This is a preprint of an article published in [Sondheimer, Rachel Milstein, and Donald P. Green. 2010. Using Experiments to Estimate the Effects of Education on Voter Turnout. American Journal of Political Science 54(1): 174-189]. Pagination of this preprint may differ slightly from the published version.

Using Experiments to Estimate the Effects of Education on Voter Turnout

Rachel Milstein Sondheimer U.S. Military Academy, West Point
Donald P. Green Yale University

The powerful relationship between education and voter turnout is arguably the most well-documented and robust finding in American survey research. Yet the causal interpretation of this relationship remains controversial, with many authors suggesting that the apparent link between education and turnout is spurious. In contrast to previous work, which has relied on observational data to assess the effect of education on voter turnout, this article analyzes two randomized experiments and one quasi-experiment in which educational attainment was altered exogenously. We track the children in these experiments over the long term, examining their voting rates as adults. In all three studies, we find that exogenously induced changes in high school graduation rates have powerful effects on voter turnout rates. These results imply that the correlation between education and turnout is indeed causal. We discuss some of the pathways by which education may transmit its influence.

The relationship between education and voter turnout ranks among the most extensively documented correlations in American survey research. From the early work of Merriam and Gosnell (1924) to today, literally thousands of cross-sectional surveys have indicated that turnout rates climb with years of formal schooling. Using Current Population Survey data, Table 1 illustrates this powerful relationship for the general population and for African Americans, who figure prominently in the analysis below. This pattern holds regardless of whether a given survey uses self-reports or public records to measure voting (Abramson and Aldrich 1982; Katosh and Traugott 1981; Wolfinger and Rosenstone 1980), whether the election in question is local or national (Hogan 1999; Rosenstone and Hansen 1993; Verba and Nie 1972; Verba, Schlozman, and Brady 1995), and whether the relationship is gauged without control variables or after an extensive set of demographic attributes has been held constant (Converse 1972; Verba, Schloz-man, and Brady 1995; Verba et al. 1993; Wolfinger and Rosenstone 1980). Although there is some dispute about the strength of this relationship in surveys conducted outside the United States (Franklin 2004; Milligan, Moretti, and Oreopoulos 2004; Powell 1986), within the United States this correlation obtains with law-like regularity.

What causal interpretation, if any, should scholars attach to this powerful and robust correlation? Perhaps none. As Richard Brody (1978) famously observed three decades ago, the microlevel relationship between education and voter turnout seems to be in tension with macrolevel patterns. The United States, like other Western-style democracies, has experienced dramatic gains in the average educational attainment of its population. As shown in Figure 1, the proportion of adults age 25 or older who completed high school was 48.0% in 1964, as compared to 73.3% in 1984 and 85.2% in 2004. Yet, the proportion of the voting-eligible population (McDonald and Popkin 2001) who participated in the 1964, 1984, and 2004 elections was 62.8%, 55.3%, and 60.3%, respectively. In contrast to other education-driven trends

Rachel Milstein Sondheimer is Assistant Professor of Political Science, U.S. Military Academy, West Point, Department of Social Sciences, 607 Cullum Road, West Point, NY 10996 (Rachel.Sondheimer@usma.edu). Donald P. Green is A. Whitney Griswold Professor of Political Science, Yale University, Institution for Social and Policy Studies, 77 Prospect Street, P.O. Box 208209, New Haven, CT 06520-8209 (Donald.Green@yale.edu).

We thank Alan Gerber for guidance throughout this project. We also thank Charles Achilles, Jayne Boyd-Zaharias, and Jeff Lewis for help with the Project STAR analysis, Lawrence Schweinhart and the High/Scope Educational Research Foundation for sharing their Perry data, Annette Taylor, Lori Canova, and the "I Have a Dream" Foundation for help and support, and the Yale Institution for Social and Policy Studies for funding. The views expressed in this article are solely those of the authors and do not represent the views of the United States Military Academy, the Department of the Army, and/or the Department of Defense.

TABLE 1 Rates of Voter Turnout by High School Graduation, by Year

	Full Sample			African American Subsample		
	2000	2002	2004	2000	2002	2004
Non–high school graduates	43.2	31.2	45.0	50.0	34.6	54.8
High school graduates	70.3	54.7	75.5	69.5	53.1	76.1
N	52,918	51,617	63,052	1,180	953	837

Source: Current Population Survey, November Supplements. Weighted using final weights from Bureau of Labor Statistics.

