - \bar{x}^i$ where \bar{x}^i equals the fraction of x^i agents in neighborhood n , Bénabou contrasts the equilibrium allocation of families with the allocation that maximizes total human capital, which may be represented as

(18)

subject to the constraint

(19)

which simply means that all agents are allocated to one of the neighborhoods. His analysis shows how the efficiency or inefficiency of the equilibrium level of stratification depends on the interplay of several factors. One factor that matters is the cross partial derivative $\frac{\partial^2 U}{\partial q \partial \theta}$; when this is negative, it implies that the marginal benefit to an

average family from a good neighborhood is higher than that of an average family.¹² This negative cross partial contributes to inefficiency in the equilibrium level of stratification.

A second factor that Bénabou identifies is $\frac{\partial^2 U}{\partial q^2}$; when this second derivative is negative, the marginal value to an individual of higher average quality is decreasing in the level of the average quality, so that inefficiency in the equilibrium allocation of neighborhood memberships can occur for a broader range of (other) features of the model.

Another approach to analyzing the efficiency of equilibrium stratification is via an analysis of willingness to pay. Becker and Murphy (2000) and de Bartolome (1990) provide particularly lucid discussions in the context of a model where the equilibrium allocation of families leads to a mix of high and low quality families in each neighborhood. They show that the equilibrium level of stratification is inefficient if the willingness to

¹²Observe that the sign of this cross-partial derivative also matters for the equilibrium level of stratification in Proposition 4, and indicates why it is possible for stratification to be efficient. Intuitively, even though the effects of changes in neighborhood composition are not internalized, if agents with high θ values exhibit

pay function is concave in τ . Intuitively, when this function is concave, the conditions for the competitive allocation of families imply that the transposition of an τ family from the worse neighborhood and an τ family from the better neighborhood produces a net increase in the willingness to pay when families are aggregated across neighborhoods.

The Bénabou (1993,1996a), Becker and Murphy (2000) and deBartolome (1990) results reflect the externalities implicit in neighborhood effects. No markets exist to compensate families for the positive neighborhood effects they engender; since the effects are not internalized, there is no presumption that the equilibrium allocation of families across neighborhoods is efficient, although it may be, as discussed in footnote 9 above.

The argument as to why inefficient levels of segregation may emerge does not account for the effect of contemporaneous neighborhood allocations on future levels of individual characteristics, which is one of the reasons given by Durlauf and Seshadri (2003) that Becker's findings on the efficiency of assortative matching may break down in some cases. One clear channel where such considerations matter is human capital, which is the focus of an important analysis by Bénabou (1996b). This paper studies efficiency related issues in the context of neighborhoods, human capital accumulation and growth. In Bénabou's model, human capital accumulation is determined in local communities via local finance of education and possible neighborhood-specific spillover effects. From the perspective of individual families, human capital for a child is maximized as the distribution of incomes in the community is shifted to the right. At the same time, the productivity of human capital for adults is determined by the aggregate human capital among all adults in the population. The efficiency of stratification versus integration thus depends on the nature of the dynamic effects of integration versus integration on the distribution of human capital. When economy wide spillovers are strong enough, all families are better off under integration, as affluent families sacrifice some family specific human capital formation for a better economy-wide distribution. This possibility is studied by Cooper (1998) who argues such effects can explain why

greater willingness to pay for high τ neighborhoods because of the production

affluent families are willing to redistribute tax revenues to less affluent districts, providing a positive political economy of educational subsidies.

3. Econometrics

While the theoretical literature on neighborhood effects is now quite extensively developed, far less work has been devoted to the econometric issues that arise in identifying such influences. Nevertheless, the econometrics of neighborhood effects does provide a number of important results for conducting and interpreting empirical work. In this section and in the empirical section, I will focus exclusively on the estimation of neighborhood effects. While many studies examined the cross-section and dynamic properties of neighborhood configurations, the focus of the great majority of these studies has been on developing appropriate ways to measure levels of and changes in economic stratification and racial segregation.¹³ As such, the causal mechanisms have generally not been explicitly addressed, and so I omit a review of the methodological issues and empirical findings in that literature.¹⁴

i. Identification

The basic econometric issues that arise in the study of neighborhood effects may be understood by considering the cross section-regression

(20)

complementarity that is described, then the stratification that is induced is efficient.

¹³Jargowsky (1997) is an excellent overview of economic stratification of communities; Massey and Denton (1993) is the standard reference on residential racial segregation in the United States.

¹⁴Two recent studies, Bajari and Kahn (2003) and Bayer, McMillan, and Rueben (2002), employ behavioral models to study segregation patterns and constitute very promising exceptions to my general assessment of the empirical segregation literature.

where, consistent with earlier notation, \mathbf{z}_i denotes an r -length vector of observable individual characteristics, \mathbf{c}_i denotes an s -length vector of contextual effects and \bar{y}_k denotes the expected value of y_i for members of neighborhood k .¹⁵ This equation is often referred to as the linear-in-means model. An important feature of this model is the assumption that all endogenous effects work through expectations. This assumption is most appropriate when neighborhoods are relatively large, for small neighborhoods, additional complications arise because of the effect of y_i 's behavior on others; see Kooreman and Soetevant (2002) for discussion of estimation for small groups. This model was initially studied by Manski (1993) in a foundational paper and has subsequently been analyzed by Moffitt (2001), and Brock and Durlauf (2001b,2003). I first consider the case where $\mathbf{c}_i = \mathbf{c}_k$ in order to focus on questions of identification that are intrinsic to neighborhood effects as opposed to identification issues that arise because of the endogeneity of neighborhoods.

To understand why identification conditions arise in this model, observe that when beliefs are rational,

$$\bar{y}_k = \mathbf{z}_k \boldsymbol{\beta} + \mathbf{c}_k \boldsymbol{\gamma} \tag{21}$$

In this expression, \bar{y}_k equals the average of the y_i in neighborhood k and appears in the regression because this average is one of the determinants of y_i . Substituting (21) into (20), the individual choices may be expressed in terms of observables via

$$y_i = \mathbf{z}_i \boldsymbol{\beta} + \mathbf{c}_i \boldsymbol{\gamma} + \mathbf{u}_i \tag{22}$$

¹⁵ Relative to earlier notation, I employ the subscript k rather than i since observations in a cross-section will be drawn from different neighborhoods, although more than one observation may be drawn from the same neighborhood.

Equation (22) summarizes the empirical implications of the linear-in-means model. The identification problem may thus be thought of as asking whether one can recover the structural parameters in (20) from the coefficients in (22).

Since (22) contains K regressors (and associated coefficients) whereas there are only K coefficients in (20), it appears that one can recover the structural parameters from a regression of y onto the various regressors, in fact the parameters of (20) are overidentified. However, this conclusion fails to account for possible collinearity between the components of (22); collinearity may potentially arise because of the presence of γ and δ in the equation. For example, following the case originally studied in Manski (1993), suppose that $\gamma = \delta$. In this case, the modeler has no basis for distinguishing between contextual and individual effects. When this condition holds, then there are only $K-1$ linearly independent regressors in (22), the associated coefficients for these linearly independent regressors are identified, but they cannot be uniquely mapped back into the K structural coefficients in (20); identification of the structural parameters in (20) thus fails. Manski (1993) has termed this failure of identification the reflection problem, to capture the intuition that the identification problem relates to distinguishing the direct effect of δ on an individual versus its indirect effect as “reflected” through the endogenous effect generated by γ .

The reflection problem as originally formulated in Manski (1993) assumes that there is a one to one relationship between individual effects and contextual effects. As such, the reflection problems may be understood as describing identification limits when there is no prior information available to allow one to identify individual and contextual influences that are distinct from one another. Such information can allow for identification.¹⁶ Brock and Durlauf (2001b) provide the necessary conditions for identification with the following theorem.

