Web Appendix
Unpacking Neighborhood Influences on Education Outcomes:
Setting the Stage for Future Research

Table of Contents

In this web appendix we include a more detailed description of effect heterogeneity and an expanded discussion of our illustrative example introduced in Section V.

A. Effect Heterogeneity

The theoretical perspectives in Section II of our chapter provide broad outlines of how neighborhood effects on educational outcomes might work, but none of these frameworks adequately considers how the daily experiences of youth differ within neighborhoods, i.e. how they spend their time, where they spend their time, with whom they spend their time, and how such exposure influences attitudes, frames, expectations, etc. These differences are a potentially important but largely uninvestigated source of effect heterogeneity. Here we provide further discussion to motivate the importance of effect heterogeneity in a conceptual framework as well as in the design of studies and subsequent empirical analysis.

We begin with differences across individuals in social networks as one possible source of neighborhood effect heterogeneity. Though most neighborhood effects theories implicitly assume that neighborhoods play some role in structuring the social networks of their residents, we actually know little about whether—or more importantly for whom—

this is the case, particularly among youth. Social networks are one of the key conduits through which information and cultural frames or scripts are transmitted, but by no means the only one. Social networks of youth of a similar age ("peer networks") have received considerable attention in the literature (e.g. Anderson 1999). Such peer networks may play important roles as cultural conduits, as most theories of peer effects assume, but our theories need to be more specific about who those peers are, which peer attachments are more common among young people in poor neighborhoods, and what is transmitted through peer networks. Harding (2009, 2010) argues that older adolescents and young adults on the street in poor violent neighborhoods have considerable cultural power and play an important role in socializing younger adolescents by exposing them to local cultural frames and scripts regarding schooling and sexual behavior.

A second source of effect heterogeneity is different behavioral adaptations to the challenges of daily life in poor neighborhoods. A focus on behavioral adaptations explicitly considers the individual as an actor who can adapt in different ways to mitigate or overcome challenges faced in different neighborhoods. The distinction developed by Sharkey (2006) between "imposed" environments (everything present in the neighborhood where an individual lives) and "selected" environments (the people and institutions with whom he or she interacts) highlights the idea that youth living in the same neighborhood may choose very different social environments for themselves. Different choices or adaptations can have different consequences. For example, violent neighborhoods provide particular challenges to adolescents. In order to feel safe, some adolescents may engage in behaviors such as demonstrating toughness, altering daily travel routines, forming strong bonds of mutual protection with friends, or relying on

older individuals for protection in order to avoid victimization (Anderson 1999, Harding 2009, 2010). These same behaviors may have unintended educational consequences because they can be interpreted as resistant or disruptive by teachers (Dance 2002). Another example is provided by Carter (2005) who argues that "cultural authenticity" (in the form of speech styles, clothing, music and other tastes) among ethno-racial minority groups can have positive payoffs in terms of group membership and solidarity (what Carter calls "non-dominant cultural capital") but also can be misinterpreted by white middle-class teachers as oppositional or resistant. Parents may also adapt their parenting practices in response to the neighborhood environment, for example by limiting their children's interactions with neighbors (Furstenberg et al. 1999, Jarrett 1997ab).

A third potential source of neighborhood effect heterogeneity is variation in family characteristics and the interaction between family characteristics and the properties of social settings. Here effect heterogeneity is driven less by differences in social interactions and more by differences between individuals and families. We illustrate some of the complexities inherent in this type of effect heterogeneity through the example in Appendix Table 1. This relatively simple example focuses on one type of neighborhood mechanism, a neighborhood resource, and two sources of effect heterogeneity: family capacity and access to non-neighborhood opportunities.