FIGURE 1 Aggregate Levels of Educational Attainment and Voter Turnout, 1962–2006

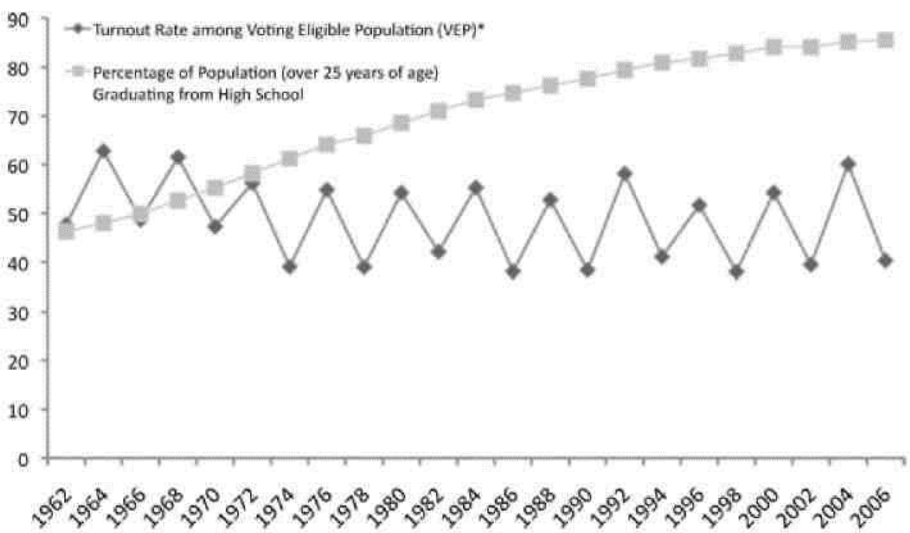

*VEP Rate is calculated using the number of people in any given election cycle in following arder.
Turnout rate using the VEP from 1962 to 1978 drawn from McDonald and Popkin (2001).
Turnout rate using the VEP from 1980 to 2008 drawn from McDonald's web site: http://elections.gmu.edu.
Educational attainment data from the Current Population Survey obtained from http://www.census.gov.

anticipated by survey researchers, such as diminished support for explicitly racially discriminatory policies (Hyman and Sheatsley 1956, 1964; Schuman, Steeh, and Bobo 1985), education-driven increases in voter turnout never materialized.

Since Brody's observation, scholars have often expressed skepticism about whether, at the microlevel, educational attainment causes voter turnout. These challenges may be grouped into two broad categories, the first of which makes the claim that the relationship observed in survey data is spurious. The difficulty of disentangling the effects of schooling from the confounding effects of family background and innate cognitive ability is a widespread concern among those who study the effects of education (Card 1999; Rosenzweig and Wolpin 2000). Even scholars such as Wolfinger and Rosenstone, who argue for a causal interpretation of this correlation, nonetheless concede that the education effects they find may be partly spurious:

> Level of education indicates not only the skills and duties learned in school but characteristics of the individual that are unrelated to school . . . years of schooling reflect family

> background more than any other demographic characteristic does. People who have gone to college are more likely to have educated and/or affluent parents. As a result, they are more likely to come from homes where books, newspapers, and magazines were read and where politics was discussed. By virtue of this socialization, those who have been to college have grown up exposed to politics and experienced in dealing with information about it. (1980, 18–20)

These reservations are bolstered by multivariate analyses indicating that the apparent effect of education diminishes after one controls for various measures of intelligence (Luskin 1990; Neuman 1986; Nie, Junn, and Stehlik-Barry 1996; Verba, Schlozman, and Brady 1995). Herrnstein and Murray go so far as to argue that "education predicts political involvement in America because it is primarily a proxy for cognitive ability" (1994, 253).

The question of whether the education-turnout relationship is spurious has set in motion a series of recent statistical investigations. Milligan, Moretti, and Oreopoulos (2004) examine turnout rates in American National Election Study data using over-time and across-state variation in compulsory schooling laws as an instrumental variable. Dee (2004) analyzes the 1980–92 High School and Beyond panel study, using local availability of two-year colleges as an instrumental variable, and the General Social Survey from 1972 to 2000, using the child labor laws to which people were exposed when they were 16 years old as an instrumental variable. Tenn (2007) analyzes 18–24-year-old respondents to the Current Population Survey panel study, tracking whether year-to-year changes in education coincide with increases in voting. Kam and Palmer (2008) use propensity score matching to gauge the effects of college attendance on an index of eight "participatory acts" measured in the Youth-Parent Socialization panel study. These studies, however, have generated conflicting results. Milligan, Moretti, and Oreopoulos (2004) and Dee (2004) conclude that educational attainment strongly influences voter turnout, but the more recent studies by Tenn (2007) and Kam and Palmer (2008) conclude that schooling has no effect.

The second challenge holds that trends in education are not what they appear. For example, Delli Carpini and Keeter (1996) argue that increases in average years of schooling over time do not translate into increased political knowledge. Years of schooling have increased on average, but the education that students receive imparts less awareness of and concern with politics. The implication is that educational attainment does not translate into increased political participation if these causal pathways do not increase as well. Another argument that calls into question the premise of rising educational attainment is the claim that education functions merely as a sorting device, differentiating high- and low-status individuals (Nie, Junn, and Stehlik-Barry 1996). If education functions as a marker of one's relative status, it is the status-associated costs and benefits of political participation that encourage those at the upper end of the distribution to participate and discourage those at the lower end. Therefore, gains in the average level of education would fail to lift turnout rates. Again, the implication is that education's effects are either spurious or contingent on other factors.

Given the centrality of education in both behavioral research and political theories of democratic citizenship, a great deal hinges on whether educational attainment exerts a causal influence on political participation. If no causal relationship exists, educational institutions may be charged with failing to "develop democratic habits" (Levinson 1999), a core objective in the minds of authors from Dewey (1916) to Gutmann (1987). On the other hand, if educational attainment should be found to promote political participation, behavioral researchers and theorists will be directed to ask more refined questions about the causal pathways through which education influences participation.

