



TERRY
SANFORD INSTITUTE
OF PUBLIC POLICY

D U K E

*Assigning Deviant Youths
to Minimize Total Harm*

Philip J. Cook
Jens Ludwig

Working Papers Series
SAN05-04

May 2005



Assigning Deviant Youths to Minimize Total Harm

Philip J. Cook, Duke University
and
Jens Ludwi, Georgetown University

May 17, 2005

Correspondence:
Philip J. Cook
Sanford Institute of Public Policy
Duke University
Box 90245
Durham, NC 27708-0245
pcook@duke.edu

Acknowledgements: The authors prepared this paper for the Executive Session on Deviant Peer Contagion, organized by Ken Dodge and Tom Dishion. We thank Al Blumstein, Ken Dodge, and other members of the Executive Session for very helpful comments. We also benefited from discussions with Allen Buchanan.

Assigning Deviant Youths to Minimize Total Harm

Philip J. Cook and Jens Ludwig

Abstract

A common practice in the fields of education, mental health, and juvenile justice is to segregate problem youths in groups with deviant peers. Assignments of this sort, which concentrate deviant youths, may facilitate deviant peer influence and lead to perverse outcomes. This possibility adds to the list of arguments in support of “mainstreaming” whenever possible. But there are other concerns that help justify segregated-group assignments, including efficiency of service delivery and protection of the public. Our analysis organizes the discussion about the relevant tradeoffs. First, the number of deviant youths (relative to the size of the relevant population, or to the number of assignment locations) affects whether the harm-minimizing assignment calls for diffusion or segregation. Second, the nature of the problematic behavior is relevant; behavior that has a direct, detrimental effect on others who share the assignment makes a stronger case for segregation. Third, the capacity for behavior control matters, and may make the difference in a choice between segregation and integration.

We briefly discuss the empirical literature, which with some exceptions is inadequate to the task of providing clear guidance about harm-minimizing assignment strategies.

Finally, we reflect briefly on the medical-practice principle “first do no harm,” and contrast it with the claims of potential victims of deviants.

Introduction

A common practice in the fields of education, mental health, and juvenile justice is to segregate problem youths in groups with deviant peers. In education, disruptive or delinquent youths may be assigned to self-contained classrooms for in-school detention, or even to alternative schools. In mental health interventions, conduct-disordered youths and drug abusers may be assigned to group therapy or placed with similar youths in group homes and day-treatment programs. In the juvenile justice system youths are placed with other offenders in camps, training schools, and detention centers. It is entirely possible that assignments of this sort, which concentrate deviant youths, may lead to perverse outcomes. Instead of helping their transition to a healthier, more productive trajectory, these assignments may actually facilitate a greater commitment to deviant behavior. The result may be an increase in the rate or seriousness of problem behaviors, either during the course of the placement or in subsequent years (McCord, 2003; Dishion, McCord and Poulin, 2002).

This possibility of an iatrogenic effect from interventions with troubled and troublesome youths is relevant in designing policy, and adds to the list of arguments in support of “mainstreaming” or diffusing them among other youths whenever possible. But there are other concerns that help justify segregated-group assignments. Delivery of treatment through specialized resources may be accomplished more efficiently in settings dedicated to that purpose, and that may also be true in monitoring and controlling the behavior of conduct-disordered youths. In any event, segregating such youths has the desirable effect of insulating other youths and adults from them. And the *threat* of

assignment to in-school detention or a training school may have some deterrent effect on bad behavior. In considering alternative assignment policies, the validity and strength of these rationales should be weighed against the possible perverse results from concentrating deviant youths.

Our analysis is intended to help organize the discussion about these tradeoffs. First, we show that the number of deviant youths (relative to the size of the relevant population, or to the number of assignment options) may well determine whether the harm-minimizing assignment calls for diffusion, segregation, or some of both. Second, the way in which individual behavior cumulates to collective harm can be of considerable relevance in characterizing the harm-minimizing assignment; in cases where bad behavior has a direct, detrimental effect on others who share the assignment (including youths who are not deviant), then there is a stronger case for segregation than when harm occurs outside of the assignment location. Third, the capacity for behavior control matters, and may make the difference in a choice between segregation and integration; an innovation that enhances control over disruptive youths may make integration more attractive.

We briefly discuss the empirical literature, which is typically inadequate to the task of providing clear guidance about harm-minimizing assignment strategies. Empirical investigations often focus just on the behavior of deviant youths under different assignments, without considering the effects on others or the overall social costs of the assignments. There are some exceptional cases where the evidence, while limited, may nonetheless be sufficient.

Given the inherent complexity of the assignment problem and the limited evidence, decisions will have to be made under uncertainty about costs and consequences. Furthermore, alternative assignment possibilities are likely to bring into conflict the interests of the deviant youths, other youths, and the taxpayers. Under these circumstances, is there an ethical obligation to give priority to the interests of the group that is to be directly acted on – the deviants? We reflect briefly on the medical-practice principle “first do no harm,” and contrast it with the claims of potential victims of deviants.

The conceptual and methodological matters discussed here are also relevant to most any policy dilemma involving sorting and mixing of different sorts of people across locations. In what follows we do not always use deviant youth as the exemplar, but refer to examples involving race and residential location of families in poverty.

Conceptual Issues

Much of the relevant literature has been concerned with establishing whether the decision of how to distribute deviant youths across locations (assignments) has an effect on their behavior. That inference has proven quite challenging in practice, as we discuss in a subsequent section. But even when this empirical effort produces reliable results, it may leave us far short of having the information necessary to make a normative judgment. To make that judgment, it is also necessary to make inferences about how the assignment affects the quantity and distribution of total harm to anyone affected by the assignment, including bystanders.

One important distinction is whether the harm is limited to the individuals judged to be at risk, or whether their behavior victimizes others as well. If there are victims, then it is important to distinguish between cases where the victims are in the same assigned location, or are outsiders. These distinctions suggest three logical cases: no other victims, victims elsewhere, victims in the assigned location. In what follows we combine the first two cases and contrast it with the third. The distinction, then, is whether the behavior harms those who share the assigned location or not. The two types are labeled “D” and “S.”

Type D encompasses those behaviors that have an effect on others who share the assigned location, such as classroom disruption (hence the “D”) and bullying.

Type S behaviors are those that do not have a direct impact on others who share the assignment. Included here are a variety of harms to the deviant youths themselves both during and after the particular assignment in question. Examples include smoking (the “S”), alcohol and drug abuse, and inattention to schoolwork. But this category is not limited to “victimless” behaviors. Also included are gang involvement and other criminal activity, as long as their impact is outside of the assigned location.