Appendix Table 1: Effect Heterogeneity by Family Capacity for Neighborhood Resource

	High Family Capacity		Low Family Capacity	
	High Access to Outside Opportunities	Low Access to Outside Opportunities	High Access to Outside Opportunities	Low Access to Outside Opportunities

High Neighborhood Resources	+	+	?	–
Low Neighborhood Resources	+	-	?	–

The two rows in the above table represent different types of neighborhoods – neighborhoods with either high or low resources. Without loss of generality, neighborhood resources might include different types of social services and non-school educational programs. Examples would be a neighborhood health center, an after-school program, or summer day camp. Such resources can remove youth from the physical and social dangers of the street by providing safe venues as well as positive socialization, for example through well-designed and supervised community centers with an array of age-appropriate social activities. We assume, however, that these resources are not abundant enough to serve all children in the neighborhood. We also assume that such resources might be acquired elsewhere, should an individual or family be able to access non-local resources. Alternatively, this model could be developed to consider how families of different types were affected by neighborhood deficits, e.g. a lack of safe streets. Presumably, high capacity families would be more capable of minimizing the effects of deficiencies.

The four columns represent different types of families. Particular differences are discussed below. The cell in each column represents a prediction of whether the outcome on some generic variable of interest would be positive or negative. Examples might be educational attainment, presence of behavioral problems, adolescent pregnancy or gang

problems. In this scheme, a neighborhood effect only occurs in a particular column (that is for a particular family type) when the outcome differs depending on whether the neighborhood has a high versus low level of resources.

Our table distinguishes families along two dimensions. First is the capacity of the family. By this we simply mean the ability of a family to take advantage of resources if they are available to them. High capacity families are those that will capitalize on the opportunities available to their children. Low capacity families will not. The level of capacity might be the result of family structure, parental education, family income, the number of children, or parental health.

The second dimension distinguishes families in terms of their access to opportunities outside their neighborhood. These opportunities might have to do with family members or friends. A single mother might send her child each weekend to play with her sister's children in a safer neighborhood. Alternatively, because a neighborhood has good access to public transportation or a family has a car, a parent might well be able to access resources for their children outside the neighborhood that are not available in their own neighborhood, such as a higher quality school.

The above typology results in four types of families: high capacity-high outside opportunities, high capacity-low outside opportunities, low capacity-high outside opportunities, low capacity-low outside opportunities families. Now consider why the level of resources in a neighborhood does or does not affect a particular type of family:

High/High	Families are able to provide necessary resources for their children whether or not those resources are available in their neighborhood. Example: Wealthy family living in the downtown of a large city whose children attend private schools.

High/Low	Because these families have low access to outside opportunities, they only can provide those opportunities to their children if they are available in their neighborhood. Example: A high capacity, but relatively poor family living in a neighborhood with many social services. This is the one family type where neighborhood resources make a difference. Example: A low-income family taking the initiative to enroll their child in a specially tailored school program, like that offered in the Harlem Children's Zone.
Low/High	What the prediction should be here is unclear, though the low capacity of the family means that they will not take advantage of neighborhood resources even when they are present. What is unclear is whether the presence of outside opportunities results in a positive outcome. For example, as Carol Stack argued in *All Our Kin* (1974), the presence of high capacity extended kin members results in positive outcomes for the children of a low capacity parent.
Low/Low	These are families that are low capacity and do not have good outside opportunities. Even if their neighborhood has considerable resources, they are unable to take advantage of them. Example: a single parent addicted to drugs who has alienated her extended family.

It is important to remember that this model is meant to represent a set of ideal types. Obviously, it is quite simplistic. Despite its simplicity, it provides at least two important insights. First, of the four types of families, the resources of the neighborhood only make a difference for one type of family—the high capacity family that does not have good access to opportunities outside the neighborhood. All other families are not affected by neighborhood resources because either they do not need them (they have high access to opportunities outside the neighborhood) or they are unable to use them (they are low capacity).

Second, our schema, if correct, would suggest that policy interventions that seek to enhance place-based resources should focus on three aims. First is providing good neighborhood resources to high capacity families. This might mean either investing in

their neighborhoods or moving them from a low to a high resource neighborhood. Second, policy interventions should be developed to help low capacity families access the resources in their neighborhoods. This means getting local social service organizations to reach out to the most troubled families as opposed to simply helping those who actively seek assistance. Third, our model demonstrates the appreciable difficulties in improving the lives of low capacity families in low resource neighborhoods. To help these families, one needs to get them access to resources generally and, just as importantly, to help them to take full advantage of those resources.