¹⁶This breakdown of identification in the absence of prior information that restrict what variables directly influence behavior is also found in the rational expectations literature where the analogous problem concerns differentiation of the effects on a given variable of expectations of future variables from the direct effects of various current and lagged variables. Wallis (1980) provides a general treatment and Sargent (1976) provides a classic example of how the lack of prior information renders very different

Proposition 5. Identification in linear individual-level models with neighborhood effects

Identification of the parameters in the linear-in-means model (20) requires

- i.* The dimension of the linear space spanned by elements of \mathbb{R}^n is n .
- ii.* The dimension of the linear space spanned by the elements of \mathbb{R}^n is at least n .

The identification problem as developed here is in certain respects unique to linear models. Identification breaks down when \mathbb{R}^n is linearly dependent on the other regressors in (20); Manski's nonidentification result specifically occurs because \mathbb{R}^n is a linear combination of 1 and \mathbb{R}^n . Linear dependence of this type will typically not arise when individual behaviors depend on other moments of the neighborhood behavior. This is most easily seen if for a nonlinear-in-means neighborhood model of the form:

$$\mathbb{R}^n = \mathbb{R}^n + \mathbb{R}^n \tag{23}$$

where \mathbb{R}^n is invertible and \mathbb{R}^n . The self-consistent expected average choice in a neighborhood is determined by

macroeconomic theories observationally equivalent. As suggested by Binder and Pesaran (2001) and Brock and Durlauf (2001b), the timing of neighborhood effects has important implications for identification; Brock and Durlauf (2001b), for example show how identification may hold for a dynamic version of (21) if contextual effects occur with a one period lag, i.e. \mathbb{R}^n depends on \mathbb{R}^n and \mathbb{R}^n . However, this possibility has yet to be systematically explored and, of course, this timing assumption needs to be justified.

where $\hat{\beta}$. As is clear from (24), the collinearity problem that can arise in the linear in means model cannot arise here, except for hairline cases. The logic of this example is in fact quite general. Brock and Durlauf (2001a) show that in the space of twice differentiable functions $\hat{\beta}$, identification will fail only for a measure zero set of $\hat{\beta}$. More important for empirical work, this argument also implies that identification will hold for nonlinear probability models of choices, for example binary or multinomial choice models of the type described above or duration data models with neighborhood effects. Brock and Durlauf (2001b,2003) provide a set of results on these cases.

One additional difference between linear and nonlinear models of neighborhood effects concerns the interpretation of estimated models that ignore endogenous effects. In the case of linear models, it is possible to interpret linear models without endogenous effects as reduced forms, if $\hat{\beta}$. (If this is not the case, then $\hat{\beta}$ will represent a vector of omitted variables in the reduced form.) As pointed out by Manski (1993), this means that nonzero parameters associated with $\hat{\beta}$ are necessary for these contextual effects to be present. In contrast, for nonlinear models such as the binary choice model, this will not be the case since the omission of the endogenous effect does not produce an associated reduced form.

ii. Self-selection

The assumption that $\hat{\beta}$ is unappealing. The reason for this is immediate given our earlier discussion of equilibrium neighborhood configurations: one does not think of residential neighborhoods as exogenously determined; hence, it is natural to expect that there is a relationship between

neighborhood choice, i.e. \hat{y}_i and unobserved heterogeneity embodied in ϵ_i . This is, of course, a form of self-selection bias whose presence has motivated a vast econometric literature; see Heckman (2001) and Manski (1995) for valuable overviews.

How may self-selection be addressed and how does self-selection affect identification? To answer these questions, it is useful to rewrite the behavioral equation as

$$y_i = \beta_0 + \beta_1 x_i + \epsilon_i \quad (25)$$

where ϵ_i by construction. Following the classic approach to selection developed by James Heckman (cf. Heckman (1979)), consistent estimation of (25) requires constructing a consistent estimate of ϵ_i and including this estimate as an additional regressor in (25); a key insight of Heckman (1979) is that once this is done, (25) may be estimated by ordinary least squares.

To see how this may be applied to the neighborhood context, suppose that for individual i , a choice has been made across J possible neighborhoods. Neighborhood j is associated with a latent “quality” measure z_{ij} which is determined by

$$z_{ij} = \gamma_0 + \gamma_1 x_i + \eta_{ij} \quad (26)$$

where x_i is a vector of observable characteristics of i that influence the quality assigned to neighborhood j and η_{ij} denotes an unobservable individual-specific quality term. Individual i is assumed to choose to reside in the neighborhood with the highest z_{ij} . Assume that $\eta_{ij} | x_i \sim \text{IID}(0, \sigma^2)$ and $\eta_{ij} | x_i \perp \eta_{ik} | x_i$. Then, for certain parametric assumptions on the densities of η_{ij} and η_{ik} , there exist analytic expressions for \hat{z}_{ij} that may be shown to be proportional to some $\pi_j(x_i)$ where the function π_j is determined by the parametric error assumptions and

the parameter vector β is estimable from a multinomial choice model of neighborhoods; denote this factor of proportionality as λ . A behavioral equation with neighborhood effects may thus be constructed as

$$\ln \lambda = \beta'x + \gamma + \epsilon \quad (27)$$

and so it is possible to estimate neighborhood effects in the presence of endogenous neighborhood choice. Explicit examples of this are given in Brock and Durlauf (2003) and Ioannides and Zabel (2002b); the appendix to this chapter describes the Brock-Durlauf approach in detail.

Self-selection corrections turn out to have important implications for identification. To see this, consider two cases. First, suppose that the decision to join a neighborhood only depends upon x . In this case (27) is now a nonlinear in means model (since λ is almost certainly nonlinear given the fact that the neighborhood choice decision is made among a set of discrete alternatives) and is thus identified outside of pathological cases, as shown in Brock and Durlauf (2001b). Second, suppose that x consists of elements of x_1 and x_2 . Since λ is nonlinear, $\ln \lambda$ will be linearly independent of x_1 even if x_2 is linearly dependent on x_1 . As such $\ln \lambda$ is an additional individual level regressor whose group level analog does not appear in the behavioral equation (20). This means that by Proposition 5, identification may be achieved.

This analysis of self-selection as a road to identification for neighborhood effects suffers from its dependence on parametric assumptions concerning the distribution of errors in (20) and (26). The particular parametric assumption made here is not essential. What appears to be more difficult is the development of a way of employing semiparametric selection corrections to facilitate identification. The analysis of identification without strong parametric assumptions is an important next step in this research. One new approach to dealing with self-selection has been proposed by Krauth (2002). In this approach, one uses information on the degree of self-selections into

groups based on observables to construct bounds on the magnitude of neighborhood effects by using this information to bound the degree of self-selection on unobservables.

iii. Unobservables and sibling data

The estimation of neighborhood effects in the presence of self-selection is an example of the more general problem of accounting for unobservable individual and neighborhood controls. Brock and Durlauf (2001b) propose using panel data methods to eliminate neighborhood level fixed effects; identification conditions are analyzed in great generality by Graham and Hahn (2003). Bayer and Timmins (2002) propose a strategy for accounting for unobservables when more than two neighborhoods are present based on techniques for dealing with unobservable attributes in models of consumer demand which, although not yet formally incorporated into an identification analysis, appears promising.

The most active area of work that has attempted to control for unobservables has focused on the use of siblings data. Aaronson (1998) and Plotnick and Hoffman (1999) do this in a regression context. Following the discussion in Aaronson (1998), consider a pair of siblings i in a given family f . Ignoring endogenous effects, individual behavior may be described by

$$y_{it} = \alpha + \beta x_{it} + \gamma + \epsilon_{it} \tag{28}$$

In this expression, γ denotes an unobservable family effect. Aaronson proposes eliminating this unobserved family characteristic by differencing outcomes between siblings and estimating

$$y_{it} - y_{jt} = \beta(x_{it} - x_{jt}) + \epsilon_{it} - \epsilon_{jt} \tag{29}$$

As such, this approach exploits the standard technique in panel data studies of eliminating fixed effects through differencing. One limitation to this method is that it does not

address the self selection issue. The reason for this, following the logic associated with (26), is that the component of the regression error that is associated with self-selection can differ across siblings if the siblings are raised in the same neighborhood at different times; the information revealed by the parent's choice of neighborhood will differ according to the parent characteristics as well as the characteristics of all neighborhoods in the choice set.

