

Experimental and Quasi-Experimental Analysis of Peer Effects: Two Steps Forward?

Bruce Sacerdote^{1,2}

¹Department of Economics, Dartmouth College, Hanover, New Hampshire 03755;
email: bruce.sacerdote@dartmouth.edu

²National Bureau of Economic Research, Cambridge, Massachusetts 02138

Annu. Rev. Econ. 2014. 6:253–72

First published online as a Review in Advance on
February 5, 2014

The *Annual Review of Economics* is online at
economics.annualreviews.org

This article's doi:
10.1146/annurev-economics-071813-104217

Copyright © 2014 by Annual Reviews.
All rights reserved

JEL codes: J00, I2

Keywords

social interactions in education

Abstract

In the past 10 years, there has been an explosion of well-identified studies that measure peer effects across many settings and for many outcomes. The emphasis on natural experiments and randomization is a highly useful one; in more standard observational studies, the self-selection of people into peer groups can make the measurement of peer effects extremely difficult. In the absence of exogenous variation, knowing that people have similar outcomes as their friends, classmates, and coworkers may tell us little about peer effects. I examine the successes, failures, and findings of experimental analyses of peer effects. I draw three broad conclusions. First, even more than in other areas of social science, the size and nature of peer effects estimated are highly context specific; peer effects in student test scores and grades are prominent in some cases and absent in others. That said, there is a pattern across studies suggesting that social outcomes (e.g., crime, drinking behavior) and career choices show larger peer influences than do test scores. Second, researchers have shown that the linear-in-means model of peer effects is often not a good description of the world, although we do not yet have an agreed-upon model to replace it. Third, despite potential temptation, we have not reached the point at which we can reliably use knowledge of peer effects to implement policies that improve outcomes for students and other human subjects.

1. INTRODUCTION AND OVERVIEW

Understanding and measuring peer effects are often viewed as a Holy Grail of social science and the key to understanding many social problems and opportunities. The Coleman report (Coleman et al. 1966) famously identifies peer inputs as a major challenge for disadvantaged students. Harris (2009) argues that parents can only influence their children's outcomes via affecting the set of peers for that student. Austen-Smith & Fryer (2005) develop a formal model of acting white, in which peer influences cause a group of minority students to lower their educational aspirations to be accepted by the group. And Glaeser et al. (1996) suggest that social interactions can explain why some locations are filled with crime and other locations are relatively safe.

Measuring the importance and nature of peer effects has proven notoriously difficult. People sort into schools, neighborhoods, and places of work in ways that make it difficult to disentangle self-selection effects from actual (causal) peer effects. If we ran a regression of one's own SAT scores on classmate SAT scores for students in high schools across the US states, we would undoubtedly find a large positive correlation. Yet few people would interpret this correlation as the effect that high school classmates have on each other.

One of the most-cited papers on peer effects is by Manski (1993), who points out that peers may have similar outcomes for any of three reasons: selection into peer groups (correlated effects), effects that emanate from a peer's background (exogenous effects), and effects from a peer's outcome (endogenous effects). Even if the only channel for peer effects to operate were the endogenous channel, there is still the reflection or endogeneity problem of identifying the causal effect of one peer on another. The traditional approach of looking for peer effects by regressing my outcomes on the average of my friends' outcomes can be terribly misleading if the formation of our friendship group is based on unobserved characteristics that we all have in common.

In the *New England Journal of Medicine*, Christakis & Fowler (2007) look at a network of friends and spouses within the Framingham Heart Study. They find that a person's chances of becoming obese are 57% greater if he or she has a friend who becomes obese. This is a powerful correlation, but how much of it results from peer effects and how much results from the tendency of people to form friendships with people from similar ethnic and socioeconomic backgrounds and with similar exercise habits? Cohen-Cole & Fletcher (2008) find that attempts to control for selection into peer groups greatly reduce the estimated peer effect on obesity.

In survey data, Case & Katz (1991) find a strong connection between one's own youth crime and neighborhood peer behavior. They are somewhat agnostic as to how much of this connection is causal. In contrast, Kling et al. (2005) examine families randomly assigned to a program designed to get some families to move to lower-poverty neighborhoods. There, the experimentally measured effects of neighborhoods on criminal behavior are more nuanced. Young males moving to lower-poverty neighborhoods see a drop in arrests for violent crime but increases in property crime and behavioral problems at school.

In the absence of experimentally induced variation, one approach to solving the thorny selection issues is to impose mathematical structure on the problem. If we simplify the selection issues and various types of peer effects into a defined series of equations, we can in theory estimate all the parameters of interest. One example is provided by Bramoullé et al. (2009). The authors use data from the National Longitudinal Study of Adolescent Health. The data contain many covariates and outcomes for a set of middle school and high school students and their friends.

Bramoullé et al. (2009) assume that peer effects take a particular form, namely a linear-in-means model (see Equation 1 for more detail). The authors achieve identification via imposing the constraint that all peer effects work strictly through the linear parameters in the model. For example, the model posits that my friends affect me directly, whereas the friends of my friends

affect me only indirectly. Thus my friends' friends outcomes or background provide an instrument to identify the degree to which my friend's outcomes affect me directly. Such solutions can be elegant and can be applied to a wide range of observed data. However, the downside to such an approach is that the estimates could be unconvincing or could leave the reader wondering whether the results are driven by the data or by the strong assumptions imposed. Would the parameter signs be reversed if we allowed friends of friends to have a modest direct effect on me? What if people are influenced not by the group mean but by their least-able friend?

Economists and other social scientists have been able to address some of these identification problems by finding (or creating) contexts in which peers are exogenously assigned to one another (e.g., thrown together at random). All manner of interesting examples have been proposed, including (a) Hoxby's (2000a) observation that schools exhibit cohort-to-cohort variation in the gender mix, (b) school admission policies that provide a different set of peers for students who are right at the margin of acceptance or denial to a selective school (Jackson 2013, Abdulkadiroglu et al. 2014), (c) the use of randomly assigned college roommates or school section mates (Sacerdote 2001, Zimmerman 2003, Duncan et al. 2005, Carrell et al. 2009, Shue 2013), (d) court-ordered desegregation of schools (Angrist & Lang 2004, Billings et al. 2014), and (e) the study of political or hurricane refugees who are resettled across a country or state in an exogenous way (Gould et al. 2009, Damm & Dustman 2012, Imberman et al. 2012).

Hoxby's (2000a) cohort-to-cohort variation is based on the insight that small classes of 15–25 students will contain significant variation in the fraction of girls. This variation stems from random or at least exogenous changes in the demographics of a school's catchment area. This variation is quite useful given that girls have on average higher reading scores and lower rates of class disruption. Thus Hoxby is able to examine peer effects in a case with random variation in the characteristics of one's peers.

In the Jackson (2013) and Abdulkadiroglu et al. (2014) examples, the authors examine shocks to peer groups that occur because some students are right at the margin of admission or non-admission to a selective school. Jackson attempts to distinguish between general school quality effects and short-run changes in peer effects by relying on cohort-to-cohort shocks in peer composition.

