Labour Economics 30 (2014) 98–108

Contents lists available at ScienceDirect

Labour Economics

journal homepage: www.elsevier.com/locate/labeco

The perils of peer effects☆

Joshua D. Angrist *

MIT Department of Economics, 77 Massachusetts Avenue, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Received 8 January 2014
Received in revised form 22 May 2014
Accepted 22 May 2014
Available online 5 June 2014

JEL classification:
C18
C31
C36
I21
31

Keywords:
Causality
Social returns
Instrumental variables

ABSTRACT

Individual outcomes are highly correlated with group average outcomes, a fact often interpreted as a causal peer effect. Without covariates, however, outcome-on-outcome peer effects are vacuous, either unity or, if the average is defined as a leave-out mean, determined by a generic intraclass correlation coefficient. When pre-determined peer characteristics are introduced as covariates in a model linking individual outcomes with group averages, the question of whether peer effects or social spillovers exist is econometrically identical to that of whether a 2SLS estimator using group dummies to instrument individual characteristics differs from OLS estimates of the effect of these characteristics. The interpretation of results from models that rely solely on chance variation in peer groups is therefore complicated by bias from weak instruments. With systematic variation in group composition, the weak IV issue falls away, but the resulting 2SLS estimates can be expected to exceed the corresponding OLS estimates as a result of measurement error and for other reasons unrelated to social effects. Research designs that manipulate peer characteristics in a manner unrelated to individual characteristics provide the most compelling evidence on the nature of social spillovers. As an empirical matter, designs of this sort have mostly uncovered little in the way of socially significant causal effects.

1. Introduction

In a regression rite of passage, social scientists around the world link students' achievement to the average ability of their schoolmates. A typical regression in this context puts individual test scores on the left side, with some measure of peer achievement on the right. These regressions reveal a strong association between the performance of students and their peers, a fact documented in Sacerdote's (2011) recent survey of education peer effects. Peer effects are not limited to education and schools; evidence abounds for associations between citizens and neighbors in every domain, including health, body weight, work, and consumption, to name a few (A volume edited by Durlauf and Young (2001) points to some of the literature.). Most people have a powerful intuition that "peers matter," so behavioral interpretations of the positive association between the achievement of students and their classmates or the labor force status of citizens and their neighbors ring true.

I argue here that although correlation among peers is a reliable descriptive fact, the scope for incorrect or misleading attributions of causality in peer analysis is extraordinarily wide. Many others have made this point (see, especially, Deaton, 1990; Manski, 1993; Boozer and Cacciola, 2001; Moffitt, 2001; Hanushek et al., 2003). Nevertheless, I believe there's value in a restatement and synthesis of the many perils of econometrically estimated peer effects. Because both peer analysis and instrumental variables (IV) estimates involve statistical correlations between group means, I find it especially useful to link econometric models of peer effects with the behavior of IV estimators.

The link with IV shows that models which assign a role to group averages in the prediction of individual outcomes should often be expected to produce findings that look like a peer effect, even in a world where behavioral influences between peers are absent. The vacuous nature of many econometric peer effects is not an identification problem; the parameters of the models I discuss are identified. More often than not, however, these parameters reveal little about human behavior or what we should expect from policy-induced changes in group composition. If the group average in question involves the dependent variable, the estimated peer effect is a mechanical phenomenon, either affirming an identify in the algebra of expectations or providing a measure of group clustering devoid of behavioral content. If the model in question includes individual covariates, putative peer effects are a test for the

☆ Presented as the Adam Smith Lecture at the European Association of Labor Economists annual meeting, September 2013, in Torino. This research was partially funded by the Institute for Education Sciences. Gaston Illanes and Gabriel Kreindler provided expert research assistance. I thank seminar participants at EALE, Maryland, Warwick, and Queens for helpful comments. Special thanks go to Bruce Sacerdote, who patiently walked me through his earlier analyses and graciously supplied new results, and to Steve Pischke, for extensive discussions and feedback repeatedly along the way. Thanks also go to an anonymous referee, the editor, and many of my peers for helpful discussions and comments, especially Daron Acemoglu, Bryan Graham, Andrea Ichino, Guido Imbens, Patrick Kline, Guido Kuersteiner, Steven Lehrer, Robert Moffitt, Victor Lavy, Parag Pathak, and Rob Townsend. The effects of their interventions were modest, but that's entirely my fault.
* Tel.: +1 857 225 0879.
E-mail address: angrist@mit.edu.

http://dx.doi.org/10.1016/j.labeco.2014.05.008
0927-5371/

equality of two-stage least squares (2SLS) and OLS estimates of the effect of these covariates on outcomes. There are many reasons why 2SLS estimates might differ from the corresponding OLS estimates. While peer effects are on the list of causes behind such divergence, they should not usually be at the top of it.

2. Peer theory

Like many in my cohort, I smoked marijuana repeatedly throughout the day in high school. Most of my friends smoked a lot of dope too. Ten years later, my youngest brother went to the same high school, but he didn't smoke nearly as much dope as my friends and I did, something that worried me at the time. My brother's friends also smoked little. In fact, by the time my brother went to our high school, nobody smoked as much dope as we did in 1975. That must be why my brother smoked so much less than me.

This youthful story bears econometric investigation. Let $\bar{s}_j$ be the smoke-alotta-dope rate among students attending high school j, that is, the school average of s_{ij}, a dummy for whether student i at school j smokes. Is there a school-level dope-smoking peer effect? It's tempting to explore the peer effects hypothesis by estimating this regression:

$$s_{ij} = \alpha + \beta \bar{s}_j + \xi_{ij}, \tag{1}$$

a model that seems to quantify the essence of my story.