Of course it is quite possible that D and S behaviors are both affected by the assignment.

To gain leverage on the assignment problem, we stipulate several simplifying assumptions and refer specifically to the school official’s task of assigning students to classrooms. The students are of two sorts, “non-deviants” (the As) and “deviants” (the Bs); their identities are pre-determined and known to the official. There are a fixed

number of classrooms and teachers. The considerations relevant to the official's decision will depend on the nature of the problem behavior.

The consequences within the classroom of Type D behavior are obviously relevant in assigning students. Other considerations may also come to bear for both types of students. Most interesting is the possibility of social influence among deviant peers. The actual behavior of a deviant youth (either D or S) may well depend on whether there are other deviant youths in the same classroom to serve as models, guides, accomplices, or an appreciative audience.

Also relevant is the capacity of teachers and other resources to influence bad behavior in the classroom or outside it. If such influence is possible, whether exercised through discipline or some sort of constructive (therapeutic) programming, then the actual behavior of deviant youths will depend on the level of resources relative to the "load" in the classroom. Resource considerations may lend weight to a prescription for either diffusion or segregation of deviant youths, depending on what might be deemed the "technology of control."

This conceptual setup provides the basis for developing some guidance for the assignment problem. Of the various assumptions, the one that may seem most artificial is that there are only two types of youths (A and B). One generalization is to assume that youths form a continuum with respect to their propensity to engage in deviant behavior, ranging from those who will initiate it under most any circumstance to those who require a powerful stimulus, such as a complete breakdown of order. While this generalization is more flexible and realistic, we stick with the more tractable assumption of a dichotomy for now. Whether dichotomy or continuum, the mix of propensities can be understood as

producing aggregate behavior that is conditioned on peer influence and external social control.

Another note on the basic setup: While our discussion is for the most part presented with reference to the problem of distributing students among classrooms, or among schools, the logic is applicable to a variety of other assignment problems. For example, the state of North Carolina is in the process of demolishing its large juvenile corrections institutions and is building a series of smaller units. Among the relevant questions are how large these new units should be and whether the protocol for assigning juveniles among them call for concentrating the most serious delinquents or mixing them in with the rest of the population. Other examples from the school context include those facing a district school system: whether to start an alternative school for deviant students, whether to locate ninth graders in middle school or high school, and whether to retain large numbers of failing students for a repeat in grade.¹ More broadly, it is relevant to any scheme that re-sorts youths across neighborhoods or schools, including housing relocation programs and school voucher programs. We return to these examples in the empirical section.

Case I: Assignment when there is no contagion and no control

If deviant peers do not influence others' behavior, and the issue of concern is limiting Type S behavior (smoking, crime outside the classroom, etc.), then the assignment problem appears trivial: The amount of Type S behavior would not be affected by how students are sorted among classrooms. Two exceptions should be noted: First, in some cases the assignment to a segregated classroom is justified as punishment,

¹ Thanks to Al Blumstein for suggesting these examples.

as in the case of in-school detention. One rationale is that the threat of such an assignment will deter bad behavior. Second is the possibility that deviant behavior is influenced by the capacity of the teacher and other resources to influence such behavior. A discussion of that possibility is left for the next section.

For disruptive (Type D) behavior, the assignment matters even if there is no social influence. In a recent theoretical article that analyzes this circumstance, Edward Lazear (2001) provides a rationale for segregation. He postulates that any one student can disrupt the productive teaching and learning activity in the classroom for a spell, and that some students are more prone to being disruptive than others. Lazear sets up the analysis by assuming two types of students: “A” students are non-disruptive all but a fraction of the time given by $(1-p_A)$, while the “B” students are disruptive a larger fraction $(1-p_B)$ of the time. Under the assumption that episodes of disruption are uncorrelated random events, the total amount of productive, non-disrupted time is given by the product $p_A^a p_B^b$, where the exponents a and b signify the number of A and B students respectively.

The cumulative harm (assumed proportional to unproductive time) due to Type D behavior in any one classroom is then

$$1 - p_A^a p_B^b$$

and total harm associated with any given assignment of students to classrooms is the sum over all classrooms.

Under this setup, consider the marginal cost to the students in a classroom (in terms of additional time lost to disruption) of replacing an A student with a B student in that classroom. That marginal cost will be

$$p_A^a p_B^b - p_A^{a-1} p_B^{b+1} = p_A^a p_B^b [1 - p_A^{-1} p_B],$$

which is a positive number because by assumption $p_A > p_B$. This marginal cost declines as the number of B students, b , grows. The intuition here is simply that the marginal cost of an extra B student is in the form of a proportional reduction in non-disrupted time. As B increases, that proportion is applied to a shrinking base of productive teaching time, $p_A^a p_B^b$. One more disruptive kid added to a classroom that is already mostly out of control makes little difference to the amount of teaching and learning that takes place there.

The mathematical result under Lazear's assumptions is that an assignment that separates all the A students from the Bs, placing them in different classrooms, maximizes the total amount of non-disrupted time, and hence minimizes total harm.² If that is the goal, then segregation is the answer. Of course, that segregated assignment may raise an equity concern, since the B students will end up with a more meager educational opportunity than the A students. That concern may be more compelling if the "disruption" in question takes the form of epileptic seizures or failure to follow directions due to ADHD, rather than to clowning or getting in fights; if the latter behaviors are viewed (rightly or wrongly) as volitional and the former not, then we might be more hesitant to assign the former group to an inferior placement. The question of whose interests should have priority is a deep one, to which we return at the end of the paper.

Other assumptions may lead to still different conclusions. The robust lesson from this analysis is that social contagion is not the only mechanism that should be considered in making assignments – the influence that youths' behavior has on the productivity of the classroom, or (more generally) the wellbeing of other students, may also be relevant.

² If there are not enough B students to fill a classroom, then they should still be assigned to the same room, filling in with As.

That set of concerns raises issues of both fairness and overall harm to the affected population.

Case II: Assignment when resources matter

In Case I, we assumed that the youths' behavior is exogenous, not influenced by the classroom context. A more realistic assumption is that the amount and quality of behavior depends on the circumstances. A "deviant youth" can then be identified as one that has relatively high potential for exhibiting deviant behavior, in the sense that it doesn't take much of a stimulus to set them off. In particular, the presence of other deviant youths may facilitate this potential. Here we consider one mechanism by which that may happen, the dilution of authority. The next section addresses another mechanism of this sort, direct social influence.

Assume each classroom has a teacher who exerts some control over the behavior of the students. The focus here is again on disruptive (D) behavior. A teacher may be able to "handle" one or two disruptive kids in her classroom, constraining their tendency to disrupt the proceedings, but at some point the "load" on her will become too great. Discipline will begin to break down, and the deviant (B) students' disruptive tendencies will be given greater scope.