Abdulkadiroglu et al. (2014) combine the possible peer effects of attending an exam school in Boston or New York with all the other effects (e.g., teacher or resource effects) of attending the school. Given that these authors find no effect on average from attending an elite exam school, they conclude that the upgrade in peer ability that occurs from attending the exam school likely has no effect on outcomes.

Sacerdote (2001) and Carrell et al. (2009), for example, use college or university settings to examine the relationships of roommates and dormmates. They examine specific contexts in which first-year roommates (or hallmates or squadron mates in the case of military academies) are randomly assigned by the housing office. This creates exogenous variation in one's peer group, which is then used to ask how much peers matter, which peers matter, and for what outcomes.

Court-ordered desegregations provide yet another potentially exogenous shock to peer groups. Angrist & Lang (2004) examine students in the receiving schools in Boston's METCO program. The idea is that a group of lower-achieving students is ordered to attend a suburban or, in some cases, a different urban school. The authors' goal is to ask whether incumbent students in the receiving schools see differential test-score growth. Papers using natural disasters such as Hurricane Katrina rely on a similar logic. Hundreds of schools in Texas and Louisiana played host to Katrina evacuees. Imberman et al. (2012) ask whether the quantity and academic ability (as measured by prestorm test scores) of the evacuees matter for incumbent student achievement.

The purpose of this current review is to broadly examine the contribution of all these experimental and quasi-experimental analyses of peer effects. I draw out common themes that have emerged to emphasize what we do not yet know, despite the recent focus on convincing identification.

The rest of the review proceeds as follows. Section 2 presents a simple linear model of peer effects and discusses how this might be expanded to handle nonlinear peer effects. I discuss a few structural or econometric approaches to identification that have been used, noting the benefits and drawbacks of each. I provide two examples of peer effects that have been estimated with and without the benefit of exogenous assignment of peers to groups. Section 3 contains the bulk of the article. I present and compare results from four different sources of exogenous variation that have been used to identify peer effects (social interactions). I compare results estimating peer effects in similar realms of student outcomes or human behavior. I examine attempts to use our knowledge of peer effects in making policy and or arranging classrooms, peer groups, and housing groups. Section 4 tries to draw broad conclusions from the results thus far and suggests possible future directions for research.

My bottom lines are as follows. (a) In the absence of some form of quasi-experiment, the endogenous selection of peers into groups makes the credible identification of peer effects difficult. (b) The size and nature of peer effects vary tremendously by outcome, age, location, and the precision with which we can define the peer group. That said, effects on test scores are more modest than effects on criminal behaviors, career choices, attitudes, and more social outcomes (e.g., joining a club, playing sports, binge drinking). (c) Finally, a combination of endogenous peer choice and our lack of understanding of the underlying model makes it difficult to proceed from merely demonstrating that peer effects exist to creating policy that uses peer effects to change outcomes in a desired direction.

2. HOW SHALL WE MODEL PEER EFFECTS?

2.1. The Linear-in-Means Model

The most basic and commonly used model of peer effects is the linear-in-means model, in which my outcome depends on the average of my peers' outcomes, my background characteristics, and the average of my peers' characteristics. Formally, this can be written as

$$Y_i = \alpha + \beta_1 \bar{Y}_{-i} + \gamma_1 \mathbf{X}_i + \gamma_2 \bar{\mathbf{X}}_{-i} + \varepsilon_i, \quad (1)$$

where Y_i represents my own outcome, \bar{Y}_{-i} represents my peers' average outcome, \mathbf{X}_i is a vector of my background characteristics, and $\bar{\mathbf{X}}_{-i}$ is a vector of my peers' average background characteristics. This model includes both what Manski (1993) terms exogenous effects from my peers' background (γ_2) and endogenous effects (β_1) from my peers' outcomes. Manski also points out that unobserved selection into peer groups will create correlated effects that will cause one's own and peers' outcomes to be correlated for reasons that are not a causal effect.

Clearly, having a single source of exogenous variation (e.g., random variation in one's peers) will not answer all three of Manski's challenges to identification. Within the experimental and quasi-experimental literature, the typical case involves some exogenous variation in the composition of classmates or roommates. In this standard case, we cannot identify both β_1 and γ_2 . The most common way for authors to proceed is to then estimate the reduced-form effects of peer background on one's own outcomes, without recovering the parameters β_1 and γ_2 .¹ For some examples of this approach, readers are referred to Zimmerman (2003), Angrist & Lang (2004), Carrell et al. (2009), and Imberman et al. (2012).

¹In other words, we know that peer background affects my outcomes, but we do not attempt to say what portion of the effect from peer background is channeled through peer outcomes.

A different way to proceed is to assume away the presence of exogenous effects and declare all peer effects to work directly through peer outcomes (i.e., assume that $\gamma_2 = 0$). One can then instrument for peer outcomes using (exogenous) peer background characteristics as the instrument and directly estimate β_1 , the endogenous peer effect.

Of course, there are numerous other assumptions, or additional sources of identification, that one might bring to bear on the problem, and many authors have done so.² One such solution is proposed by De Giorgi et al. (2010). They propose that, although my (randomly assigned) classmates may have peer effects that stem both from outcomes and from their backgrounds, the friends of my friends only affect me via my friends' outcomes. This is a clever approach, even if it is not completely clear that the assumption will be valid in all contexts.

A more general solution is proposed by Moffitt (2001) and is often called the partial population approach. The idea is to consider a random assignment treatment (e.g., paying people to visit the gym) and allow the fraction of group members who are given the treatment to vary. Not only will this allow identification of the direct effect of the treatment on one's own behavior, but the treatment also provides an instrument for my peers' behavior, which should allow identification of β_1 .

A nice application of this idea is Babcock & Hartman's (2010) experiment in encouraging University of California, Santa Barbara, students to exercise at the university gym. Prior to implementing the experiment, the authors elicited information about the friendship network. Random assignment led to subjects having different fractions of their friends being treated. Thus the authors can estimate the effect of having treated friends for both treatment and control subjects and are able to extrapolate these results to treating entire dorms or friendship networks.

Kuhn et al. (2011) provide a nice partial population study of how consumption varies when one's peers win the lottery. The Dutch Postcode Lottery is arranged so that all ticket holders within a winning neighborhood receive a prize. It turns out that both winning and nonwinning households in the neighborhood are more likely to buy a new car relative to people in a neighborhood in which no one won a lottery prize.

It is worth noting that the linear-in-means model has been popularized not necessarily because it is the best description of the world, but because it is a convenient and parsimonious formulation. In one of the original peer effects papers, Summers & Wolfe (1977) model student test scores as a function of peer mean outcomes because the group mean is the only peer characteristic available. Yet when people discuss peer effects mechanisms or relate anecdotes, they rarely think of the peer mean or class mean as being the relevant influence. Rather, the peer effects that we observe day to day tend to stem from a friend or close set of friends introducing us to a new sport or a new career possibility or encouraging us to work hard (or not) at some endeavor.