A second approach to employing sibling data to uncover neighborhood effects is proposed by Solon, Page, and Duncan (2000). In this approach, one compares the covariance of outcomes for siblings in the same community and a pair of unrelated individuals in the same community and uses the values to draw inferences about the role of neighborhoods versus families as determinants of individual outcomes. To compute the relationship, one may work with (28) with ϵ_{it} omitted (the correlation approach has the virtue that it does not require ϵ_{it} or ϵ_{jt} to be observable); I work with covariances rather than correlations for ease of exposition. Two assumptions are imposed on this model. First, members of a common family i possess identical individual effects, i.e. $\alpha_{it} = \alpha_{jt}$. Second, no account is made of endogenous effects. Under these assumptions, the covariance of two siblings in a common neighborhood is

$$\text{Cov}(y_{it}, y_{jt}) = \sigma_{\epsilon}^2 + \sigma_{\alpha}^2 \quad (30)$$

whereas the covariance of two unrelated neighbors is

$$\text{Cov}(y_{it}, y_{jt}) = \sigma_{\epsilon}^2 \quad (31)$$

These two equations provide a way of testing the null hypothesis that neighborhood effects are zero. If $\sigma_{\alpha}^2 = 0$, then it must be the case that $\text{Cov}(y_{it}, y_{jt}) = \sigma_{\epsilon}^2$; this converts to a correlations inequality when the variances of the different types are equal.

The correlations approach to neighborhood effects is an interesting alternative to regression approaches. However, there are two limitations to this approach. First, the comparison of correlations provides relatively little information about the nature and magnitude of neighborhood effects. While the inequality $r_{ij} \leq r_{ij}$ is implied by the absence of neighborhood effects, it can also hold when neighborhood effects are present. Hence, comparisons of correlations provide only a weak testing framework for neighborhood effects. Further, one cannot make firm statements about the magnitude or nature of neighborhood effects and so the assessment of policies to change neighborhood composition is not possible.

A second limitation of the method is that it imposes strong assumptions on the determinants of behaviors within families. The assumption that all individual-specific characteristics are identical within a family, i.e. $\epsilon_{ij} = \epsilon_{ik}$, is quite strong and would appear to rule out family background variables such as income. Perhaps more important, the model assumes that differential treatment of siblings within families may be modeled by an uncorrelated and identically distributed error term. However, this assumption is problematic as it ignores such possibilities as birth order effects in childrearing, etc. Assumptions of this type have proven to be a serious problem in other contexts. For example, the use of twins data versus other siblings in heritability studies has been strongly faulted for failing to account for childrearing differences; for a discussion of this issue and many other interpretation problems with studies that are based on correlations of this type see Goldberger and Kamin (2002). Hence, it seems important to investigate whether the error assumptions in these models are empirically appropriate.

These limitations do not mean that correlation analyses are without interest; they are clearly useful as data summaries and suggestive of the relative importance of neighborhoods in variation across individuals. At a minimum, the correlations approach is promising as a way of providing data summaries which can help guide theoretical modeling.

4. Empirical studies

i. Ethnography

One important source of evidence on neighborhood effects is ethnographic studies, which are an important empirical tradition in the sociology literature. This type of evidence is underutilized in economics. While such studies obviously cannot be subjected to the sorts of empirical criteria one associates with econometrically-based empirical work, such evidence is nevertheless corroborative of quantitative studies. Further, to the extent that quantitative studies require identification assumptions of the type described in Section 3, it is possible that support for such assumptions may be taken from this source.

In the 1960's, ethnographic studies produced early, albeit controversial evidence of neighborhood effects. Lewis (1966) explicitly describes how a culture of poverty existed in poor Puerto Rican communities and ascribes a number of social ills to the norms that exist in those places. Important recent ethnographic work includes Anderson (1990,1999) who has documented how poor inner city communities have developed codes of conduct that are conducive to high levels of violence; Liebow (1967) draws related conclusions in the context of African Americans. Duneier (1992) studies the social relationships that develop at a neighborhood restaurant and illustrates the importance of social interactions even in this very particular context.

From the evidentiary standards of economics, ethnographies are generally regarded as suspect because of their particularity and because of fears that the observer's prejudices are determining the findings. While these are certainly valid concerns, they do not imply that such studies are without insight. In particular, it seems that such studies can facilitate decisions on how to define neighborhoods in more quantitative studies and may also provide some insight into the choice of control variables.

ii. Experiments

A second source of evidence on neighborhood effects is controlled experiments. The social psychology literature is filled with studies of how groups affect individual

perceptions; see Aronson (199) for a survey. The value of this literature is reflected in a classic experiment due to Sherif et al (1961) known as the Robbers Cave experiment. Sherif and coworkers brought a group of middle class teenage boys to a camp at Robbers Cave Oklahoma. For the first two weeks of the camp, interactions between the boys were not guided by the experimenters. After two weeks, the boys were randomly divided into two groups, the Eagles and the Rattlers. (The one exception to randomness was that friendships observed by the experimenters were broken up.) After these groups were formed, competitive activities such as games were organized between the groups. Sherif et al (1961) document how these groups quickly became the source of strong feelings of identity. Members of each group developed negative stereotypes about members of the other group in terms of intelligence and honesty. This experiment makes clear how neighborhoods, even under random assignment, can influence cognition and behavior towards others.

The new experimental economics has conducted many experiments that have addressed social influences on behavior, although relatively few have focused on the sorts of influences that have motivated the neighborhood effects literature. One exception is a recent study by Falk and Ichino (2003). In this experiment, individuals are organized randomly into groups of different sizes and assigned to fill envelopes, with the compensation for the activity set independent of the group's output. In some cases, individuals worked alone. In others, individuals worked side by side. For some groups, members were informed of the different productivity levels ascribed to previous groups. In others, effort was made visible, i.e. each worker could see how hard others were working. Falk and Ichino (2003) found that there was a strong correlation in effort within pairs of workers and that the effort level of groups was strongly influenced by information about the performance of others.

While it is difficult to translate these types of findings into implications for the types of neighborhood effects that are conventionally studied, experimental evidence does strongly buttress the general claim that social influences matter.

iii. Econometric studies with observational data

Since the pioneering work of Datcher (1982) there is now a rich empirical literature designed to assess neighborhood effects. Table 2 summarizes 25 empirical studies taken from this very large literature; the selection of studies reflects an effort to explore the range of empirical analyses that have appeared as well as some bias towards more recent work.¹⁷

These studies cover a wide range of individual outcomes, definitions of neighborhoods and neighborhood effects, as well as a range of econometric methodologies. Nevertheless, a few general conclusions may be drawn.

First, the bulk of empirical studies of neighborhood effects find evidence of their presence. Of course, this may reflect a publication bias against negative results. However, a researcher with a strong prior that neighborhood effects are present would not have these beliefs strongly changed by the body of observational studies.

Second, the neighborhoods effects literature is highly unsystematic in its choice of neighborhood variables by which to measure effects. One finds variables such as the median income in neighborhood, the percentage of professional and managerial workers among all workers, behaviors of neighbors, etc. Further, the choice of neighborhood variables is rarely motivated by theory. (This is less true for studies that explore information transmission such as Bertrand, Luttmer, and Mullainathan (2000) or Drewianka (2003) which construct neighborhood variables to reflect contact probabilities, but even here the variables are not derived from explicit search models.). Taken together, the evidence of neighborhood effects in this literature is largely a black box, i.e. it is difficult to translate the findings of the papers into specific microeconomic mechanisms. To be clear, this probably largely reflects the failure of the theoretical neighborhoods literature to provide more guidance on the generative mechanisms that produce neighborhood effects. As discussed in Section 1, neighborhoods models typically assume certain neighborhood variables appear in the specification of preferences, technologies or beliefs and analyze the consequences.

