

Neighborhoods and Adolescent Development:

How Can We Determine the Links?

Greg J. Duncan

Northwestern University

Stephen W. Raudenbush

University of Michigan

October 16, 1998

Prepared for the symposium “Does It Take a Village? Community Effects on Children, Adolescents, and Families,” Pennsylvania State University, November 5-6, 1998. Portions of this paper are drawn from Greg J. Duncan and Stephen W. Raudenbush “Assessing the Effects of Context in Studies of Child and Youth Development,” Educational Psychologist, forthcoming. It has benefited from discussions with fellow members of the MacArthur Foundation Methodology Working Group – Robert Sampson, Helena Kraemer, Ron Kessler, John Nesselroade, and, especially, Tom Cook, – as well as comments from Johanne Boisjoly, Dorothy Duncan, Rachel Dunifon, Kathleen Harris, Lori Kowaleski-Jones, Daniel Levy, Jens Ludwig, Shelia Murray, Marianne Page and Gary Solon. We are grateful to the Family and Child Well-being Research Network of the National Institute of Child Health and Human Development for supporting this research

Neighborhoods and Adolescent Development:

How Can We Determine the Links?

I. Introduction and Summary

Despite ample theoretical reasons to suspect that neighborhood conditions influence adolescent development and behavior, the task of securing precise, robust and unbiased estimates of neighborhood effects has proved remarkably difficult. This paper provides an assessment of the conceptual and, especially, methodological issues involved as well as guidance on the most promising research designs for obtaining an unbiased understanding of the nature of neighborhood effects.

Key methodological issues include: i) obtaining neighborhood-level measures that approximate the theoretical constructs of interest; ii) allowing for the possibility of simultaneous influences between youth and their contexts; iii) avoiding bias from unobservable characteristics of parents that influence both choice of neighborhood and child outcomes; iv) consideration of ways in which families mediate and moderate neighborhood influences; and v) using samples with sufficient variability in neighborhood conditions.

We argue that: i) studies that draw their samples from only a handful of different neighborhoods have little chance of distinguishing among the many theoretical ways in which neighborhoods may influence youth; ii) neighborhood-effects estimates from studies that measure neighborhood characteristics from youth or parental self-reports or by aggregating responses of youth or their parents are likely to be biased, especially when the youth outcomes themselves are based on youth or parent reports; iii) neighborhood data drawn from independent samples of residents or by more economical systematic social observation methods are more promising for addressing some of hypotheses of interest; iv) a simple but informative method of estimating upper bounds on the scope of potential neighborhood effects is to estimate outcome correlations for pairs of youth who live close to one another; and v) quasi- and random-assignment experimental studies represent our best hope for discovering the scope, if not nature, of neighborhood influences.

II. Why Neighborhood Conditions Might Matter

Why might extrafamilial contexts -- neighborhoods, communities, schools and peers -- affect an adolescent's behavior? The literature is filled with answers to this question, some but not all of which argue that higher-SES environments are better for children. Because this literature is reviewed more completely in other papers in this symposium, we provide in this section an exceedingly brief and selective review of theories of contextual -- especially neighborhood -- effects, with an eye toward motivating our methodological discussion.

Jencks and Mayer (1990) develop a taxonomy of theoretical ways in which neighborhoods may affect child development. They distinguish:

- “epidemic” theories, based primarily on the power of peer influences to spread problem behavior;
- theories of “collective socialization,” in which neighborhood role models and monitoring are important ingredients in a child's socialization;
- “institutional” models, in which the neighborhood's institutions (e.g., schools, police protection) rather than neighbors per se make the difference;
- “competition” models, in which neighbors (including classmates) compete for scarce neighborhood resources; and
- models of “relative deprivation,” in which individuals evaluate their situation or relative standing vis-à-vis their neighbors (or classmates).

The first three of these explanations predict that “better” environments promote positive development. The last two predict that some youth may be negatively affected by exposure to higher-SES environments.

Since adolescents typically spend a good deal of time away from their homes, explanations of neighborhood influences based on peers, role models, schools and other neighborhood-based resources would appear to be more relevant for them than for younger children. However, it is possible that neighborhood influences begin long before adolescence. A substantial minority of 3- and 4-year-olds are enrolled in center-based daycare or preschool (Hofferth and Chaplin, 1994). Physically dangerous neighborhoods may force mothers to be isolated in their homes and thus restrict opportunities for their children's interactions with peers and adults (Furstenberg,

1993). Parks, libraries and children's programs provide more enriching opportunities in relatively affluent neighborhoods than are available in resource-poor neighborhoods. Parents of high socioeconomic status may be observed to resort less frequently to corporal punishment and to engage more frequently in learning-related play. Thus, there are many ways in which neighborhood conditions might affect both children and adolescents (Chase-Lansdale et al., 1997).

Social disorganization theory identifies key elements of collective socialization and institutional forces likely to influence child and adolescent development. Following Shaw and McKay (1942), Sampson and his colleagues have argued that a high degree of ethnic heterogeneity and residential instability leads to an erosion of adult friendship networks and undermines a values consensus in the neighborhood (Sampson & Lauritsen 1994), which in turn means that problem behavior among young people is not controlled as effectively as in more socially-organized neighborhoods.

Sampson et al. (1997) argue for the importance of the concept of "collective efficacy," which combines social cohesion (the extent to which neighbors trust each other and share common values) with informal social control (the extent to which neighbors can count on each other to monitor and supervise youth and protect public order). Thus it represents the capacity for collective action by neighbors. Sampson et al. (1997) find that collective efficacy so defined relates strongly to neighborhood levels of violence, personal victimization and homicide in Chicago, after controlling for prior crime and for social composition as measured by census variables.

Wilson's (1987) explanation of inner-city poverty in Chicago relies on a more complicated model in which massive changes in the economic structure, when combined with residential mobility among more advantaged blacks, results in homogeneously impoverished neighborhoods that provide neither resources nor positive role models for their children and adolescents.

Furstenberg (1993) and Furstenberg et al. (1998) argue for the importance of family-management practices in understanding neighborhood effects. Basing their work on both ethnographic and a survey-based study, they point out that families formulate different strategies for raising children in high-risk neighborhoods, ranging from extreme protection and insulation

to an active role in developing community-based “social capital” networks that can help children at key points in their academic or labor-market careers. This work highlights the need to consider family-neighbor interaction effects in neighborhood research.

III. Methodological challenges to “getting context right”

Distinguishing empirically among these complementary and, in some cases, competing theories is not an easy task. Collectively, the theories suggest many possible mechanisms, most of which are not easily measured. But measurement issues are only part of a collection of conceptual problems that await the aspiring neighborhood-effects researcher.

Building on Manski (1993), Moffitt’s very useful review paper (1998) distinguishes among: i) the simultaneity problem; ii) the omitted-context-variables problem; and iii) the endogenous membership problem. To this list we would add: iv) consideration of ways in which families mediate and moderate neighborhood influences; as well as: v) the more practical problem of selecting samples with sufficient contextual variability.

To frame the methodological issues, consider a model in which adolescent i ’s achievement or problem behavior (y) is an additive function of i ’s family (FAM) and extra-familial contextual (CON) influences:

$$(1) \quad y_{ic} = A' \text{FAM}_i + B' \text{CON}_{ic} + e_{ic}.$$

For the moment, we assume one child per family and C ($c = 1 \dots C$) distinct contexts. Our interest is in obtaining unbiased estimates of B' , the effect of context on the youth outcome. Nonadditive interactions between FAM and CON are considered below and do not invalidate our discussion based on model (1).¹

The simultaneity problem

A first possible problem is that of simultaneous causation – that contextual conditions themselves may be a function of y ’s behavior. In this case we have a two-equation system:

$$(2) \quad y_{ic} = A' \text{FAM}_i + B' \text{CON}_{ic} + e_{ic}.$$

$$(3) \quad \text{CON}_c = C' y_{ic} + D' Z_c + w_c$$

where Z_c is a vector of other determinants of contextual conditions.

The idea that children are not only shaped by, but also shape their family environments is a familiar one to developmentalists and a key element of “transactional” models of development (Sameroff and Chandler, 1975). That two-way “transactions” may play a role in extrafamilial contexts is best seen in the case of best friends or peer groups. In the case of best friends, CON might be the behavior of i’s best friend. Equation (2) then reflects the assumption that i’s behavior is causally linked to the behavior of his or her best friend. But equation (3) then reflects the assumption that i’s best friend’s behavior is also causally dependent on i’s own behavior. Identification of the Bs and Cs in a two-equation system such as (2)-(3) is a difficult task.

Less intuitive but not implausible are simultaneity problems involving adolescents and their neighborhood-based contexts. Suppose, as do Sampson, Raudenbush, and Earls (1997), that the “collective efficacy” of the adults in a neighborhood is a forceful deterrent to the problem behavior of the neighborhood’s youth. It is quite possible that a neighborhood’s collective sense of efficacy is itself determined by youth behavior and that misbehavior of even one youth, if sufficiently serious, could affect the context (in this case the collective efficacy in the neighborhood) of that youth.

Addressing the simultaneity problem for peer contexts is particularly difficult, since it is virtually impossible to find Z_c determinants of the behavior of peers that are not also determinants of i’s own behavior. Moffitt (1998) and Manski (1993) provide a more complete analysis of the peer case, assuming that i’s choice of group members (e.g., i’s best friend) is exogenous and distinguishing between the effects on i’s behavior of i’s best friend’s: i) behavior and ii) demographic characteristics. They term the former, behavioral, effects “endogenous social interactions” and the latter, demographic-characteristic, effects “exogenous social interactions.”

