

DOES PEER ABILITY AFFECT STUDENT ACHIEVEMENT?

ERIC A. HANUSHEK,^{a*} JOHN F. KAIN,^b JACOB M. MARKMAN^c AND STEVEN G. RIVKIN^d

^a *Stanford University, National Bureau of Economic Research and University of Texas at Dallas. Hoover Institution, 434 Galvez Mall, Stanford University, Stanford, CA 94305-6010, USA*

^b *University of Texas at Dallas, Dallas, TX, USA (Deceased)*

^c *Department of Economics, Converse Hall, Amherst College, Amherst, MA 01002, USA*

^d *Amherst College, National Bureau of Economic Research and University of Texas at Dallas. Department of Economics, Converse Hall, Amherst College, Amherst, MA 01002, USA*

SUMMARY

Empirical analysis of peer effects on student achievement has been open to question because of the difficulties of separating peer effects from other confounding influences. While most econometric attention has been directed at issues of simultaneous determination of peer interactions, we argue that issues of omitted and mismeasured variables are likely to be more important. We control for the most important determinants of achievement that will confound peer estimates by removing student and school-by-grade fixed effects in addition to observable family and school characteristics. The analysis also addresses the reciprocal nature of peer interactions and the interpretation of estimates based upon models using past achievement as the measure of peer group quality. The results indicate that peer achievement has a positive effect on achievement growth. Moreover, students throughout the school test score distribution appear to benefit from higher achieving schoolmates. On the other hand, the variance in achievement appears to have no systematic effect. Copyright © 2003 John Wiley & Sons, Ltd.

1. INTRODUCTION

The peer group composition of schools is undeniably important in the minds of parents as well as policy makers at the local, state and federal level. Residential location decisions of families, various state and federal laws, and court interpretations of school district policies frequently have an implicit if not explicit peer group component. There have nonetheless been relatively few direct investigations of the impact of peer groups on student performance and what evidence exists has been open to widely varying interpretations. This paper pursues a unique identification strategy based on small perturbations in cohort composition to extract the causal impacts of peer group characteristics on achievement.

Peer group effects have played a prominent role in a number of policy debates including ability tracking, anti-poverty programmes in both rural areas and urban ghettos, and school desegregation. In addition, opposition to the growing demand for expanded school choice or the provision of education vouchers often focuses on the importance of peers and potential for greater economic and social isolation of disadvantaged students. At the same time advocates of choice often tout its potential for reducing the impact of neighbourhood economic and racial segregation.

The role of peers also has entered increasingly into theoretical analyses of school choice. Starting from the observation that many people express concern about other students, a variety of analyses

* Correspondence to: Eric A. Hanushek, Stanford University, National Bureau of Economic Research and University of Texas at Dallas, Hoover Institution, 434 Galvez Mall, Stanford University, Stanford, CA 94305-6010, USA.
E-mail: hanushek@stanford.edu

(e.g. Benabou, 1993, 1996; Caucutt, 2002; de Bartolome, 1990; Epple and Romano, 1998) have examined the equilibrium properties of choice and peer group effects on student achievement.

This attention to peer effects has taken place largely in the absence of compelling empirical evidence on the impact of peer group characteristics on a variety of academic, social and labour market outcomes. As Brock and Durlauf (2001), Manski (1993) and Moffitt (2001) point out, the empirical analysis of peer influences has been inhibited by both conceptual and data problems—problems that raise serious questions about interpretation of the existing studies, even those that use more sophisticated econometric techniques including instrumental variables. These critiques, in part precipitated by parallel analyses of neighbourhood poverty (e.g. Mayer and Jencks, 1989; O'Regan and Quigley, 1999; Rosenbaum and Popkin, 1991), point to a number of potentially severe empirical problems that are at least partially present in the recent set of randomized housing experiments aimed at understanding neighbourhood effects (e.g. Rosenbaum, 1995; Katz *et al.*, 2001; Ludwig *et al.*, 2001).¹

This paper makes use of a unique matched panel data set on students and schools to identify the impacts of specific peer group characteristics on academic achievement. It directly confronts the central specification issues that impinge on our ability to estimate the achievement effects of peer group composition. These include the confounding influences of unobserved or badly measured student, family and school factors, and of the reciprocal nature of peer group interactions.

The basic strategy involves the successive elimination of the components of individual student achievement growth that are most likely to lead to confusion of family and school influences with peer group effects. While controls for observable characteristics are used, the ability to control for fixed individual, school and school-by-grade effects on test score gains permits the clearest identification of peer effects. Ultimately, we identify these effects by considering the impact of small differences in peer group characteristics for successive cohorts of students in a given school.² This panel data approach is robust to most of the commonly cited estimation dangers.

One problem not addressed by the fixed effects framework is the reciprocal nature of peer interactions that likely introduces simultaneous equations bias when contemporaneous peer behaviour is proxied by current average achievement. In the fixed effects framework, specifications based on lagged peer achievement eliminate the problem of simultaneous equations bias and capture the systematic predetermined aspects of peer interactions. However, this approach also ignores the impact of current peer behaviour not captured by lagged achievement and may therefore lead to an underestimate of peer influences. This issue is addressed at length in Section 4 below.

Our basic estimation of elementary school achievement growth indicates that the achievement level of peers has a positive effect on achievement that is roughly constant across quartiles of the school achievement distribution. In contrast, the variance in achievement appears to have no systematic influence.

Despite the fixed effects methodology, important limitations remain in terms of both data structure and the available measures of peers and outcomes. The role of peers can be complex. Influences may come from friends or role models, or peer group composition may alter the nature of instruction in the classroom. We do not have information at the classroom level or about friends but instead rely on aggregations of students within the same school and grade. We also have limited information on attributes of peers, though we do include measures for the three most

¹ For a review and critique of these studies, see Moffitt (2001).

² This methodology is similar to that used by Hoxby (2000) in the estimation of class size and racial composition effects for students in Connecticut.

commonly expressed peer characteristics: achievement, race and socio-economic status. Finally, due to limited outcome measures available for elementary school students, all empirical work examines academic achievement. Many of the policy discussions and parental concerns focus on other outcomes including teen pregnancy, drug use, high school attrition, attitudes towards minorities and employment, to name but a few.

2. CONCEPTUAL ISSUES IN THE ESTIMATION OF PEER GROUP EFFECTS

The identification of specific social interaction effects is a daunting task. Not only must the analysis address the endogenous choice of neighbourhoods and schools, but it must also separate peer influences from the effects of other school characteristics and account for the fact that student and peer achievement are determined simultaneously. In this section we outline an empirical framework with which to examine peer influences, following closely the work of Brock and Durlauf (2001) and Moffitt (2001), with special emphasis on the educational context.