When authority is subject to dilution, then the best assignment of students to classrooms will depend on the capacity of teachers – what "load" can they "handle" – and the number of Type B youths relative to the total number of classrooms. For example, if each teacher is capable of fully constraining one B student at most, and there are fewer B students than classrooms, then the harm-minimizing assignment is to disperse the B

students so there is no more than one anywhere. Note that that assignment is not random, but rather requires an explicit policy of identifying and assigning the Bs with an eye to the load per teacher.³

If there are more B students than classrooms, then the nature of the harm-minimizing assignment may still be dispersion, or it may become segregation. In general the answer will depend on how behavior and control capacity interact.

The addition of another B into a classroom has both a direct and indirect effect on total disruption in that classroom. The direct effect is whatever that new youth adds to the total, given the circumstances. The indirect effect stems from the fact that the addition of that youth changes the circumstances by diluting control capacity, with the result that the behavior of all Bs in that classroom may deteriorate.

This is a familiar dynamic in community youth programs, mental health placements, and, most obviously, the criminal-justice area. If, for example, a new criminal gang forms, its crimes will add to the load on the local police department and court system. Unless there is excess capacity, the result is likely to be a reduction in the probability and severity of punishment for any given crime. With the sanction threat reduced, criminal activity will become more profitable, with the result of further increases in the crime rate and further dilution of criminal-justice resources. The new “equilibrium” may be characterized by a much higher crime rate than could be explained in a direct sense by the formation of that new gang. This type of vicious cycle has been

³ While systematic dispersion of Bs is ideal, it may not be feasible. In that case it is of interest to ask whether random assignment would be better than segregation. The answer will depend on the shape of the function relating concentration to disruption.

used, for example, to explain the extraordinarily high rates of homicide in Columbia (Gaviria, 2000).⁴

While it is certainly possible to adapt the Lazear model (discussed under Case I) to this new circumstance, there is not much to be learned, since the implications for the harm-minimizing assignment will be sensitive to the specific assumptions about how behavior responds to dilution of control. Some qualitative statements can be made if we stipulate a sharp threshold phenomenon with respect to “load.” First, if there are too many Bs to keep every classroom inside this threshold, then some segregated classrooms may be part of the harm-minimizing solution so that the other classrooms can remain relatively calm. That is to say, the harm-minimizing solution will entail a combination of diffusion and segregation. Second, if teachers differ in their capacity to handle Bs, then the assignment should be adjusted accordingly. Similarly, if there are extra control resources that can be allocated, then they should be assigned together with the students with an eye to staying within threshold in as many classrooms as possible.

If there is no threshold, but rather a pattern of steadily increasing disruption per B as the number of Bs increases, then it might appear that equalizing the load across classrooms would minimize harm. But the Lazear model makes clear that that is not necessarily the case. The question becomes just how the behavior in question contributes to total harm. If the concern (as in that model) is classroom time available for productive activities, then there is a limit to just how bad things can be – namely, zero productive

⁴ A related dynamic is in reference to the informal social control exerted by a community over its miscreants. That capacity is also limited. One mechanism that may be relevant is the attenuation of norms condemning bad behavior. As such behavior becomes more common, the result may be an implicit downward shift in the cultural definition of deviance. Daniel Patrick Moynihan popularized this idea as “defining deviance down.”

time. So even an explosive growth in disruptiveness (resulting from the addition of Bs to the classroom) will result in a muted growth in harm, asymptotic to the limit.

In sum, if actual behavior depends on the relationship of load to control capacity, then the total harm resulting from a given number of Bs may be quite sensitive to their assignment among classrooms. The harm-minimizing assignment will depend on the number of Bs relative to the number of classrooms and the capacity of ordinary (unspecialized) teachers to assert control.⁵

The capacity of a teacher to maintain control may be expanded through training and adoption of effective innovations, which in turn could affect the harm-minimizing assignment. One important example is the Good Behavior Game, developed in the 1960s by Montrose Wolf, one of the founders of behavior analysis, with two graduate students (Barrish, Saunders, and Wolf, 1969; Embry, 2002). This game adopts a simple behavioral strategy to help maintain order in a classroom. The teacher divides the class into teams and for pre-set intervals keeps a visible scorecard of “fouls” committed by each team. Teams that behave well (fewer fouls than some pre-set limit) receive a prize. This method has been subjected to extensive testing, including a large randomized trial with first graders, the Baltimore Prevention Project, in 1985-6. The game was demonstrated effective at reducing classroom disruption and increasing pedagogic productivity; what’s more, and more surprising, it appears to have caused long-term improvements in behavior, so that even in sixth grade the experimental students were better behaved, less likely to smoke, and so forth (Kellam et al., 1998). Dennis Embry

⁵ This summary suggests that another policy margin relevant to the segregation-integration choice is classroom size. Smaller, more numerous classrooms provide an opportunity to separate deviant youths and expand control capacity. General evidence that smaller classroom size is helpful to student learning is provided, for example, by Alan Krueger in his analysis of Tennessee’s experiment with reduced classroom size for students in kindergarten through third grade (Krueger, 1999).

(2002) has nominated the Good Behavior Game as a “behavioral vaccine,” for its power to reduce impulsive, disruptive behaviors and set deviant youths on a healthier course.

Thus the Good Behavior Game provides teachers with the capacity to reduce both Type D and Type S behaviors. The harm associated with having deviant youths in the classroom is thereby reduced. If the game were difficult to implement (as suggested by the fact that it is not currently used much in ordinary classrooms), it would provide an impetus to segregate deviant youths into classrooms with teachers that do use the method. On the other hand, if it were widely adopted, as suggested by advocates, then it would reduce the costs of mainstreaming deviant youths.

Case III: Assignment when deviant peer influence is important

As in Case II we assume that the behavior of deviant youths depends on the circumstances, but now consider a different mechanism: youths may influence each other directly, in addition to whatever indirect influence may occur via the dilution of authority. Deviant peer influence of this sort is akin to the spread of a contagious disease, and is often referred to as social contagion (Cook and Goss, 1996).