These distinctions have important policy implications. If context effects on, say, delinquency operate through peer behavior, then public policymakers might be able to stem local epidemics of teen problem behavior by focusing on prevention among a key set of high-risk adolescents. But if it is the neighbor or peer characteristics rather than behavior that matter, then more costly (in terms of resources and political capital) programs such as school and residential mobility programs become more important.

Estimating the reduced-form version of (2)-(3) in which i 's behavior is regressed on FAM and i 's best friend's demographic characteristics (but not behavior) identifies the existence of social interactions but does not provide distinct estimates of exogenous and endogenous social interactions. Identifying the distinct role of endogenous and exogenous effects of peers is all but impossible, even in this simplistic framework in which they have assumed that i 's choice of group members (e.g., i 's best friend) is exogenous.

The identification problem is somewhat less serious in the case of neighborhood contexts, since there is some hope for finding Z_c determinants of neighborhood structure (e.g., regional changes in economic structure) that are not also determinants of i 's own behavior.

The omitted-context variables problem

Distinct from the simultaneity problem is the more conventional problem of omitted variables – in this case context-level variables.² Regressing i 's achievement or behavior on his or her family and contextual characteristics will bias estimates of B' if important characteristics of i 's context are omitted from the regression. To illustrate this, we can add to (1) a component of the error term that is assumed constant within contexts:

$$(4) \quad y_{ic} = A' \text{FAM}_i + B' \text{CON}_{ic} + u_c + e_{ic}.$$

This model leads to the familiar omitted-variables problem and a biased estimation of B' (the effect of CON on y) if other contextual conditions (represented in (4) by u_c): i) are omitted from the estimation of equation (4); ii) affect y ; and iii) are correlated with CON. Attempts to measure key contextual constructs generally adopt administrative-data approaches, which, we shall argue, are limited in scope; or survey-based approaches, which often suffer from measurement error bias. Another measurement approach, explained in a subsequent section, is that of systematic social observation (SSO).

Administrative-data approaches. It is easy to argue that many existing studies of neighborhood context suffer from omitted-variables bias. Most draw their data from the decennial census. Every ten years, the Census Bureau collects information that can be used to construct demographic-based neighborhood measures such as the fraction of individuals who are poor, the fraction of adults with a college degree, and the fraction of adult men without jobs. Subsets of such data are available for Census blocks and block groups; complete Census data are available for “tracts” (geographic areas encompassing 4,000 to 6,000 individuals with boundaries

drawn to approximate neighborhood areas), ZIPCodes, cities, counties, metropolitan areas, labor market areas and states. Other administrative data bases can be used for measuring certain physical characteristics of neighborhoods and schools as well as certain ecological risk factors (such as crime and infant mortality rates of neighborhoods), although these data are often not uniformly measured across contexts, nor available for geographic areas as small as tracts.

There are two problems with studies that rely on Census-based sources. Some use only a single neighborhood measure such as the poverty or welfare-receipt rate in the Census tract of residence or a single (often factor-analysis-derived) index of such measures. Significant coefficients on the census-based measure are taken as evidence of neighborhood effects. In some cases where only one census variable is used (e.g., tract poverty or welfare rate), the interpretation is further extended (erroneously in most cases) to suggest that it is neighborhood poverty or welfare dependence as such, as opposed to many other dimensions of neighborhoods correlated with rates of poverty or welfare receipt, that is behind the neighborhood effect.

The multitude of theoretical ways in which neighborhood processes operate suggest that many different kinds of measures, even if all of them can be derived from Census-based sources, are needed to capture the different kinds of neighborhood effects. For example, epidemic models focus on the presence of “problematic” peers and have often been implemented with measures of neighborhood poverty or low-SES job structure (Crane, 1991; Clark, 1992). In contrast, social control and institutional models focus more on the presence of higher-SES neighbors than the presence or absence of low-SES neighbors. This distinction is subtle, but easily conceived if SES is thought to have at least three strata - say, low, medium and high levels of SES.

The diversity of U.S. neighborhoods produces different combinations of these three strata, which enables researchers to distinguish empirically among their effects on developmental outcomes. Brooks-Gunn, Duncan, Klebanov, and Sealand (1993) and authors of several chapters in Brooks-Gunn, Duncan, and Aber (1997) find that it is the presence or absence of affluent, high-SES neighbors rather than the presence or absence of poor neighbors that relates most strongly to child and adolescent outcomes. Wilson’s focus on male joblessness adds yet another correlated but theoretically distinct dimension of neighborhood structure and places more demands on the contextual dispersion of the sample needed to distinguish between Wilson’s and other models.

Much more problematic from an omitted-variables point of view are tests of theoretical approaches based on neighborhood influences not well captured by administrative data. Institutional, social disorganization and family-process models are examples, since the required measures are not readily available from census-based sources. Absent from Census forms and either absent from or inconsistently measured in other administrative sources are data about schools (e.g., number, type and quality of schools in the area); law enforcement (e.g., number of police, number and type of crimes, percent of crimes reported that are cleared by arrest, various characteristics of local police practices); access to transportation (e.g., distance to freeway entrances and public transportation); drugs and gang activity; neighborhood collective efficacy; churches and other community institutions.

Survey-based approaches. In attempting to go beyond the constraints imposed by Census or administrative-based sources, some studies have sought to use youth or their parents as informants about the characteristics of their neighborhoods or schools. A major problem with this strategy is that measurement errors in these assessments are likely to be correlated with the measurement errors in the youth-based outcomes. For example, a depressed mother may give overly-pessimistic assessments of both neighborhood conditions and her children's social behavior. If the mother's mental health is not controlled in the regression analysis, then the estimated relationship between neighborhood conditions and child behavior will be overstated. Another example of bias induced by self-reports is when an adolescent's report of his or her peers' attitudes is spuriously correlated with his or her self-reports of attitudes and behavior. Note that this correlated-errors problem can be eliminated if the youth outcomes are based on administrative (e.g., test scores, school attendance, arrest records) rather than survey data, a strategy adopted in Cook et al. (1998).

There is less reason to suspect measurement-error-driven bias from context measures formed by aggregating demographic characteristics such as ethnicity, sex, or social class to construct segregation indices or other measures of social composition. For example, Lee and Bryk (1989) constructed measures of the social and ethnic composition of US high schools from student survey data and used those measures to predict the same students' academic achievement. There is a small risk that a student's report of demographic background is influenced by his or her achievement. In contrast, an aggregated measure of perceived instructional quality would quite plausibly reflect the achievement of the student reporters, and it

would therefore be inadvisable to use such a measure of instructional quality as a predictor of student achievement.

A more satisfying, if expensive, strategy is to obtain an independent sample of capable informants about a context and pool their reports to create context-level measures. This approach has been successfully used in national data on school climate (c.f., Raudenbush, Rowan, & Kang, 1991), with multiple teachers surveyed about their degree of control, collaboration, and supportive administrative leadership; and, as described below, in data assessing the social cohesion, informal social control, and collective efficacy of neighbors in Chicago (Sampson et al., 1997). In both cases, 15-30 informants per context were required to obtain reliable contextual-level measures. Clearly the expense of this measurement strategy grows rapidly with the number of contexts sampled, and it will increase during the course of a longitudinal study as mobility creates greater dispersion of participants across contexts and, hence, produces more contexts to be assessed.

Endogenous membership

The contexts in which children develop are not allocated by a random process, which leads to the third – endogenous membership - problem. This is most clearly seen in the case of selection of best friends and peer groups, where decisions rest almost entirely with the adolescent.³ A youth's immediate neighborhood and, to a somewhat smaller extent, school also have an element of choice, in this case on the part of the parent. The propensity of children to live in better or worse neighborhoods or attend better or worse schools depends on parental background characteristics and current circumstances, not all of which can be easily measured.

As with the omitted-context-variables problem, the endogenous membership problem involves omitted variables, but in this case at the level of the individual (in the case of youth i 's choice of best friends or peers) or family (in the case of youth i 's parent's choice of neighborhood or school). Presuming the latter, family-based, source of endogenous-membership bias, we can illustrate the problem and potential solutions by adding to our model a subscript (f) denoting family and a error component (v_f) denoting family-specific influences on choice of context.

$$(5) \quad y_{ifc} = A' \text{FAM}_{if} + B' \text{CON}_{ic} + v_f + u_c + e_{ic}.$$

Here the omission of explicit measures of v_f will bias B' to the joint extent: i) v_f is an important determinant of y ; and ii) v_f is correlated with CON. A parallel argument holds for omitted individual-level influences on choice of context.

The direction of endogenous-membership bias in estimates of (5) is uncertain. Suppose parents choose between: i) holding two jobs and using the extra income to buy a better neighborhood; or ii) having a single earner and living in a poorer neighborhood. Suppose further that parents who live in poorer neighborhoods and/or send their children to worse schools make up for the deficiencies of the neighborhood or school through the additional time that parents spend with their children. Neighborhood or school conditions matter in this scenario, but an empirical analysis will show this to be the case only if it adjusts for differences in parental time use. Failure to adjust for parental employment will cause conventional regression-based approaches to understate neighborhood or school effects. In terms of equation (5), failure to include parental employment as part of FAM will likely bias the estimate of B' toward zero.

Another scenario, also leading to an understatement of neighborhood or school effects, is one in which parents well equipped to resist the effects of bad neighborhoods choose to live in them to take advantage of cheaper housing or perhaps shorter commuting times. Unless measures of parental competence are included in the model, the estimated effects of bad neighborhoods or schools on youth outcomes will be smaller than if parents were randomly allocated across neighborhoods.