Attempts to estimate peer effects on educational achievement directly have been relatively limited. Hanushek (1972, 1992) finds no peer achievement effects, while Henderson *et al.* (1976), Summers and Wolfe (1977) and Zimmer and Toma (2000) report positive influences of higher achieving peers, at least for some students. Consideration of ability tracking in schools likewise has yielded mixed results, even though policy has presumed that tracking is generally bad for achievement (e.g. see Oakes, 1992; Argys *et al.*, 1996). The evidence on achievement effects of racial composition has been much more voluminous, although the results are no easier to summarize or interpret (*cf.* Armor, 1995).

In general there has been limited attention given to the mechanisms through which peers affect outcomes. The most common perspective is that peers, like families, are sources of motivation, aspiration and direct interactions in learning. Moreover, peers may affect the classroom process—aiding learning through questions and answers, contributing to the pace of instruction, or hindering learning through disruptive behaviour à la Lazear (2001).

Most analysis has focused on the identification of the 'reduced form' relationship between outcomes and specific measures of peer group quality, typically ignoring the precise structure of the underlying causal relationship. An outcome measure is regressed on peer group characteristics that are usually constructed as school aggregates of family background variables or achievement. These measures are readily available and, if they adequately capture the influences of families, would seem appropriate for peers. While we follow in that tradition, it is important to note that ambiguities about the correct measurement of family backgrounds exist, and these naturally transfer to the measurement of peer influences.³

The attention to understanding the full causal structure in the case of peers reflects the fact that peer composition is a product of both parental choices of neighbourhood and school and school policy makers' decisions on attendance rules and classroom placement of students. One important theme, that follows the interpretation of Moffitt (2001) and motivates our work here, is that

³ The empirical analysis of family background has generally relied just on readily available measures of socio-economic status of families to proxy for the learning environment in the home without much attention to the details of the structure. This lack of attention to detail partially reflects the fact that little consideration is given to policies directed at changing the characteristics of families, so the details of the causal structure have been less important. When considered directly, however, analyses have largely questioned whether such things as family income or parental education are the driving forces (e.g. Mayer, 1997 on family income or Behrman and Rosenzweig, 2002 on mother's education).

existing peer results are very sensitive to the measurement and specification of various influences on achievement. A central aspect of our analysis is the replication of alternative specifications within a consistent database so that elements of the previous inconsistency of findings can be disentangled.

The key issue in the identification of peer group effects on achievement is the separation of the effects of peers from other confounding influences. Two potential problems have pervaded the peer literature. First, measures of peer attributes may be good proxies for omitted or mismeasured factors that affect individual achievement, leading to biased results that quite generally exaggerate the importance of peers. Second, because of the simultaneous nature of peer interactions—a student both affects her peers and is affected by peers—separating the causal impacts is extraordinarily difficult, at least in the most general form. The formal theoretical literature has concentrated most attention on the latter issue, while we believe the former is much more important in the practical estimation of peer effects in schools. This section begins with a general model of educational achievement and then develops these two aspects of estimating peer effects.

2.1. A General Model of Achievement

In its abstract form, we begin with the commonly held view that today's achievement is influenced not just by current family, school and peer interactions but also by those of the past that establish the base for any current learning. Following Brock and Durlauf (2001) and Manski (1993), we separate peer influences into endogenous (behavioural) effects and exogenous or predetermined (contextual) effects. The first category refers to the contemporaneous and reciprocal influences of peer achievement on schoolmates, reflecting the fact that the achievement of peers is governed by similar achievement relationships. The second category includes measures of peers that are unaffected by current behaviour, such as socio-economic status or race. We follow past convention and represent the endogenous peer variable as $\bar{A}_{(-i)gs}$ (average peer achievement in a grade) and the predetermined peer variables as $\bar{P}_{(-i)gs}$ (where the subscript $(-i)$ means that these variables represent the average values computed over all students in the school and grade other than student i).⁴ Equation (1) describes achievement (A) for student i in grade G and in school s :

$$\begin{aligned}
 A_{iGs} = & \underbrace{X_{iGs}\beta_G + S_{Gs}\delta_G + \bar{P}_{(-i)Gs}\lambda_G + \bar{A}_{(-i)Gs}\gamma_G}_{\text{current inputs}} \\
 & + \underbrace{\sum_{g=1}^{G-1} X_{igs}\beta_g + \sum_{g=1}^{G-1} S_{gs}\delta_g + \sum_{g=1}^{G-1} \bar{P}_{(-i)gs}\lambda_g + \sum_{g=1}^{G-1} \bar{A}_{(-i)gs}\gamma_g + e_{igs}}_{\text{cumulative past inputs}} \quad (1)
 \end{aligned}$$

where X and S are vectors of family background and school variables, respectively. Because it is useful for developing the estimation issues, this representation separates current and past influences.⁵

⁴ For expositional purposes, we write this model as having a single endogenous peer measure. The analysis is easily generalized, and the general results are unchanged by adding other contemporaneous achievement measures.

⁵ Even though we present achievement just in terms of school experiences, leaving out preschool experiences is done solely for expositional ease. Given the estimation strategy, it has no effect on the results.

2.2. The Importance of Measurement in Peer Effect Estimation

In reality, researchers virtually never possess the entire history of the relevant inputs. Consequently specifications based on equation (1) are rarely if ever directly estimated. The most common alternative, lacking historical information, bases estimation solely on measures of the current values of X , S , $\bar{P}_{(-i)}$ and $\bar{A}_{(-i)}$. But, estimation of specifications of this form offers little hope of providing consistent estimates of the peer parameters (λ and γ). The main issue—one that is not specific to peer effect estimation—is that current characteristics will generally be correlated with unobserved past determinants of achievement, introducing the standard problem of omitted variables bias.

In peer estimation, ignoring history has a stronger impact. Because members of peer groups tend to have similar experiences over time through systematic neighbourhood and school choice, many omitted historical factors will be common to the peer group. Perhaps even more relevant, many left out or poorly measured contemporaneous inputs will also tend to be common to the group. Common past and current omitted factors that affect both individual i 's achievement and peer achievement will induce a correlation of contemporaneous peer factors and the individual error term (e), making peer effects appear important even when they have no true impact, i.e. even when λ and γ are identically zero.⁶

Such problems have been widely discussed within the general achievement literature. For example, in discussions of the interpretations of peer influences contained in the 1966 Coleman Report (Coleman *et al.*, 1966), the possible interactions of model misspecification and peer group measurement entered into early critiques (Hanushek and Kain, 1972; Smith, 1972). In other words, the nature and measurement of peer group factors implies that common model misspecification is particularly damaging to inferences about the importance of peers.