The only way to break this impasse is to marshal new evidence regarding the causal role of schooling. Unlike previous studies, which use observational data and therefore rely on strong assumptions about the exogeneity of educational attainment, this essay presents the results of two randomized experiments and one quasi-experiment in which the educational attainment of treatment and control groups was exogenously manipulated. Like most experiments, the three studies involve idiosyncratic settings and populations, but all share a common feature: the random or near-random interventions led to increased high school graduation rates among children assigned to the treatment groups. If the education-as-cause hypothesis is correct, the children assigned to the treatment group should, years later, vote at higher rates than their counterparts in the control group. This article reports the results of years of detective work tracking down the subjects in these studies, during which care was taken to use equivalent methods to locate members of the treatment and control groups. The three studies jointly indicate that exogenous increases in schooling induce substantially higher rates of voter participation.

The presentation of our research is structured as follows. We begin by laying out a Neyman-Rubin model of the local average treatment effect of education on voter turnout and indicating what assumptions must be invoked in order to estimate the parameters of this model using experimental data. Next, we describe the three studies and present some initial findings comparing treatment and control group outcomes. Bivariate probit, which allows explicitly for the possibility that unobserved correlates of education may also influence the vote, is applied to the three studies. Finding that educational attainment strongly influences voter turnout, we conclude by discussing possible microlevel mechanisms accounting for the education effect.

Neyman-Rubin Causal Model

The logic underlying randomized experiments (and research designs that attempt to approximate randomized clinical trials) is often explicated in terms of a notational system that has its origins in the work of Neyman ([1923] 1990) and is sometimes termed the "Rubin Causal Model" after Rubin (1978, 1990). The advantage of this notational system lies in the fact that it illuminates the core assumptions on which causal inference depends, regardless of whether inferences are to be drawn from experimental or observational data. Using the logic of the Rubin Causal Model to explicate what Green and Gerber (2002) have termed "downstream experimentation," we show how experiments designed to increase educational attainment can reveal the effects of educational attainment on voter turnout.

Definitions. The model consists of three variables. For each individual i, let $Z_i = 1$ refer to assignment to the treatment group, and let $Z_i = 0$ refer to assignment to the control group. Let $\mathbf{Z}$ be the N-dimensional vector of random assignments with elements Z_i, and let $X_i(\mathbf{Z})$ denote whether an individual i graduates from high school given the random vector of treatment assignments $\mathbf{Z}$. If everyone assigned to the treatment group graduates and no one from the control group graduates, $X_i(\mathbf{Z})$ would equal Z_i for all i. Of course, that situation is unlikely in practice because many factors beyond treatment assignment affect educational attainment.

Just as $X_i(\mathbf{Z})$ represents the *potential* educational outcomes associated with a set of experimental assignments, $Y_i(\mathbf{Z}, \mathbf{X})$ is defined as the potential electoral participation of an individual to a vector of random assignments and educational outcomes. By imposing what is termed the "stable unit treatment value assumption" (Rubin 1978), which holds that the potential outcomes for a given person are unrelated to the treatment status of other individuals, we may write $X_i(\mathbf{Z})$ as $X_i(Z_i)$ and $Y_i(\mathbf{Z}, \mathbf{X})$ as $Y_i(Z_i, X_i)$. Substantively, this assumption implies that the treatment of one child had no effect on the outcomes of other children, which seems plausible in our application, particularly since voting by individual i is unlikely to have been affected by the educational attainment or treatment assignment of individual j.

For individual i, the causal effect of treatment assignment on high school graduation is $X_i(Z_i = 1) - X_i(Z_i = 0)$. A basic assumption in the analysis that follows is that the expected value of this effect across the N observations is nonzero. This assumption makes intuitive sense: in order to estimate the downstream effects of a random intervention, the intervention itself must have an effect on graduation rates. We further assume what Imbens and Angrist (1994) term "monotonicity": $X_i(Z_i = 1) \geq X_i(Z_i = 0)$ for all i, which means being assigned to the experimental treatment makes each individual more likely to graduate. In the applications described below, it is hard to imagine preschool, small classes, or college scholarships diminishing a subject's propensity to graduate from high school.

The causal effect of primary interest is the influence of high school graduation on voter turnout. The key assumption on which the identification of this causal parameter rests is the exclusion restriction, the stipulation that the potential voting outcome is a function of graduation rates alone, and that once one takes graduation rates into account, treatment assignment exerts no direct influence on voting (i.e., $\mathbf{Y}(\mathbf{Z}, \mathbf{X}) = \mathbf{Y}(\mathbf{X})$). Substantively, this assumption means that voting is affected not by the program that induced additional schooling but rather by the schooling itself and enables us to write the causal effect of X on Y as $Y_i(X_i = 1) - Y_i(X_i = 0)$. In other words, the causal effect of education on turnout is the difference between two states of the world, one in which the individual graduates and another in which he or she does not.