We initially assume that contagious transmission for deviant behavior is limited to those with a propensity for that behavior. Three recent studies from widely differing contexts are among those that offer general support for this assumption. First, in the Baltimore Prevention Project cited above, Kellam et al. (1998) found that first graders with aggressive tendencies were the ones whose behavior was most affected by the presence of other aggressive children in the classroom. Second an analysis of the influence of randomly assigned roommates for first-year students at a university found

that males who had been drinkers in high school had a positive effect on the quantity of drinking by their college room mates, but only if that room mate had also been a drinker in high school (Boisjoly et al., 2003). Third, an analysis of the post-release behavior of delinquents in Florida training schools found that the likelihood of post-release recidivism was positively related to the concentration of youths who happened to be in that reformatory with similar criminal histories; for example, auto thieves were more likely to steal again following release if they had been locked up with a relatively high number of other auto thieves (Bayer et al., 2003).

In evaluating alternative assignment strategies, contagion is analytically similar to the problem of authority dilution. The addition of a badly behaved individual to a classroom adds to the total harm directly (by his own behavior), and also indirectly through his influence on other deviants in the room. This process may exhibit a threshold phenomenon. In the analysis of disease epidemics, a critical point occurs where there is enough contact between infected contagious individuals and “susceptible” individuals so that each case of infection results in more than one additional case -- initiating explosive growth. The concentration of deviants required to produce this result is akin to the capacity of a teacher to handle a load of disruption. And the two mechanisms may be related substantively, as well as by analogy: the teacher’s ability to respond effectively to disruption may be instrumental in stopping the spread of a contagious process.

The possibility of social influence is not limited to Type D (disruption) behavior, but also includes Type S. As with auto thieves in the training-school study mentioned above, the behavior of a deviant youth outside of the immediate assignment may be influenced by who was in the same location with him. But there is an important

difference between S and D behavior in evaluating alternative assignments. Much of the harm from D-behavior may fall on non-deviant classmates, but by assumption the harm from S-behavior is outside of that setting. Thus one of the main justifications for segregation is absent in the case of S behavior.

If S-behavior is the issue and social contagion a real possibility, then the best assignment depends (as usual) on the number of deviants relative to classrooms, and the shape of what might be called the “social-contagion function.” If there are fewer deviants than classrooms, then no more than one per classroom is the harm-minimizing solution. If there is more than one per classroom, it is important to know something about how the behavior of Bs responds to the concentration of other deviants. Here are two possibilities:

- If social influence is characterized by a threshold phenomenon, such that social amplification of S behavior does not begin until there are at least $n+1$ deviants in the classroom, then the key issue is how the ratio of deviants to classrooms relates to that threshold. If there are 10 classrooms and fewer than $10n$ deviants, then systematic diffusion is the harm-minimizing solution. If social influence is characterized by a threshold phenomenon, but there are more than $10n$ deviants to be divided among the 10 classrooms, then a mixed approach may be best, with some Ss assigned to “dilute” classrooms (n deviants) and others in classrooms that are more concentrated. In some cases the best assignment may include a few segregated classrooms with all Bs, and the rest integrated. But it may also work out that equal division is the harm-minimizing solution. Finding the best answer

depends on knowing the shape of the relationship between harm and deviant concentration.

- Second is the possibility that there is contagion without a threshold, so that deviant influence grows stronger as the concentration of Bs in the classroom increases. In other words, suppose that the addition of another B to the classroom tends to make each of the others more committed to deviant behavior, and that effect increases monotonically with the number of Bs. In this case, the harm-minimizing solution is (once again) equal division of Bs among available classrooms.

In the introduction, we noted that our analysis of sorting and mixing is relevant to other domains besides deviant behavior. Consider a quite different domain, taken from the influential Coleman Report of 1966. That report analyzed the effect of racial segregation in schools. The empirical analysis was based on a national sample of American schoolchildren. James S. Coleman and collaborators estimated the effect of the racial makeup of the student's school (percentage black) on his or her achievement test score. His estimation procedure assumed that the effect was linear, and possibly different for black than non-black students.

To preserve as much of the previous notation as possible, we define the following variables and analyze a representative school:

a = number of non-black students in the school (type A)

b = number of black students in the school (type B)

k = total number of students in the school, $k=a+b$

Y_i = achievement test score of student of type i ($i=A$ or B)

p_i = constant associated with student of type i

β_i = parameter indicating the effect of racial concentration

We follow Coleman in assuming a linear relationship which may differ between groups, as follows:

$$Y_i = p_i + \beta_i b/k$$

If the goal is to maximize the sum of achievement-test scores for the classroom, then the maximand is:

$$ap_a + a\beta_a b/k + bp_b + b\beta_b b/k$$

or

$$ap_a + \beta_a b(k-b)/k + bp_b + \beta_b b^2/k$$

A couple of interesting cases come out of this setup. If the effect of the race mix is the same on blacks and non-blacks, so that $\beta_a = \beta_b = \beta$, then the expression simplifies to

$$ap_a + bp_b + \beta b.$$

In this case, the number of black students in the school has a linear effect on total achievement. If we then sum across all schools, all that is relevant is the total number of black and non-black students – their distribution across schools has no effect. So in this case, the decision of how to mix different types of students matters in the small (at the level of the individual school) but not at the systems level. This point has been noted as well by other analysts of peer and neighborhood effects (Jencks and Mayer, 1990; Galster, Quercia and Cortes, 2000).

Another possibility, made interesting by the fact that Coleman concluded that it best fit the data, is that the black concentration in a school affects black scores but not white scores. In that case, we have the total score in the school equal to:

$$ap_a + bp_b + \beta_b b^2/k$$

If we sum over all schools in the system, it turns out, assuming $\beta_b < 0$ (as found by Coleman), that uniform integration is the policy that maximizes total scores.

Summary

One lesson is that an evaluation of the effects of an assignment policy for deviant youths should not be limited to the behavior of the youths in question. Also required is an assessment of the harm to all who are affected by the assignment, including the bystanders.

A slight generalization of our analysis serves as a useful summary. The harm in any one classroom is a function of the number of deviant youths. Total harm in the system is the sum over the classrooms. Thus:

$$\text{Harm in classroom } c = H_c = H(b_c, a_c)$$

$$\text{Total harm} = \sum_c H_c$$

If H_c is a linear function of b_c , then total harm is determined only by the total number of Bs in the system. But under a variety of circumstances harm is decidedly not linear in b_c and the assignment does matter. We reviewed cases in which there was a limited capacity to maintain order, existence of deviant peer influence, or harm that is not proportional to bad behavior (as in the case of classroom disruption).

A full analysis would also take account of the budget impact on the relevant agencies, and other mechanisms that may be relevant – especially the deterrent effect of the threat of assignment to a segregated classroom (or training school).