Strong modeling assumptions (using linear in means or other formulations) often allow one to separate the direct effect from the reflection (multiplier effect) and hence identify β_1 . For example, in an early version of the Dartmouth College roommates paper (Sacerdote 2000), I take the linear-in-means model quite literally and note that with exactly two roommates, we have two equations:

$$Y_1 = \alpha + \beta_1 Y_2 + \gamma_1 X_1 + \gamma_2 X_2 + \varepsilon_1, \quad (2a)$$

$$Y_2 = \alpha + \beta_1 Y_1 + \gamma_1 X_2 + \gamma_2 X_1 + \varepsilon_2. \quad (2b)$$

Given these two equations, it is possible to substitute Equation 2b into Equation 2a. I then solve for a reduced form, which expresses Y_1 in terms of X_1 , X_2 , and the four parameters α , β_1 , γ_1 , and

²Economists are nothing if not clever and prolific.

γ_2 . I write down the formulae for the residual of the reduced form and the covariance of the reduced-form residuals across roommates, expressing them in terms of the variance of ε and β_1 (see Sacerdote 2000 for details). Armed with this algebra, I then run the reduced form to obtain reduced-form parameters and the reduced-form residual variance and covariance across roommates. Finally, I can back out the structural parameters from the reduced-form estimates.

A different structural approach to the problem is Graham's (2008) modeling of peer effects in the Tennessee STAR experiment. Graham imposes restrictions on the various components of the variance of performance across classrooms. He notes that classroom variance could stem from variance in teacher effects, variance in individual student-level ability, and finally "excess" variance that is caused by peer effects. By assuming the linear-in-means model, he can obtain exact expressions for each of these variances, obtain reduced-form estimates, and finally back out β_1 . With either of these structural approaches (Sacerdote 2001 or Graham 2008), one needs to ask how much the results (and the appearance of precise estimates) are being driven by assuming a strong model.

2.2. Going Beyond the Linear-in-Means Model

The linear-in-means model has the virtue of simplicity but does suffer from the disadvantage of not being a convincing description of the world. In fact, as Hoxby & Weingarth (2005) point out, peer effects with nonlinearities may be much more interesting because nonlinearities open up the possibility that some people (or students) could be helped by a change in peers without making other people worse off.

Many attempts to test the linear-in-means model reject the model in favor of more complex alternatives. For example, Carrell et al. (2009) find that higher-ability peers at the US Air Force Academy provide greater positive peer effects for lower-ability students than for middle-ability students. In examining peer effects from Katrina evacuees, Imberman et al. (2012) find something of the reverse: Higher-ability peers help higher-ability students more than lower-ability students.

Hoxby & Weingarth (2005) provide a coherent framework for estimating peer effects that go beyond the linear-in-means model. Their suggestion is to calculate the fraction of peers within different quantiles (e.g., deciles) of peer ability. The fraction of peers in each ability level can then be interacted with a dummy variable for the quantile of one's own ability. In the case of deciles, this would of course lead to 100 interaction terms, and the results can be presented graphically rather than as coefficients.

Hoxby & Weingarth (2005) then provide a classification of different models of peer effects that can be tested using these coefficients. For example, the boutique model suggests that students do well in a classroom that contains students similar to themselves, whereas the rainbow model implies that students benefit from having peers of diverse abilities.

There are of course numerous ways in which Equation 1 could be generalized; the Hoxby-Weingarth quantiles and interactions idea is just one approach. However, this specification has the advantage of being fairly flexible and transparent. As a result, many nonlinear peer effects estimations follow a version of this specification.

2.3. The Importance of Exogenous Variation in Peer Group Composition or Attributes

The process by which neighbors, students, or coworkers clump together to form a neighborhood, firm, or classroom is anything but random. Indeed, parents spend a great deal of time choosing a place to live or a school or advocating for a particular teacher so that their child will have a set of

peers, a classroom, and a teacher that are advantageous for the child. High-productivity firms may be particularly attractive to other high-productivity workers, and so on.

Numerous studies show, for example, that youth in high-crime neighborhoods are themselves much more likely (than other youth) to engage in criminal and antisocial activity (see, e.g., Case & Katz 1991). But it is difficult to know what portion of this relationship is causal. In contrast, Kling et al. (2005) have exogenous variation in neighborhood choice via the Moving to Opportunity (MTO) experiment. They can measure the impact of neighborhood (broadly defined to include the associated shift in school choice) on arrests and self-reported crime.

Using data from the Framingham Heart Study, Christakis & Fowler (2007) find that a person is 57% more likely to be obese if he or she has a friend who is obese. The obvious question is, how much of this is causal, and are there unobserved characteristics that lead two people with similar health trajectories to form a friendship? I would argue that we do not really know and that it may be that only a small part of the coefficient may be causal. Cohen-Cole & Fletcher (2008) find that attempts to control for selection into peer groups greatly reduce the estimated peer effect on obesity. Running the same regression with randomly assigned college roommates would not necessarily solve the debate either because college roommates (and their connection and health status) provide such a different context than friendships among middle-aged adults.

3. STRATEGIES FOR IDENTIFICATION AND WHAT WE HAVE LEARNED

3.1. Exogenous Movements of People

Social scientists have exploited all manner of natural and controlled experiments that cause an exogenous shock to a location (or school) and hence to a peer group. One can think about the peer effects on the people who are moving or the peer effects that the migrating people (students) have on incumbents in the new location. Typically, it is easier to claim that we are isolating peer effects for the latter type of analysis, but in principle, both types of studies are interesting.

Court-ordered (or voluntary) desegregation of schools is one event that causes a potentially exogenous shift in peer characteristics for students in the receiving schools. Angrist & Lang (2004) consider the effects of busing students from Boston to Brookline public schools on students who live in Brookline. They rely on variation in the percent of METCO students across different years and different grades within a school. They find no effects on test scores for local students.

The one exception to this finding is that African-American students see significant declines in third grade test scores from adding METCO students. A 5% increase in the percent of METCO students for a third grade cohort reduces the test scores of African-American students by 0.3 standard deviations. Angrist & Lang (2004) point out that with a point estimate this large, it is unlikely that such an effect could be working through test-score peer effects alone; such an increase in the percentage of METCO students reduces average peer test scores by only 1.25 percentage points.

This study suggests no peer effects on average, but there are hints that there may be big peer effects for some subgroups. Imberman et al. (2012) strike a similar note. These authors look at variation in peer composition that stems from the arrival of Hurricane Katrina evacuees into schools in Houston, Texas, and throughout Louisiana. There is not much evidence that on average the arrival of Katrina evacuees lowered the achievement of incumbent students in the receiving schools.

However, when Imberman et al. (2012) interact one's own achievement (by quintile) with the percent of evacuees in each quintile of the score distribution, they find stronger evidence of peer effects. A 10–percentage point increase in the fraction of “evacuees from the lowest quartile of the statewide distribution” is associated with a reduction in test scores for top-quartile incumbents of 0.17 standard deviations. Conversely, a 10–percentage point increase in the fraction of evacuees

from the top quartile of the statewide distribution helps top-quartile incumbents by 0.09 standard deviations.