Estimation of Eq. (1) is superfluous, of course. Any regression of s_{ij} on $\bar{s}_j$ produces a coefficient of unity:

$$\frac{\sum_j \sum_i s_{ij}\left(\bar{s}_j - \bar{s}\right)}{\sum_j n_j \left(\bar{s}_j - \bar{s}\right)^2} = \frac{\sum_j \left(\bar{s}_j - \bar{s}\right)\left(n_j \bar{s}_j\right)}{\sum_j n_j \left(\bar{s}_j - \bar{s}\right)^2} = 1$$

In fact, the properties of Eq. (1) emerge without algebra: The group average on the right hand side is a fitted value from a regression of the left hand side on dummies indicating groups (high schools, in this case). The covariance between any variable and a corresponding set of regression fitted values for this variable is equal to the variance of the fits, producing the result that covariance over variance equals one.

The tautological nature of the relationship between individual data and group averages is not a story about samples. Let β denote the population regression coefficient from a regression of (mean zero) y on $\mu_{y|z} = E[y|z]$, for any random variables, y and z. The scenario I have in mind is that z indexes peer-referent groups (like high schools). For any z, we can be sure that

$$\beta \equiv \frac{E\left[y\mu_{y|z}\right]}{V\left[\mu_{y|z}\right]} = 1, \tag{2}$$

a relation that follows by iterating expectations:

$$\begin{aligned} E\left[y\mu_{y|z}\right] &= E\left(E\left[y|z, \mu_{y|z}\right] \times \mu_{y|z}\right) = E\left(E[y|z] \times \mu_{y|z}\right) \\ &= E\left[\mu_{y|z}^2\right] = V\left[\mu_{y|z}\right]. \end{aligned}$$

Others have commented on the vacuous nature of regressions of individual outcomes on group mean outcomes. Manski (1993) described the problem this way: "… observed behavior is *always* consistent with the hypothesis that individual behavior reflects mean reference-group behavior" (italics mine). Manski's extended discussion, however, suggests that the tautological nature of Eq. (2) is a kind of troubling special case, one that can in principle be avoided given sufficient ex ante information on how individuals choose their peer reference groups. In the same spirit, Brock and Durlauf (2001) and Jackson (2010), among others, describe regressions like Eq. (2) as posing an identification problem, one for which we might, with suitable econometric ingenuity, find a solution. Yet, the coefficient in my simple regression of individual outcomes on high school mean outcomes is identified in a technical sense, by which I mean, *Stata* (or even *SAS*) should have no trouble finding it.

Econometric models of endogenous peer effects are typically more elaborate than the one I've used to describe the Angrist brothers' smoking habits. Discussing peer effects in the Tennessee STAR class size experiment, Boozer and Cacciola (2001, p.46) observed: "Of course, since the setup just discussed delivers a coefficient of exactly 1, it is improbable a researcher would not realize his error, and opt for a different estimation strategy." Elaboration, however, need not produce a coherent causal framework. In a more recent analysis of the STAR data, for example, Graham (2008) models achievement in STAR classrooms as satisfying this equation:

$$y_{ci} = \alpha_c + (\gamma - 1)\bar{\varepsilon}_c + \varepsilon_{ci}, \tag{3}$$

where α_c is a class or teacher effect and $\gamma > 1$ captures social interactions. The residual ε_{ci} is a kind of placeholder for individual heterogeneity, but not otherwise specified.

As in many discussions of peer effects, Graham (2008)'s narrative imbues Eq. (3) with a causal interpretation: "Consider the effect of replacing a low-ε with high-ε … mean achievement increases for purely compositional reasons and … because … a high-ε raises peer quality" (p. 646). Graham (2008)'s subsequent discussion introduces covariates that might be causally linked with changes in α_c. On it's own, however, Eq. (3) is a weak foundation for causal inference. I can fit this model perfectly as follows: set α_c equal to the group average, $\bar{y}_c$, and $\varepsilon_{ci} = y_{ci} - \bar{y}_c$. Since $\bar{\varepsilon}_c = 0$ in this specification, any γ will do. My proposal, which identifies α_c with the only conditional mean function that can be constructed given information on individuals and groups and nothing else, satisfies Eq. (3) under any sample design or data generating process, including those with random assignment to groups and groups of differing or even infinite size. Eq. (3) therefore seems no more useful than the tautological relation described by Eq. (2).

2.1. Control yourself

Many econometric models of peer effects build on a theoretical framework that explains behavior as a function of both individual and group characteristic. Townsend (1994), for example, hypothesized that, controlling for household demographic structure, individual household consumption responds to village average consumption in a theoretical relationship generated by risk sharing. Bertrand et al. (2000) described spillovers in welfare use that emerge as a result of ethnic networks – these are parameterized as acting through neighborhood and ethnicity group averages, controlling for individual characteristics. With individual covariates included as controls, a regression of y on group average y typically does not produce a coefficient of unity. This feature notwithstanding, I don't believe that the coefficient on group averages in a multivariate model of endogenous peer effects reveals the action of social forces.

I interpret covariate-controlled endogenous peer relationships here using a model for the population expectation of outcomes conditional on individual characteristics and peer group membership. My discussion focuses on a specification from Manski (1993), who notes that the following conditional expectation function (CEF) is typical of econometric research on peer effects:

$$E[y|x, z] = \beta\mu_{y|z} + \gamma x. \tag{4}$$