It is perhaps more likely that parents especially ill-equipped to handle bad neighborhoods or schools are most likely to live in them, because these parents lack the (partly unmeasured) wherewithal to move to better neighborhoods. In this case, the coincidence of a poor neighborhood or school and the poor developmental outcomes of their children results from their inability to avoid either, thus leading to an overestimation of the effects of current neighborhood conditions. Conversely, parents who are effective in promoting the developmental success of their children may find their neighborhood choices dominated by considerations of developmental consequences. If this parental capacity is not captured in measured parental characteristics, then the coincidence of positive developmental outcomes for their children and living in a better neighborhood would be misattributed to current neighborhood conditions and thus lead to an overestimation of neighborhood effects. Here again, in terms of regression model

(5), the omitted v_f factors are unobserved characteristics of the parents (e.g., concern for their children's development) that affect developmental outcomes. Regressions of contextual effects that do not control for all relevant parenting variables will produce biased estimates of the effects of those contextual factors.

Random assignment. There are three approaches for addressing the endogenous membership problem. The best is to rely on data in which families are randomly assigned to neighborhood contexts. In terms of (5), the random assignment of CON effectively eliminates the correlation between family-specific omitted factors (v_f) and context (CON), and thus eliminates the possibility of omitted-variables bias. As described below, HUD's Moving to Opportunity program contains such experimental data on neighborhood context. Second-best solutions to the nonrandom context selection problem are to rely on quasi-experimental data such as those generated by the Gautreaux program, analyses of which are also summarized below.

Measure the unmeasured. The best non-experimental approach to the endogenous membership problem is to locate data that measure the crucial family and individual-level omitted variables. For example, some child-development data sets contain fairly sophisticated measures of parenting characteristics and parental mental health. Controls for such measures in regression-based analyses can help reduce the endogenous-membership bias to the extent that those measures capture the determinants of the process of contextual choice. Of course, it is impossible to demonstrate that all relevant variables have been included in a model, which suggests that this measurement-based strategy for the omitted-variables problem should be accompanied by others.

Instrumental variables. Another nonexperimental approach to the bias problem is to replace the contextual measure (CON) in equation (5) with a predicted value of CON that is purged of CON's spurious correlation with unobserved parenting or other family or individual-level measures. The instrumental variables approach is often implemented as a two-step procedure (Greene, 1993). In the first step, the contextual measure (CON) is the dependent variable and is predicted by other variables in (5) plus exogenous variables that are not themselves determinants of y .⁴ In the second stage, equation (5) is estimated replacing CON with the predicted value of CON obtained in the first stage.

Evans, Oates and Schwab (1992) adopt this strategy to adjust for endogenous membership problems in a school-based study, although they rely on dubious instrumental variables. Their dependent variables of interest are high-school completion and out-of-wedlock teen childbearing. Their contextual variable is the SES of the student body. When they ignore endogeneity issues and regress their outcomes on student-body SES and family-level controls, they find highly significant, beneficial effects of high student-body SES. However, when they estimate a two-equation model, with the first equation regressing student-body SES on characteristics of the metropolitan area in which the student resides and the second regressing the developmental outcomes on predicted student-body SES and family-level controls, the effects of student-body SES disappear.

Dubious in the procedures of Evans et al. (1992) is the assumption that metropolitan-level characteristics do not influence youth outcomes such as high-school completion or fertility. It is easy to imagine that both labor-market conditions and metropolitan-specific norms might well influence the cost-benefit calculus behind these decisions, which would invalidate their approach. As described below, a more promising approach for identifying an instrumental-variables model is to rely on data from true or quasi-randomized residential mobility designs.

Sibling models. Yet another approach to the endogenous membership problem is to use sibling-based fixed-effects models to eliminate the biasing influence of omitted persistent, unmeasured parental characteristics. In fixed-effects models, each sibling's score on the dependent and independent variables is subtracted from the average value of all siblings in the family. In the special case of two siblings per family, the deviation-from-means model becomes a sibling difference model. If we replace the subscript i in (5) with 1 (for sibling 1) and 2 (for sibling 2), and, assuming that there is sufficient cross-sibling variability in contextual conditions to reference CON with the sibling subscripts, the sibling difference model takes the following form:

$$(6) \quad y_2 - y_1 = A' (FAM_2 - FAM_1) + B' (CON_2 - CON_1) + (v_{f2} - v_{f1}) + (u_{c2} - u_{c1}) + (e_2 - e_1)$$

In terms of measured variables, this amounts to estimating a regression in which sibling differences in the outcome of interest are regressed on sibling differences in family and contextual characteristics. If observed (FAM) and unobserved (v_f) family factors are constant

across siblings, then they drop out of (6), thus eliminating the omitted-variable bias caused by family-determined endogenous group membership. Even if unobserved family factors differ across siblings, it is often reasonable to assume a low correlation between sibling differences in those family factors and sibling differences in context, in which case even unmeasured sibling-specific family factors will not impart much bias to estimates of neighborhood effects (B').

Aaronson (1997) demonstrates the feasibility of this approach using data on Panel Study of Income Dynamics adolescents. He uses family residential changes as a source of neighborhood background variation within families to estimate sibling-based neighborhood effects that are substantially free of family-specific heterogeneity biases associated with neighborhood selection. Using a sample of multiple-child PSID families where the adolescent siblings are separated in age by at least three years, he estimates sibling-difference models of children's educational outcomes and finds evidence of neighborhood effects. In fact, his family fixed-effect regressions using the neighborhood poverty rate as the measure of neighborhood conditions show even larger neighborhood effects on high school graduation and grades completed than conventional OLS models.

Sibling models are not without problems, however. They require multiple-child families, which introduces a potential source of sample selection bias. And Griliches (1979) points out that differencing between siblings reduces but does not eliminate endogenous variation in neighborhood regressors, since parental decisions to change neighborhoods may be motivated in part by their child-specific developmental consequences. At the same time, sibling-difference models filter out much of the exogenous variation in contextual conditions. And finally, family moves are often motivated by events which may themselves may affect youth development.

Families as mediators and moderators of neighborhood effects

Thus far, our discussion of family influences on youth outcomes has assumed that they play a confounding role in attempts to gauge neighborhood influences. We have concluded that researchers who fail to measure them (especially the part of them that determines neighborhood choice) risk bias in their assessments of neighborhood influences. It is important to recognize the possibility of two other, more substantial, roles for families in neighborhood studies: i) mediators that help account for the "reduced form" effects of neighborhood conditions on youth outcomes; and ii) moderators in which families and neighborhoods jointly influence youth outcomes.

Families as mediators. Models presented thus far assume that neighborhood conditions affect children directly. But it is also likely that characteristics of families such as income, living arrangements, parenting and decision-making and parental mental health are shaped by social and physical contexts like neighborhoods (Duncan, Connell and Klebanov, 1997). If neighborhoods affect parents in ways that in turn affect children, then mere adjustments for family differences as in equation (1) will understate the total effects of neighborhoods on children. Better to conceive of these relationships recursively, with family conditions playing the role of mediators.

It is easy to imagine how neighborhoods might affect parents, especially low-SES parents. Persistent residence in a neighborhood with high levels of crime, low levels of economic opportunity, weak marriage pools, and poor transportation can erode and eventually dissipate the competence and commitment of single mothers to seek employment in that neighborhood, to marry, or to move to a neighborhood where they can work and provide safe activities outside the home for their children (Korbin and Coulton, 1997).

Assessments of neighborhood influences on families face many of the same methodological problems as studies of neighborhood influences on children. Particularly vexing is the endogenous membership problem – how much of an association between, say, bad neighborhood conditions and low parental income reflects neighborhood conditions causing employment problems and how much is the spurious result of omitted factors (e.g., mental health) causing both employment problems and residence in a bad neighborhood? An important area for future research is to secure unbiased estimates of the links between neighborhood characteristics and family conditions.

Families as moderators. Cook et al. (1998) argue for interactions between family- and neighborhood-level conditions. Their Philadelphia-based analysis (as well as more general one presented in other chapters of Furstenberg et al., forthcoming) illustrates the subtlety of the relationships among neighborhood conditions, family management practices and youth outcomes.

Surprisingly, they found only modest differences in family management practices across their diverse sample of neighborhoods. Management styles were more restrictive in disadvantaged neighborhoods, but the relationship was not a very strong one. Families’

institutional connections mattered for youth academic outcomes, but only in more disadvantaged neighborhoods. Parental restrictiveness reduced the involvement of children in potentially beneficial activities, but only in higher-risk neighborhoods.

Variability in contextual characteristics

Estimating models of how neighborhood characteristics affect youth achievement and behavior with survey-based, nonexperimental data requires youth samples that are dispersed across a variety of contexts. If the researcher hopes to go beyond assessments of a single “good” vs. “bad” neighborhood dimension and begin to distinguish among competing neighborhood-effects theories (e.g., based on male joblessness vs. resources vs. collective socialization), then the needed sample dispersion across different kinds of neighborhoods is considerable. But since costs increase at least linearly with the number of sampled neighborhoods, it may be impossible to conduct worthy studies of neighborhood effects with small research budgets.

A discussion of the tradeoffs and options rests on the nature of the contextual data to be analyzed. One option is with administrative (e.g., Census-based) data collected by geographic area and associated with sampled youth through address matching or some similar procedure. A second is by aggregating characteristics or survey responses of sampled youth or their parents. Our earlier discussion warned of instances when aggregation procedures are likely to bias estimates of contextual effects.

Is it desirable to cluster sample observations by context? To save costs, most surveys “cluster” their samples by selecting multiple households within a limited number of neighborhoods or multiple students within a limited number of classrooms. Whether this clustering hinders or helps attempts to model context effects depends on the source of the contextual data.