B. Further Random Assignment Design Research Considerations

In Section V of our chapter we very briefly sketched a study design to examine one potential neighborhood effects mechanisms, neighborhood violence. The goal of the hypothetical study was to estimate the effects of neighborhood violence on youth educational outcomes. Here we discuss features of the research design in greater detail. Our aim is not to present a complete research design (or even one that is clearly feasible without vast resources), but rather to discuss aspects of the research design that illustrate ways to apply the ideas discussed in the main text.

Consider first the various approaches to identifying the effect of neighborhood violence. Handled through a standard multivariate regression analysis, any estimate of neighborhood violence on youth educational achievement might be biased by confounding characteristics of the youth, his or her family, and other aspects of the neighborhood environment that influence school performance and that are also correlated with neighborhood violence. This situation leaves researchers in a dilemma: Where can a

researcher find exogenous sources of variation to identify the effects of a neighborhood characteristic like exposure to violence? We need one or more variables that produce variation in exposure to violence but do not otherwise affect the outcome of interest. We can think about this by considering two sources of plausible exogenous variation: variation that is planned and variation emerging from a natural experiment. Planned variation can arise from random assignment or experimental designs. Randomization ensures that individuals (or whatever unit is randomized) assigned to the treatment and control groups are the same across both observed and unobserved characteristics, except for chance variation. One strategy is for the researcher to indirectly manipulate neighborhood conditions by encouraging residential mobility, for example, by offering a housing voucher (such as in the Moving to Opportunity mobility experiment). A second strategy is to directly alter neighborhood characteristics through a place-based intervention. In a placed-based intervention, a random subset of neighborhoods would experience a direct intervention such as newly built after school clubs or the addition of street workers to reduce gang violence, and a random subset of otherwise comparable neighborhoods would not receive the intervention. A third strategy, the one we use in this example, is to randomly manipulate exposure to the neighborhood characteristic.

Well-conceived and properly implemented random assignment studies have high potential for unraveling causal effects, but researchers are faced with trade-offs. Practical and ethical considerations may make the manipulation of certain processes difficult or undesirable no matter how important these processes may be from an explanatory perspective. However, mechanisms that involve, for example, gangs, drugs, or violence can still be examined through experimental manipulations as long as the manipulation

focuses on increasing the safety of the treated group and leaves the control group no worse off. Further, a number of circumstances can dilute the intended treatment effects (often described as issues of noncompliance). For example, participants may not take up the offer of a new opportunity, service or program and ethical considerations often preclude requiring or mandating participation. Researchers also cannot prevent study participants from accessing a similar service on their own, or unintentionally receiving the experimental treatment even if initially assigned to the control group. Second, depending on the mechanisms and outcomes of interest, experiments can be expensive and may require a very long follow-up before meaningful outcomes are measurable. The costs of experiments may limit sample sizes and statistical power to detect effects that are small or moderate in magnitude. Third, a series of practical constraints must be carefully weighed: implementation and noncompliance may mean that not everyone in the treatment group receives the intended intervention, the study population may be narrowly defined and thus results may not be generalizable to full-scale or mandatory programs, and well-done experiments can only focus on a limited number of interventions. As a result, while experiments are useful for identifying causal effects once leading hypotheses have been identified, they may not be as useful for generating hypotheses. In fact, the power of randomized designs to make contributions to the literature may not be fully leveraged until after researchers have analyzed observational and/or qualitative data to generate a narrow set of hypotheses. For more on random assignment study designs, we refer readers to Orr (1999).