Billings et al. (2014) provide a nice analysis of effects stemming from the ending of court-ordered desegregation in Charlotte-Mecklenburg County schools. In 2001, the school system was prevented from using race to assign schools. As a result, schools experienced a rapid change in composition and accompanying resegregation. The authors use a study design that looks at the effects of a differing new school assignment within neighborhood*old school assignment cells. In other words, the school catchment areas were redrawn such that there are often two students in the same neighborhood who used to attend the same school but are now assigned two different schools with very different demographic profiles.

Billings et al. (2014) use “percent minority” as the key right-hand side variable. A 10–percentage point increase in the share of minority students lowers high school test scores by approximately 0.02 standard deviations for students who were in high school at the time of resegregation. Nonminority students in these same cohorts also see decreased high school graduation rates from an increase in the percent of minority students at their assigned school. Minority students do not experience lower graduation rates from being assigned to a higher-percent minority school.

Interestingly, those students who were in middle school at the time of resegregation do not experience statistically significant negative effects on high school test scores. That suggests that over time, the negative effects (of assignment to a higher-percent minority school) dissipate, possibly because schools learn to serve their new population or possibly because the county was able to direct resources where they were most needed.

Clearly, it would be a stretch to claim that the effects found in the Billings et al. (2014) all, or even mostly, stem from peer effects. Having a high percentage of minority peers may be correlated with having teachers with lower test-score value added or schools with fewer resources.

Perhaps most interestingly (from a peer effects point of view), Billings et al. (2014) find large and persistent effects on criminal activity. A 10% increase in “percent minority” at the assigned school leads to a 1.4–percentage point increase in the likelihood “ever arrested” for male minority students.

The MTO experiment has given rise to an equally fascinating set of papers also connected (broadly speaking) to the peer effects literature. The experimental group was given a housing voucher that could only be used to relocate in a US Census tract that had a poverty rate of 10% or less. The control group was not given such a voucher, and a third group was given a Section 8 housing voucher that could be used in any location.

Kling et al. (2005) find the striking result that assignment to the MTO experimental group significantly lowers arrest rates for girls and raises property arrests for boys. For example, the treatment-on-the-treated estimate for girls on the number of total arrests is -0.54 relative to the control mean of 0.61. With regard to the number of property arrests for boys, the treatment-on-the-treated estimate is 0.36 relative to a control mean of 0.47. As with the Charlotte-Mecklenburg resegregation study, these experimental effects need not be peer effects per se, but peer effects may be playing a large role.³

Several papers estimate peer effects on natives from the arrival of political refugees or immigrants. Gould et al. (2009) use the arrival of large numbers of immigrants into Israel in the 1990s.

³In related work, Kling et al. (2005) and Ludwig et al. (2001) find that MTO increases various measures of well-being and reduces obesity rates by 4.5 percentage points. There is no evidence that MTO raises test scores or graduation rates for the youth in the experimental group families.

To control for endogenous choice of location, the authors rely on variation in the fraction of immigrants across cohorts and within schools.

There seems to be little effect on average from being in a school*cohort that has a higher fraction of immigrants. However, when Gould et al. (2009) focus on native students with less-educated mothers, the peer effects from increases in the fraction of immigrants are economically and statistically significant. A 10–percentage point increase in the fraction of immigrants leads to a 2–percentage point increase in the high school dropout rate (against a mean of 8 percentage points).

3.2. Random Variation Across Cohorts

The concept of gaining identification from within-school cohort-to-cohort variation in peer characteristics need not be limited to instances of immigration, natural disasters, or political asylum. In fact, one of the earliest applications of random cohort-to-cohort variation in peer composition is by Hoxby (2000a). Her insight is that elementary school cohorts are small enough that they have meaningful variation in the fraction of girls and the racial mix. Girls tend to have higher reading scores and fewer behavioral infractions than boys. Therefore, naturally occurring variation in the gender (or racial) mix of an elementary school cohort can provide exogenous variation in peer characteristics and peer test scores, in particular.⁴

Hoxby (2000a) finds that both boys' and girls' reading scores rise approximately 0.04 test-score points when their cohort is 10–percentage points more female. (Average test scores are roughly 35 with a standard deviation of 2.5.) The effects on math scores are a similar size. The result for math scores leads to an interesting point. Hoxby notes that it is unlikely that average peer test scores are the sole channel for the peer effects she observes. Girls only have slightly higher math scores than boys (i.e., a 10–percentage point increase in the fraction of girls only slightly moves average math scores). If the effect on math scores of having more girl peers worked solely through an increase in average math scores, then the reduced-form effect of a 1 point rise in peer scores would yield an increase in one's own score of 1.7–6.8 points.

Lavy & Schlosser (2011) are able to take this identification strategy one step further and look at the mechanisms behind the peer effects from a more-female cohort. They use cohort-to-cohort gender variation within schools in Israel. Having 10% more girls in one's cohort raises the probability of a high school student obtaining a matriculation diploma by approximately 0.5% for boys and 1% for girls. The matriculation diploma is needed to continue on to university study, and approximately 50% of boys and 60% of girls achieve this milestone.

Lavy & Schlosser (2011) also find statistically significant effects on the number of credits taken and on enrollments in math and science courses. But what is particularly interesting about this paper is the authors' ability to examine effects of the percentage of girls on student-reported measures of the classroom environment, in-school disruptions, and student effort. Having a larger fraction of girls greatly reduces classroom disruptions and violence. This arguably is an important mechanism through which gender variation affects student achievement.

Anderson & Lu (2012) take the literature in an interesting direction by asking how elementary students benefit by having a larger fraction of girls seated directly around them in the classroom. A girl who is surrounded by all girl peers experiences a 0.20 standard test-score rise relative to a girl who is surrounded by all boys.

⁴Hoxby (2000b) uses cohort-to-cohort variation to estimate the effect of class size on student achievement.

3.3. Using Test-Score Discontinuities

Jackson (2013) combines cohort-level variation plus test-score cutoffs for school admission in Trinidad to separately identify both the effects of attending a highly selective school and the importance of peer effects within those schools. For the typical student, peer effects are a small portion of the total school effect. However, for the highest-achieving students, peer effects are nearly one-third of total school value added.

Abdulkadiroglu et al. (2014) measure the effects of attending elite exam schools in Boston and New York. Admission to these schools is largely based on an entrance exam, with a cutoff used to determine which students are admitted and which are denied. This sets up a nice regression discontinuity as one can compare outcomes for students with scores just above the admissions cutoff with the outcomes for students whose scores fall just below.

Abdulkadiroglu et al. (2014) motivate their paper as a study of peer effects, as successful applicants to these exam schools are surrounded by some of the highest-achieving peers in the country. The authors find no impacts of attending an exam school on SAT scores or on state standardized exams. From this, they conclude that having a higher-achieving peer group does not help one's own test-score outcomes. They also find no effects on the probability of attending a selective college.