Suppose that administrative data are to be used to measure context. Suppose further that neither the costs of collecting administrative data (as is often the case) nor the costs of collecting survey data (as is rarely the case) depend on the number of sampled contexts. Under these conditions, clustering samples by context would be undesirable. Such clustering creates a (statistically) inefficient dependence across observations. In the absence of cost savings in collecting interview or contextual information, the optimal design would sample one subject per context.⁵

In most instances, interviewing costs associated with additional subjects per context are substantially lower than costs associated with subjects drawn from different contexts. The tradeoff between interviewing costs and statistical efficiency has long been a concern of sampling statisticians and, in the case of typical household surveys, leads to designs with relatively modest (e.g., 4 to 10) subjects per cluster. Studies that measure context by aggregating characteristics, perceptions or behavior of respondents provide an additional rationale for clustering. Careful consideration of the costs and benefits (e.g., as with the Chicago Neighborhood Study) leads to larger cluster sizes (i.e., between 15 and 30 subjects per cluster).

Contextual variability. The cost of such clustering is that there is less dispersion across contexts than if the same field budget were spent on a less-clustered sample. How limiting is it to restrict the geographic variability of the sample? Unfortunately, the relatively limited variability in neighborhood conditions found in sections of cities or even in entire geographic areas of many cities poses difficult tradeoffs for study design. Duncan and Raudenbush (forthcoming) illustrate the scope of the problem by drawing tract-based data from the 1980 decennial census. They formed subsets of tracts to approximate typical study designs: i) all tracts in the United States (to approximate national samples; ii) all tracts in the city of Chicago (to approximate a large study in a single but diverse city); iii) all Chicago tracts with a 30+% poverty rate (to approximate an “underclass” study in a large city); iv) all tracts in the city of Atlanta (to approximate a large study in a less diverse large city); v) all Atlanta tracts with a 30+% poverty rate (to approximate an “underclass” study in a less diverse large city); and vi) all tracts in the city of Rochester, NY (to approximate a study in a medium-sized city).

They drew from the census files seven tract-level demographic measures often used in neighborhood-based research, each of which is associated with a distinct neighborhood-level theoretical process: *race* - the percentage of individuals in the tract who are black; *female headship* - the percentage of households headed by women; *welfare* - the percentage of households receiving public assistance; *poverty* - the percentage of non-elderly individuals with below-poverty household incomes; *high educational level* - the percentage of adults with college degrees; *neighborhood stability* - the percentage of households who had lived in the same dwelling five years before; and *joblessness* - the percentage of adult males who worked fewer than 26 weeks in 1979.

To assess potential multicollinearity problems using these measures, Duncan and Raudenbush (forthcoming) took the various collections of tracts and regressed each of the neighborhood measures on the remaining six neighborhood measures. Not surprisingly, the resulting R^2 s indicated a great deal more multicollinearity in the city-specific samples than for the national set of tracts. For example, only 29% of the variation in the fraction of college-graduate adults could be accounted for by the other six measures in the national sets of tracts. In the city-specific samples, the squared multiple correlations ranged from .31 to .75 and averaged .50 – nearly twice the degree of multicollinearity in the national set of tracts. Overall, the average extent to which the city-based squared correlations exceeded those for all U.S. tracts ranged from .07 to .25 across the seven measures.

The experiences of the authors contributing chapters to Brooks-Gunn et al. (1997, Volume I) illustrate the same point in a different way. All sought a coordinated analysis of neighborhood effects by matching census-tract-based neighborhood measures (of the kind listed above) to the addresses of children and youth from a variety of samples. In the case of data from two national samples – the PSID and Children of the National Longitudinal Survey of Youth – the sample dispersion provided ample degrees of freedom to support estimates of the effects of a number of theoretically distinct but empirically correlated neighborhood dimensions. This was also the case for data from the eight-city Infant Health and Development Program. However, data from samples of three more specialized studies – children from a single urban school district in upstate New York, youth from high-poverty neighborhoods in New York, Baltimore and Washington, D.C., and children from high-poverty neighborhoods in a large Southern city – only supported estimates of the effects of a single “good” vs. “bad” neighborhood dimension. Research designs that support only one neighborhood dimension are obviously incapable of identifying the nature of neighborhood processes.

V. Some Promising Approaches

Rather than conclude with this depressing list of problems facing analysts of contextual effects, we prefer to draw uplifting examples from recent work that solve at least some of these problems. Our examples include: i) the Project on Human Development in Chicago Neighborhoods, which incorporates a well-conceived design for obtaining measures of

neighborhood constructs; ii) correlation-based approaches to estimating upper bounds on the influences of neighborhood effects; iii) findings from the quasi-experimental Gautreaux project; and iv) preliminary findings from one site in HUD's Moving to Opportunity random-assignment experiment.

The Project on Human Development in Chicago Neighborhoods

Sampson et al. (1997) is a promising approach to measuring context in a way that corresponds closely to theoretical constructs. As part of a study of the delinquent behavior of youth in a sample of Chicago neighborhoods, they measure the "collective efficacy" of neighborhoods by conducting a survey of adult residents in sampled neighborhoods rather than relying exclusively on decennial-census measures.

Few researchers have the resources to conduct independent surveys of neighborhood residents to obtain their contextual measures. The Chicago Project also incorporates systematic social observation ("SSO" Reiss, 1988) as an alternative source of contextual information. Some of its measurement strategies (e.g., videotaping and coding) are cumbersome and expensive. However, in one strategy, trained observers can fairly quickly assess aspects of a neighborhood such as its degree of social and physical disorder. Interviewers dispatched to conduct interviews can also be used to conduct such observations at a cost far less than that of conducting an independent survey of residents.

The Chicago Project implemented its "deluxe" version of SSO by having a van drive five miles an hour down every street within 80 target neighborhood clusters. Videotape recorders on both sides of the van captured physical characteristics of the streets and buildings on each side of the street as well as visible aspects of social interaction. Trained observers then coded the videotapes, noting the status of buildings (residential versus commercial, detached homes or apartments, whether vacant or burnt out, their general condition, presence of security precautions such as bars or grates, etc.), presence of garbage, litter, graffiti, drug paraphernalia, broken bottles, abandoned cars, and other aspects of the physical environment.

The driver and a second rider in the van, trained to observe social interactions, also recorded their observations via audiotape. Social interactions included, for example, adults drinking in public, drug sales, children playing in the street, and apparent gang activity. Scales tapping social and physical disorder, housing conditions, and other aspects of the neighborhood

environment showed high internal consistency across face blocks within neighborhood clusters and reasonably high construct validity as indicated by correlations with theoretically linked constructs measured by an independent community survey, by the census, and by official crime data. Analyses now underway are estimating the value added by the videotapes, above the information gleaned from the audiotapes. Generally, the videotaped data are far more expensive than the audiotaped data. It is feasible to use the audiotape strategy even when samples are not highly clustered because data collection per blockface is comparatively cheap.

SSO has substantial promise for efficient collection of data on the social organization of neighborhoods -- data not available from administrative records. However, some of the constructs that can be captured through interviews, such as “collective efficacy” in Sampson et al. (1997), are not accessible via observational methods. Given the expense of interviewing residents in unclustered samples, researchers interested in neighborhood effects face difficult tradeoffs, discussed below. Similar tradeoffs face school researchers, who might opt for observational measures (c.f., Mortimore et al., 1988) as an alternative to survey methods designed to capture school organization and climate.

Correlations as upper-bound estimates of the influence of context

An alternative approach to the problem of unobservable neighborhood variables relies on correlations between children who are neighbors or classmates or, equivalently, on the explained variance of neighborhoods, schools or classmates to provide an upper bound on the possible effect of these contexts. To motivate the logic of this approach, we first note that many studies have used sibling correlations to estimate the importance of shared family and other environmental experiences. For example, sibling correlations for years of completed schooling are quite high – around .55 – indicating that there are important elements of the genes, family environments, neighborhoods, schools and other aspects of the shared environments of siblings that make siblings much more alike in terms of completed schooling than two individuals drawn at random from the population.

Solon, Page, and Duncan (1997) argue that an analogous correlation for children growing up in the same neighborhood but not in the same family indicates how much of what is important in the shared environments of siblings lies outside the immediate family. A high completed-schooling correlation for unrelated neighbor children, for example, is consistent with a strong

neighborhood effect and would imply that shared neighborhood conditions are an important component of the sibling correlations. (An alternative interpretation is that the extra-familial correlations are driven by the often-similar family backgrounds of children in neighboring families.) Neighbor correlations close to zero would suggest that the scope for pure (i.e., extrafamilial) neighborhood effects is quite small.

The beauty of sibling and neighbor correlations is that they provide an upper bound on the estimated impact of both measurable and unmeasurable aspects of the environments shared by siblings and neighbors. In particular, neighbor correlations address the omitted-context-variables problem because the strength of the correlations does not depend on whether the contextual factors driving them can be measured. Of course, a limitation of these correlations is that they reveal nothing about the process by which familial and extra-familial influences operate to make siblings and neighbors more alike than two individuals drawn at random from the population. Nor do the neighborhood correlations address either the simultaneity or endogenous-membership problem.

Solon et al. (1997) formalize these arguments in terms of the same kind of additive model given in (1):

$$(7) \quad y_{sfc} = A'FAM_{fc} + B'CON_c + e_{sfc}.$$

As before, c references the extra-familial context (e.g., school class, neighborhood, peer group) and f references families. Siblings within families are referenced with s . FAM_{fc} is a vector of shared family influences for all siblings within the same family; CON_c is a vector of extra-familial contextual influences (henceforth called “neighborhood” but applicable to other contexts) shared by all siblings and neighboring children; and e_{sfc} is an error term.