Empirical Issues

In choosing an assignment policy, it would be useful to know the outcomes for all youths in the system for each of the options under consideration. To keep things simple, suppose that only two schemes are under consideration, one that segregates deviant youth ($T=1$) and the other that “mainstreams” such youth into classrooms with non-deviant youth ($T=0$). In a randomized experiment in which entire school systems are randomly assigned to the segregated or integrated conditions, the net effect of the difference in assignment schemes can be estimated by regressing outcomes for each youth on an indicator for the system’s treatment assignment. In the equation below, a limited number of parameters are to be estimated, allowing for different intercepts and response coefficients for non-deviant (A) and deviant (B) youths. The residual terms ε_c and ε_i represent other characteristics of classrooms and individuals respectively that affect the outcome; because of random assignment, the distribution of these characteristics will be balanced across treatment conditions (that is, orthogonal to the assignment indicator T). In this equation, the variable I_i is an indicator for whether the individual is type A ($I_i=1$) or B ($I_i=0$).

$$Y_i = \alpha + \beta I_i + \beta_A I_i * T_i + \beta_B (1 - I_i) * T_i + \varepsilon_c + \varepsilon_i$$

In this setup, β_A provides an estimate of the average effect on As of segregated classrooms, while β_B provides the same information for the Bs. A weighted average of these two coefficients would provide an overall estimate. The relationship between these estimates and “harm” depends on the nature of the outcome measure. Positive β s indicate segregation is harmful if the outcome measure is crime, injury, absenteeism, or smoking. If the outcome is an achievement test score, then a positive β indicates that segregation is beneficial.

In practice there may be a number of outcome measures that are deemed relevant. A complete analysis would require combining them in some fashion to produce an comprehensive measure of welfare (Nagin, 2001).

This group assignment experiment highlights conditions under which analysts could reliably compare two (or more) assignment practices. This ideal is rarely achieved in practice. Among the limitations of available research literature are inadequate measurement and uncertain interpretation.

Measurement. Often, the outcome measures are limited to one or several measures of the behavior of the target youth (usually the deviant youths). That approach is inadequate if the behavior of other youths may also be affected by the assignment. Furthermore, data on the behavior measures, even if collected for all youths, may not adequately capture the effects on welfare, as in the case of classroom disruption, bullying, and so forth.

Adding up. A comprehensive assessment should consider the system-wide effects, although that is not standard practice. For example, a finding that the behavior of Bs tends to deteriorate in the presence of a high concentration of Bs may be presented as

evidence that Bs should be integrated. But that conclusion does not necessarily follow, since it only looks at the costs of concentration, and not the possible benefits to those classrooms and students that avoid contact with Bs in a concentrated assignment.

Self-selection. The Coleman Report cited above provides one well-known example of estimating peer effects on youth outcomes. Coleman's estimates, like most estimates of peer effects in the literature, were based on natural variation rather than experimentally induced variation. In "nature," students in integrated classrooms are likely to differ in relevant ways from students in segregated classrooms. Coleman's analysis attempts to adjust for this problem by controlling for measures of family background and other student characteristics available to him in his data set (the EEOC). But those measures are far from comprehensive. Inevitably there remains the question: Why do two observationally equivalent minority students wind up in very different types of schools?

Suppose, for example, that parents who choose an integrated school for their children tend to have different attitudes toward race and education on average than those who choose a more segregated school. ("Choice" in this case may be exercised by choosing where to live or choosing between private and public schools.) If African-American parents who are most committed to education are the ones who manage to navigate their child's way into an integrated setting, and parent competence and attitudes cannot be "controlled for" (because they are not measured adequately), then Coleman's analysis and similar non-experimental estimates will confound the effects of unmeasured parent attributes with the effects of school racial composition. Put differently, the self-selection of families and youth into different educational or other social settings may

compromise the internal validity of estimates for peer effects on youth outcomes, causing biased estimates.

Self-selection could also affect the external validity (generalizability) of estimates of peer effects if peer influences differ among settings. For example, suppose that in the mid 1960's, when the EEOS was collected, minority families who sent their children to integrated school settings only did so in areas where the local white population was relatively progressive and open to the idea of school desegregation. In this case the effect of racial desegregation on whites in desegregated schools may be quite different from the hypothetical effect on whites of integrating all-white schools. Introducing minority students to all-white schools may have more negative effects on the whites in those schools (and the newly introduced black students as well) than the estimates in the Coleman Report would suggest. (For more on the issue of "local average treatment effects," see Imbens and Angrist, 1994, and Angrist, Imbens and Rubin, 1996).

Classrooms, Control or Contagion? Even if, despite the self-selection problem, the parameter estimates are valid, there remains the question of just what can be learned from them. The parameters measuring peer context in a standard observational study do not identify the causal mechanism that is responsible for the estimated effect. Possible mechanisms include social contagion, the effects of "load" on control, and unmeasured or poorly measured aspects of the classroom environment itself (Manski, 1993). Understanding the mechanism is usually important in designing an appropriate assignment and related policies.

This discussion suggests that standard practice within the empirical literature on peer effects is unlikely to be informative about the net impact of different assignment

processes within our public-school or juvenile-justice systems. But there is a growing experimental literature that eliminates most of the self-selection problem and is helpful in other respects as well. Here we review a small sample of such experiments.

First-year dormitory assignments

Many residential colleges and universities assign some first-year students to dormitory rooms more-or-less at random. That assignment scheme provides a natural experiment for studying peer effects. Usually membership in social groupings is influenced in part by individual choice, which makes it difficult to distinguish the effects of the group from the characteristics of the people who tend to join that group. Random assignment eliminates that problem. In one study, randomly assigned roommates at a large mid- western university were found to influence each others drinking under some circumstances (Boisjoly et al., 2003). The effect was only found for males, and only when both roommates had been drinkers in high school – drinking in college by those who had been abstainers in high school was not affected by first-year roommate assignment. But when two drinkers were assigned to the same room, the effect (in the male dorms) was to amplify their drinking.

If the university's goal were to reduce drinking by undergraduates, then these results are sufficient to point the way. It would be advisable in making dormitory assignments to match high-school abstainers with drinkers. That conclusion is strongest if drinking is primarily a Type S behavior rather than a Type D behavior. If on the other hand a roommate's drinking tends to negatively affect the quality of life in the dorm room (a Type D behavior), then the normative circumstance is more complicated. It

could be argued that the abstainers should be spared the disruptive influence of a drinking roommate, even if the total prevalence of alcohol problems would be reduced by forced integration.

Baltimore Prevention Project

In the Baltimore Prevention Project, introduced above, first graders were assigned to classrooms on a random basis. Sheppard Kellam and colleagues (Kellam et al., 1998) evaluated the effects on behavior of this assignment (and the application of the Good Behavior Game) years later, when the children were in sixth grade. One finding was that the initial assignment had a long-term effect on the behavior of the most aggressive boys. Those who were identified as aggressive in first grade behaved worse five years later if they happened (by chance) to end up in a first-grade classroom with a relatively high concentration of other aggressive boys.