¹⁷Jencks and Mayer (1990) is a standard survey of earlier empirical work on neighborhood effects. See Dietz (2002) for a useful recent overview.

Third, there has been relatively little systematic attention to questions of model uncertainty. While many empirical studies of neighborhood effects check for the robustness of results along a few dimensions, this has not been done systematically. One exception is Ginther, Haveman, and Wolfe (2000), who explore the robustness of estimated neighborhood effects for alternative choices of individual level controls. They find that the magnitude and statistical significance of neighborhoods effects is very sensitive to the choice of individual level controls. One problem with this conclusion is that it does not address the question of what information is contained when one aggregates across model specifications. Put differently, one is not so much interested in the distribution of neighborhoods effects across model specifications but in estimates of neighborhood effects that do not condition on a particular model when model uncertainty is present. Brock, Durlauf and West (2003) address this general problem drawing on Draper (1995) and other recent research in statistics. Sirakaya (2003) is unique in applying model averaging ideas to the analysis of neighborhood effects. Her analysis finds strong evidence of endogenous neighborhood effects that accounts for the information contained in each of a large set of potential models of recidivism by ex-felons on probation. This work has important implications for future work on neighborhood effects.

Fourth, relatively little systematic attention has been given to the identification problems that exist between endogenous and contextual effects. This is true at two levels. First, most studies fail to address the differences between these effects and the attendant implications for the models that are estimated. Second, to the extent that both effects are included in the same model, one does not see discussion of what assumptions are being implicitly made to allow for identification. However, there are some recent exceptions to this; examples include Drewianka (2003), Ioannides and Zabel (2002a,b), Minkin (2002), and Sirakaya (2003). These studies have been able to parse the two types of effects and typically find both are present.

Fifth, efforts to control for unobserved heterogeneity have had varying effects on the analysis of neighborhood effects. From the perspective of unobserved family fixed effects, findings of neighborhood effects appear to be mixed, see for example Aaronson (1998) who finds that estimates of neighborhood effects are robust to allowing for family

effects in siblings data whereas Plotnick and Hoffman (1999) do not. More recently, Aizer and Currie (2002) find evidence of neighborhood effects on the utilization of publicly funded prenatal care that is robust to the incorporation of spatial fixed effects. My own reading of the literature is that this form of unobserved heterogeneity has not been shown to be of first order concern in interpreting the existing empirical studies.

Sixth, the robustness for neighborhood effects findings to controls for self-selection appears to depend on the method employed. Evans, Oates and Schwab (1992) appear to be the first to use instrumental variables to control for self selection. This study employed a measure of the percentage of students in a school who are disadvantaged to evaluate how neighborhood effects influence high school drop out and teen fertility rates. The measure of school level disadvantage is instrumented with metropolitan area unemployment, college completion and poverty rates and median income. Evans, Oates and Schwab (1992) find that although neighborhood effects are statistically significant when treated as exogenous, the coefficient is statistically insignificant (with a change of sign) when it is instrumented. Foster and McLanahan (1996) employ a similar instrumental variables strategy and also find that instrumental variables estimates fail to find evidence of neighborhood effects. These results contain an important cautionary message. However, it is unclear exactly how to interpret them. It is unclear that metropolitan area instruments can account for neighborhood effects that occur conditional on a metropolitan area; a point explicitly acknowledged by Evans, Oates and Schwab. For example, the effect of neighborhood characteristics on aspirations may depend on the college completion rate among adults in a neighborhood relative to the metropolitan area. There is also a question as to the validity of the instruments. It is unclear why individual behavioral decisions do not directly depend on metropolitan area characteristics and hence why they are excluded from the original behavioral equation.

It is also possible to find cases where instrumental variables estimates of neighborhood effects do not reduce estimated magnitudes. Rivkin (2001) uses similar instruments to Evans, Oates and Schwab (1992) in a study of high school-based effects and, unlike Evans, Oates and Schwab (1992) finds that coefficient estimates using instrumental variables are larger than those in a baseline model. Rivkin (2001) does not interpret this as evidence of neighborhood effects, but argues this indicates that the types

of instruments employed may, if anything, exacerbate endogeneity bias. However, there is no reason why this must be so and his results are equally consistent with the interpretation that neighborhood effects matter.

In contrast, evidence of neighborhood effects is not diminished in the one study that employs an explicit self-selection correction of the type discussed in Section 3, Ioannides and Zabel (2002b). This paper finds strong evidence of neighborhood effects on housing demand even after explicitly modeling the neighborhood choice decision. While this result does require parametric assumptions on various model errors, it is suggestive that if self-selection is used as a source of information, it may prove valuable. Overall, this paper may be regarded as the best example of the integration of econometric methods into empirical work that has appeared in the neighborhoods literature.

Taken as a whole, these six observations illustrate that the empirical literature based on statistical analysis of observational data provides only limited support for the importance of neighborhood effects.

iii. Correlation studies and sibling data

Several authors have used the correlation approach described in Section 3 to uncover neighborhood effects using sibling data. These studies have generally found little evidence of neighborhood effects. Solon, Page and Duncan (2000) analyze correlations in educational attainment using geocode data in the Panel Study of Income Dynamics (PSID) and find after controlling for some basic background characteristics, the residual correlation in neighboring children's outcomes is on the order of .1. Page and Solon (2001) study adult incomes of females and find that the correlation between neighboring girls is only 1/3 that of sisters. They further find that much of the correlation between neighbors is driven by income differentials between urban and non-urban areas. Duncan, Boisjoly, and Harris (2001) study correlations between siblings, friends, schoolmates (members of a common grade at a school), and neighbors for measures of high school achievement and delinquency. They find that sibling correlations are much higher than the others, and that friendship correlations are much larger than neighbor and

schoolmate effects. While these findings suggest a larger role for family background than neighborhood characteristics as determinants of behavior, they do not directly address the economic significance of neighborhood effects nor do they provide much insight into the effects that policies that alter neighborhood membership would have, for reasons described in Section 3.

iv. Quasi-experiments

An important alternative to the use of observational data such as the Panel Study of Income Dynamics is the use of data in which government interventions into the residential choices of individuals are used to assess the effects of neighborhoods. Such interventions are examples of what in economics are known as “quasi-experiments,” the idea being that the intervention at least partially defines groups of individuals who have or have not randomly received a treatment (drawing an analogy from biostatistics), in this case, a new group membership, thereby allowing for the measurement of group effects.

One example of such an intervention is the Gautreaux program. In 1967, Dorothy Gautreaux led a group of plaintiffs to sue the Chicago Housing Authority, claiming that placement of poor families in public housing in poor neighborhoods constituted a form of discrimination. A consent decree between the plaintiffs and the CHA resolved the case and produced a housing program that in essence assigned one group of families to other parts of Chicago and another to suburban communities outside the city. Sociologist James Rosenbaum has organized and conducted interviews with families that had participated in the program in order to determine the effects of living in suburban communities on poor families. In a series of studies (cf, Rosenbaum and Popkin (1991), Rosenbaum (1995)), he showed that families living in suburbs experienced substantially better socioeconomic outcomes along a number of dimensions. As described in Rosenbaum (1995, pg. 242), these differences are particularly pronounced with respect to outcomes for children. For example, the percentage of college attendees among children whose families moved to suburbs was 54% whereas the percentage for children whose moves kept them in the city of Chicago was 21%; when one considers only 4-year

colleges the attendance rates are 27% versus 4%. While these data suffer from some self-selection problems that render their causal interpretation problematic (an issue well understood by Rosenbaum), they are extremely suggestive and have greatly helped to stimulate research on neighborhood effects.¹⁸