Solon et al. (1997) show that the sibling covariance in y_{sfc} can be expressed as:

$$(8) \quad Cov(y_{sfc}, y_{sfc'}) = Var(A'FAM_{fc}) + Var(B'CON_c) + 2Cov(A'FAM_{fc}, B'CON_c),$$

i.e., the sum of shared family variance, shared neighborhood variance and twice the covariance between family and neighborhood factors.

The covariance between neighboring children from different families is:

$$(9) \quad Cov(y_{sfc}, y_{s'fc'}) = Var(B'CON_c) + 2Cov(A'FAM_{fc}, B'CON_c) + 2Cov(A'FAM_{fc}, A'FAM_{fc'}),$$

i.e., the sum of the shared neighborhood variance and twice the covariance between family and neighborhood factors plus twice the covariance in family backgrounds among neighboring children. In comparing sibling (in (8)) and neighbor (in (9)) covariances, it can be seen that shared neighborhood variance and covariance between family and neighborhood factors are common to both. The shared family variance is obviously missing from the non-family neighbor covariance, while the family-background covariance of neighboring children is missing from the sibling covariance but is a part of the non-family neighbor covariance.

National surveys such as the PSID and NLSY draw their samples from a set of tightly clustered neighborhood areas that often encompass only one or two blocks. Thus these clusters approximate neighborhood areas and, using anonymous cluster identification, it is possible to calculate both sibling and neighbor correlations for various outcomes of interest. Solon et al. (1997) calculate such sibling and neighbor correlations with a representative PSID sample consisting of individuals age 8-16 in 1968. For their outcome measure – years of completed schooling -- the sibling correlation (.54) is much higher than the estimated neighbor correlation (.19), suggesting a rather limited scope for the effects of extra-familial contexts. After removing effects of easily observed socioeconomic characteristics of families (race, family income, family structure and maternal schooling) shared by children living in the same neighborhood, the neighbor correlation dropped further (to .10), suggesting an even more limited scope for unique neighborhood influences. All in all, the results suggest that neighborhood influences are much more limited than family influences in accounting for individual differences in completed schooling.

Correlations from Add Health. Duncan, Boisjoly and Harris (1998) draw data from the National Longitudinal Survey of Adolescent Health (Add Health) to generalize this approach.⁶ They use correlations between siblings within a family, between grademates within a school, between schoolmates residing in the same Census block group, and between peers as defined by a set of “best friend” nominations as respective upper-bound estimates of the potential influence of family, school, neighbors and peers.⁷ Add Health is a nationally representative study of adolescents in grades 7 through 12 in the U.S. in 1995, the vast majority of whom responded to an in-school, self-administered questionnaire, and a systematically chosen subset of whom responded to two waves of at-home personal interviews.

The data are uniquely able to characterize sibling, neighborhood, peer and school environments of sample members. The sample itself is clustered within 134 schools drawn from a school-based sampling frame. In-home interviews administered to a random subset of students from each school provide representative samples of schoolmates that can be used to generate schoolmate-based correlations. The in-school questionnaires administered to all students asked each adolescent respondent to name his or her five best male friends and five best female friends, providing data for best-friend correlations. We distinguish instances where best friend nominations are and are not mutual. In-home data can also be aggregated across schoolmates residing in the same neighborhood, with neighborhood residence established through address matching to Census block group.⁸ The design provides data on friendship and neighbor pairs that can span school grades, but must be from sampled schools. All twin pairs found in the schools were included in the in-home interviewing, providing substantial sample sizes for monozygotic twin, dizygotic twin and non-twin siblings.

Two key outcomes available in the data are the Add Health Peabody Picture Vocabulary Test, an achievement/ability measure based on receptive vocabulary, and a delinquency scale. Both are continuous measures, the first with a nearly normal distribution. The delinquency scale consists of items on painting graffiti, damaging property, shoplifting, getting into a serious physical fight, stealing a car, stealing something worth less than \$50, stealing something worth \$50 or more, burglarizing a building and selling drugs.

Table 1 presents correlations in age-adjusted PPVT and delinquency scores for various groups of male siblings, friends, neighborhoods and grademates. Brothers are divided into MZ and DZ twins, non-twin full siblings close and not close in age, and half-siblings. Best-friends, neighbor and grademate correlations are computed both on scores that incorporate a simple age adjustment and on scores that also adjust for a handful of readily observed measures of family SES – income, parental education and family structure. Furthermore, grademate correlations are computed separately for middle- and high school grademates to test the hypothesis that context-driven correlations strengthen as children move from early to later adolescence.

PPVT correlations across family groups have a similar pattern to those found in the voluminous literature on full-scale IQ correlations.⁹ MZ twin correlations (.80) are considerably higher than DZ twin correlations (.56). Simple behavioral genetics models produce 48% and

32% estimates of heritability and shared environmental influences, respectively.¹⁰ As with the Solon et al. (1997) neighbor correlations of completed schooling, grademate and neighbor correlations are much smaller than sibling correlations, and drop even more in the presence of adjustments for family SES. The results imply an upper-bound explained variance of 4% for neighborhoods effects and 5% for school effects on achievement. At .46 (and, SES-adjusted, .33), mutual best-friend correlations are closer to sibling correlations. Correlations for non-mutual best friends (i.e., where person A named B as best friend but B named as best friend someone other than A) are considerably lower. Unfortunately, it is impossible with correlation-based method to determine upper bounds on either the exogenous or endogenous component of peer effects.

Correlations for delinquency are not as high as those for the test scores, and are much less affected by adjustments for family SES. In this case, simple behavioral genetics models imply 30% and 15% estimates of heredity and shared environmental influences on delinquency.¹¹ As with test scores, the delinquency correlations among neighbors (.02) and grademates (.00 to .04) suggest at most a very modest scope for neighborhood and school effects. That grademate correlations are smaller for older as compared with younger adolescents is surprising, and indicates that school-based contextual influences may be stronger at the earlier ages. At .28, best friend correlations are higher than all but the twin correlations.

Caveats regarding correlations. There are a number of important qualifications for using correlations among classmates, peers and neighbors as upper-bound estimates of the importance of extra-familial contexts. First, the transitory nature of neighborhoods, schools, and, especially, peer groups (Urberg, et al., 1995) in the lives of children may impart a measurement-error-induced downward bias to the correlations. For example, since residential mobility is quite common in the United States, especially among younger children, children sharing a neighborhood at any given point may have quite different residential histories. However, this bias may not be large since residential moves typically occur between similar neighborhoods (Solon et al., 1997). Peer “mobility” also typically occurs among individuals with similar characteristics (Urberg, et al., 1995).

Second, the endogenous nature of context, particularly peer groups, will cause best-friend correlations to overstate, perhaps dramatically, the causal role played by contexts. Third, the

differential reliability of our various outcome measures will impart correspondingly differential bias to our correlations.

Third, there is no easy way with the correlation method of allowing for nonlinear neighborhood effects, so this technique says little about how living in an extremely disadvantaged (or, for that matter, advantages) neighborhood might affect outcomes.

Finally, and perhaps most important from a policy perspective, effect sizes that program evaluators commonly view as moderate or even large translate into small proportions of variance in individual outcomes “explained” by neighborhood membership (Duncan and Raudenbush, forthcoming; Cain and Watts, 1972; Rosenthal and Rubin, 1982) and into small intra-neighborhood correlations. Duncan and Raudenbush (forthcoming) consider standardized effect sizes -- that is, standardized mean differences between a set of experimental neighborhoods and an equal number of control neighborhoods, commonly viewed as small ($d = .2$ of a standard deviations, medium ($d = .4$), large ($d = .6$) or very large ($d = .8$). These effect sizes would give rise to the intra-neighborhood correlations given below:¹²

Standardized mean difference	Approximate intra-neighborhood correlation (ICC)
$d = .20$	ICC = .01
$d = .40$	ICC = .04
$d = .60$	ICC = .08
$d = .80$	ICC = .14

Thus, even very large effect sizes translate into correlations generally regarded as small. A small correlation between neighbors does not rule out a large effect size associated with a measured difference between neighborhoods.¹³

Gautreaux as a quasi-experimental design

Rosenbaum (1991) was able to circumvent endogenous membership bias by using data from an unusual quasi-experiment involving low-income black families from public-housing

projects in Chicago. As part of the Gautreaux court case, nearly 4,000 families volunteered to participate in a subsidized program that arranged for private housing, much of it in predominantly white Chicago suburbs, but some of it in predominantly white sections of the city of Chicago itself.

The program's procedures create a quasi-experimental design with respect to the initial neighborhoods in which participants are placed. While all participants come from the same low-income black city neighborhoods (usually public housing projects), some move to middle-income white suburbs, while others move to white and black urban neighborhoods. Until 1990, participants were assigned to city or suburb locations in a quasi-random manner. Apartment availability was determined by housing agents who do not deal directly with clients, and availability was unrelated to client interest. Counselors offered clients units as they became available according to their position on the waiting list, regardless of clients' locational preference. Until 1990, counselors did not honor clients' preferences because it was feared that others would demand similar treatment, creating bottlenecks and conflicts. Although clients can refuse an offer, very few have done so, since they risk not getting another in the six-month period of their eligibility. Consequently, participants' preferences for city or suburbs have virtually nothing to do with where they end up moving.

The quasi-random assignment of Gautreaux families to their (in this case new) contexts provides statistical leverage against the endogenous membership problems by all but eliminating the correlation between family characteristics (both measurable and not) and context. A disadvantage is that inferences regarding neighborhood effects from these data are limited to low-SES families willing to volunteer for such programs.