An approach to the general problem of estimating achievement relationships, which we follow below, begins by taking the first difference of equation (1). The value added specification reduces the data requirements to the inputs relevant for grade G , since all of the historical influences on the current achievement level drop out, as in equation (2):

$$\Delta A_{iGs}^c = X_{iGs}^c \beta + S_{Gs}^c \delta + \bar{P}_{(-i)Gs}^c \lambda + \bar{A}_{(-i)Gs}^c \gamma + v_{iGs}^c \quad (2)$$

where ΔA_{iGs}^c is the achievement gain (difference between current grade and previous grade test scores) for student i in grade G in school s in cohort c .⁷ Student achievement growth is related to the contemporaneous inputs (which are the flows of these factors over the observed time period), and the generic problems of omitted historical variables are circumvented.

Even with the value added form, consideration of peer influences complicates the estimation of achievement models. The problems of poorly measured individual school and background factors (either because of omitted or error-prone measures) have the same extra impact in the value added models because of the 'strong proxy' nature of peer measures. One important and relevant example is systematic but unmeasured elements of teacher quality. As a simple illustration, assume that the error term in equation (2) omits an aspect of teacher quality (κ_G) that, while uncorrelated with X_G

⁶ See Brock and Durlauf (2001), Manski (1993) or Moffitt (2001) for a formal development of the general cases of issues raised in this section.

⁷ An alternative estimation approach is to add a measure of prior achievement to the right-hand side. This approach does not constrain the parameter on prior achievement to be one but does add other complications with estimation (see Hanushek, 1979; Rivkin *et al.*, 2001). The identification of separate cohorts at this point facilitates development of the subsequent estimation strategy.

and the measured school aspects of S_G , is common to the achievement of peers. Even in the case where current peer achievement is irrelevant (i.e. $\gamma = 0$), the estimation of the effects of $\bar{A}_{(-i)}$ will yield an estimator of peer effects, $\hat{\gamma}$, with upward bias that systematically makes peers look more important. The magnitude of the bias is directly related to the importance of the omitted factor in determining achievement.⁸

In part to circumvent such problems of mismeasured current inputs commonly affecting all peers, a number of studies have simply dropped $\bar{A}_{(-i)}$ from the specification and included its lagged value in its place. Unfortunately, by itself this introduces a series of statistical and interpretative problems (depending on the precise nature of the underlying behaviour). For example, the lagged average achievement score is likely to remain correlated with the error term because of the serial correlation in unobserved teacher, school and individual factors. Thus, the simultaneous equations and omitted variables biases in estimating peer effects, while altered in form, are not eliminated. Additionally, the substitution of lagged achievement introduces another type of bias that we discuss in detail below.

Our primary strategy for dealing with these general issues begins by extracting fixed components of individuals and schools to deal not only with the most significant omitted variables problems but also the key elements of neighbourhood and school selectivity. Importantly, the substitution of lagged values of peer achievement in place of the current value within such a fixed effects framework is not subject to the problems of simultaneous equations and omitted variables biases, because the fixed effects eliminate the systematic family and school influences that are correlated over time.

From the starting point of equation (2), equation (3) decomposes the error, ν , into a series of components that highlights those factors most likely to contaminate the peer estimates:

$$\nu_{iGs}^c = \omega_i + \omega_s + \omega_{Gs} + \tau_G^c + \theta_{Gs}^c + \varepsilon_{iGs}^c \quad (3)$$

The first three terms capture time invariant individual (ω_i), school (ω_s) and school-by-grade effects on achievement (ω_{Gs}); the fourth factor (τ_G^c) captures cohort-by-year differences in the testing regime; the fifth component (θ_{Gs}^c) captures school-by-grade effects that vary from cohort to cohort, most notably the quality of teaching; and the final factor (ε) is a random error capturing individual shocks that vary over time.

Our approach makes use of matched panel data to remove explicitly the first four components: fixed individual, school, school-by-grade and cohort-by-grade effects. Notice how these fixed effects account for the primary systematic but unobserved differences in students and schools. The student fixed effects (in the gains formulation of equation (2)) account for all student and family factors that do not vary over the period of achievement observation and that affect the rate of learning—including ability differences, family child rearing practices, general material inputs, consistent motivational influences, and parental attitudes towards schools and peers. This approach thus directly deals with many of the most difficult issues of potential bias in the peer estimates arising from omitted and mismeasured individual and family factors.

⁸ Consider in this simple example that a common omitted factor, κ_G , enters linearly into the achievement of i and of all peers with a coefficient of β . The bias simplifies to $\beta^2 \text{var}(\kappa_G)/\text{var}(\bar{A}_{(-i)G})$. Intuitively, because the common omitted factor appears both in i 's achievement and that of peers, a positive bias on peer achievement is introduced.

Moffitt (2001) discusses a variety of possible sources of such errors including the potential impact of measurement problems.

Next consider any fixed differences in schools that are not perfectly correlated with the student fixed effects or included covariates (S and X). While these are typically correlated with peer group composition through school and neighbourhood choice, they are accounted for by school fixed effects. Finally, even systematic within-school changes in achievement gains across grades can be accounted for through the use of school-by-grade fixed effects.

The importance of the multiple cohorts should not be underestimated. For example, consider the possibility that achievement for students in some schools tends to decline as the students age due to factors other than peer achievement (e.g., adolescence may be more disruptive for economically disadvantaged students). If only fixed individual and school effects were removed (as is possible with panel data for a single cohort) in the estimation, the resulting positive peer effect estimate would suggest that students were responding to peers when in fact other factors had introduced a spurious relationship between the achievement gains of all students in a school. On the other hand, if fixed student and school-by-grade effects are removed—as is possible with data for multiple cohorts—such systematic changes in specific schools cannot drive the results.

The estimation of peer effects along with the fixed individual and school-by-grade effects intuitively relies on perturbations in the pattern of peers across grades and cohorts, i.e. the estimates are identified by small within-school and grade differences in peer group characteristics between cohorts. Such differences emanate from two sources: mobility into or out of the school and, less importantly, changes in student circumstances (e.g. income or achievement). The large annual mobility of students, averaging greater than 20% per year in the Texas public schools, accounts for much of the differences across grades and cohorts.⁹

One significant concern of course is the possibility that the observed changes in peers simply act as proxies for other changes in family or school inputs. Whether our estimation strategy can generate consistent estimates hinges upon whether the two time-varying components of the error terms (θ and ε in equation (3)) are orthogonal to the included peer variables. Three possibilities seem most important. First, the small observed changes in peer circumstances may be related to changes in family conditions that could bias the estimates. An increase or decrease in peer average income or achievement may result from similar changes in own family income that precipitate a school transfer and exert a direct effect on outcomes. Alternatively, shifts in local labour market conditions may cause changes in both own family and peer group average income, making it difficult to disentangle the influences of peers and family. Second, changes in school characteristics may affect both own achievement and that of peers. For example, the funding and availability of compensatory education programmes is linked to school average income, possibly building in a correlation between peer average income and programmatic effects. Or teacher differences in a specific grade may vary with peer characteristics. Finally, school selection by other families may be driven by attributes of the school, and the effects of such attributes may be confused with peer effects. Consider a school that is becoming dysfunctional, say because of an ineffective principal, and finds that all of its upper income families flee over time. In such a case, achievement of the remaining students could fall along with the incomes of peers, erroneously suggesting that peer income affects achievement even when there is no such relationship.