The Gautreaux findings are important as they represent an early effort to provide evidence of neighborhood effects based on an external intervention into neighborhood configurations. However, the various ways in which the allocation of families across new locations in Chicago versus suburbs was determined by unobserved characteristics of the families in the program, clouded any policy inferences one could draw from the program. One consequence of this is an extended effort by the Department of Housing and Urban Development to conduct an experiment in altering neighborhood memberships that is more conducive to causal inferences. This program, the Moving to Opportunity demonstration (MTO) has been underway in five cities, Baltimore, Boston, Chicago, Los Angeles and New York since 1994. The demonstration provides housing vouchers to a randomly selected group of families; within this subsidized group, families in turn were randomly allocated between unrestricted vouchers (users are known as the Section 8 group) and vouchers that could only be used in census tracts with poverty rates below 10% (whose users are the Experimental group).¹⁹

Recent evaluations of the effects of the vouchers include Hanratty, McLanahan, and Pettitt (2001), Katz, Kling and Liebman (2001), Leventhal and Brooks-Gunn (2001,2002), Ludwig, Duncan, and Hirschfield (2001) and Rosenbaum and Harris (2001). These assessments reveal a number of interesting findings. For children, there

¹⁸Rosenbaum's analyses compare families that were moved to alternate public housing in Chicago to families that stayed in the suburbs; those that moved and then returned to Chicago are not included. This means the sample of suburban families differs from a random selection of families in that it consists of those families who were willing to forgo the benefits of the city (proximity to family and friends, etc.). Such families might well tend to have parents who place an unusually high value on economic achievement, so the success of their offspring, for example, might be due to this latent variable and not the suburban environment per se. While the differences in outcomes may be due to neighborhood effects rather than the self-selection of more "ambitious" families into suburbs, one simply cannot determine this from the data.

¹⁹See Goering (1999) for a detailed description of the MTO demonstration and Goering, Feins, and Richardson (2002) for an overview of MTO findings.

appear to have been impressive gains for both Section 8 and MTO movers along several dimensions. Katz, Kling, and Liebman (2001) conclude in the case of Boston area families that children in Section 8 and Experimental groups exhibited substantial reduction in behavioral problems (the index of behavioral problems they use is about 30% lower for the two groups than those in the control group), and that Experimental group children exhibit lower incidences of asthma attacks (approximately 50% lower than either Section 8 or control children) and injuries requiring medical attention (Experimental group children exhibited injury rates that are over 50% lower than the control group and about 30% lower than the Section 8 group). Leventhal and Brooks-Gunn (2002) also find that behavioral problems were reduced for children in the Section 8 and control groups for New York City families; however, they failed to find evidence of improved health outcomes. Leventhal and Brooks-Gunn (2002) also found some evidence that mental health for boys was improved by moving to better neighborhoods. Ludwig, Duncan and Hirscheffeld (2001) find evidence using data for families in Baltimore that neighborhood moves reduce incidents of juvenile crime, finding that moves from high to low poverty neighborhoods reduce juvenile arrests for violent crimes by a factor on the order of 30% to 50%. Rosenbaum and Harris (2001) find for the Chicago demonstration that economic benefits for household heads, with employment rates for Section 8 and MTO families rising from 29.3% and 24.5% to 42.9% and 46.3% respectively.

In contrast to children's outcomes, the effects of the MTO demonstration on adults are more mixed. Hanratty, McLanahan, and Pettitt (2001) find evidence that hours worked among Los Angeles families rose substantially for Section 8 and Experimental families; interestingly, the increase is over 35% larger for the Section 8 families. Leventhal and Brooks-Gunn (2001) find that substantial decreases in depression (more specifically, depressive behaviors) for Experimental families versus control families; Section 8 families exhibit no improvement. In contrast, Katz, Kling, and Liebman (2001) find little effect of either type of voucher on adult economic outcomes. Goering, Feins, and Richardson (2002) also report evidence from a cross-site survey that finds no evidence of voucher effects on adult welfare use or labor market activity.

As important as the MTO demonstration is, there are limitations to the information it has provided. First, the evidence thus far only describes how the vouchers have benefited those who have employed them. Relatively high percentages of eligible families have failed to use the vouchers; for the Experimental group, the percentage of eligible families using the vouchers ranges from 34% to 61% across cities (Goering, Feins, and Richardson (2002)). At best (and to be clear this is very carefully discussed by researchers involved with MTO), one cannot extrapolate the findings to the broader population of the poor. Second, one needs to recognize that much of the benefits of the programs may be attributable to the increase in income associated with voucher eligibility as opposed to the shift in neighborhoods per se. The improvements one observed between families that employed vouchers with neighborhood poverty restrictions are much less dramatic when compared with families who were given unrestricted vouchers (which is unsurprising, of course since agents with more options should over all be better off) as opposed to those who did not receive vouchers.²⁰ Third, it is impossible to determine what aspects of the different neighborhoods led to improved outcomes. To give one example (one that is discussed by Katz, Kling and Liebman (2001)) the reductions in asthma rates may be due to improvements in housing quality (asthma is strongly associated with rat infestations) and nothing about the neighborhood per se. The divergence in the impact of vouchers on children versus adults is further suggestive that one needs to be very careful in drawing causal inferences on particular neighborhood effects from the MTO studies. Finally, there is a question of generalizability. Moving large numbers of poor families to more affluent communities will induce general equilibrium effects in terms of the location decisions of other families, the ability of schools in these neighborhoods to provide needed services, etc. One can easily imagine that the commitment of affluent families to public schools would be ended by a massive influx of poor families into their communities. Hence, one cannot simply assert that the effects of this program will be replicated if it is implemented on a wide scale; a point forcefully made in Sobel (2002). For these reasons, one cannot blithely use the MTO

²⁰Rosenbaum and Harris (2001, pg. 336) find, for example how among MTO movers, the percentage that said the condition of their housing is good or excellent increased from 33.9% to 80.6% after moving.

evidence to advocate large scale housing relocation programs as an antipoverty policy, an error one finds in Fiss (2000), for example.

While the Gautreaux program and MTO demonstration directly address the sorts of neighborhood effects that have motivated the neighborhood literature, other studies of quasi-experiments have focused more on uncovering evidence of social interactions in more restricted contexts. One context is that of college roommates, where it is straightforward to identify cases of random assignment. Sacerdote (2001) examines the effects of freshmen roommate assignments and finds they have a substantial effect on academic effort. Kremer and Levy (2003) find significant peer group effects on alcohol use. These studies employ very “clean“ data compared to the larger quasi-experiments; for example, there is no issue of the take up rates that appears in residential neighborhood experiments. On the other hand, there is some question as to the applicability of the findings in roommate contexts to broader neighborhood notions since rooming contexts impose especially high levels of contact. Further, roommates create “interference” in individual choices; one can easily imagine that one low effort roommate can make it harder for others to study.

v. Aggregate studies

A final group of studies of neighborhood effects has focused on aggregate level data. Glaeser, Sacerdote and Scheinkman (1996) find that crime rates across cities and precincts are both far more variable than would be predicted if individual decisions were interdependent. Topa (2001) studies interactions across physically contiguous neighborhoods in Chicago. Specifically, he estimates a nonlinear regression model in which the unemployment rate in one neighborhood is allowed to depend on the unemployment rates of adjacent communities; the nonlinear relationship is structural in the sense that it is derived from a stochastic process designed to model information transmission across neighborhoods. Topa (2001) finds spatial interdependences in neighborhood level unemployment rates are present and quantitatively important.

One problem with this approach to studying neighborhood effects is that it equates evidence of such effects with correlation in behavior across agents. In the case of Glaeser, Sacerdote, and Scheinkman (1996) correlation in the unobserved components to individual behavior will increase the variance of sample means; for Topa (2001) correlated unobservables in local neighborhoods can produce correlation in the levels of unemployment. This sort of correlation does not necessarily affect the interpretation of individual level regressions.