Table 1 shows that several researchers find positive linear-in-means peer effects in test scores for elementary and secondary students. It reports estimates of the effects of peer background score on one's own score. Hoxby's (2000a) estimates are at the higher end of the range, although again she emphasizes that effects from the fraction of girls or the fraction of nonwhite students need not work solely through test scores. This point likely applies equally well to most other studies. Angrist & Lang (2004) find only limited evidence for peer effects of METCO students on test scores of non-METCO students.

Dahl et al. (2012) use variation in which new parents have peers who are subject to a new (more generous) paternity leave policy in Finland. Some subjects have coworkers (peers) who have newborn children just before the policy takes effect, and some peers have a child born just after the policy goes into effect. The authors find substantial peer effects in take-up of the program. Having a peer take up the program makes one 11 percentage points more likely to take up the program. Furthermore, the authors add structure to the problem to estimate the snowball effect as take-up spreads throughout a firm.

3.4. Random Assignment of Dormmates and Roommates

As an alternative to cohort-to-cohort variation, many authors rely on the use of randomly assigned roommates, dormmates, or classmates within a college setting. Although such studies are highly specialized to a particular context, they do have the advantage of being able to identify specific groupings of peers who interact closely with one another. In other words, it is possible to closely define the peer group.

For instance, in a series of papers, Carrell et al. (2009) and Lyle (2009) make use of randomly assigned squadrons at the US Air Force Academy and the US Military Academy at West Point, respectively. The US Air Force Academy provides an ideal laboratory for testing peer effects for at least two reasons. First, incoming students are randomly assigned to squadrons with approximately 30 other first-year students in the squadron. Other than class time, there are only limited opportunities for interactions with students outside of one's squadron. This means that the relevant peer group is better defined than in other settings. Second, all students in the entire cohort are also randomly assigned to course sections. All first-year students take the same set of core courses, which are graded using the same exams on a common curve.

Table 1 Peer effects across a variety of experimental and quasi-experimental studies^a

Paper	Sample	What is estimated	Identification	Standardized effect of increase in peers' background	Effect of peers' outcome on one's own outcome
Peer effects in test scores					
Lavy & Schlosser (2011)	High school students in Israel	Effect of the change in the fraction of lower-ability students on matriculation test scores	Cohort variation within school	0.012 to 0.036	
Hoxby (2000a)	Elementary school students in Texas	Effect of the fraction of female students and fraction of nonwhite students on one's own test score and the implied effect of peer test scores on one's own test score	Cohort variation within school	Girls' reading score: 0.34 to 0.52 Boys' reading score: 0.31 to 0.50	
Imberman et al. (2012)	Elementary school students in Louisiana	Effect of peer test scores on one's own test score	Hurricane-induced shocks to peers	Math: 0.33 (0.15) Reading: 0.00 (0.270)	
Imberman et al. (2012)	Middle and high school students in Louisiana	Effect of peer test scores on one's own test score	Hurricane-induced shocks to peers	Math: 0.15 (0.08) Reading: 0.08 (0.08)	
Peer effects in GPA					
Sacerdote (2001)	Undergraduate students at Dartmouth College	Effect of roommate background (Academic Index) and roommate GPA on one's own GPA	Randomly assigned roommates	−0.030; not statistically significant	0.068 (0.029)
Zimmerman (2003)	Undergraduate students at Williams College	Effect of roommate background (SAT) on one's own GPA	Randomly assigned roommates	0.104; not statistically significant	
Stinebrickner & Stinebrickner (2006)	Female undergraduate students at Berea College	Effect of roommate background (SAT) on one's own GPA	Randomly assigned roommates	0.017; not statistically significant	

(Continued)

Table 1 (Continued)

Paper	Sample	What is estimated	Identification	Standardized effect of increase in peers' background	Effect of peers' outcome on one's own outcome
Stinebrickner & Stinebrickner (2006)	Female undergraduate students at Berea College	Effect of roommate background (high school GPA) on one's own GPA	Randomly assigned roommates	0.100	
Carrell et al. (2009)	Undergraduate students at the US Air Force Academy	Effect of roommate background (SAT verbal score) on one's own GPA	Randomly assigned roommates	0.072	
Peer effects in other educational outcomes					
Lavy & Schlosser (2011)	High school students in Israel	Effect of the change in the fraction of lower-ability students on the probability of obtaining matriculation diploma		0.015 to 0.036	
Jackson (2013)	High school students in Trinidad	Effect of changing peers' incoming test scores on the number of tests passed	Cohort variation	0.03 to 0.09	
Peer effects in social or career outcomes					
Dahl et al. (2012)	Norwegian fathers	Effect of peer take-up of paternity leave on one's own likelihood	Peers' timing of child's birth		0.153 (0.054)
Kuhn et al. (2011)	Dutch neighbors	Effect of neighbor winning lottery on the probability of buying a new car	Postcode Lottery		0.068 (0.031)
De Giorgi et al. (2010)	Undergraduate students at Bocconi University	Effect of peer choosing to major in economics on one's own choice to major in economics	Randomly assigned course sections		0.04
Duncan et al. (2005)	Undergraduate students	Effect of roommate's high school binge drinking on one's own college binge drinking	Randomly assigned roommates	0.19; roommate drinking in high school leads to two more episodes	

(Continued)

Table 1 (Continued)

Paper	Sample	What is estimated	Identification	Standardized effect of increase in peers' background	Effect of peers' outcome on one's own outcome
Marmaros & Sacerdote (2002)	Undergraduate students at Dartmouth College	Hallmates' effect on one's own likelihood of taking a high-paying job	Randomly assigned hallmates/dormmates		0.236 (0.091)
Sacerdote (2001)	Undergraduate students at Dartmouth College	Effect of roommate joining a fraternity/sorority on one's own likelihood	Randomly assigned roommates		0.078 (0.038)

^aStandard errors are in parentheses. If the authors did not report the standardized unit effect of a background characteristic, I converted to standardized units by multiplying the key peer coefficient by an individual-level standard deviation of the right-hand-side variable and dividing by the individual-level standard deviation of the left-hand-side variable. In such cases, rather than attempting to convert standard errors, I note whether the reported peer coefficient was not statistically significant. Alternatively, for studies for which I report a range, these are ranges given by authors in summarizing their own results.

In this highly controlled setting, Carrell et al. (2009) find evidence for peer effects at the squadron level. A one-standard deviation increase in squadron average peer SAT scores raises one's own first-year grade point average (GPA) by 0.05. Zimmerman (2003) finds somewhat similar results for randomly assigned freshman roommates at Williams College. Having a roommate with one-standard deviation higher SAT verbal scores raises one's own GPA by 0.02 to 0.03 points. In Sacerdote (2001), I do not find statistically significant effects in this linear-in-means model. However, when using a more flexible model, I find that all student types benefit significantly from having a roommate in the top 25% of incoming scores, with a benefit of approximately 0.06 GPA points. In the most-flexible model in which I allow peer effects to vary by roommate type and one's own type, higher-ability roommates provide the most benefit to higher-ability students. Lower-ability roommates cause the most harm to lower-ability students.