⁹ Individual cohorts will differ through random factors such as age patterns of children within a given school district. The estimation, however, relies on both changes across grades for a given cohort and differences among cohorts. Hanushek *et al.* (forthcoming) describe and analyse student mobility. The overall mobility rates cited also exclude any school changes related to natural movements from elementary to middle schools, even though such moves will generally affect school peers.

The severity of these potential problems depends in part on the ability to control for changes in families and schools. In this analysis (described below) we include time-varying measures of family income, school characteristics, compensatory programme status and overall school transfer behaviour. Perhaps more important given the controls for student and school-by-grade fixed effects, the potential severity also depends upon the speed with which families relocate in response to school conditions. The concern is simply that families adjust to changes in school quality (including peer composition) and thereby might induce bias through equilibrium selection behaviour. But, because residential moving is a costly process that undoubtedly includes some slow adjustment, movement due to parental selectivity of schools is almost certainly much slower than the movement of peer characteristics found in exogenous year-to-year variations. The assumption that families also react slowly (i.e. not in the current year) to specific variations in teacher quality seems natural, implying that there is no reason to believe that the choice behaviour of parents to current changes in teachers leads to any presumption about correlation of peer factors and annual variations in teacher quality. It seems plausible that differences captured by school or school-by-grade quality provide the prime motivation for any family selection of schools. Any remaining variations in annual teacher quality, even if large, must be orthogonal to the school-by-grade estimates of quality. Particularly because the average family has more than one child, mobility reactions to current shocks to teacher quality are likely to be minimal.

In sum, the choice of neighbourhood and school will tend to bias upwards the estimated effect of peer achievement unless an exogenous source of variation in peer achievement can be identified. Our estimation strategy, which relies on small changes over time and grades in peer characteristics, will provide consistent estimates of the underlying peer parameters unless systematic changes in the contemporaneous innovations to achievement (θ and ε) are correlated with the predetermined peer effects. Given the available measures of year-to-year changes in family and school characteristics and the structure of the data allowing for the removal of student and school-by-grade fixed effects, such correlations are likely to be of a very low order of importance.

2.3. The Reflection Problem

Before completing this discussion, an additional issue of peer influences must be introduced. The most vexing estimation problem, formulated in detail by Manski (1993), is the possible simultaneous determination of achievement for all classmates, with high achievement by one student directly improving the achievement of classmates and *vice versa*. This possibility, which has also received the most theoretical attention, is captured in equation (2) by the inclusion of $\bar{A}_{(-i)G}$, the average achievement of peers. If the achievement of each peer is also governed by equation (2), we would have $\bar{A}_{(-i)G}$ directly related to A_{iGs} through individual i 's influence on the others in the class. As Moffitt (2001) shows, this situation can be thought of as a standard simultaneous equations problem, where the induced correlation of $\bar{A}_{(-i)G}$ and v_{iG} leads to inconsistent estimation of the peer effect parameter. The reflection problem (in the terminology of Manski, 1993) presents a conundrum, because it is extremely difficult to identify the separate structural peer parameters (λ and γ) of equation (2) through standard exclusion restrictions.¹⁰ Without imposing functional

¹⁰ Necessary conditions for identification can be found in Brock and Durlauf (2001). An alternative estimation approach that relies on randomization of peers helps with problems of common omitted factors and the estimation of reduced form relationships but still leaves the basic simultaneous equations problem (see Moffitt, 2001 and the estimation in Sacerdote, 2001).

form restrictions, one needs to find aggregate peer factors that do not have an individual analogue in the achievement relationship, something that is difficult given the underlying conceptual basis that portrays peers as essentially extended families. Nor does randomization help, because current behaviour of the individual and peers will still be important.

The estimation and interpretation issue in this framework is whether the contemporaneous behaviour of peers is important or whether any peer relationship is essentially captured by the underlying characteristics including prior achievement. Understanding the dimensions of this issue requires more detailed consideration of the peer components in equation (2). At the outset, it is important to note that the standard terminology in the reflection problem—distinguishing between behavioural and contextual factors—can be confusing in the case of peers and achievement. Similar to measures of family background, the predetermined measures of peers, such as aggregate parental education levels or racial composition of classmates, are in part proxies for attitudes, behavioural patterns and learning related activities that systematically enter into the behaviour and learning of each student.

Of course the endogenous peer component represented here by current aggregate achievement is distinguished from the other factors mainly because of the reciprocal nature of the determination of peer achievement. It is 'behavioural' in the sense that each student's actions directly affect the rest of the class. It is this issue of simultaneity that severely complicates the estimation of γ , not the fact that exogenous characteristics are unrelated to peer group behaviour.

Our estimation concentrates on models that employ lagged peer achievement instead of the contemporaneous value. Interpretation of estimation built on lagged peer achievement depends on the relationship between lagged and current behaviour.¹¹ If lagged achievement captures all of the relevant variation in current peer behaviour (i.e. there are no year-to-year shocks in current behaviour), there is no bias. Of course in the fixed effects framework there would be no need to substitute for current peer achievement if this were the case. More realistically, lagged peer achievement is likely to be an imperfect proxy for the current value. If the difference between current and lagged measures of peer achievement is random (e.g. the probability of a family shock in grade G such as divorce is randomly distributed), the estimated effect of peer achievement will generally be biased towards zero in a normal proxy variable effect. Even if the current innovation to peer behaviour is correlated with lagged peer achievement, under most conceivable circumstances the estimated effect will still be downward biased. Therefore the estimated effect of lagged achievement should provide a lower bound estimate of γ , and we find little reason to believe, at least based on past estimation and descriptions of classroom behaviour and interactions, that the changes in individual behaviour in a particular grade are especially important when compared to the underlying systematic differences captured by lagged achievement.