When assessing the validity of this exclusion restriction, one should keep in mind that X may be a coarse proxy for educational attainment. Whether one graduates from high school is a less discriminating measure than whether one graduates from high school, college, or neither. And the highest degree obtained is a less refined measure of educational attainment than subject tests that gauge the breadth and depth of one's schooling. It should be stressed, however, that the particular X one chooses to employ in an empirical analysis primarily serves as a scaling device, so that the causal effect can be interpreted based on the units that characterize X (e.g., "This is the effect of graduating from high school."). The instrumental variables estimator described below provides consistent estimates of causal effects even when X is measured coarsely. In effect, the exclusion restriction says that the program to which people were randomly assigned (Z) has no direct effect on voting net of one's latent level of educational attainment, for which X is a proxy.

A related issue is how to regard factors such as interest in politics, which may be related to educational attainment. If interest in politics is caused by educational attainment, the model presented above will correctly recover the "total" effect of X, including the causal influence that is transmitted through interest in politics and other mediating factors. The model breaks down if political interest causes voting and the randomly assigned program Z influences interest in politics directly, not via the mediating variable of educational attainment. In this instance, our causal inference about the effects of education would be biased, as part of what we are attributing to education is really due to political interest. When reading the descriptions of the experimental interventions below, one should reflect on whether it is plausible to think that their effects were transmitted through channels other than educational attainment.

Suppose that one wanted to remain as agnostic as possible about the mechanisms through which these interventions influenced outcomes. One could simply focus on the so-called intent-to-treat effects, the difference between $Y_i(Z_i = 1)$ and $Y_i(Z_i = 0)$. As noted below, the instrumental variables estimator is simply this quantity divided by $X_i(Z_i = 1) - X_i(Z_i = 0)$. In effect, the IV estimator takes the intent-to-treat effect and rescales it by the differences in educational attainment for the treatment and control groups.[1] Ultimately, the causal inferences we draw may be traced to the intent-to-treat relationship, that is, the relationship between randomly assigned treatment groups and voting rates.

Estimation. It is impossible to observe causal effects directly, because we do not see a given person simultaneously in his or her treated and untreated states. In order to estimate causal effects, we use the fact that random assignment generates groups whose expected responses to treated and untreated states are the same. Under random assignment, the expected outcome for the untreated control group is the expected outcome that would have been obtained by the treatment group had it gone untreated; conversely, had the control group been treated, its expected outcome would be the same as the expected outcome of the (treated) treatment group.

In our application, the parameter of interest is what is often termed a "local average treatment effect." Here the local average treatment effect is the causal effect of high school graduation for the subset of individuals whom Angrist, Imbens, and Rubin (1996) call "compliers"—those who graduate from high school if and only if assigned to the treatment group. Just who the compliers are will depend on what the random treatment is. Thus, estimation of a local average treatment effect inevitably raises questions of external validity, because there is no guarantee that treatment effects will be the same across different populations of compliers. That is why we sought out an array of different experiments. The fact that we obtain similar results across these experiments suggests that the local average treatment effect has greater generality than might be suggested by the modifier "local."

Angrist, Imbens, and Rubin (1996) prove that under the five assumptions discussed above (random assignment, stable unit treatment value, monotonicity, nonzero causal effect of Z on X, and no effect of Z on Y net of X), instrumental variables (IV) regression is a consistent estimator of the local average treatment effect of X on Y, even when X is thought to be related to unobserved causes of Y. Note that this property is what sets apart the instrumental variables approach used here from more conventional regression approaches. Ordinary least squares (OLS) will be inconsistent when educational attainment is correlated with unobserved factors, such as intelligence. Instrumental variables regression leverages random assignment to produce consistent estimates even in the presence of con-founders.

The instrumental variables estimator suggested by Angrist, Imbens, and Rubin (1996) has a very simple mathematical form, which allows for transparent presentation of results. The estimator is the ratio of two regression coefficients: the estimated effect of Z on Y (in path-analytic terms, the "total effect" of random assignment on voting) over the estimated effect of Z on X (the apparent effect of random assignment on schooling). In the sections that follow we preview the statistical results by presenting these two simple quantities.

Bivariate probit. Our parameter of interest is the effect of high school graduation on voter turnout, and we seek to estimate it while allowing for the possibility that high school graduation is correlated with unobserved causes of turnout. As noted above, the proper approach involves the logic of instrumental variables, whereby random assignment (or near-random assignment) serves as an instrumental variable predicting graduation rates. In our application, a complication arises because high school graduation and voting are both dichotomous variables, and linear models do not confine predicted probabilities of voting to the [0,1] interval. We therefore depart from the nonparametric framework of the Angrist, Imbens,

[1]An analogous adjustment occurs when bivariate probit (see below) is used, but this nonlinear model imposes additional restrictions that also affect the precision of the resulting estimates. Whereas instrumental variables estimation produces treatment estimates that have significance levels that are very close to those obtained from a reduced form regression of Y_i on Z_i , bivariate probit tends to produce more significant treatment estimates than a probit regression of Y_i on Z_i.

and Rubin (1996) model and turn to a nonlinear estimation approach, bivariate probit (Wooldridge 2002), which has been shown to perform well in Monte Carlo simulations and empirical applications involving dichotomous treatments and outcomes (Terza, Bradford, and Dismuke 2008), particularly when the analysis focuses solely on Y_i, X_i, and Z_i, and excludes covariates (Angrist 2001).