Foster (2006) does not find evidence of roommate effects at the University of Maryland. Siegfried & Gleason (2006) examine randomly assigned first-year roommates at Vanderbilt University. They find no benefit to one's own GPA stemming from roommates' SAT scores or high school GPA, but there is a positive effect from having a roommate who completed five or more Advanced Placement courses. Additionally, Siegfried & Gleason find some evidence for non-linearities; higher-ability roommates provide the largest benefit to other higher-ability students. Hoel et al. (2006) find that having students with higher incoming scores in one's own college classroom does not raise one's academic performance.

Stinebrickner & Stinebrickner (2005) find that a student's GPA suffers when (randomly assigned) roommates bring a video game to school. Stinebrickner & Stinebrickner (2006) find that students benefit from being assigned roommates who had higher high school GPAs. The authors hypothesize that high school GPA is a better proxy for study effort than are ACT or SAT scores.

The most interesting peer effects may go beyond test scores and grades. Carrell et al. (2011) find that squadron mates' incoming physical fitness scores have a large influence on one's own physical fitness scores at the US Air Force Academy. Having squadron mates who average a person-level standard deviation higher on the incoming score raises one's own physical fitness by 0.165 standard deviations. Moreover, Carrell et al. (2008) find that adding to a group one additional peer who cheated on high school exams creates an additional 0.33 to 0.47 cheaters in college.

In Sacerdote (2001), I find that an additional 10% of one's dormmates joining a Greek organization is associated with a 3% increase in one's likelihood of joining. Marmaros & Sacerdote (2002) regress one's likelihood of taking a high-paying job on the average outcome of one's randomly assigned freshman hallmates and find a statistically significant coefficient of 0.24. In contrast, Arcidiacono & Nicholson (2005) find no evidence of peer effects on specialty choice within a medical school class.

Duncan et al. (2005) examine how my predisposition to drink alcohol interacts with that of my randomly assigned roommates and thereby affects my binge drinking in college. For males, if I binge drank in high school, the assignment of a roommate who also did so will raise my number of (college) binge-drinking episodes per month by approximately four episodes. This essentially doubles my number of binge-drinking episodes. The peer effect exists only for men (not women) who have a predisposition to drink. These authors demonstrate a very large, but also highly plausible peer effect in the health of college students.

Boisjoly et al. (2006) also show that randomly assigned roommates have a large influence on attitudes and perceptions. Being randomly assigned a black roommate increases white students' support for affirmative action policies, even one to three years after their freshman year. Being assigned a minority roommate raises a white student's number of interactions with minority students. Finally, white students with black roommates during their freshman year report being more comfortable interacting with people outside their own racial group.

Despite the lack of conformity in these findings from roommate studies, some important tentative conclusions can be reached. First, linear-in-means models for academic outcomes find modest or, in a few cases, zero peer effects from college roommates and dormmates. Second, making the model more flexible (and nonlinear) often increases the size and statistical significance of estimated peer effects. Carrell et al. (2013) make this point for the US Air Force Academy data. This finding is also present in the Dartmouth roommates data and elsewhere.

Third, there are certain nonacademic outcomes that display much larger peer effects. Binge drinking, participation in the Greek system, physical fitness, and attitudes toward minority students all display large peer effects. Peers may be a major determinant of these outcomes, at least in the short run.

This dichotomy between peer effects in undergraduate GPA and peer effects in more social outcomes is evident in **Table 1**. Of the five studies listed that look at GPA, three find small and not statistically significant results in a linear-in-means model. In fact, there are several additional studies cited above (e.g., Foster 2006) that find no peer effects when a linear-in-means model is imposed. In contrast, as shown in the lower half of **Table 1**, the effects on roommate drinking, occupational choice, or membership in a Greek organization appear to be strong and statistically significant.

Recently, a series of fascinating studies has emerged that look at long-run peer effects among randomly assigned business school section mates. Shue (2013) finds that Harvard Business School (HBS) section mates who become CEOs or CFOs influence each other's compensation levels and propensity to make an acquisition. CEOs are 20 percentage points more likely to make an acquisition if all their sections mates do so as well. The effects are particularly strong following a reunion year for the relevant HBS class.

In a similar vein, Lerner & Malmendier (2013) look for peer effects in entrepreneurship among HBS section mates. Interestingly, they find that a peer's prior entrepreneurial experience can reduce unsuccessful entrepreneurship among classmates. Finally, Ahern et al. (2012) find negative peer effects in altruism among business school classmates and positive peer effects in risk-taking behavior.

3.5. Explicit Experiments

Of course, not all well-identified peer effects studies stem from the use of natural experiments. The search for knowledge on peer effects has led to a series of controlled experiments. In many ways, the Moffitt (2001) partial population approach is ideal. Babcock & Hartman (2010) implement this setup in the context of gym workouts. Rather than randomly assigning peers, these authors give their random intervention (encouragement to exercise) to differing fractions of peers.

This is an intriguing approach. Opportunities to randomly assign peers are rare in a real-world setting, whereas opportunities to randomly implement a treatment are much more common. In addition, the partial population approach can take advantage of an existing, potentially quite robust peer network. In contrast, randomly assigning roommates or squadron mates will not guarantee that the imposed peer groups will interact.

In fact, Carrell et al. (2013) suggest that artificially manipulated peer groups may fall apart or fail to interact in the ways hoped for by policy makers or predicted by observational data. They attempt to improve academic performance for the lowest third of incoming cadets at the US Air Force Academy. The randomized pre-experimental data suggest that lower-ability students benefit strongly from being placed with squadron mates who have high SAT verbal scores.

Carrell et al.'s (2013) optimized squadrons consist of groups comprising 50% lower-ability cadets and 50% cadets with high SAT verbal scores, leaving the middle-ability cadets separated

into homogeneous squadrons by themselves. Performance for the lower-ability students actually fell rather than increasing as predicted (relative to lower-ability students in the randomly assigned control group). Our hypothesis to explain this perverse finding is that in the new bifurcated treatment squadrons, the lower-ability students largely interacted with each other rather than with the higher-ability students.

Carrell et al.'s (2013) survey data on study partners and friends confirm this hypothesis. Lower-ability students in the treatment group were much more likely to choose each other as study partners rather than choose a higher-ability study partner. Not only did we have a mechanical impact from making more lower-ability students available (in the treatment squadrons), we also observed higher same-ability attraction of study partners, above and beyond the effect of availability.

The lesson drawn from the experiment is that endogenous peer choice can make the manipulation of peer effects (by a social planner) a challenging job, and such manipulation raises the possibility for serious and unintended consequences.

Duflo & Saez (2003) show that there are statistically significant peer effects in the decision to enroll in a retirement plan. A set of employees was randomly encouraged to attend a benefits fair to learn about the 403(b) supplemental retirement plan. Eleven months later, the treated employees were 1.4 percentage points more likely to be enrolled in the plan, relative to employees in departments with no treated individuals. Interestingly, untreated employees who had treated peers showed the same increase in enrollment. Moreover, the untreated peers in treated departments were much more likely to attend the benefits fair than employees in untreated departments. A natural interpretation of these findings is that information about the fair and the retirement plan is flowing across peers within a department.