Rosenbaum and his colleagues analyze Gautreaux youth outcomes using data from interviews conducted with these children and their mothers in 1989, at which point the age of the children averaged 18 (Kaufman and Rosenbaum, 1992). Among their findings:

- More city movers dropped out of high school than did suburban movers (20% in the city vs. less than 5% in the suburbs).
- Although test scores were not available for individual respondents, they found that suburban movers had virtually the same grades as city movers (a C+ average in city and suburbs). Since suburban students usually get about a half-grade lower than city students

with the same achievement test scores, the grade parity of the two samples implies a higher achievement level of suburban movers.

- Although research finds that blacks are underrepresented in college tracks in racially integrated schools (Coleman, et al., 1966; Rosenbaum and Presser, 1978), the Gautreaux results showed that suburban movers were more often in college tracks than city movers (40% vs. 23%).
- While the higher suburban standards might be a barrier to college attendance by these youths, we found that suburban movers had significantly higher college enrollment than city movers (54% vs. 21%).
- Among the Gautreaux youth attending college, almost 50% of the suburban movers were in four-year institutions, whereas only 20% of the city movers were. Of those not attending four-year institutions, two-thirds of the suburban movers were working toward an associate's degree, while just half of the city movers were.
- For youth not attending college, a significantly higher proportion of the suburban youth had full-time jobs than city youth (75% vs. 41%). Suburban youth were also four times as likely to earn over \$6.50 an hour as city youth (21% vs. 5%). The suburban jobs were significantly more likely to offer benefits than city jobs (55% vs. 23%).

Crucial questions for reconciling the large effects found by Rosenbaum with the more modest ones found in the nonexperimental literature are: to what extent his use of quasi-experimental data better addresses the endogenous membership problem; whether large neighborhood effects exist for underclass blacks but not for other population groups; and whether the volunteer nature of his sample produces larger effects than would be the case for a more general sample of low-income, inner-city blacks.

It is also important to note that quasi-experiments such as Gautreaux can help avoid bias problems in assessing the effects of neighborhood conditions on family conditions – the “family as mediator” issue raised earlier. Popkin, Rosenbaum, and Meaden (1993) compared city vs. suburban movers on employment outcomes for mothers and found substantially greater employment (but not higher wage rates) for mothers assigned to suburban as opposed to city locations.

The Moving to Opportunity experiment

With funding for 10 years, MTO is randomly assigning housing-project residents in five of the nation's largest cities to one of three groups: i) a group receiving housing subsidies to move into low-poverty neighborhoods (called here the *experimental group*); ii) a comparison group receiving conventional Section 8 housing assistance but not constrained in their locations (the *Section 8 group*); iii) a second comparison group receiving no special assistance (the *control group*).

Ludwig et al. (1998) use the experimental data from the Baltimore site to: i) evaluate the effects of the two program components on the frequency of criminal activity among adolescents, as reflected in the Maryland Department of Juvenile Justice's criminal-offender records; and ii) use the experimental variation as part of an instrumental-variables procedure to estimate a model of the effects of neighborhood poverty on such criminal activity.

By way of background, eligibility for the Baltimore MTO program was limited to families with children who lived in the five poorest census tracts (average poverty rate of 67% in the 1990 Census) in Baltimore City. Virtually all of the families who volunteered for the program were African-American and headed by females. As with families at all other MTO sites, baseline surveys in Baltimore reveal that escaping from gangs and drugs was the most important stated reason for participating in the MTO program.

Not all families randomly assigned to experimental and Section 8 groups relocated during their six-month eligibility periods. Indeed, only about half of the experimentals moved, with 90% of this group moving to census tracts with poverty rates under 10%. Three-quarters of the Section 8 families moved, with 16%, 27% and 19% of these moving to census tracts with poverty rates under 10%, 10-20% and 20-30%, respectively. Thus, the program's randomization applies to the "intention to treat" volunteer families with the three treatments, but not to the actual neighborhood conditions chosen by those families.

Ludwig et al.'s (1998) analysis sample is restricted to 358 13-17-year-old children who continued to live in the state of Maryland for at least one year following baseline random assignment to the experimental (n=157), Section 8 (n=94) and control (n=107) groups. Juvenile justice records reveal that 15% of the 358 had been arrested for a violent offense (rape, robbery, assault) prior to baseline randomization, 12% had been arrested for a property offense (burglary,

auto theft or theft/larceny), and 11% had been arrested for a collection of “other” offenses (e.g., disorderly conduct, weapons, drugs).

Table 2 shows key results from a comparison of the experimental and Section-8 program groups relative to the controls on pre- to post-program changes in violent, property and other crime arrests for the 358 MTO children who were ages 13 to 17 for at least one year during the post-program.¹⁴ For girls, there are no significant differences in crime rates across either of the program groups.

For boys, the data reveal a number of interesting differences. First comparing experimental and controls, there is a sizable and statistically significant reduction in the proportion who are arrested for violent (17 percentage points) and “other” (13 points) offenses. The point estimate of program effects on property crime was negative but not statistically significant. Effect sizes for the Section 8 vs. control group comparison were smaller for violent crime but similar for the property and “other” crime categories.

These “intention to treat” estimates of the effects of the MTO program offer follow directly from the random-assignment nature of the program. But it is also possible to go beyond direct program assessments and use random assignment to identify more general models of the relationship between neighborhood characteristics and juvenile crime. Simply comparing experimental relocators with control families will produce biased estimates of the effects of neighborhood poverty on juvenile crime, since families are randomized with respect to treatment groups rather than where they actual residential locations. However, the MTO random assignment can be used as an instrumental variable in a model of neighborhood effects since treatment assignments influence relocation outcomes by changing the “price” of relocation, yet by construction are uncorrelated with the unobservable determinants of juvenile crime.

Ludwig et al. (1998) use random assignment to instrument neighborhood poverty rates as part of a model of the effects of neighborhood poverty on juvenile criminal outcomes. It is important to note that the analysis does not identify the specific attributes of the neighborhoods that are responsible for these effects, so they could not distinguish between the effect of neighborhood poverty itself from dimensions of neighborhood quality that are correlated with poverty.

The first stage of the instrumental-variables estimation strategy consists of estimating the likelihood of residence in a very low- or medium-poverty neighborhood during the post-program period. With these predicted values, one can estimate a model of crime prevalence. Not surprisingly, first-stage estimates showed that treatment group assignments are powerful predictors of families' post-program neighborhood poverty rates. The results of estimating second-stage, crime, equations are presented in Table 3. For males, living in a low-(<10%) rather than high- (>40%) poverty neighborhood is estimated to have large effects on the prevalence of arrests for violent (-36 percentage points) and "other" (-41 percentage points) crimes. Both of these effects are statistically significant at conventional levels, although the rather large standard errors associated with these estimates show that they are not estimated very precisely. Compared with high poverty areas, living in a neighborhood with moderate poverty rates (10 to 39 percent) has no clear effects on violent crimes, but may for males have some effects on participation in "other" crimes. For females, in all but one case (an anomalous positive effect of residence in moderate-poverty neighborhoods on property offenses), the results show no effects of neighborhood conditions on crime.

All in all, Ludwig et al. (1998) find that the experimental MTO treatment in particular, and neighborhood poverty (and/or its correlates) more generally, appear to have sizable effects on juvenile arrests. Not surprisingly, these effects are concentrated among the subgroup of teens who tend to be most criminally active: males, and those in the highest-crime years of 13 to 17.

V. Conclusion

Correlational studies based on youth from general-population samples indicate that the family (and possibly the genetic component of the family) accounts for much more of the variation in youth achievement and behavior than neighborhood and school contexts. Although small in size, the degree of neighborhood-based "action" may still be large enough to be consistent with cost-effective neighborhood-based interventions. Unaddressed in these correlational studies are assessments of why context matters and whether context matters much more for disadvantaged than for general-population youth.

We have argued that regression-based approaches to estimating models of contextual effects on youth outcomes face three daunting methodological challenges: simultaneous

causation, omitted (contextual) variables, and endogenous membership. In the case of neighborhood context, the second and third of these are particularly problematic. Also problematic in these kinds of efforts are modeling the role of the family and ensuring adequate variability in contextual conditions across the youth sample.

We conclude that convincing quantitative assessments of the effects of neighborhood conditions on youth achievement and behavior require either: i) nonexperimental data, drawn from geographically-dispersed samples, containing theoretically-motivated contextual measures, and estimated with models that address problems of simultaneity, omitted context variables and endogenous membership; or ii) experimental or quasi-experimental data, based on theoretically interesting samples and experimental conditions, that are estimated with models that address problems of simultaneity and omitted-context variables.

Very few existing studies are in a position to provide the needed data. In the case of nonexperimental data, we have argued that the design of the Project on Human Development in Chicago Neighborhoods is an important advancement toward developing measures of context that correspond closely to theoretical constructs. Neighborhood studies conducted with representative youth samples in Philadelphia, Prince Georges County (MD) and Chicago by the MacArthur Network on Youth in High-Risk Settings may also solve this problem. It remains to be seen whether analyses of these data can address simultaneity and endogenous membership problems.

It is conceivable that data from national-sample studies such as the PSID, NLSY and Add Health could support convincing studies of neighborhood effects, but only if such studies address our list of model-based concerns through innovative statistical modeling. PSID and NLSY-based assessments of measured neighborhood characteristics are confined to information available in administrative data sources – an important constraint. Context measurement is not as problematic in the Add Health design, which provides a wealth of interview-based assessments about youth, parents and schools from a highly clustered yet geographically dispersed sample. Here again it remains to be seen whether analyses of these data can address the simultaneity and endogenous membership problems as well as avoid the correlated-error problem endemic to studies that construct context measures by aggregating respondent characteristics.