The two-equation system in our analysis includes an equation for voter turnout (Y_i) and an equation for high school graduation (X_i). The model specifies a linear model of the *latent* propensity to vote (Y_i*) and to graduate (X*) as functions of the realized outcomes of high school graduation (X_i) and random assignment (Z_i).

$$Y_i^* = \beta_0 + \beta_1 X_i + \varepsilon_Y \quad (1)$$

$$X_i^* = \alpha_0 + \alpha_1 Z_i + \varepsilon_X \quad (2)$$

The translation of latent propensities into observed outcomes is as follows: $Y_i = 1$ if $Y_i^* > 0$; otherwise, $Y_i = 0$. Similarly, $X_i = 1$ if $X_i^* > 0$; otherwise, $X_i = 0$. The bivariate probit specification rests on the assumption that (ε_Y, ε_X) are independent of the randomly assigned Z_i and distributed bivariate normal with mean zero and unit variance. The use of bivariate probit estimation, as opposed to single equation probit, is justified by concern that the correlation between ε_Y and ε_X is nonzero. Correlation between these disturbance terms will bias conventional probit regressions of Y on X. Bivariate probit is also superior to other techniques that involve a linear approximation of either equation (1) or (2) and therefore risk bias (see Wooldridge 2002).

The main parameter of interest in the bivariate probit model is β_1, the effect—in probits, or standard deviations in a standard normal distribution—of high school graduation on voter turnout. Expressed in terms of percentage-point changes, the local average treatment effect of high school graduation on voting is calculated as $\Phi(\beta_0 + \beta_1) - \Phi(\beta_0)$, where $\Phi(\bullet)$ refers to the cumulative normal distribution function, and β_0 and β_1 are maximum likelihood estimates calculated using STATA statistical software.

Overview of Three Studies

In an effort to measure the downstream effects of education on voter turnout, we searched the education literature for randomized interventions that generated significant increases in schooling. The pool of extant experiments proved to be surprisingly small. As Boruch, De Moya, and Snyder (2002) point out, only 10% of the research projects funded by the Department of Education's Planning and Evaluation Services involved any form of randomized field trials, and of the 1,200 articles printed in the *American Educational Research Journal* since its inception in 1964, just 35 used randomized designs to study curriculum interventions. After narrowing the candidate list of experiments to those for which data were available, we were left with two randomized studies, the Perry Preschool Experiment and the Tennessee STAR experiment. We were unable to obtain data from existing quasi-experimental studies of the "I Have a Dream" (IHAD) scholarship program, but we were able to gain the cooperation of one IHAD site, which assisted us in the collection of data on its students. This section describes the three programs and provides preliminary evidence of their causal effects.

Perry Preschool Experiment

The High/Scope Perry Preschool program is a longitudinal evaluation of the effects of intensive preschool on graduation from high school and future life outcomes. The Perry program is widely cited as a model evaluation both for its methodology and its lasting effects on low-income youth (Barnett 1985). The experiment's central hypothesis is that "good preschool programs can help children in poverty make a better start in their transition from home to community and thereby set more of them on paths to becoming economically self-sufficient, socially responsible adults" (Schweinhart, Barnes, and Weikart 1993, 3).

The study drew its participants from preschool-age children who were expected to attend Perry Elementary School in Ypsilanti. Because the program's broader aim was fighting poverty, the subject population was drawn from a low-income subgroup of this residential area. Researchers first identified families of low socioeconomic status using a score based on three factors: parents' educational levels, parents' occupational levels, and the number of rooms in the family household.[2] Within this group of low SES families, researchers then sought to identify young children with relatively low levels of intelligence.

[2]Based on these inclusion criteria, families eligible for the study were more economically disadvantaged than the national African American population. For example, half of the fathers and mothers of the children in the study had dropped out of school by eighth and ninth grade, respectively. Overall, 20% of the mothers and 10% of the fathers of children in the study graduated from high school. The national graduation rate at this time was a bit over 50% and approximately 33% for African Americans.

Before randomly assigning the experimental groups, researchers matched students into pairs based on IQ scores from the Stanford-Binet test.[3] A member of each pair was randomly assigned into one of two unmarked groups. Next, some of these ranked pairs were swapped between groups so that both groups would have similar mean socioeconomic status, mean intellectual performance, and gender ratios. A coin flip determined which of these two groups was assigned to the program (treatment) condition and which to the no-program (control) condition. Finally, researchers moved siblings of those students chosen for the treatment condition into this condition as well so as to eliminate spillover effects from the intervention. At the outset, 64 students were each assigned to the treatment and control groups. However, two students were subsequently moved from the treatment to control because their single mothers were employed outside of the home and thus unable to participate in the home visits and classes necessary for the treatment regimen. For purposes of analysis, we consider these two children as members of the treatment group, as placed during their *original* random assignment.[4] In all, 100 families participated in the study, 47 in the treatment cohort and 53 in the control cohort. Four students moved away during the preschool intervention, and one other passed away shortly into the program. Thus the Schweinhart analysis focuses on the remaining 123 students who either received the preschool treatment or would have received the treatment had they been selected for the experimental cohort. The final numbers of the Schweinhart sample are 58 students in the treatment group (33 males and 25 females) and 65 students in the control group (39 males and 26 females). All subjects are African American.[5]

Students entered the study in five waves based on their ages beginning in 1962. Four of the waves received two years of preschool at ages 3 and 4, but the first wave received only one year of preschool because it began the program at age 4. The preschool treatment consisted of

TABLE 2 Perry Experiment: Effects of Random Assignment on Graduation Rates and Voter Turnout

	Control	Treatment
Graduated from high school	44.4%	65.0%
Voted in 2000 or 2002	12.7%	18.3%
N	63	60

Source: Voting records obtained from Voter Contact Services.