3. THE UTD TEXAS SCHOOLS PROJECT MICRODATA

The cornerstone of this research is the analysis of a unique matched panel data set of school operations constructed by the UTD Texas Schools Project, a project conceived of and directed by John Kain. The data track the universe of three successive cohorts of Texas public elementary

¹¹ The more customary interpretation of equation (2) as a structural representation of achievement would lead to the estimation based on predetermined peer achievement being viewed as a reduced form relationship. We avoid this because we think of $\bar{A}_{(-i)}$ as itself a proxy for the current behavioural interactions that are important.

students as they progress through school, beginning with students who attended third grade in 1992. For each cohort there are over 200,000 students in over 3000 public schools. Unlike many data sets that sample only small numbers from each school, these data enable us to create quite accurate measures of peer group characteristics.

3.1. Sample Characteristics

We use data for grades three through six for the three successive cohorts. Only black, Hispanic and white students are included; the relatively few Asians are excluded in order to simplify the models.

The student data contain a limited number of student, family and programme characteristics including race, ethnicity, gender and eligibility for a free or reduced price lunch (the measure of economic disadvantage) and compensatory education services, but the panel feature can be exploited to account implicitly for time-invariant individual effects on achievement gains. Importantly, students who switch schools can be followed as long as they remain in a Texas public school.

Beginning in 1993, the Texas Assessment of Academic Skills (TAAS) was administered each spring to eligible students enrolled in grades three through eight. The criteria referenced tests evaluate student mastery of grade-specific subject matter. Unique IDs link the student records with the test data. This paper presents results for mathematics, although the results are qualitatively quite similar for reading achievement. Consistent with the findings of our previous work on Texas, schools appear to exert a much larger impact on math than reading in grades four through six (see Hanushek *et al.*, 2002; Rivkin *et al.*, 2001). Each math test contains approximately 50 questions. Because the number of questions and average percent right varies across time and grades, we transform all test results into standardized scores with a mean of zero and variance equal to one. The regression results are robust to a number of transformations including the raw percentage correct. In order to avoid complications associated with classification as limited English proficient (LEP) or disabled, all LEP and special education students are dropped from the analysis.¹²

Importantly, the student database can be linked to information on teachers and schools through the school IDs. The school data contain detailed information on individual teachers including grade and subject taught, class size, years of experience and student population served. While individual student–teacher matches are not possible, students and teachers can be uniquely related to a grade on each campus. Each student is assigned the school average class size and the distribution of teacher experience for teachers in regular classrooms for the appropriate grade and school year.

3.2. Family Background and School Variables

The fixed effects capture all stable student, family and grade specific school effects on achievement growth. Time-varying student and family factors (X) include indicator variables identifying eligibility for a free or reduced price lunch, school transfer and participation in the Federal Title 1 compensatory education programme for low income children. The vector of time-varying teacher

¹² Our analysis of special education (Hanushek *et al.*, 2002) suggests that a higher proportion of special education students in a grade raises the achievement of regular education students. This finding is, however, impervious to the mix of mainstreamed and pull out instruction and leads us to believe that it is not peer group composition per se but potentially aspects of classroom management and/or fiscal effects from added special education funding.

characteristics (S) includes average class size, percent of teachers with zero years of experience, and percent of teachers with one year of experience. (The relevant set of teacher variables is based on prior work on a generalized achievement model in Rivkin *et al.*, 2001).

The final important issue is the construction of the peer group characteristics. Variables are calculated from information on schoolmates by grade (own information is excluded from the calculations). Proportion black, proportion Hispanic and proportion eligible for a free or reduced price lunch use current information. Their construction is straightforward, though proportion eligible for a reduced price lunch is likely to be a noisy measure of peer economic circumstances.¹³

As previously discussed, the construction of measures of peer achievement is much more problematic. We concentrate entirely on predetermined measures of achievement in order to capture the effects of pre-existing differences in the level of human capital that may influence peers through social interactions. By measuring peer achievement with peer test scores from two grades earlier (but for the current classmates), we avoid building in a mechanical peer correlation.¹⁴ We explicitly consider both the mean and standard deviation of peer achievement and always exclude a student's own achievement from the calculations.

The inability to assign students to classrooms also means that peer group characteristics are computed by grade rather than by classroom. While such aggregation reduces problems introduced by the non-random division of students into classes, it also eliminates all within-grade and year variation in peer group characteristics and the possibility of examining the effects of ability tracking. As a consequence, the coefficient on the variance of peer achievement does not present a straightforward interpretation, because the schoolwide variance of peer achievement is not simply the average variance of achievement in classrooms as is the case with the level of peer achievement and percentage eligible for a reduced price lunch. Additional assumptions concerning changes over time in the division of students into classes are required in order to draw causal inferences from this coefficient.

4. THE EFFECTS OF PEERS

Baseline level and value added specifications and more complicated fixed effects models generate a series of estimated peer group effects. Table I reports results from level and value added specifications that do not remove either student or school fixed effects. These preliminary specifications are similar to the bulk of existing work and provide a baseline from which to compare the fixed effect estimates. Table II reports results from student, school and school-by-grade fixed effects specifications that attempt to account for the endogeneity of the choices made by families and schools.¹⁵ Using the full fixed effect model, Table III reports results from specifications that permit peer group effects to vary by a student's ranking in the school test score distribution.

¹³ The division of students into two family income categories misses substantial within-category variation. In addition, student cooperation is required to be classified and students may become more reluctant as they age, though the school-by-grade fixed effects should address this problem. Unfortunately, there is no additional information on family income, so that this often-used variable is the sole indicator of economic circumstances.

¹⁴ The problem with achievement in the previous grade (G-1) is that the dependent variable is the test score gain. A particularly good teacher who substantially increases achievement in grade G-1 might reduce the expected gains in grade G, given that the grade G-1 test score provides the baseline with which to measure grade G achievement gains. School specific non-random measurement error in the grade G-1 score may also be negatively correlated with grade G gains.

¹⁵ We do not report estimates from specifications in which only subsets of the peer characteristics were included. These estimates were quite similar to those for specifications that included all peer characteristics.

Table I. Estimated effects of peer group characteristics on mathematics achievement level and achievement gains (absolute value of Huber–White adjusted *t*-statistics in parentheses)

	Level		Gain	
	Without school characteristics	With school characteristics	Without school characteristics	With school characteristics
Average math score in grade G-2	0.42 (22.73)	0.42 (22.91)	-0.07 (6.12)	-0.07 (6.18)
Standard deviation of scores in grade G-2	-0.07 (1.67)	-0.06 (1.35)	0.08 (3.36)	0.09 (3.64)
Proportion eligible for reduced price lunch	0.16 (4.82)	0.13 (3.67)	-0.08 (4.78)	-0.10 (5.24)
Sample size	526,546		1,028,162	

Note: All specifications include percent black and percent Hispanic, interactions of percent black and percent Hispanic with own race, dummy variables for reduced price lunch eligibility, school transfer, Title I programme eligibility, gender, black, Hispanic and cohort-by-grade-by-year indicators.