Because the initial assignment was made randomly, it can be concluded with some confidence that Kellam's finding reflects a causal process involving peer effects. Perhaps the first-grade exposure to other aggressive boys exacerbated their tendencies, or perhaps the first grade provided an opportunity for the aggressive boys to find each other and become troublemaking companions thereafter.

What does this finding imply about assigning aggressive children to first grade classrooms? Assuming we only care about their subsequent behavior, and that behavior translates into social harm in a straightforward (additive) manner, then the answer could potentially be derived from the analysis of these data. Given an estimated relationship between first-grade class composition and subsequent behavior, simulations could be run

to determine the total amount of trouble associated with different hypothetical assignments. The best assignment may well be an even distribution of these boys among classes, but not necessarily – if they are sufficiently prevalent, it may be best to concentrate some of them and diffuse the rest. In any event, the “right” answer would require that the objectives be identified explicitly.

Moving to Opportunity

The best available evidence to date on the existence of “neighborhood effects” on youth outcomes derives from the U.S. Department of Housing and Urban Development’s Moving to Opportunity (MTO) housing-voucher experiment. MTO overcame the self-selection problem that plagued previous studies of neighborhood effects by randomly assigning public housing families who volunteered for the demonstration into different mobility treatment groups, two of which were offered housing vouchers to move to lower-poverty areas and one of which – the control group – received no additional services under the program.

Short-term findings from MTO suggested that moving from high- to low-poverty neighborhoods reduces youth involvement in problem behavior and may increase achievement test scores (Kling, Katz and Liebman, 2001; Ludwig, Duncan and Hirschfield, 2001; Ludwig, Ladd and Duncan, 2001; Goering and Feins, 2003). Additional years’ worth of post-randomization data suggest that moving to a neighborhood with less poverty produces lasting behavioral benefits among adolescent girls but not boys (Kling and Liebman, 2004; Kling, Ludwig and Katz, 2004).

Most relevant for our purposes, MTO tells us nothing about the effects of moving youth across neighborhoods on those who reside in either the origin or destination neighborhoods of MTO participants. The harms to neighbors of Type D behavior are ignored. For that reason it is not possible to assess the net social benefit of the mobility induced by the experimental intervention. Even more problematic is projecting what would happen if housing vouchers were offered to *all* public-housing families under the same terms as the MTO experimental group.

School Vouchers

As another example, consider the case of school vouchers, which, like housing vouchers, help move youths to new social settings. Much of the recent evidence comes from local randomized housing-voucher experiments. The first and perhaps best-known program comes from Milwaukee. While students in Milwaukee were not actually randomly assigned to “voucher” or “no voucher” groups, as in a classical randomized experiment, private schools in the program that had more applicants than slots were required to admit students on the basis of a random lottery. In this sense the Milwaukee voucher program provides a particularly credible natural experiment; the best available study of Milwaukee suggests that attending private rather than public school produces gains in math but not in reading (Rouse, 1998).

More recent voucher programs in New York City, Washington, DC and Dayton, Ohio use formal randomized experimental designs and yield suggestive evidence of positive private-school effects, at least on some students in some grades, although the experimental designs behind these programs is undermined somewhat by response rates

to the follow-up surveys used to measure student outcomes that are lower than one might wish (Howell and Peterson, 2002; Howell, Wolf, Campbell and Peterson, 2002).⁶

In all of these school-voucher experiments, the causal mechanism is not well identified. It could be that the private schools offer better instruction on average, or that the peer environment is more conducive to academic achievement. If the latter mechanism is part of the mix, then the system-wide effects of a voucher program may be quite different than the effects on those who take up the vouchers. For example, suppose that those who volunteer for these voucher experiments are among the best students in their old public schools but now among the weakest students in their new private schools. In this case their take-up of the voucher offer leads to a reduction in average student quality in both the origin and destination schools. If individual achievement is positively influenced by the average quality of peers, then the result would be that the voucher students would benefit (by association with better students) while all others – both public and private – would lose. So once again we need a more comprehensive measure.

System-level studies

The case of school vouchers provides a useful example of the distinction between “mover” and “systems-level” studies. What would happen to the average achievement of all American youth if our country adopted a large-scale school voucher program?

⁶ If we were certain that the characteristics of non-respondents in the voucher treatment and control groups were on average identical then relatively low response rates would reduce the statistical bias of the analysis but not introduce bias. However as an incentive to participate in follow-up surveys, control group families who responded to these surveys were entered into a new voucher lottery (Howell, Wolf, Campbell and Peterson, 2002, p. 196). This raises the possibility that the voucher treatment group respondents may have been the families who were most satisfied with their voucher experiences while the control responders consisted disproportionately of those families who were least satisfied with their public school experience. This response pattern would lead the analysis to overstate the effects of voucher receipt on achievement.

Because no such system has ever been implemented in the United States, we cannot answer this question directly. But intriguing evidence on this point comes from Chile, which has offered students private-school vouchers for the past two decades.

The number of new private schools that operate in Chile has increased substantially over this time, with the growth in the size of the new private sector varying considerably across metropolitan areas within Chile over time. Chang-Tai Hsieh and Miguel Urquiola (2002) estimate the effects of this resorting program by comparing trends in average student outcomes across metropolitan areas within Chile that experienced differential growth in the local private school market.

Their “difference-in-difference” analysis suggests that Chile’s school voucher program has redistributed students from public to private schools. As a result, average test scores have declined in the public schools and increased in private schools. That pattern may in principle be the result of the concurrent effects of several mechanisms: the direct effects and peer effects of changes in the quality of the public and private student enrollees, as well as productivity changes in private and public schools. It is telling in this respect that the aggregate effect of the growth in the private school market appears to be nil – when public and private students are considered together, areas that had more voucher-induced movement into private schools did no better than students in areas with less movement. These results suggest that whatever gains the “switchers” enjoyed were compensated by losses to those left behind in the public schools.

The Chilean experience highlights the tradeoff between mover experiments and systems-level studies. The former provide highly credible evidence on the effects of assignment policies on a sub-set of youth within the social system of interest, but provide

limited information about systems-level effects. Studies conducted at the systems level, on the other hand, capture the overall effects of changing assignment procedures within a system. However these systems-level studies come at the cost of losing the experimental variation that drives differences in peer settings in mover studies. In the case of Chile, the private school sector developed more rapidly in large, urban, wealthy metropolitan areas. Whether changes over time in student outcomes within smaller, rural areas of Chile provides a good control group for large wealthy areas is not clear.