2.5 hours in the classroom on weekday mornings as well as a weekly 90-minute teacher visit to the mother and student at the family's place of residence during a weekday afternoon. The program began in mid-October and lasted through May of each year. The curriculum was designed based on the ideas of Jean Piaget in that teachers emphasized "active learning" and students were encouraged to plan, carry out, and finally review the activities given to them. There was a focus on "open-ended questions" in an effort to engage students in conversations both with adults and other students. Teachers also attempted to maintain a daily routine that promoted responsibility and independence. Following the preschool intervention, subjects entered kindergarten as scheduled by local school regulations based on date of birth. No further treatment occurred after the conclusion of the initial intervention.

Outcome measures. The principal investigators of the Perry study shared with us follow-up data on the high school graduation status of 121 of 123 subjects for whom they have longitudinal information. Voter turnout data were gathered from public records, using data supplied by the firm Voter Contact Services, based on each participant's name, address, and birthday.[6] Twenty of the original 123 subjects now live outside of Michigan in 11 different states and the District of Columbia. Ninety-six subjects still live in Michigan, and seven are deceased.

Table 2 presents an overview of the Perry results. The randomly assigned groups, as noted above, had different educational outcomes. High school graduation rates were 44.4% in the control group and 65.0% in the treatment group. Consistent with the hypothesis that schooling exerts a causal influence on electoral participation, voter

[3]Only three families identified through this process declined to participate in the study. This decision was made prior to assignment into the control and treatment groups and thus has no effect on the internal validity of the study. Note that analysis using a matched-pair design is unnecessary because each member of every pair had the identical chance of being selected for the treatment and control groups.

[4]One could thus consider the treatment group an "intent-to-treat" group since not all subjects originally assigned to the preschool group received the treatment.

[5]Throughout the 1960s, the residents of the South Side of Ypsilanti, where Perry Elementary is located, tended to be African American. Although there was no formal segregation at the time of the experiment, housing patterns resulted in a predominantly African American student population.

[6]Voter Contact Services, a commercial vendor, gathers voter registration and turnout information from public agencies charged with maintaining this data at the county level. Both Voter Contact Services and Polimetrix (see below) were instructed to indicate whether subjects had voted in federal elections. Federal elections, rather than state or municipal elections, are employed in this analysis so that all subjects had the same opportunity to vote, regardless of where they resided.

turnout was higher in the treatment group (18%) than the control group (13%).

I Have a Dream Natural Experiment

The "I Have a Dream" (IHAD) program is a comprehensive scholarship program aimed at increasing high school graduation and postsecondary matriculation rates of at-risk youth. While IHAD began as a means of financing college education for low-income students, it has since evolved into a wide-ranging program, which in the words of one sponsor, "aims to provide a middle class experience" for students from disadvantaged backgrounds. In addition to the promise of college scholarships, students in the program receive tutoring in various academic subjects, work with mentors on establishing and achieving educational and career-oriented goals, and participate in extracurricular activities together from the time of their initial selection in elementary school through high school and sometimes beyond.

This study focuses on a particular intervention conducted by the "I Have a Dream" Foundation of Boulder County in Colorado. At the outset of this project, the national IHAD foundation mailed a letter concerning our research to all IHAD programs that had students in tenth grade or older. The Boulder project distinguished itself from the group of interested program sites for two reasons. Their foundation was in the beginning stages of an independently undertaken project to update contact information for their participants. Second, due to its relationship with the University of Colorado, the Boulder Valley School District has a history of being amenable to educational evaluations of its students.

In 1992, enrollment in IHAD was offered to and accepted by all 79 fifth-grade students in Lafayette, Colorado, who qualified for the free or reduced lunch program at three elementary schools. Drawing upon the extensive quasi-experimental literature evaluating IHAD's effects (Aron and Barnow 1994; Higgins et al. 1991; Kahne and Bailey 1999; Kuboyama 2000; McGrath and Hayman 1998; Shoemaker and Sims 1997; Strusinski 1997), we formed a control group consisting of those students on free or reduced lunch at the same schools but attending fourth or sixth grade in 1992. This design creates treatment and control cohorts with the similar ethnicity, age, and socioeconomic backgrounds. In contrast to the Perry study, approximately two-thirds of the students are non-Hispanic whites, and fewer than 5% are African Americans.