Table II. Estimated effects of peer group characteristics on mathematics test score gains, controlling for student, school or school-by-grade fixed effects (absolute value of Huber–White adjusted *t*-statistics in parentheses)

	Student fixed effects	School fixed effects	School-by-grade fixed effects	Stayers	Movers
Proportion eligible for reduced price lunch	0.18 (3.33)	0.12 (1.74)	0.10 (1.82)	0.19 (1.38)	0.16 (2.73)
Average math score in grade G-2	0.17 (7.42)	0.16 (6.21)	0.15 (6.03)	0.24 (6.26)	0.14 (5.20)
Standard deviation of scores in grade G-2	0.06 (1.26)	0.05 (1.08)	0.02 (0.53)	0.12 (1.72)	0.05 (0.88)
Sample size	1,028,162	1,028,162	1,028,162	299,730	728,432

Note: All specifications also include individual fixed effects, percent black and percent Hispanic, interactions of percent black and percent Hispanic with own race, dummy variables for reduced price lunch eligibility, school transfer, Title I programme eligibility and cohort-by-grade-by-year indicators, and average class size, proportion of teachers with zero years of experience and proportion of teachers with one year of experience. Stayers are students who do not change school between fifth and sixth grade, while movers are those who do.

Specifically, separate peer group effects are estimated for each of the four quartiles of the school test score distribution (based on scores in third grade). Variable means and standard deviations are reported in Appendix Table AI.

All specifications include a number of variables whose coefficients are not reported in addition to average achievement, the variance in achievement and proportion eligible for a subsidized lunch. These include percent black and percent Hispanic along with interaction terms between black and percent black and Hispanic and percent Hispanic in order to permit the effects to differ by student race and ethnicity, dummy variables for reduced price lunch eligibility, school transfer and Title I programme eligibility, and cohort-by-grade-by-year indicators.

Table III. Estimated effects of peer group characteristics on mathematics test score gains by quartile of each school's test score distribution, controlling for student and school-by-grade fixed effects (absolute value of Huber–White adjusted *t*-statistics in parentheses)

	Bottom quartile	Second quartile	Third quartile	Top quartile
Proportion eligible for reduced price lunch	0.10 (1.44)	0.11 (1.91)	0.09 (1.73)	−0.03 (0.62)
Average math score in grade G-2	0.12 (4.85)	0.12 (5.20)	0.12 (5.64)	0.08 (3.48)
Standard deviation of scores in grade G-2	−0.10 (2.82)	−0.05 (1.50)	−0.01 (0.45)	0.01 (0.42)
Sample size	1,028,162			

Note: The specification also includes individual and school-by-grade fixed effects, percent black and percent Hispanic, interactions of percent black and percent Hispanic with own race, dummy variables for reduced price lunch eligibility, school transfer, Title 1 programme eligibility and cohort-by-grade-by-year indicators, average class size, proportion of teachers with zero years of experience and proportion of teachers with one year of experience.

4.1. Baseline Models

Table I presents basic models of the level and growth of achievement. Odd number columns exclude and even number columns include school characteristics. Because the estimates are quite insensitive to their inclusion, the remaining tables report only specifications that include the school variables. All specifications in Table I include indicator variables for race and ethnicity (which are subsumed in the individual fixed effects in subsequent tables).

Not surprisingly, there is a very strong positive relationship between math achievement level and the average achievement of peers in the levels specifications (columns 1 and 2). However, this relationship disappears or is reversed for the value added specifications that examine growth in achievement (columns 3 and 4). As discussed, coefficients from the levels specifications almost certainly confound peer effects with omitted family characteristics. Though achievement growth models substantially reduce the problem of omitted variables, these specifications may also be subject to contamination from a number of sources.¹⁶

In contrast to the levels specifications, the value added estimates conform to prior expectations for the income variable. A higher proportion of schoolmates eligible for reduced price lunches significantly reduces achievement gains. Finally, achievement gains are not negatively related to the standard deviation of student achievement as might be expected if heterogeneity reduces the effectiveness of classroom instruction. The negative relationship observed in the levels specifications disappears once value added models are introduced.

¹⁶ One possible explanation for the negative value added estimates is related to the test score instrument used in Texas. The test does a poor job of capturing gains in knowledge at the upper end of the distribution. To the extent that lower achieving students are catching up to others in Texas in terms of basic skills and the average peer achievement variable is a proxy for initial achievement level, the estimate of the average peer achievement effect would be subject to the type of negative bias observed in Table I. While our continuing work is investigating the possibilities of non-linearities, the results below on quintiles of the distribution do not suggest any simple explanation of this type.

4.2. Fixed Effect Estimates

Neither the simple value added nor levels specifications in Table I provide defensible estimates of peer group effects. Both are contaminated by problems of omitted variables bias that no doubt contribute to the unexpected directions of various peer effects. In this section we exploit the power of the stacked panel data sets to isolate the independent effects of peers. Coefficients from five specifications are reported. The first controls for student fixed effects (in achievement growth), the second controls for both student and school fixed effects, the third controls for student and school-by-grade fixed effects, the fourth restricts the sample to those in the same campus in both years (essentially removing a student specific school fixed effect), and the fifth restricts the sample to campus switchers. In our opinion, it is the variation in peer group variables that remains after controlling for both student and school-by-grade fixed effects that offers the most convincing identification of the true effects of peers on mathematics achievement.

Table II demonstrates that the overall estimates are quite sensitive to the error specification. Student fixed effects produce quite different estimates than the simple value added models in Table I, and the addition of school fixed effects leads to further changes for most variables.

Consider the pattern of estimates for the (lagged) peer average math score variable. Column 1 shows that the removal of student fixed effects leads to a highly significant coefficient of 0.17 that is only slightly reduced by the additional controls for school and school-by-grade fixed effects. Controlling for school differences by restricting the sample to only non-movers produces a somewhat higher estimate of 0.24 (column 4).¹⁷ The larger effect for non-movers is consistent with peer influences that are stronger for students who have been in the group for a longer period.

The difference between the school-by-grade fixed effect estimates and those for non-movers may also reflect an artifact of the variable construction.¹⁸ Despite the fact that own test score is excluded from the calculation of peer average achievement, a potential problem remains. With the fixed effects methodology the removal of the school-by-grade mean uses own test score as a part of the calculation (own achievement score contributes to schoolmates' peer average score).¹⁹ Because differences between cohorts far exceed the within cohort differences that result from the omission of each student's score from her peer group calculation and because own score enters with a two year lag, however, we do not believe that this mechanical result could be entirely responsible for the differences in estimates.²⁰

The range of 0.15 to 0.24 produced by the full sample school-by-grade and non-mover specifications provides the best estimate of the actual peer test score effect, though this should probably be considered a lower bound. The lagged test scores, even constructed from information for all schoolmates, provide a noisy measure of current peer achievement because some students undergo substantial changes as they progress through school.²¹

¹⁷ The separate estimates are restricted to students who remain in the same campus for both fifth and sixth grade, where the removal of individual fixed effects also removes an individual specific school fixed effect that is not contaminated by own prior score.