An intriguing recent development in using aggregate data to uncover neighborhood effects is due to Glaeser, Sacerdote and Scheinkman (2002) and exploits a social multiplier property of peer effects. A social multiplier exists in models with endogenous interactions because feedback effects between agents imply that the effect on an individual of a change in a private variable will be lower than the effect on an individual if every member of his neighborhood experiences the same change. Glaeser, Sacerdote and Scheinkman (2002) exploit this to consider how the response of an outcome variable to some control changes at different levels of aggregation; specifically, they argue that the elasticity of an individual's behavior to a change in an individual specific control should be smaller than the elasticity of a change in the group average behavior to the average of the control. This argument is used to show the presence of social interaction effects for cases ranging from the grade point averages of Dartmouth students considered by rooming group and dormitory and crime rates at the city, state and national level. This appears to be a very promising approach.

vi. Identifying neighborhoods

All of the empirical studies that I have discussed take a particular neighborhood structure as known *ex ante*. This common assumption is necessitated by data limitations, i.e. the definition of neighborhood is determined by the information available in the data set and not by any substantive criteria. Important data sets such as the General Social Survey (GSS) link individual data to metropolitan areas, which clearly are far broader than the notions of neighborhood suggested by theories of social interactions; even a data

set such as the PSID has limitations as one cannot identify candidates for neighborhoods than census tracts. I am unaware of any systematic evaluation of the question of the appropriate units for measuring neighborhoods. Such an analysis is clearly important if one wants to engage in policy evaluation or other analyses which require the ability to translate regression coefficients into model parameters.

While the size of units has not been addressed, there has been some research on the appropriate social notion of a neighborhood. As Akerlof (1997) has argued, theories of social influences naturally lead one to contemplate a social space in which individuals are situated. As emphasized there and in Akerlof and Kranton (2000), the ways in which individuals are influenced by others is strongly influenced by the set of self-perceptions that constitute identity, so that two individuals in the same physical space may experience very different social interactions.

Despite the importance of the question of neighborhood definition, there has been relatively little research on this question. Aizer and Currie (2002) is an unusual exception to this in that their analysis of the role of neighborhood effects allows ethnicity and residential proximity to jointly define the relevant interaction group for individuals (in this case less educated mothers and their use of publicly funded prenatal programs). The one systematic study of the social space in which to define neighborhoods is Conley and Topa (2002).

Conley and Topa (2002) focus on correlations in unemployment across neighborhoods in Chicago. Their analysis focuses on 75 "Community Areas" in Chicago which previous researchers have constructed from the 866 Chicago census tracts with the explicit intent of identifying areas with a common sense of community, etc. Conley and Topa use unemployment data from these community areas and construct spatial correlation functions to understand how the neighborhoods covary. However, rather than simply employ Euclidean distance to construct these functions, they construct four different notions of distance between neighborhoods: 1) physical distance, which refers to the distance between centroids of neighborhoods, 2) travel time distance, which refers to the time necessary to travel via public transportation from the center of one neighborhood to another, 3) racial and ethnicity distance, which is a measure of the similarity in ethnic composition between two neighborhoods based on 9 categories, 4) occupational

distance, which is a measure of occupational similarity between two neighborhoods using 13 occupational categories. Conley and Topa find that the measure of ethnic distance seems to be the most salient dimension along which neighborhoods exhibit spatial correlation. Once one controls for racial and ethnic distance, one finds little additional spatial correlation. However, none of the metrics appear to explain much once one accounts for racial and occupational composition within a neighborhood. Hence, Conley and Topa conclude that it is likely that social interactions, if any, occur at a lower level of aggregation. While their results are not decisive, their methodology is an important advance.

5. Additional evidence on neighborhood effects

Evidence of neighborhood effects may also be found in a number of related literatures.

i. Classroom effects

A range of authors have studied the effects of classroom composition on educational outcomes. These studies are of interest both from the perspective of the “neighborhoods” defined by classrooms and because one of the reasons why the allocation of families across residential neighborhoods may matter is because of such classroom effects. An additional virtue of these studies is that in certain respects classroom composition is more amenable to quasi-experimentation, as will become apparent below.

Within the economics literature, one of the early and still most important studies is Henderson, Mieszkowski, and Sauvageau (1978). This study employed an unusually detailed data set comprised of French speaking students in Montreal for whom panel data was collected to measure language and mathematics skills; the data set allowed for a range of controls for family, teacher, and schools effects. In this study, mean IQ of

classmates was employed as a measure of peer group effects. Henderson, Mieszkowski and Sauvageau (1978) found that peer effects are clearly present, but that these effects are concave in the sense that the marginal effect of an increase in mean classroom IQ is decreasing in the level of the mean IQ. This result is important, as Henderson, Mieszkowski, and Sauvageau (1978) argue, because it suggests that tracking classes by IQ is inefficient if the objective of the school is to maximize average educational achievement. This work also makes clear how there are distributional consequences from classroom mixing as more able students are hurt and less able students are helped.

Hanushek, Kain, Markman, and Rivkin (2001) conduct a similar analysis using data from the University of Texas Dallas Schools Project, which has compiled a complete data set of Texas students based on tracking third grade students in 1992 for four years. Measuring peer effects by past test score performance by students in the same grade, Hanushek, Kain, Markman, and Rivkin (2001) find that mean test score performance by others in the same grade improves student performance; unlike Henderson, Mieszkowski, and Sauvageau (1978), some evidence of nonlinearity appears.

Other studies have focused on the identification of peer effects through government programs that affect classroom composition. Angrist and Lang (2002) study the effects of the Metco program in Boston, a desegregation program that sends (primarily) African American students to suburban schools. Focusing on Brookline Massachusetts, they find that there is little evidence of any adverse peer effects induced by the transfer of lower achieving inner city students into the Brookline schools. The study does not address the effects on the inner city students themselves. Boozer and Cacciola (2001) study the peer effects using data from Project Star, a state of Tennessee program that was originally designed to assess the effects of smaller class sizes on performance. Boozer and Cacciola study the effects of the percentage of students who were previously enrolled in these small classes on the performance of their classmates in subsequent years. They find these peer effects to be nontrivial and conclude that much of the net benefit of lower class sizes is due to spillover effects. Minkin (2002) reanalyzes the Project Star data under the assumption that the strength of peer influences differs between classmates who were classmates the previous year and those who were not.

Interestingly, he finds much smaller peer effects for all combinations of students than do Boozer and Cacciola.

ii. Social capital

One literature that is closely related to the study of neighborhood effects is the literature on social capital. While social capital does not possess a precise definition, the set of ideas various authors have tried to capture with the term is well summarized in Ostrom (2000):

“Social capital is the shared knowledge, understanding, norms, rules, and expectations about patterns of interactions that groups of individuals bring to a recurrent activity...When they face social dilemmas or collective-action situations...participants must find ways of creating mutually reinforcing expectations and trust to overcome the perverse short-run temptations they face.” (pg. 176)

As this definition makes clear, there are close connections between the sorts of nonmarket interactions that have been discussed in the neighborhoods effects literature and those that appear in the social capital literature.

Appeals to social capital as an important determinant of individual and aggregate behavior have become very common throughout the social sciences. Much of this work is difficult to interpret due to variation in the definition of social capital across studies, a failure to explicitly deal with identification, and a tendency for these studies to conflate any empirical correlations between a group variable with an individual outcome with a causal role for social capital; critiques of the empirical literature include Durlauf (2002) and Durlauf and Fafchamps (2003). Nevertheless, a number of social capital studies do provide empirical evidence that neighborhood characteristics help predict certain individual outcomes and so are a useful empirical corroboration of findings in the empirical literature. Durlauf and Fafchamps (2003) provide a broad overview of the empirical social capital literature; the discussion here identifies a few studies that are particularly interesting from the perspective of neighborhood effects.