Members of the treatment and control cohorts were contacted via a telephone survey. In order to maintain confidentiality, the school district provided names and contact information to IHAD, which coordinated the interviews.[7] It should be stressed that exactly the same district-supplied contact information was used to track the treatment and control cohorts. In order to identify which students had participated in free or reduced lunch programs as elementary students, the survey asked respondents to reminisce about their elementary and high school experiences in the Boulder Valley. Respondents were asked about their clubs and after-school activities as well as their memories of their school lunch experiences. They were asked various questions including whether they brought their lunches from home or purchased them at school, with whom they sat at lunch, as well as whether or not they remember participating in the free and reduced lunch program. Respondents who recalled participating in the lunch program were then asked if they would be willing to participate in a follow-up survey. This filter is intended to produce similar sampling biases for both the treatment and comparison cohorts.

Using this methodology, our research team administered the initial survey to 246 individuals, 58 of whom acknowledged participation in the free and reduced lunch program and were willing to take the in-depth survey.[8] Given that the study comprises two sets of control group respondents (who attended fifth grade during the year preceding and following the IHAD cohort), we should expect twice as many subjects in the control versus experimental cohort. The survey numbers bear this out as 32.8% or 19 of the 58 respondents who recall participation in the free and reduced lunch program also acknowledged participation in IHAD.

[7]Boulder Valley School District provided the IHAD Foundation of Boulder County with a list of the 1,020 students who entered the fifth grade in the fall of 1991 and 1992 and the fourth grade in 1992 at the three elementary schools from the initial intervention as well as their most recent addresses and phone numbers on file with the district. This list of students was then sent to a firm that attempted to find updated contact information for each subject using public records. Of the 1,020 initial subjects, phone numbers were located for 992 individuals. Once the phone list was compiled, three persons hired through the IHAD Foundation attempted to contact each person to administer the initial survey. These lists were rotated through the interviewers three times, with each attempting to call at different times throughout the day and during the week. To ensure equal likelihood of contacting subjects in the treatment and control cohorts, those in the IHAD program were not sought out using more up-to-date information from the Boulder IHAD Foundation. To preserve the symmetry between treatment and comparison groups, only information found through the aforementioned process was used.

[8]There was some attrition here as four respondents qualified for the free or reduced lunch program but declined to take the in-depth survey and three respondents agreed to take the in-depth survey at a later date but were unable to be reached at their scheduled appointment times.

TABLE 3 IHAD Natural Experiment: Effects of Near-Random Assignment on Graduation Rates and Voter Turnout

	Control	Treatment
Graduated from high school	61.5%	79.0%
Voted in any election through 2004	33.3%	42.1%
N	39	19

Source: Voting records obtained from Voter Contact Services.

Although an earlier IHAD survey included questions concerning civic and political participation, we opted to gather verified information on turnout from registration rolls rather than rely on self-reports and the biases they may introduce (Bernstein, Chadha, and Montjoy 2001; Brady 1999; Hyman 1944; Silver, Anderson, and Abramson 1986).[9]

Table 3 reports the results of the IHAD intervention on graduation and voting rates. Consistent with previous evaluations of IHAD programs, our data show that IHAD increased the graduation rate from 61.5% to 79.0%. The IHAD intervention is also associated with increased voting rates. Voting records indicate that 42% of those receiving the IHAD treatment cast ballots, as compared to 33% of the control group.

STAR Experiment

In 1985, the Tennessee state legislature allocated $12 million in funding over four years to statewide research concerning the effects of class size reduction on educational outcomes. The Student Teacher Achievement Ratio (STAR) experiment sought to test the effects of small classes on students in kindergarten through third grade. The study (described in detail by Achilles 1999; Word et al. 1990) represents the largest randomized experiment to test the effects of small classes on student achievement in the United States. In contrast to the Perry and IHAD studies, STAR subjects were drawn from a broad socioeconomic spectrum, with African Americans comprising approximately one-third of the sample (Krueger 1999; Krueger and Whitmore 2001; Mosteller 1999).[10] Close to half of the subjects were eligible for free or reduced lunch (Nye, Hedges, and Konstantopoulos 1999).

Students and their teachers were randomly assigned to one of three experimental conditions: small classes of 13–17 students, regular classes with 22–25 students, and regular classes with a full-time teacher's aide.[11] Once assigned to a particular type of classroom condition, students ostensibly remained in those conditions through the remainder of the experiment. Critics such as Hanushek (1999a, 1999b) have pointed out that approximately 108 of the 6,505 participants in the STAR database moved between experimental groups after kindergarten, apparently due to lobbying by parents seeking smaller classes. For this reason, we restrict our attention to 1,576 kindergarteners in the subset of 18 schools for which the original random assignment information is available.[12] In keeping with the principle of maintaining the integrity of the original random assignments, we ignore which treatments students actually received and use only the randomly *assigned* class size as an instrumental variable.

Outcome measures. The STAR intervention lasted until fourth grade, whereupon students entered regularly sized classes. Kindergarteners who participated in Project STAR were expected to graduate from high school in 1998. Follow-up studies indicate that subjects assigned to attend small classes subsequently exhibited higher levels of academic performance (Finn and Achilles 1999; Krueger 1999) and higher rates of high school graduation in comparison to subjects in regular-size classes (Finn, Gerber, and Boyd-Zaharias 2005). Subjects assigned to small classes were also more likely to have taken college entrance exams, specifically the SAT or ACT, although there were no significant differences in these scores (Krueger and Whitmore 2001).