¹⁸ We thank an anonymous referee for pointing out this possibility.

¹⁹ This is of course fundamentally different from including the contemporaneous outcome measure in the peer group calculation (which would clearly raise the reflection issue).

²⁰ The non-mover estimates are generated from a specification that does not remove school-by-grade fixed effects, though the similarity of the school and school-by-grade fixed effect estimates in columns 2 and 3 suggests that this is unlikely to introduce a serious problem.

²¹ One other possibility that we explore is that students are more heavily influenced by own race/ethnic peers, possibly due to segregation of social interactions within schools. However, the estimates (not reported) do not support this hypothesis.

Though peer average achievement exerts a significant effect on achievement, there is little or no evidence in any of the specifications that changes in the heterogeneity of students (measured by variation in peer achievement) affects the rate of achievement growth. This finding suggests that ability grouping per se may have minimal effects on average achievement. Note, however, that we have information only on variance for the grade as a whole and not variance within classes. If schools alter their ability grouping policies in response to cohort differences in achievement variance, the aggregate measure of variance will not capture classroom differences across cohorts, and the estimates will not be informative.

Finally, the estimates in Table II do not support the view that lower income peers harm achievement. Because percent eligible for a free or reduced price lunch is a very noisy measure of peer income, this result is not that surprising. Moreover, the fact that the variable confounds actual income differences with differences in school efforts to classify children as disadvantaged introduces additional complications, perhaps contributing to the positive coefficient on percent low income in all specifications.

A potential problem with the fixed effects approach is that the removal of all between school variation reduces the ratio of signal to noise by leaving too little actual variation in peer group characteristics. Appendix Table AII reports the residual variances of the peer group variables following the removal of student and school or school-by-grade fixed effects. In the case of the achievement variable, slightly less than 10% of the original variance remains, meaning that a one standard deviation change in the residual roughly equals 0.1 standard deviations of the original test score distribution. This is far from the tip or edge of the distribution, and the pattern of estimates do not support a simple measurement error explanation. A similar reduction in variance occurs for the standard deviation of test score and the percentage receiving a subsidized lunch.²²

The strong conclusion that comes from this analysis is that the achievement of peers has a strong and direct influence on learning. While the exact causal mechanism remains ambiguous—because we cannot rule out the importance of current peer behaviour as opposed to simple skill differences—the estimates provide clear evidence of peer effects.

4.3. Differences by Quartile

The results in Table II reveal significant effects of peer average achievement for all students, but peer influences may affect some students more than others depending on their initial position in the school achievement distribution. To examine this possibility, we interacted all peer group variables with indicators for the student's position in each school's achievement distribution based on test score in third grade.²³ The coefficients in Table III are generated from a single regression.

The results reveal little variation by school achievement quartile with the exception that students in the top quartile may be somewhat less responsive to peer achievement.²⁴ There is no support for the belief that students further below the median are differentially affected than others closer

Rather the uniformly smaller (though still highly significant) estimates suggest that the use of only own demographic group schoolmates in the construction of the peer characteristics introduces measurement error.

²² Though variation in percent eligible for a subsidized lunch emanates from both changes in the student body and changes in classification as students age, it is student mobility that accounts for the bulk of the variance. In fact calculations of variance (not reported) showed that fixing student eligibility at the fourth grade response leads to virtually no change in the residual variances.

²³ Third grade test score was used to avoid a direct link with the dependent variable.

²⁴ An alternative explanation for the lower estimate for students in the top quartile is that ceiling effects in the test attenuate the estimates for these students.

to the centre of the distribution. Finally, preliminary work (not shown) found no evidence of non-linearities from specifications that ignored a student's position in the test score distribution but included quadratic terms for all peer characteristics.

5. CONCLUSIONS

The difficulties of isolating school and peer group effects have been well documented. We have attempted to overcome problems of omitted variables and simultaneous equations biases through the use of a fixed effects framework and lagged measures of peer achievement. The results strongly support the view that standard specifications are subject to biases, as the sequential introduction of student, school and school-by-grade fixed effects led to substantial changes in the magnitude and often the direction of peer effect estimates. We believe that the variation in peer group characteristics that remains after controlling for student and school-by-grade fixed effects in the rate of achievement growth and a number of time-varying student, family and school characteristics provides a valid source of identification for the estimation of peer group effects.

The results themselves provide little evidence that average income or the heterogeneity of peers in terms of variation in achievement levels affect growth in mathematics achievement. These results should be qualified by the fact that proportion eligible for a reduced price lunch is a noisy measure of income and by the fact we use grade rather than classroom level data. While it is possible that schools may act to counter adverse effects of heterogeneity, the year-to-year changes used in the fixed effects models suggest that given the structure of schools, an increase in the variance of achievement does not have a significant negative effect on math learning.

Perhaps the most important finding is that peer average achievement has a highly significant effect on learning across the test score distribution. A 0.1 standard deviation increase in peer average achievement leads to a roughly 0.02 increase in achievement. Given that a one standard deviation change in peer average achievement is 0.35 of a standard deviation of the student test score distribution and that the use of lagged test score introduces error into the measure of peer achievement, the point estimate suggests that differences in peer characteristics have a substantial effect on the distribution of achievement when cumulated over the entire school career.

One important drawback of the analysis is the exclusion of current peer behaviour that is uncorrelated with past peer achievement. If innovations to behaviour form an important avenue through which peers affect outcomes, the inability to capture such behaviour might lead to a serious underestimation of peer influences. Unfortunately, the identification of current behavioural effects presents serious obstacles in concept as well as practice. However, the well-documented persistence of a student's performance over time suggests that systematic differences among students account for much more of the variation in peer group quality than student variations from year-to-year.

In terms of public policy, the fact that the effects are similar across the test score distribution suggests that a reallocation of students will have little impact on the overall state or school average. Rather it will affect only the distribution of achievement across schools; winners from having more able peers are balanced by losers with less able peers. The findings also imply that there will be additional external benefits to improving student performance through special programmes, tutoring and the like. While such benefits are likely to be small in comparison to the main effect for the student receiving any treatment, it is clear that student outcomes are intertwined in important ways. Finally, much more must be learned about the effects of peers on other social

and economic outcomes, and classroom level data are needed to learn more about the impact of ability grouping.