One theme in the empirical social capital literature concerns the effect of family moves on children's outcomes. Specifically, a standard argument in the social capital literature is that the strength of attachment to communities is a function of stability of residents; in other words, individuals who frequently move will benefit less from the social support structures provided by neighborhoods than those who are long time residents. A number of studies have found evidence that more frequent moves are associated with a number of undesirable outcomes. In a widely cited study Hagan, MacMillan, and Wheaton (1996) find, using Canadian data, that adverse effects of moves on a range of educational attainment measures is mediated by parental involvement as the effects of moves are far greater when parents appear to be less engaged with offspring (as measured by perceptions of the children). Sandefur, Meier, Hernandez (1999) find that high school completion and post-secondary enrollment are both negatively associated with the number of moves an individual experiences while growing up. Two caveats should be kept in mind when assessing these studies. First, the evidence of a statistical relationship between the number of moves and various outcomes is not uniformly strong; see for example Furstenberg and Hughes (1995) who find only weak evidence of such a relationship. Second, these studies do a poor job of accounting for the endogeneity of family moves. Clearly, family moves will be correlated with unobserved family characteristics such as parental interest in offspring. Hence, for reasons parallel to the discussion of self-selection into neighborhoods, one cannot interpret these studies causally.

From the perspective of neighborhood effects, perhaps the most important work in the social capital literature revolves around efforts to identify detailed characteristics of neighborhoods and their relationship to neighborhood quality. In this context, much of the focus has been on schools as the relevant social category as opposed to residential community. Morgan and Sorensen (1999a) is a good example in this regard. This paper studies the relationship between gains in mathematics achievement between the 10th and 12th grades for students in the National Educational Longitudinal Study of 1988. Morgan and Sorensen (1999a) find that educational gains are positively associated with the density of friendship networks (measured as the percentage of five closest friends attending the same school as the individual in the study) and negatively associated with

the density of parental networks (measured as knowledge by a parent of the parents of his friends.) This finding has been regarded as controversial in the sociology literature (cf. comments by Carbonaro (1999) and Hallinan and Kubitschek (1999) and the rejoinder by Morgan and Sorensen (1999b)). These controversies relate more to the interpretation of Morgan and Sorensen's findings in the context of particular social capital theories than to the findings of how different measures of social structure predict educational achievement. One important lesson from the controversy is that even very precise measures of social structure may prove difficult to map back into specific theories of neighborhood effects.

The most important work in detailing neighborhood characteristics and associated outcomes has been produced by the Project on Human Development in Chicago Neighborhoods (PHDCN). The PHDCN is an extremely detailed data collection project that covers several hundred neighborhoods in Chicago. As described in Sampson, Morenoff, and Earls (1999 pg. 639), the available data include responses to questions such as "About how often do you and people in your neighborhood do favors for each other?" and the likelihood that one's neighbors would intervene if one's child were observed skipping school.

Sampson, Morenoff, and Earls (1999) use the PHDCN to study a range of social aspects of neighborhoods. In particular, they distinguish the social capital of a neighborhood as "the resource potential of personal and organizational networks" (pg.635) from the collective efficacy of a neighborhood, "a task-specific construct that relates to the shared expectations and mutual engagement by adults in the active support and social control of children." (pg. 635). The purpose of this distinction is to differentiate general notions of the levels of neighborhood social resources from the use of these resources. By delineating how neighborhood members help one another, for example through monitoring one another's children, Sampson, Morenoff, and Earls (1999) give a rich portrait of how neighborhoods benefit their members, illustrating how help in childrearing or trust among neighbors are important mediating variables in understanding why poor neighborhoods have adverse effects on their members.

iii. Segregation

A number of recent studies have explored the effects of segregation on individual outcomes. Cutler and Glaeser (1997) consider a set of individual regression of individual outcomes such as high school graduation, college graduation, employment status, and nonmarital fertility and explore whether these outcomes are explained, for African Americans, by the levels of racial segregation in the metropolitan areas in which they live. Evidence for this is determined via the sign and statistical significance for a variable that consists of the cross product of a dummy variable for race and a measure of segregation in the metropolitan area of residence. This variable is consistently significant and implies that segregation lowers black outcomes. The estimated magnitudes are also quite large; according to the point estimates, 2/3's of the black/white difference in single motherhood is explained by segregation. This paper also finds that racial segregation has substantial explanatory power beyond that associated with income segregation. Mayer (2002) studies the effect of economic segregation on years of schooling. She finds that increased variance in family incomes across census tracts reduced years of schooling among children in less affluent families; within census tract variance does not appear to matter.

Other studies have found that within-neighborhood ethnic concentrations matter for understanding individual outcomes. Borjas (1995) has argued that for the United States, within-neighborhood levels of "ethnic capital," defined as average education within an ethnic group, matters for intergenerational mobility. Clark and Drinkwater (2002) find that employment outcomes are strongly related to ethnic group percentages for a range of minorities in England and Wales.

While suggestive of neighborhood effects, mapping the findings of segregation studies into evidence on neighborhood effects is complicated. For example, the Cutler and Glaeser (1997) findings on the effects of racial segregation on African Americans may reflect differences in discrimination across different metropolitan areas; similar reasoning applies to studies of the effect of ethnic concentrations on individual outcomes. This problem is not addressed by the use of instrumental variables such as number of rivers in a metropolitan area to account for the endogeneity of segregation in a

metropolitan area as such instruments may proxy for regional and historical differences in the treatment of blacks. In the case of economic segregation, interpretation difficulties arise, for example, in that it is unclear why the variance across census tracts should be informative with respect to intraneighborhood interactions.

iv. Social attitudes

The importance of neighborhoods may also be seen through the analysis of attitudes. One idea that permeates the social capital literature is that the social structure of neighborhoods can affect a range of attitudes and beliefs about other individuals and groups. This idea provides a causal mechanism as to why certain behaviors may be predicted from group characteristics. For example, to the extent that certain neighborhood characteristics induce feelings of trust and reciprocity towards others, the transfer of information about job opportunities may be facilitated. A recent study by Alesina and La Ferrara (2002) evaluates this idea using data from the GSS. Individual data on various attitudes related to trust is analyzed using data on the characteristics of the metropolitan area of respondents in the survey. Alesina and La Ferrara (2002) find that levels of trust are higher among individuals who reside in metropolitan areas with higher degrees of economic inequality and racial heterogeneity. A useful advance in this type of work would be the collection of attitudinal data for smaller units than metropolitan areas, which seem too large to well approximate the relevant neighborhoods for social interactions.

One question raised by the analysis of neighborhoods and social attitudes concerns how attitudinal data of this type translate into actual behavior. At an individual level, experimental evidence in Glaeser, Laibson, Scheinkman and Soutter (2000) suggests that expressions of trust in answering the GSS questions are better understood as predicting trustworthy behavior by the answerer than actual feelings of trust. At an aggregate level, this question has been indirectly addressed by a recent literature that has addressed the effect of neighborhood composition on government policies. Alesina, Baqir, and Easterly (1999) find that the level of racial heterogeneity in local political

jurisdictions is negatively associated with the level of government services. Alesina and La Ferrara (2000) find that racial heterogeneity is associated with lower rates of participation in community social activities, an effect that is particularly pronounced among those who possess racially prejudiced views. However, while these studies are certainly consistent with evidence that neighborhood composition affects social attitudes, one can easily imagine alternative explanations for the observed behaviors. And of course, racial composition of a community is endogenous, as usual making causal claims problematic.

v. Home ownership and individual behavior

A number of authors have studied the relationship between aspects of housing and various behaviors. One argument that has been made is that home ownership is causally associated with various socially desirable behaviors. One reason for such a relationship concerns investment in public goods. DiPasquale and Glaeser (1999) argue, for example, that homeownership increases individual incentives to invest in both local amenities as well as forms of social capital, interpreted as something that increases the degree of connectedness to others. DiPasquale and Glaeser (1999) find that for the US, homeowners exhibit higher levels of participation in nonprofessional organizations, greater knowledge of and participation in local politics, etc. They also find similar results for Germany, although the magnitudes of the effects are substantially smaller. Green and White (1997) find that children of homeowners are less likely to drop out of high school or to experience pregnancy as teenagers. While Green and White are quite circumspect in interpreting their results, one can apply the same causal arguments as made in DiPasquale and Glaeser. Both of these studies attempt to address the endogeneity of homeownership; DiPasquale and Glaeser (1999) use the average homeownership rate of an individual's socioeconomic class whereas Green and White use a bivariate probit specification to jointly model outcomes and homeownership, both studies acknowledge that their corrections are far from ideal.