Beshears et al. (2011) add further nuance to this literature. They run an experiment in which employees are given different sets of information about their peers' participation in 401(k) plans. For some employees, they find discouragement effects, in which learning that a high fraction of peers participating in the plan can reduce one's own levels of participation. This is a somewhat unexpected finding and again suggests that interventions involving peers do not always have the intended effect.

Falk & Ichino (2006) run an experiment to measure peer effects in worker productivity. Their setting involves paying high school students to stuff envelopes. One group of students worked in isolation, whereas students in the other group were paired. The authors find strong positive peer effects in productivity in the paired treatment, and average output is higher when peers are present.

The assigned task in Falk & Ichino's (2006) experiment is realistic, although it is a short-run job that does not require much human capital. At the other end of the spectrum, Waldinger (2012) looks at the productivity of German scientists before and after many of their colleagues are dismissed by the Nazis. Whereas there is no effect on productivity at the department level, losing a coauthor reduces one's own publications and citations by 13–18%.

4. DISCUSSION AND CONCLUSIONS

What broad conclusions can we draw from this burgeoning literature, and where might the literature be headed? A good place to start is with peer effects in elementary and secondary schools as there are a large number of studies there and the policy implications may be particularly important. Approximately half the analyses find no statistically significant peer effects from classmates' background ability. Angrist & Lang's (2004) main specification is a good example of this. Some studies find modest-sized linear-in-means peer effects at the classroom or school*grade level. For example, in some specifications, Imberman et al. (2012) use a linear-in-means model and instrument

for peer average ability with the percent of Katrina evacuees. They find that peer average ability (in standard deviation units) affects one's own test score with a coefficient of 0.33.

Relying on gender variation across cohorts (within a school), Hoxby (2000a) finds that a 1 point increase in peer average reading score leads to a 0.3 increase in one's own reading score. From this, she concludes both that peer effects may be quite substantial and that it is likely that gender mix may work through more channels than simply peer average test scores.

The bottom line with the linear-in-means models within elementary and secondary education is that approximately half the studies find either modest or large effects on test scores. Half the studies do not find evidence of peer effects in test scores.

Perhaps the more interesting finding (and my second bottom line) is that the estimated peer effects become larger and more statistically significant when the assumptions of the linear-in-means model are relaxed. Imberman et al. (2012) see no peer effects on average from the arrival of Katrina evacuees into Louisiana schools. But robust peer effects are evident once we allow the effect to vary by own position in the test-score distribution and by the ability of the incoming evacuees. Hoxby & Weingarth (2005) have similar findings. Students tend to perform better as more students of their type are added to the classroom.

A third broad conclusion is that social and non-test-score outcomes show very large peer effects. Perhaps not surprisingly, the amount male college students binge drink is greatly influenced by their peers (Duncan et al. 2005). Studies from the US Air Force Academy (and West Point and Annapolis) show large peer effects in student cheating and in physical fitness scores.

Lerner & Malmendier (2013) find that HBS peers with entrepreneurial experience reduce one's own amount of unsuccessful entrepreneurship. Shue (2013) finds that randomly assigned HBS section mates who become CEOs and CFOs respond to each other's compensation levels and tendency to make acquisitions.

How much are the career choices of college students influenced by their peers? This is an open question. Marmaros & Sacerdote (2002) find a strong correlation between one's own initial career choice and that of randomly assigned freshmen floormates and dormmates. Using data at Bocconi University, De Giorgi et al. (2010) find that if a strongly connected peer chooses economics as a major, one is 4% more likely to choose economics as well. This sometimes moves a student away from a previously stated preferred major in a way that reduces his or her grades and expected earnings. Somewhat contrary to these two papers, Arcidiacono & Nicholson (2005) do not find that medical school students are influenced by their peers' choices of specialty.

A fourth lesson from the literature is that we do not yet know enough about the nature of peer effects to engage in social engineering of peer groups to affect students' outcomes in a desired direction. Carrell et al. (2013) arrange squadrons at the US Air Force Academy in an effort to boost the academic performance of the students with the lowest incoming ability. The experiment had the opposite effect on lower-ability students, possibly because of the way in which endogenous peer choice responded to the newly arranged squadrons.

Marmaros & Sacerdote (2006) (who use email traffic to measure interactions) perform some simulations of how much across-race interaction would occur if nonwhite students were evenly spread across dorms or if the percentage of nonwhite students in a graduating cohort were increased. However, I now take those calculations with a pound of salt given the large and unexpected endogenous policy responses seen in Carrell et al. (2013).

Social scientists may be on safer ground when implementing policies such as those in Babcock & Hartman (2010) or Duflo & Saez (2003), which simply intervene with part of a peer network (hoping to influence untreated peers) but do not attempt to actively manipulate who is in the network or how much they interact.

Overall there is little doubt that peers are a powerful influence on human behavior. Despite the serious identification problems, we have made solid progress toward understanding the domains in which peers matter most and approximately how big the effects are.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I thank Scott Carrell, Caroline Hoxby, Doug Staiger, and many others for the positive peer effects that led to this article. All remaining omissions and errors are a unique product of the author. I thank Mahnum Shahzad for outstanding research assistance. The National Science Foundation and the Institute of Education Sciences, US Department of Education, provided generous funding.

LITERATURE CITED

- Abdulkadiroglu A, Angrist JD, Pathak PA. 2014. The elite illusion: achievement effects at Boston and New York exam schools. *Econometrica* 82:137–96
- Ahern KR, Duchin R, Shumway T. 2012. *Peer effects in economic attitudes*. Work. Pap., Univ. Mich., Ann Arbor
- Anderson M, Lu F. 2012. *Peer effects in microenvironments: the benefits of homogeneous classroom groups*. Work. Pap., Univ. Calif., Berkeley
- Angrist JD, Lang K. 2004. Does school integration generate peer effects? Evidence from Boston's METCO program. *Am. Econ. Rev.* 94:1613–34
- Arcidiacono P, Nicholson S. 2005. Peer effects in medical school. *J. Public Econ.* 89:327–50
- Austen-Smith D, Fryer RG Jr. 2005. An economic analysis of “acting white.” *Q. J. Econ.* 120:551–83
- Babcock PS, Hartman JL. 2010. *Networks and workouts: treatment size and status specific peer effects in a randomized field experiment*. NBER Work. Pap. 16581
- Beshears J, Choi JJ, Laibson D, Madrian BC, Milkman KL. 2011. *The effect of providing peer information on retirement savings decisions*. NBER Work. Pap. 17345
- Billings SB, Deming DJ, Rockoff J. 2014. School segregation, educational attainment and crime: evidence from the end of busing in Charlotte-Mecklenburg. *Q. J. Econ.* 129:435–76
- Boisjoly J, Duncan GJ, Kremer M, Levy DM, Eccles J. 2006. Empathy or antipathy? The impact of diversity. *Am. Econ. Rev.* 96:1890–905
- Bramoullé Y, Djebbari H, Fortin B. 2009. Identification of peer effects through social networks. *J. Econom.* 150:41–55
- Carrell SE, Fullerton RL, West JE. 2009. Does your cohort matter? Measuring peer effects in college achievement. *J. Labor Econ.* 27:439–64
- Carrell SE, Hoekstra M, West JE. 2011. Is poor fitness contagious? Evidence from randomly assigned friends. *J. Public Econ.* 23:657–63
- Carrell SE, Malmstrom FV, West JE. 2008. Peer effects in academic cheating. *J. Hum. Resour.* 43:173–207
- Carrell SE, Sacerdote BI, West JE. 2013. From natural variation to optimal policy? The importance of endogenous peer group formation. *Econometrica* 81:855–82
- Case AC, Katz LF. 1991. *The company you keep: the effects of family and neighborhood on disadvantaged youths*. NBER Work. Pap. 3705
- Christakis NA, Fowler JH. 2007. The spread of obesity in a large social network over 32 years. *N. Engl. J. Med.* 357:370–79
- Cohen-Cole E, Fletcher JM. 2008. Is obesity contagious? Social networks vs. environmental factors in the obesity epidemic. *J. Health Econ.* 27:1382–87