Data from quasi- and randomized experimental studies such as Gautreaux and, especially, Moving to Opportunity provide convincing “intention to treat” assessments of their respective mobility programs as well as crucial leverage for implementing an instrumental-variables-based strategy against the vexing endogenous membership problem. An added advantage from a policy perspective is that these assessments are drawn for samples of disadvantaged, inner-city youth. Disadvantages include the facts that they do little to identify exactly what aspect of context matters the most and that they are often based on samples of families that volunteer for the programs.

Unless they aspire to very expensive “big science” data collections, researchers interested in generating yet more sources of data to assess neighborhood effects should avoid nonexperimental studies that, given resource constraints, are confined to local samples. Secondary analyses of the emerging underanalyzed studies mentioned in this paper (especially Add Health), using statistically-innovative methods, are one potentially fruitful avenue for their efforts.

Researchers intent on new data collections are advised to consider clever ways of taking advantage of Gautreaux-type “natural experiments” that provide exogenous sources of contextual variation. One example is of the public housing relocations associated with the court-order desegregation of public housing in Yonkers. Jeanne Brooks-Gunn is leading an effort to compare families moving to the new integrated public housing site with families who applied for the new housing but, by losing the housing lottery, did not have the opportunity to move.

At a more micro level, studies of peer interactions might be able to circumvent the endogenous membership problem by taking advantage of some natural experiments such as universities’ random assignment of freshman roommates in some dormitories.¹⁵ Comparisons of attitudes, behavior and achievement between randomly-assigned and self-selected roommates provide an indication of the size of the endogeneity problem.

At a more macro level, there is value in before-after comparisons of the effects of a beneficial economic “shock” on neighborhoods and families from natural experiments such as legalized gambling. Sites such as Tupelo, Mississippi are especially interesting since a riverboat gambling industry appeared there almost overnight and restrictive laws in neighboring states have maintained Tupelo’s monopoly on the industry.

Natural experiments such as these have their limitations. But the quasi-experimental variation in context they offer researchers enriches their analytic value to a point far beyond that of most nonexperimental studies.

- Aaronson, D. (1997). Sibling estimates of neighborhood effects. In Brooks-Gunn, J., G.J. Duncan, & L. Aber, Eds. Neighborhood Poverty: Policy Implications in Studying Neighborhoods, Vol II. (pp. 80-93). New York: Russell Sage Foundation.
- Billy, J., and Udry, J. (1985). Patterns of adolescent friendship and effects on sexual behavior. Social Psychology Quarterly, 48, 27-41.
- Bouchard, T., and McGue, M. (1981). Familial studies of intelligence: A review. Science, 212, 1055-1059.
- Brooks-Gunn, J., Duncan, G. and Aber, J. L. (1997). Neighborhood poverty: Context and consequences for children, New York: Russell Sage.
- Brooks-Gunn, J., Duncan, G., Klebanov, P. and Sealand, N. (1993). Do neighborhoods affect child and adolescent development? American Journal of Sociology, 99(2), 353-395.
- Bryk, A., & Raudenbush, S. W. (1992). Hierarchical linear models for social and behavioral research: Applications and data analysis methods. Newbury Park, CA: Sage.
- Cain, G. and Watts, H. (1972). Problems in making policy inferences from the Coleman Report. American Sociological Review, 35(2), 228-252.
- Chase-Lansdale, L., Gordon, R., Brooks-Gunn, J. and Klebanov, P. 1997. "Neighborhood and Family Influences on the Intellectual and Behavioral Competence of Preschool and Early School-age Children," in J. Brooks-Gunn, G. Duncan and L. Aber, Neighborhood Poverty: Context and Consequences for Children, New York: Russell Sage.
- Clark, R. (1992). Neighborhood effects on dropping out of school among teenage boys. Mimeo. Washington, D.C.: Urban Institute.
- Coleman, J. (1966). Equality of educational opportunity. Washington, D.C.: GPO.
- Cook, T., Kim, Jeong-Ran, Chan, Wing-Shing and Settersten, R. (1998). How do neighborhoods matter? In F. Furstenberg, Jr. et al. (Eds.), Managing to make it: Urban families in high risk neighborhoods. Chicago: University of Chicago Press.
- Crane, J. (1991). The epidemic theory of ghettos and neighborhood effects on dropping out and teenage childbearing. American Journal of Sociology, 96(5), 1126-1159.

- Duncan, G. and Raudenbush, S. (Forthcoming). Assessing the effects of context in studies of child and youth development. Educational Psychology.
- Duncan, G., Boisjoly, J. and Harris, K. (1998) Sibling, peer and schoolmate correlations as indicators of the importance of context for adolescent development. Paper presented at the biennial meetings of the Society for Research on Adolescence, San Diego, February, 1998.
- Duncan, G., Connell J. and Klebanov, P. (1997) Conceptual and methodological issues in estimating causal effects of neighborhoods and family conditions on individual development, in Brooks-Gunn, J., Duncan, G. and Aber, J. L. (eds). Neighborhood poverty: Context and consequences for children, Volume I, New York: Russell Sage.
- Evans, W. N., Oates, W. E., and Schwab, R. M. (1992). Measuring peer group effects: A study of teenage behavior. Journal of Political Economy, 100, 966-991.
- Falconer, D. (1981). Introduction to quantitative genetics. New York: Longman.
- Furstenberg, F. F. (1993). How families manage risk and opportunity in dangerous neighborhoods in W. J. Wilson (ed.), Sociology and the Public Agenda. Newbury Park, CA: Sage Publications.
- Furstenberg, Jr., Cook, T., Eccles, J., Elder, G. and Sameroff, A. (1998). Managing to make it: Urban families in high risk neighborhoods. Chicago: University of Chicago Press.
- Greene, W. (1993). Econometric analysis. 2nd Edition. New York: MacMillan.
- Griliches, Z. (1979). sibling models and data in economics: Beginnings of a survey. Journal of Political Economy, 87, S37-S64.
- Hausman, J. & Wise, D. (1979). "Attrition bias in experimental and panel data: The Gary Income Maintenance Experiment." Econometrica, 47(2): 455-473.
- Hofferth, S. and Chaplin, D. 1994. Caring for Young Children while Parents Work: Public Policies and Private Strategies. Washington, DC: The Urban Institute.
- Jencks, C. & M.D. Brown. (1975). Effects of high schools on their students. Harvard Education Review, 45, 273-324.

- Jencks, C. and Mayer, S. (1990). The social consequences of growing up in a poor neighborhood. In L. Lynn and M. McGahey (eds.), Inner-City Poverty in the United States (pp. 111-186). Washington, D.C.: National Academy Press.
- Kandel, D.B. (1978). Homophily, selection, and socialization in adolescent friendships. American Journal of Sociology 84, 427-436.
- Kaufman, J. & J. Rosenbaum (1992). The education and employment of low-income black youth in white suburbs. Educational Evaluation and Policy Analysis, 14(3), 229-240.
- Korbin, J. & Coulton, C. (1997). Understanding the neighborhood context for children and families: Combining epidemiological and ethnographic approaches, in Brooks-Gunn, J., Duncan, G. and Aber, J. L. (eds). Neighborhood poverty: Policy implications in studying neighborhoods, Volume II, New York: Russell Sage.
- Lee, V., & Bryk, A. (1989). A multilevel model of the social distribution of educational achievement. Sociology of Education, 62, 172-192.
- Ludwig, J., Duncan, G., & Hirschfield, P. (1998). Urban poverty and juvenile crime: Evidence from a randomized housing-mobility experiment. Mimeo. Northwestern University.
- Manski, C. (1993). Identification of endogenous social effects: The reflection problem. Review of Economic Studies, 60, 531-542.
- Manski, Charles F. (1996). Learning about treatment effects from experiments with random assignment of treatments. The Journal of Human Resources, 31, 709-733.
- Moffitt, R. (1998). Policy interventions, low-level equilibria, and social interactions. Johns Hopkins University Working Paper.
- Mortimore, P., Sammons, P., Stoll, L., Lewis, D., & Ecob, R. (1988). School matters. Los Angeles: University of California Press.
- Nichols, . (1978). Twin studies of ability, personality and interests. Homo, 29, 158-173.
- Plomin, R., J.C. DeFries, & G.E. McClearn (1990). Behavioral genetics: A primer (2nd Ed.), New York: W. H. Freeman.
- Popkin, S., J. Rosenbaum, & P. Meaden (1993). Labor market experiences of low-income black women in middle-class suburbs: evidence from a survey of Gautreaux program

- participants. Journal of Policy Analysis and Management, 12, 556-73.
- Raudenbush, S., Rowan, B., & Kang, S. (1991). A multilevel, multivariate model for studying school climate in secondary schools with estimation via the EM algorithm. Journal of Educational Statistics, 16(4), 295-330.
- Reiss, A. J. Jr. (1988). Systematic observation surveys of natural social phenomena. In Perspectives on attitude assessment: Surveys and their alternatives. (pp. 132-50). Proceedings of a Conference Held at the Bishop's Lodge, Santa Fe, New Mexico. Prepared for the Office of Naval Research. Distributed by NTIS.
- Rosenbaum, J. (1991). Black pioneers—do their moves to the suburbs increase economic opportunity for mothers and children? Housing Policy Debate, 2(4), 1179-1213.
- Rosenbaum, J.E. and S. Presser (1978). Voluntary racial integration in a magnet school. School Review, 86(2), 156-186.
- Rosenthal, R., & Rubin, D. (1982). Comparing effect sizes of independent studies. Psychology Bulletin, 92, 500-504.
- Rowe, D. (1983). Biometrical genetic models of self-reported delinquent behavior: Twin studies, Behavioral Genetics, 13, 473-489.
- Sameroff and Chandler (1975). Reproductive risk and the continuum of caretaking causality. In Review of Child Development Research, vol. 4, edited by F.D. Horowitz. Chicago: University of Chicago Press.
- Sampson, R. & Lauritsen, J.L. (1994). Violent victimization and offending: Individual, situational and community-level risk factors. In A.J. Reiss & J. Roth (Eds.), Understanding and preventing violence: Social influences, volume 3 (pp. 1-114). Washington DC: National Academy Press.
- Sampson, R., Raudenbush, S., & Earls, F. (1997). Neighborhoods and violent crime: A multilevel study of collective efficacy. Science, 277, 918-924.
- Shaw, C. and McKay, H. (1942). Juvenile delinquency and urban areas. Chicago, IL: University of Chicago Press.