Additional evidence of the existence of neighborhood effects may be found in Glaeser and Sacerdote (1999), who study the relationship between housing structure and various measures of social interaction. They find that individuals in larger apartments appear to have stronger social connections with neighbors than do house dwellers, but that apartment dwellers are less likely to participate in local politics. This finding, combined with a finding that street crime tends to be concentrated around large apartments, leads Glaeser and Sacerdote to conclude that apartment dwellers tend to be socially isolated from the larger communities in which they live. While the causality in this analysis is unclear, the suggestion that very localized differences in housing configurations alter individual behavior is quite intriguing from the perspective of neighborhood effects.

vi. Geography and social customs

A final source of evidence on neighborhood effects may be developed from recent studies that explore how regional variations in certain behaviors appear to be driven by social custom. This idea is of longstanding importance in areas such as sociolinguistics, where regional variations in syntax and pronunciation persist despite the influences of mass media, cf. Chambers (1995) and Labov (2001). Recent analyses have extended this idea to substantive economic behaviors. In one important study, Young and Burke (2001) study patterns in cropsharing contracts between landowners and tenants in Illinois. They document how cropsharing percentages are concentrated on simple divisions (50/50, etc.) even though there is no reason, given standard theories of contracts, for this to occur. They further show that there is substantial regional variation in these simple contracts with different shares predominating in different regions. These two facts suggest an important role for local norms in determining contract terms. Young and Burke (2003) extend this work to show that these norms have significant distributional consequences. Similar findings are developed by Burke, Fournier, and Prasad (2003) in the context of medical care. This paper documents strong regional differences in the use of alternative medical treatments such as choice of coronary care and shows how these

may be understood as arising from local social norms. These two applications illustrate new dimensions along which to understand how neighborhood effects matter.

6. Conclusions

The new literature on neighborhood effects covers an enormous range of theoretical, econometric and empirical issues. As such, it defies easy summary. These three components of the neighborhood effects literature are in different states of development. There now exists a reasonably well developed theoretical literature that addresses both how neighborhood effects influence aggregate behavior and how these effects influence equilibrium neighborhood formation. In contrast, the econometrics of neighborhood effects is still in a nascent stage. The literature has identified deep identification problems that exist due to Manski's (1993) reflection problem between endogenous and contextual effects. Further, the estimation problems that exist because of self-selection and other types of unobserved heterogeneity are relatively well understood. However, work is only beginning on ways to overcome these problems so that credible empirical work may proceed. Finally, there is a large empirical literature that has explored neighborhood effects over an impressive range of behaviors. Much of this work is interesting and suggestive that neighborhood effects matter. However, this work has generally not dealt with the econometric issues that arise for neighborhoods models and so cannot be regarded as meeting the evidentiary standards of the more successful literatures in economics, for example, program evaluation.

What suggestions does this assessment raise for future research? One general implication is that future empirical work should attempt to simultaneously address behaviors within neighborhoods and neighborhood configurations via structural models. Structural models will allow for a full exploration of self-selection in neighborhoods models and allow for the analysis of policy interventions in ways beyond the current literature. One example of how neighborhood choice may help elucidate the nature and magnitude of neighborhood effects is via hedonic price arguments. With the exception of Ioannides, who only addresses this indirectly, the information embedded in house prices

on neighborhood effects has yet to be exploited. Work by Nesheim (2002) and Heckman, Ekelund, and Nesheim (2002a,b) make important advances in the use of hedonic price models to uncover factors such as neighborhood effects.

One can also identify a number of specific challenges. It seems important to develop more realistic housing market models. A generally neglected issue is that of house market dynamics; issues of capital gains and the role of future neighborhood composition have essentially been ignored in the current literature. These are extremely hard problems and compound the general difficulties that exist in urban economic theory. Recent theoretical work by Ortalo-Magne and Rady (2002a,b) on housing market dynamics may prove to be important in enriching current neighborhoods models.

Finally, I believe that much more attention needs to be paid to the microeconomic foundations of neighborhood effects. One promising approach concerns the role of self-identity in behavior, which has been introduced into economics by Akerlof and Kranton (2000). This work suggests that individual preferences and beliefs are conditioned in a fundamental fashion by the type of social identity they wish to possess and to present to others. Neighborhoods are a likely source of identity. An example of how this may be important is work by Ferguson (2001) and Ogbu (2003) on African American school performance in Shaker Heights Ohio. Shaker Heights has received much attention as it is a middle class community that appears to have been relatively welcoming to African American families and yet in which there are substantial test score gaps between racial groups. Ogbu (2003) is particularly persuasive on the role of collective identity in conditioning factors ranging from study habits to aspirations. While work on identity suggests a role for how neighborhoods directly influence its members, related work on stigma, in particular the important recent book by Loury (2002). Loury suggests that neighborhoods also induce effects because of the way they influence how neighborhood members perceive others. Loury's view may be interpreted as saying stigma against certain groups arises because beliefs about groups are fundamentally underidentified. In other words, stereotypical thinking can persist in a society because the experiences and information available to individuals typically cannot falsify the stereotypes. The extent to which this is so very much depends on how individuals are organized in social space. In other words, stereotypes about the poor or about certain ethnic groups survive partially

because the lack of interactions with these groups alters the experiences and information possessed by others. Hence, one can see a role for neighborhood configurations in underpinning Loury's arguments. These approaches to understanding why neighborhoods matter, relying as they do on particular views of human cognition, makes clear the importance of greater data collection efforts at more disaggregated levels than one typically finds in neighborhoods studies.

To be clear, given its youth, the neighborhoods effects literature has made enormous strides in expanding economists' understanding of social influences and more progress will certainly be made over time.

Appendix. Selection correction for neighborhood effects regressions based on the multinomial logit model

This appendix is designed to illustrate how one can employ the Heckman-type selection corrections to account for endogeneity of neighborhood choice. The analysis is taken from Brock and Durlauf (2003). Under the assumptions that 1) ϵ_{ij} is double exponentially distributed, i.e. $f(\epsilon) = \frac{1}{2} \exp(-|\epsilon|)$ (so that neighborhood choice obeys a multinomial logit model) and 2) η is normally distributed, arguments in Lee (1983) may be used to show that

$$P_{ij} = \frac{\exp(\beta_j' X_{ij})}{\sum_k \exp(\beta_k' X_{ik})} \frac{1}{1 + \exp(\beta_j' X_{ij} - \eta)} \quad (32)$$

where

$$\eta = \frac{\int \exp(\beta_j' X_{ij}) \exp(\beta_k' X_{ik}) \exp(\eta) \exp(-\eta^2/2) d\eta}{\int \exp(\beta_j' X_{ij}) \exp(\beta_k' X_{ik}) \exp(\eta) \exp(-\eta^2/2) d\eta} \quad (33)$$

with ϕ and Φ denoting the density and distribution function of a normal (0,1) random variable and

$$\eta = \frac{\int \exp(\beta_j' X_{ij}) \exp(\beta_k' X_{ik}) \exp(\eta) \exp(-\eta^2/2) d\eta}{\int \exp(\beta_j' X_{ij}) \exp(\beta_k' X_{ik}) \exp(\eta) \exp(-\eta^2/2) d\eta} \quad (34)$$

The correction described by eqs. (33) and (34) allows for consistent estimation of the behavioral parameters in (32). The parameters β are estimated in a first stage multinomial logit analysis and used to form η ; η is simply a regression parameter in (32).