- Coleman JS, Campbell E, Hobson C, McPartland J, Mood A, et al. 1966. *Equality of educational opportunity*. Rep., US Dep. Health Educ. Welf., Off. Ed., Washington, DC
- Dahl GB, Løken KV, Mogstad M. 2012. *Peer effects in program participation*. NBER Work. Pap. 18198
- Damm AP, Dustmann C. 2012. *Does growing up in a high crime neighborhood affect youth criminal behavior?* Work. Pap., Univ. Coll. London
- De Giorgi G, Pellizzari M, Redaelli S. 2010. Identification of social interactions through partially overlapping peer groups. *Am. Econ. J. Appl. Econ.* 2:241–75
- Duflo E, Saez E. 2003. The role of information and social interactions in retirement plan decisions: evidence from a randomized experiment. *Q. J. Econ.* 118:815–42
- Duncan GJ, Boisjoly J, Kremer M, Levy DM, Eccles J. 2005. Peer effects in drug use and sex among college students. *J. Abnorm. Child Psychol.* 33:375–85
- Falk A, Ichino A. 2006. Clean evidence on peer effects. *J. Labor Econ.* 24:39–57
- Foster G. 2006. It's not your peers, and it's not your friends: some progress toward understanding the educational peer effect mechanism. *J. Public Econ.* 90:1455–75
- Glaeser EL, Sacerdote B, Scheinkman JA. 1996. Crime and social interactions. *Q. J. Econ.* 111:507–48
- Gould ED, Lavy V, Passerman MD. 2009. Does immigration affect the long-term educational outcomes of natives? Quasi-experimental evidence. *Econ. J.* 119:1243–69
- Graham BS. 2008. Identifying social interactions through conditional variance restrictions. *Econometrica* 76:643–60
- Harris JR. 2009. *The Nurture Assumption: Why Children Turn Out the Way They Do*. New York: Free
- Hoel J, Parker J, Rivenburg J. 2006. *A test for classmate peer effects in higher education*. Work. Pap., Reed Coll., Portland
- Hoxby C. 2000a. *Peer effects in the classroom: learning from gender and race variation*. NBER Work. Pap. 7867
- Hoxby CM. 2000b. The effects of class size on student achievement: new evidence from population variation. *Q. J. Econ.* 115:1239–85
- Hoxby CM, Weingarth G. 2005. *Taking race out of the equation: school reassignment and the structure of peer effects*. Work. Pap., Harvard Univ., Cambridge, MA
- Imberman SA, Kugler AD, Sacerdote BI. 2012. Katrina's children: evidence on the structure of peer effects from hurricane evacuees. *Am. Econ. Rev.* 102:2048–82
- Jackson CK. 2013. Can higher-achieving peers explain the benefits to attending selective schools? Evidence from Trinidad and Tobago. *J. Public Econ.* 108:63–77
- Kling JR, Ludwig J, Katz LF. 2005. Neighborhood effects on crime for female and male youth: evidence from a randomized housing voucher experiment. *Q. J. Econ.* 120:87–130
- Kuhn P, Kooreman P, Soetevent A, Kapteyn A. 2011. The effects of lottery prizes on winners and their neighbors: evidence from the Dutch Postcode Lottery. *Am. Econ. Rev.* 101:2226–47
- Lavy V, Schlosser A. 2011. Mechanisms and impacts of gender peer effects at school. *Am. Econ. J. Appl. Econ.* 3:1–33
- Lerner J, Malmendier U. 2013. With a little help from my (random) friends: success and failure in post-business school entrepreneurship. *Rev. Financ. Stud.* 26:2411–52
- Ludwig J, Duncan GJ, Hirschfield P. 2001. Urban poverty and juvenile crime: evidence from a randomized housing-mobility experiment. *Q. J. Econ.* 116:655–79
- Lyle DS. 2009. The effects of peer group heterogeneity on the production of human capital at West Point. *Am. Econ. J. Appl. Econ.* 1:69–84
- Manski CF. 1993. Identification of endogenous social effects: the reflection problem. *Rev. Econ. Stud.* 60:531–42
- Marmaros D, Sacerdote B. 2002. Peer and social networks in job search. *Eur. Econ. Rev.* 46:870–79
- Marmaros D, Sacerdote B. 2006. How do friendships form? *Q. J. Econ.* 121:79–119
- Moffitt RA. 2001. Policy interventions, low-level equilibria, and social interactions. In *Social Dynamics*, ed. SN Durlauf, HP Young, pp. 45–82. Washington, DC: Brookings Inst.
- Sacerdote B. 2000. *Peer effects with random assignment: results for Dartmouth roommates*. NBER Work. Pap. 7469

- Sacerdote B. 2001. Peer effects with random assignment: results for Dartmouth roommates. *Q. J. Econ.* 116:681–704
- Shue K. 2013. Executive networks and firm policies: evidence from the random assignment of MBA peers. *Rev. Financ. Stud.* 26:1401–42
- Siegfried JJ, Gleason MA. 2006. *Academic roommate peer effects*. Unpublished manuscript, Vanderbilt Univ., Nashville
- Stinebrickner R, Stinebrickner TR 2005. How much does studying matter? *Proc. Fed. Reserve Bank Clevel.*, pp. 55–59. Cleveland, OH: Fed. Reserve Bank Clevel.
- Stinebrickner R, Stinebrickner TR. 2006. What can be learned about peer effects using college roommates? Evidence from new survey data and students from disadvantaged backgrounds. *J. Public Econ.* 90:1435–54
- Summers AA, Wolfe BL. 1977. Do schools make a difference? *Am. Econ. Rev.* 67:639–52
- Waldinger F. 2012. Peer effects in science: evidence from the dismissal of scientists in Nazi Germany. *Rev. Econ. Stud.* 79:838–61
- Zimmerman DJ. 2003. Peer effects in academic outcomes: evidence from a natural experiment. *Rev. Econ. Stat.* 85:9–23