The second type of study design that harnesses exogenous variation in the treatment is a “natural experiment” in which policy manipulations or other social or

economic processes not directly related to the outcome create exogenous variation in either residential mobility or neighborhood characteristics. Examples include the residential mobility created by demolition of public housing projects (Jacob 2004), and the use of natural boundaries that determine public school districts (Hoxby 2001). Here researchers must be opportunistic and leverage naturally occurring exogenous variation in the neighborhoods that youth experience. The challenge is that such variation must be known a priori and must be measured. Moreover, any claim that such variation is truly exogenous is based on assumptions about social and economic processes and subject to skepticism from other researchers. The neighborhood effects literature has produced few if any natural experiments whose identification strategies have not been met by significant skepticism.

In the illustrative example in Section V of the chapter, we focus on planned variation through a random assignment design. We propose manipulating the amount of time youth spend in their neighborhood under the assumption that out-of-school time use affects exposure to violence which affects educational performance. Note that there are multiple ways in which exposure to violence may affect educational outcomes (see Section II in the main text) and that many, such as stress or trauma, require long-term cumulative exposure to violence. Since in this example the change in exposure to violence is short term, we focus on outcomes that are more likely to respond immediately to a reduction in exposure to violence, such as school attendance, tardiness, and engagement in school. These responses to exposure to violence can ultimately contribute to achievement.

Our proposed study design is individual random assignment embedded within a matched-neighborhood design. The matched neighborhoods allow us to place youth into neighborhoods that are similar to their own on demographic and economic characteristics but have lower levels of violence. In the first stage, selected high violence neighborhoods are matched to neighborhoods comparable on poverty, race, and education that have lower rates of violence (see Seith et al, 2003 on the feasibility of matched-neighborhood designs). An intervention could be designed that is non-academic and that involves engagement with the neighborhood – such as working on a local clothing drive. For the second stage of the study, we could recruit youth from the neighborhoods with higher levels of neighborhood violence, so that the neighborhoods where subjects are assigned by the intervention are all less risky than their home neighborhoods. We would randomly assign these youth to teams in different locations outside of their own neighborhoods, some of which would have lower neighborhood violence than others. Some type of financial incentive might be designed to encourage participation in the intervention activity.

The first important piece of data collection is baseline information through a survey measuring the characteristics of the individuals and their families prior to the intervention (including a variety of indicators to gauge family capacity as described in the first section of this appendix) and baseline characteristics of the neighborhoods through census or comparable community-level data including crime rates. During the intervention, monitoring of program implementation verifies treatment fidelity and identifies any unintended consequences. To assess the impact of the intervention, data collection would include a follow-up survey to document youth time use (hours spent

engaged in certain activities, where and with whom); longitudinal school records to track youth attendance, disciplinary actions taken by the school, and school grades over the course of the intervention; follow-up information on the characteristics of the experimental and control neighborhoods; and open ended qualitative interviews to capture more nuanced aspects of time use and participation in the program as well as subjects' experiences with schooling and with neighborhood violence both prior to and during the intervention.

Our proposed randomized design to manipulate time use can test the hypothesis that spending time in a lower violence neighborhood results in lower exposure to violence and increases youth school attendance, reduces tardiness and school disciplinary action. Under this hypothesis, we would expect, for example, the difference between post-intervention attendance rates and pre-intervention attendance rates to be larger in groups assigned to lower violence neighborhoods. The intervention manipulates exposure to neighborhoods that differ in their levels of violence but does not directly manipulate exposure to violence itself. We can use survey data and qualitative data to measure the intervention's impact on exposure by asking youth about their experiences with violence in their day-to-day activities and traveling to and from school, and we can consider measures of anxiety or stress that can be captured through survey reports.

Stratification of the sample based on pre-intervention youth characteristics – such as measures of their vulnerability – or family characteristics -- such as family capacity -- could be used to explicitly test hypotheses about effect heterogeneity. For example, we might expect that only highly vulnerable youth in a high violence neighborhood would be affected by exposure to violence. The sample would be stratified by youth vulnerability

and then the youth in each stratum would be randomized to the control or intervention group.