Recommendations for Research

The main lesson from the analysis above is that when deviants' behavior is either influenced by social context, or victimizes bystanders, a complete evaluation may require going well beyond observations on the deviants themselves. The study of first-year dormitory assignments was especially illuminating because it included information on the college drinking behavior of roommates who were "non deviant" in that they did not drink in high school; the study may be considered deficient in that it did not include measures of disruption within the dormitory setting. An important limitation of the MTO study was that it lacked measures of the effects of those who relocated under the program on either their new neighborhoods or their old neighborhoods.

The problem, of course, is that a more comprehensive study tends to be more costly to implement. Judging whether it is worthwhile to measure contextual effects in any given evaluation project requires knowledge of when they are likely to be important. The accumulation of such knowledge should be a goal of any long-term research strategy.

Normative considerations

Empirical evidence, even when reliable, is inevitably an incomplete guide to formulating assignment policies for deviant youths. Also needed is a clear statement of values. Here we offer several observations, focusing on how to take account of the possibility of deviant peer contagion.

If the welfare of the deviant youths is the benchmark for judging policy, then negative effects from clustering such youths for the delivery of services are of great concern. The possibility that an intervention would actually do harm to the group that it was intended to help is not just ironic, but downright scandalous. That's the power of Joan McCord's findings concerning the summer camps for the Cambridge-Somerville Study's delinquent youths; she concludes that by concentrating such youths for an extended period, the camp experience not only did not set them straight, but actually served to reinforce their deviant propensities (McCord, 2003).

The Hippocratic stricture "First do no harm," applies in this case: physicians are to be especially sensitive to the possibility that their prescriptions will hurt rather than help their patients. While the operational significance of this stricture is not clear, it may be a useful check on the tendency to demonstrate professional know-how by offering some active remedy, rather than let nature take its course. Of course, even well-advised remedies may be harmful for some patients, but "evidence based medicine" is supposed to be guided by the empirical probabilities, insisting on an expectation of benefit and the informed consent of the patient or his or her guardians.

While this standard is not irrelevant to assignment policies for deviant youths, it should be noted that the choice between mainstreaming and some form of segregation

brings into play the welfare of other youths – those who would have more or less contact with the deviant youths depending on the assignment – and, of course, cost. If mainstreaming is good (on average) for the deviant youths, but harmful (on average) to the others, then the welfare of the deviants alone may be deemed an inadequate standard for making the judgment. But there is no consensus on this matter.

Indeed, since the 1970's those who favor mainstreaming children with mental and physical disability have the dominant voice in school-assignment policies. Public Law 94-192 guaranteed children with disabilities the right to an education in a mainstreamed setting with their peers, rather than in separate classrooms (Dodge, Kupersmidt, and Fontaine, n.d., p. 42). In the Willie M case beginning in 1979, this right was extended by a federal court to violent delinquents with mental health problems. Among the advocates' arguments are that first, mainstreaming is in the disabled youths' interest, and second, that their interest should be given priority over other considerations, including financial cost and classroom management problems.

Suppose that we accept the principle of a rebuttable presumption that deviant youths should be mainstreamed. The rebuttal then might take one of several forms: first, an argument that with available resources and technology, deviant youths would make better progress in a segregated setting; second, that other youths would benefit from removing certain deviant peers to segregated settings; or third, that a particular class of deviant youths did not deserve the "mainstreaming" presumption because they should be held culpable for their actions (rather than deemed "disabled"). The first two of these arguments would be based on empirical evidence.

Several observations about such arguments follow from the analysis in this paper:

Since the evidence on the consequences of alternative assignment policies for deviant youths and other youths is typically incomplete or weak, rejecting the null hypothesis (rebuttable presumption) is going to be difficult in practice. Where mainstreaming is the norm, as in public schools, the presumption will be difficult to overcome, given the difficulties of inferring causal effects from available evidence. In those cases where there is experimental evidence, it is typically incomplete. So even when the preponderance of the evidence suggests that mainstreaming would cause more harm overall than segregation, that conclusion will not be beyond reasonable doubt.

On the other hand, if there is no presumption of mainstreaming (as in public housing policy, for example, or with respect to youthful offenders who are not deemed disabled), then the preponderance of the evidence should rule, and the interests of all who are affected be taken into account. In that case the more complete analysis of harm suggested in our analysis above is relevant in principle, although still difficult in practice. A finding that concentrating deviant youths tends to facilitate negative peer influence makes a difference, but only as one of several relevant mechanisms.

The complexities introduced by uncertainty about the effects of alternative policies, and concerns about how those effects are distributed, can be illustrated by the example of a school-system assignment policy that is to be guided by projected effects on an end-of-grade achievement test. The hypothetical choice is between mainstreaming a group of deviant students, and concentrating them in special classrooms or a separate school. Evidence is available on how the choice of assignment will affect achievement

test scores for each group; the evidence yields probability distributions of results for deviant students, other students, and overall. Here are several of the possible decision rules, in order of increasing sensitivity to the welfare of the deviant group:

1. Choose segregation unless Mainstreaming causes a higher expected achievement score for the non-deviant group.

This standard might be considered if the deviant group in question were deemed to be criminals or otherwise blameworthy, and deserving of assignment to a segregated setting.

2. Choose the assignment that has the higher overall expected average achievement score (averaging deviants with others).
3. Choose Mainstreaming unless reliable evidence indicates that there is at least a 95% chance that the deviant students will have higher average scores if separated.

This standard is the implementation of the “rebuttable presumption” in favor of mainstreaming.

4. Choose the assignment that has the higher expected average achievement score for the deviant students, regardless of the effect on the other students.

The interplay of values and statistical reasoning presents an interesting challenge for policymakers in this case and similar cases in other domains.