APPENDIX

Table AI. Variable means and standard deviations

	Mean	Standard deviation
Test score gain	0.03	0.63
Peer characteristics		
Proportion eligible for reduced price lunch	0.44	0.26
Average math score in grade G-2	0.04	0.35
Standard deviation of scores in grade G-2	0.91	0.13

Table AII. Residual variance of peer characteristics following removal of fixed effects

	Fixed effects removed			
	None	Student	Student & school	Student & school by grade
Proportion eligible for reduced price lunch	0.07	0.0038	0.0030	0.0029
Average math score in grade G-2	0.12	0.0179	0.0137	0.0114
Standard deviation of scores in grade G-2	0.02	0.0035	0.0028	0.0024

ACKNOWLEDGEMENTS

We benefited from comments by William Brock, Steven Durlauf, Caroline Hoxby and participants of the Brookings Conference on Empirics of Social Interactions (January 2000). Support for this work has been provided by the Spencer Foundation, the Mellon Foundation, the Smith Richardson Foundation and the Packard Humanities Institute.

REFERENCES

- Argys LM, Rees DI, Brewer DJ. 1996. Detracking America's schools: equity at zero cost? *Journal of Policy Analysis and Management* **15**(4): 623–645.
- Armor DJ. 1995. *Forced Justice: School Desegregation and the Law*. Oxford University Press: New York.
- Behrman JR, Rosenzweig MR. 2002. Does increasing women's schooling raise the schooling of the next generation? *American Economic Review* **92**(1): 323–334.
- Benabou R. 1993. Workings of a city: location, education, and production. *Quarterly Journal of Economics* **434**(3): 619–652.
- Benabou R. 1996. Heterogeneity, stratification, and growth: macroeconomic implications of community structure and school finance. *American Economic Review* **86**(3): 584–609.
- Brock WA, Durlauf SN. 2001. Interactions-based models. In *Handbook of Econometrics*, Heckman JJ, Leamer EE (eds). North-Holland: Amsterdam; 3297–3380.
- Caucutt EM. 2002. Educational policy when there are peer group effects—size matters. *International Economic Review* **43**(1): 195–222.
- Coleman JS, Campbell EQ, Hobson CJ, McPartland J, Mood AM, Weinfeld FD, York RL. 1966. *Equality of Educational Opportunity*. U.S. Government Printing Office: Washington, DC.

- de Bartolome CAM. 1990. Equilibrium and inefficiency in a community model with peer group effects. *Journal of Political Economy* **98**(1): 110–133.
- Epple D, Romano RE. 1998. Competition between private and public schools, vouchers, and peer-group effects. *American Economic Review* **88**(1): 33–62.
- Hanushek EA. 1972. *Education and Race: An Analysis of the Educational Production Process*. Health-Lexington: Cambridge, MA.
- Hanushek EA. 1979. Conceptual and empirical issues in the estimation of educational production functions. *Journal of Human Resources* **14**(3): 351–388.
- Hanushek EA. 1992. The trade-off between child quantity and quality. *Journal of Political Economy* **100**(1): 84–117.
- Hanushek EA, Kain JF. 1972. On the value of ‘equality of educational opportunity’ as a guide to public policy. In *On Equality of Educational Opportunity*, Mosteller F, Moynihan DP (eds). Random House: New York; 116–145.
- Hanushek EA, Kain JF, Rivkin SG. Forthcoming. Disruption versus Tiebout improvement: the costs and benefits of switching schools. *Journal of Public Economics*.
- Hanushek EA, Kain JF, Rivkin SG. 2002. Inferring program effects for specialized populations: does special education raise achievement for students with disabilities? *Review of Economics and Statistics* **84**(4): 584–599.
- Henderson V, Mieszkowski P, Sauvageau Y. 1976. *Peer Group Effects and Educational Production Functions*. Economic Council of Canada: Ottawa.
- Hoxby CM. 2000. The effects of class size on student achievement: new evidence from population variation. *Quarterly Journal of Economics* **115**(3): 1239–1285.
- Katz LF, Kling JR, Liebman JB. 2001. Moving to opportunity in Boston: early results of a randomized mobility experiment. *Quarterly Journal of Economics* **116**(2): 607–654.
- Lazear EP. 2001. Educational production. *Quarterly Journal of Economics* **116**(3): 777–803.
- Ludwig J, Duncan GJ, Hirschfield P. 2001. Urban poverty and juvenile crime: evidence from a randomized housing-mobility experiment. *Quarterly Journal of Economics* **116**(2): 655–679.
- Manski CF. 1993. Identification of endogenous social effects: the reflection problem. *Review of Economic Studies* **60**(July): 531–542.
- Mayer SE. 1997. *What Money Can't Buy: Family Income and Children's Life Chances*. Harvard University Press: Cambridge, MA.
- Mayer SE, Jencks C. 1989. Growing up in poor neighborhoods: how much does it matter? *Science* **243**(March 17): 1441–1445.
- Moffitt RA. 2001. Policy interventions, low-level equilibria, and social interactions. In *Social Dynamics*, Durlauf S, Peyton Young H (eds). MIT Press: Cambridge, MA; 45–82.
- Oakes J. 1992. Can tracking research inform practice? Technical, normative, and political considerations. *Educational Researcher* **21**(4): 12–21.
- O'Regan KM, Quigley JM. 1999. Accessibility and economic opportunity. In *Essays in Transportation Economics*, Winston C, Gomez-Ibanez JA, Tye W (eds). Brookings Institution: Washington, DC; 436–468.
- Rivkin SG, Hanushek EA, Kain JF. 2001. Teachers, schools, and academic achievement. Working Paper No. 6691, National Bureau of Economic Research (revised).
- Rosenbaum JE. 1995. Changing the geography of opportunity by expanding residential choice: lessons from the Gautreaux Program. *Housing Policy Debate* **6**(1): 231–269.
- Rosenbaum JE, Popkin SJ. 1991. Employment and earnings of low-income blacks who move to middle-class suburbs. In *The Urban Underclass*, Jencks C, Peterson PE (eds). Brookings Institution: Washington, DC; 342–356.
- Sacerdote B. 2001. Peer effects with random assignment: results for Dartmouth roommates. *Quarterly Journal of Economics* **116**(2): 681–704.
- Smith M. 1972. Equality of educational opportunity: the basic findings reconsidered. In *On Equality of Educational Opportunity*, Mosteller F, Moynihan DP (eds). Random House: New York; 230–342.
- Summers AA, Wolfe BL. 1977. Do schools make a difference? *American Economic Review* **67**(4): 639–652.
- Zimmer RW, Toma EF. 2000. Peer effects in private and public schools across countries. *Journal of Policy Analysis and Management* **19**(1): 75–92.