Although this research design is appealing in its simplicity, it is important to note its limitations. The randomized experiment will not allow for testing of other neighborhood effects mechanisms unless one is willing to make very strong assumptions about selection on observables and independent mechanisms. As we discuss in the main text, examining multiple mechanisms requires multiple sources of exogenous variation. Moreover, even well-thought out simple designs can generate perplexing findings that might be due to the unanticipated effects of treatment assignment, and this sometimes necessitates a reconsideration of the initial study design. Nevertheless, when designed well, randomized experiments and natural experiments, such as the one described here focused on neighborhood violence, offer promise for uncovering mechanisms.

Appendix References

Anderson, Elijah. 1999. *Code of the Street: Decency, Violence, and the Moral Life of the Inner-city.* New York: W. W. Norton.

Angrist, Joshua, Guido Imbens, and Don Rubin. 1996. Identification of Causal Effects Using Instrumental Variables. Journal of American Statistical Association 91(343): 444-55.

Brooks-Gunn, Jeanne, Greg J. Duncan, Pamela K. Klebanov, and Naomi Sealand. 1993. "Do Neighborhoods Influence Child and Adolescent Development?" *American Journal of Sociology* 99:353–95.

Carter, Prudence L. 2005. *Keepin' It Real: School Success Beyond Black and White.* New York: Oxford University Press.

Dance, L. Janelle. 2002. *Tough Fronts: The Impact of Street Culture on Schooling.* New York: Routledge Falmer Press.

Furstenburg, Frank F. et al (eds). 1999. *Managing to Make It: Urban Families and Adolescent Success*. Chicago: University of Chicago Press.

Grannis, Rick. 1998. "The Importance of Trivial Streets: Residential Streets and Residential Segregation." *American Journal of Sociology* 103(6): 1530-64.

Harding, David J. 2009. "Violence, Older Peers, and the Socialization of Adolescent Boys in Disadvantaged Neighborhoods." *American Sociological Review* 74(3): 445-464.

Harding, David J. 2010. *Living the Drama: Community, Conflict, and Culture Among Inner-City Boys.* Chicago: University of Chicago Press.

Hoxby, C. M. 2001. "All School Finance Equalizations Are Not Created Equal." *The Quarterly Journal of Economics, 116,* 1189-1230.

Jacob, Brian 2004. "Public Housing, Housing Vouchers and Student Achievement: Evidence from Public Housing Demolitions in Chicago." *American Economic Review*. 94(1): 233-258.

Jarrett, Robin L. 1997a. "African American Family and Parenting Strategies in Impoverished Neighborhoods," *Qualitative Sociology* 20: 275-288.

Jarrett, Robin L. 1997b. "Bringing Families Back In: Neighborhood Effects on Child Development." Pp. 48-64 in Jeanne Brooks-Gunn, Greg J. Duncan, and J.

Lawrence Aber, eds. *Neighborhood Poverty*, Vol. 2: *Policy Implications in Studying Neighborhoods*. New York: Russell Sage

Krieger N, Chen JT, Waterman PD, Soobader MJ, Subramanian SV, Carson R. 2002. "Geocoding and Monitoring of US Socioeconomic Inequalities in Mortality and Cancer Incidence: Does the Choice of Area-Based Measure and Geographic Level Matter?: the Public Health Disparities Geocoding Project." *American Journal of Epidemiology.* 156:471–482.

Orr, Larry L. (1999). *Social Experiments: Evaluating Public Programs with Experimental Methods.* Thousand Oaks, CA: Sage Publications.

Sampson, Robert J. and Steve Raudenbush. 1999. "Systematic Social Observation of Public Spaces: A New Look at Disorder in Urban Neighborhoods." *American Journal of Sociology* 105: 603-651.

Seith, David C., Verma, Nandita, Bloom, Howard S., and George C. Galster. (2003). Exploring the Feasibility and Quality of Matched Neighborhood Research Designs. MDRC.

Sharkey, Patrick T. 2006. "Navigating Dangerous Streets: The Sources and Consequences of Street Efficacy." *American Sociological* Review 71: 826-846.

Stack, Carol. 1974. *All Our Kin: Strategies for Survival in a Black Community*. New York: Harper & Row.