Solon, G., Page, M. and Duncan, G. (1997). Correlations between neighboring children in their socioeconomic status as adults. Mimeo. University of Michigan.

Urberg, K., Degirmencioglu, S., and Tolson, J. (1995). The structure of adolescent peer networks. Developmental Psychology, 31, 540-547.

Wilson, W. J. (1987). The truly disadvantaged: The inner city, the underclass and public policy. Chicago: University of Chicago Press.

Table 1: Correlations in Age-Adjusted PVT and Delinquency Scores Within Various Groups of Adolescent Males

	PVT achievement test score	Delinquency score	Number of pairs
<i>Family-based</i>			
Monozygotic twin boys	.80	.45	141
Dizygotic twin boys	.56	.30	123
Non-twin full brothers <2 years apart	.55	.23	168
Non-twin full brothers >2 years apart	.64	.13	160
Half brothers	.44	.14	107
<i>Mutual best friends</i>			
Not family-SES adjusted	.46	.28	282
Family-SES adjusted	.33	.28	282
<i>Non-mutual best friends</i>			
Not family-SES adjusted	.33	.16	672
Family-SES adjusted	.12	.15	672
<i>Neighbors</i>			
Not family-SES adjusted	.18	.02	153,110
Family-SES adjusted	.04	.02	153,110
<i>7th and 8th grade grademates</i>			
Not family-SES adjusted	.14	.04	16,822
Family-SES adjusted	.05	.03	16,822
<i>10th through 12th grade grademates</i>			
Not family-SES adjusted	.19	.00	166,165
Family-SES adjusted	.05	.00	166,165

Source: Calculations by Johanne Boisjoly using data from the National Longitudinal Survey of Adolescent Health.

Table 2: Experimental and Section 8 vs. Control Regression Coefficients and Standard Errors of the Effects of MTO Program on Fraction Arrested for Juvenile Crime

	% arrested	
	Experimental vs. Controls	Section 8 vs. Controls
MALES		
Violent offenses	-17.3 (7.1)*	-5.7 (10.6)
Property offenses	-9.5 (7.5)	-9.5 (8.9)
Other offenses	-13.2 (6.2)*	-14.5 (5.9)*
FEMALES		
Violent offenses	2.9 (8.6)	-7.5 (4.2)
Property offenses	1.5 (5.2)	5.7 (10.1)
Other offenses	-3.9 (5.8)	0.6 (9.0)

These prevalence estimates of the change in arrest probabilities are derived from a “difference in difference” probit model.

*p<.10

Table 3: Instrumental Variables Coefficient and Standard Error Estimates for Effects of Neighborhood Poverty on Juvenile Crime

		% arrested	
		Residence in <10% vs. >40% poverty tract	Residence in 10%-39% vs. >40% poverty tract
MALES			
	Violent offenses	-36.3 (18.8)*	-3.6 (22.7)
	Property offenses	-17.5 (16.2)	-30.2 (22.0)
	Other offenses	-40.7 (16.3)*	-49.2 (23.4)*
FEMALES			
	Violent offenses	-2.3 (13.3)	-9.4 (15.7)
	Property offenses	8.4 (10.0)	25.5 (13.3)*
	Other offenses	-1.6 (12.1)	7.1 (19.9)

These prevalence estimates of the change in arrest probabilities are derived from a “difference in difference” probit model.

* p<.10

¹ Also of note is the potential for a nonlinear relationship between context and outcomes. Jencks and Mayer (1990) discuss the policy importance of nonlinearities, pointing out that the net impact of redistributing contextual resources from the rich to the poor could produce a net gain or loss depending on the relative sizes of gains and losses to poor and rich adolescents affected by the policies. Typically, nonlinearities can be readily handled in the context of equation 1.

² Moffitt (1998) presents a more general discussion of these problems under the heading “correlated unobservables.”

³ We emphasize that the possible endogeneity of group membership is distinct from the possible endogeneity of the social interactions themselves. As noted in our section on simultaneity, the problem of determining whether youth *j*'s behavior affects youth *i*'s behavior is distinct from the problem of determining why *i* and *j* choose to become best friends.

⁴ More formally, the identifying variable must be uncorrelated with the error term(s) in (5).

⁵ While such a design would be optimal for estimating regression coefficients, it provides no information about variation within and between contexts. To the extent it is important to gauge the magnitude of unmeasured sources of variation within and between contexts, the unclustered sample design is problematic, even when costs are ignored. Clustering observations within contexts also enables analysts to use “fixed effects” regression techniques to estimate family models that are free from bias from neighborhood factors. In these models, all families in a neighborhood area are identified and the data are transformed by subtracting each adolescent's measure from neighborhood averages. In capitalizing on intra-neighborhood variance, fixed-effects models produce no explicit estimate of neighborhood effects, but they do purge family effect estimates of neighborhood-based bias.

⁶ The cited paper uses data from Add Health's in-school survey. The data presented here are preliminary and are taken from the in-home survey but follow the same procedures.

⁷ Behavioral geneticists have used correlations in this way for nearly a century in studying family influences (Plomin et al., 1990), as have sociologists for the last quarter century when examining school effects (e.g., Jencks and Brown, 1975; Coleman, 1966) and peer effects (e.g., Kandel, 1978; Billy and Udry, 1985).

⁸ We are grateful to John Billy for supplying us with the necessary anonymized block-group identifiers.

⁹ Bouchard and McGue (1981) conducted a meta-analysis of 212 IQ correlation studies and report the following weighted average correlations for pairs reared together: .86 for MZ twins; .60 for DZ twins; .47 for non-twin siblings and .31 for half-siblings. In the case of twin studies of verbal comprehension, Nichols (1978) reports average correlations for identical and fraternal twins of .78 and .59.

¹⁰ Behavioral geneticists use the pattern of correlations among siblings and parent-child pairs with varying degrees of genetic relatedness and co-residence during childhood to estimate the role of genes and shared and unshared environmental influences (Falconer, 1981). Roughly speaking, a personality trait with a mixture of purely genetic and unshared environmental causes should produce outcome correlations twice as high in monozygotic (MZ, i.e., one-egg) twins as in either dizygotic (DZ, i.e., two-egg) twins or siblings born at different times, since MZ twins have 100% genetic relatedness while DZ and other full siblings share only 50% genetic relatedness. Under these assumptions, the extent to which the MZ correlation is less than perfect reflects the importance of a combination of unshared environmental causes

and measurement error. Allowing for the potential importance of the environments shared by twin and nontwin siblings (e.g., unchanging parenting practices, permanent family resources and, for twins, in-utero conditions) complicates inferences from these correlations. If genes were unimportant and environments similar then one would expect similar MZ, DZ and non-twin sibling correlations. Simple behavioral genetics models suggest that the heritability of a trait equals twice the difference between MZ and DZ correlations and that the role of shared environment can be expressed as twice the DZ correlation minus the MZ correlation.

¹¹ These MZ and DZ twin correlations are considerably smaller than the .71 and .47 correlations reported in Rowe (1983).

¹² Suppose we have two "treatment" groups and we compute the standardized mean difference, d , between those two groups. The proportion of variance explained by "treatment" is then R^2 , which in this case is:

$$R^2 = d^2 / (d^2 + 1 / (p * q))$$

that is, the square of the effect size divided by the sum of that squared effect size and the reciprocal of p times q where p is the proportion of subjects in treatment group 1 and $q = 1 - p$ is the proportion in treatment group 2. This relation is commonly used in meta-analysis. Setting $p = q = .50$ (a balanced design) gives

$$R^2 = d^2 / (d^2 + 4)$$

which is the formula used in the table. We note that the intra-cluster correlation (ICC) is the ratio of the between-cluster variance to the total variance (Bryk and Raudenbush, 1992, Chapter 4). In our hypothetical example, all between-cluster variance is created by the "treatment." Hence, R^2 is equivalent to the ICC. In non-experimental settings, we do not have treatment groups but rather risk groups defined by neighborhood characteristics (that is, we have a "high" and a "low" risk set of neighborhoods.) If the risk groups are not equal in size, the R^2 s for each d will be lower than those in our table. Thus, those in our table set an upper bound on R^2 for any given d , because the maximum $p * q$ occurs at $p = q = .5$ given $p + q = 1$.

¹³ Nor does this imply that interventions producing even very small effects (e.g., $d = .10$) are ill-advised. Cost-effectiveness depends upon effect sizes relative to cost, and it is quite possible for there to be socially-profitable intervention policies in the context of a small effect size.

¹⁴ The estimates are of the effects on juvenile crime of assignment into a particular MTO treatment group, known as the "intent-to-treat" effect (Manski, 1996). The form of the regression is a so-called "difference-in-difference" model following Hausman and Wise (1979), since comparisons of pre- to post-baseline changes across treatment groups produces sharper estimates of program impacts than comparing outcomes across treatment groups in the post-program period. Difference-in-difference estimates also help adjust for the fact that the prevalence of criminal activity for the experimental group was somewhat higher in the pre-program period.

¹⁵ Daniel Levy, Michael Kremer, and Richard Freeman are currently engaged in a research project based on this idea.