References

- Angrist, Joshua D., Guido W. Imbens and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91(434): 444-472.
- Bayer, Patrick, Randi Pintoff, and David Pozen. 2003. "Building Criminal Capital Behind Bars: Social Learning in Juvenile Corrections." New Haven, CT: Yale University Dept. of Economics, mimeo.
- Beard, Belle Boone. 1934. *Juvenile Probation*. New York: American Book Company.
- Boisjoly, Johanne, Greg J. Duncan, Michael Kremer, Dan M. Levy, and Jacque Eccles. 2003. "Peer Effects in Binge Drinking among College Students." Working Paper, Evanston, IL: Northwestern University, Institute of Social Research.
- Coleman, James S. and others. 1966. *Equality of Educational Opportunity*. Washington, DC: U.S. Department of Health, Education and Welfare.
- Cook, Philip J. and Kristin Goss. 1996. "A Selective Review of the Social Contagion Literature." Durham, NC: Duke University, Department of Public Policy Studies, mimeo.
- Crane, Jonathan. 1991. "The Epidemic Theory of Ghettos and Neighborhood Effects on Dropping Out and Teenage Childbearing." *American Journal of Sociology* 96(5): 1226-59.
- Dishion, Thomas J., Joan McCord, and Francois Poulin. 1999. "When Interventions Harm." *American Psychologist* 54(9): 755-764.
- Dodge, Kenneth A., Janis B. Kupersmidt, and Reid Griffith Fontaine. No Date. *Willie M: A Legacy of Legal, Social, and Policy Change on Behalf of Children*. Durham, NC: Duke University, Center for Child and Family Policy.
- Duncan, Greg J., Katherine A. Magnuson, and Jens Ludwig. 2003. "The Endogeneity Problem in Developmental Studies." Working Paper. Evanston, IL: Northwestern University, Institute for Policy on Research.
- Embry, Dennis D. 2002. "The Good Behavior Game: A Best Practice Candidate as a Universal Behavioral Vaccine." *Clinical Child and Family Psychology Review* 5(4): 273-297.
- Fagan, Jeffrey, Franklin E. Zimring, and J. Kim. 1998. "Declining Homicide in New York City: A Tale of Two Trends." *Journal of Criminal Law and Criminology* 88(4): 1277-1323.

Galster, George C., Roberto G. Quercia and Alvaro Cortes. 2000. "Identifying Neighborhood Thresholds: An Empirical Exploration." *Housing Policy Debate* 11(3): 701-732.

Gaviria, Alejandro. 2000. "Increasing Returns and the Evolution of Violent Crime: the Case of Colombia." *Journal of Development Economics* 61: 1-25.

Gladwell, Malcolm. 2000. *The Tipping Point: How Little Things Can Make a Big Difference*. Boston: Little, Brown.

Glaeser, Edward, Bruce Sacerdote, and Jose Scheinkman. 1996. "Crime and Social Interactions." *Quarterly Journal of Economics* 111: 507-48.

Guryan, Jonathan. 2001. "Desegregation and Black Dropout Rates." Cambridge, MA: National Bureau of Economic Research, Working Paper # 8345.

Hanushek, Eric A., John F. Kain and Steven G. Rivkin. 2002. "New Evidence about Brown v. Board of Education: The Complex Effects of School Racial Composition on Achievement." Working Paper. Dallas: University of Texas at Dallas.

Howell, William G. and Paul E. Peterson. 2002. *The Education Gap: Vouchers and Urban Schools*. Washington, DC: Brookings Institution Press.

Howell, William G., Patrick J. Wolf, David E. Campbell, and Paul E. Peterson. 2002. "School Vouchers and Academic Performance: Results from Three Randomized Field Trials." *Journal of Policy Analysis and Management* 21(2): 191-218.

Hsieh, Chang-Tai and Miguel Urquiola. 2002. "When Schools Compete, How Do They Compete? An Assessment of Chile's Nationwide School Voucher Program." Working Paper. New York: Columbia University.

Imbens, Guido W. and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62(2): 467-475.

Jencks, Christopher and Susan E. Mayer. 1990. "The Social Consequences of Growing Up in a Poor Neighborhood." In *Inner-City Poverty in the United States*, Laurence E. Lynn and Michael G.H. McGeary (eds.). Washington, DC: National Academy Press, pp. 111-186.

Johnson, Michael, Helen F. Ladd and Jens Ludwig. 2002. "The Benefits and Costs of Residential Mobility Programs." *Housing Studies* 17(1): 125-138.

Jones, Marshall B. and Donald R. Jones. 2000. "The Contagious Nature of Antisocial Behavior." *Criminology* 38: 25-46.

- Kellam, Sheppard G., Xiang Ling, Rolande Merisca, C. Hendricks Brown, and Nick Ialongo. 1998. "The Effect of the Level of Aggression in the First Grade Classroom on the Course and Malleability of Aggressive Behavior into Middle School." *Development and Psychopathology* 10: 165-185.
- Kleiman, Mark A.R. 1999. "Getting Deterrence Right: Applying Tipping Models and Behavioral Economics to the Problem of Crime Control" In *Perspectives on Crime and Justice: 1998-1999 Lecture Series*. Washington, DC: National Institute of Justice.
- Kling, Jeffrey, Jens Ludwig and Lawrence Katz. 2004. "Youth Criminal Behavior in the Moving to Opportunity Experiment." Princeton, NJ: Princeton University, Industrial Relations Section, Working Paper 482.
- Krueger, Alan. 1999. "Experimental Estimates of Educational Production Functions." *Quarterly Journal of Economics* 114(2): 497-532.
- Lazear, Edward P. 2001. "Educational Production." *Quarterly Journal of Economics* 116(3): 777-803.
- Ludwig, Jens, Greg Duncan, and P. Hirschfield. 2001. "Urban Poverty and Juvenile Crime: Evidence from a Randomized Housing Experiment." *Quarterly Journal of Economics* 116: 655-80.
- Ludwig, Jens, Helen F. Ladd and Greg J. Duncan. 2001. "Urban Poverty and Educational Outcomes." *Brookings-Wharton Papers on Urban Affairs*. Edited by William Gale and Janet Rothenberg Pack. Washington, DC: Brookings Institution, pp. 147-201.
- Manski, Charles F. 1995. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60: 531-542.
- McCord, Joan. 2003. "Cures That Harm: Unanticipated Outcomes of Crime Prevention Programs." *Annals, AAPSS* 587: 16-30.
- Meyer, Bruce D. 1995. "Natural and Quasi-Experiments in Economics." *Journal of Business and Economic Statistics* 13: 151-162.
- Moffitt, Robert. 2001. "Social Interactions, Policy Interventions, and Low-Level Equilibria" in Stephen N. Durlauf and H. Peyton Young (eds). *Social Dynamics* Washington DC: Brookings Institution Press: 45-82.
- Nagin, Daniel. 2001. "Measuring the Economic Benefits of Developmental Prevention Programs." in Michael Tonry, (ed.) *Crime and Justice: A Review of Research* v. 28: 247-384.

Rouse, Cecilia. 1998. "Private School Vouchers and Student Achievement: An Evaluation of the Milwaukee Parental Choice Program." *Quarterly Journal of Economics* 113(2): 553-602.

Sah, R. 1991. "Social Osmosis and Patterns of Crime." *Journal of Political Economy* 99: 1272-95.

Schelling, Thomas C. 1971. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1: 143-186.