In the data available to us, the only long-term measure of educational attainment is graduation from high school. A subject is coded as a one if records indicate that she graduated and zero otherwise.[13] Unfortunately, due

[9]A further reason for concern about the prior survey is its timing. Interviewers contacted subjects in the summer and fall of 2004, prior to the November presidential election, which would have been the first opportunity for many of the subjects to vote.

[10]Krueger and Whitmore (2001) estimate that, in comparison to the entire state of Tennessee, Project STAR schools have higher than average minority enrollment rates. Moreover, African American participation in the program is twice the national average enrollment rate.

[11]During the course of the experiment, class size was the only change made by the participating schools. Teachers did not alter their curricula, and a follow-up study confirms that teachers did not substantially alter their teaching practices based on cohort assignment (Evertson and Randolph 1989). Because schools had different numbers of pupils, the probability of receiving the treatment varies somewhat from school to school. However, adding fixed effects for school has no material effect on our results because voting rates across schools do not differ significantly by school. If anything, the ITT grows larger when controlling for school.

[12]Krueger (1999) performed the same coding using the same set of 18 rosters and finds 1,581 students. Our rosters were coded twice, producing 1,576 both times.

[13]Subjects who earned a GED were coded as not having graduated (personal correspondence with Jayne Boyd-Zaharias in 2005). In

to STAR researchers' budget constraints, this follow-up information was gathered for approximately 40% of the observations in the study (Pate-Bain, Fulton, and Boyd-Zaharias 1999). Restricting analysis to those for whom we have both the original kindergarten assignment and high school graduation data reduces the N to 811.[14] Fortunately, the analysis presented below shows the results to be robust across different sample definitions: all STAR data with known random assignments and the subset with nonmissing graduation data.

The Project STAR database contains no data on voter participation. We contracted with the firm Polimetrix, an independent polling firm that maintained a compilation of all state voter files, to determine voter registration and turnout. Polimetrix matched the names and birth dates of the Project STAR subjects to their national list of registered voters. Our information, however, was sometimes insufficient to generate a unique match. To differentiate among multiple matches, the search process also considered the locations to which subjects were likely to move. For example, individuals born in Tennessee have a higher probability of moving to Kentucky or Missouri than to Connecticut or North Dakota. Using this information, and blind to the treatment or control status of the subjects, Polimetrix searched for each subject in their national voter database and compiled a list of likely matches. A hit was determined by name, date of birth, and state of residence, in that order, and assigned a score indicating the likelihood of a correct match.

Before taking the scores into account, this process found at least one hit for 1,455 out of the 1,576 subjects in our dataset.[15] There was at least one hit for each of the 811 subjects for whom we have high school graduation data. Using this list of hits, we construct the voter turnout variable matching the most likely individual out of the possible hits. We also create a more restrictive voter turnout variable requiring that each individual meets a minimum match score.

The turnout variable is coded one for individuals who voted in any one of the following elections: the 2000 primary, the 2000 general, the 2002 primary, or the 2002 general election. The broad turnout variable is coded as a zero for subjects whose best hit did not vote in any of the aforementioned elections. The restricted turnout variable constrains hits to a minimum level match such that any subject whose best hit is below a particular threshold (indicating an unlikely match according to Polimetrix programmers) is considered to be unregistered and thus a nonvoter.

Table 4 shows the relationship between random assignment and education and voting outcomes. The right-hand panel of Table 4 compares graduation rates among those subjects for whom data are available. Students randomly assigned to the treatment group graduated at a rate of 90.1%, as compared to 85.0% among those in the control group. Regardless of whether we restrict our attention to students with graduation data or examine the entire sample, and regardless of whether we use lax or restrictive criteria for determining whether a name was matched to the voter file, we find higher voting rates among the treatment group than the control group. Evidently, random or near-random inducements that increase educational attainment also increase voter turnout rates. The next section estimates the magnitude and statistical significance of this relationship.

Statistical Results

From Tables 2, 3, and 4, we know that the intent-to-treat effects are found to be positive in all three studies. We also know that the treatment group in each study had a higher graduation rate than its control group counterpart. The task remaining is to put these two pieces of information together, using bivariate probit to estimate the causal effect of graduating from high school and the statistical uncertainty associated with this estimate.

Table 5 reports the results of bivariate probit regressions applied to each dataset. The estimates of most interest are those in the middle of the table, indicating the local average treatment effect of high school graduation on voter turnout.[16] In all three datasets, this coefficient is

a previous version of this article, earning a GED was coded as graduating from high school in the Perry and IHAD studies. Using this coding, the bivariate probit estimate for the effect of education in the Perry study was 1.177 (bootstrap SE = 1.047); for IHAD, it was 0.824 (bootstrap SE = .921). The current version of the article uses the same standard of high school graduation to maintain consistency across the three studies. Angrist and Pischke (2009, chap. 8) report that conventional standard errors are downwardly biased for the STAR experiment; our bootstrapped results impose a correction factor larger than the one they report.

[14]For reasons detailed below, the sample drops from 1,576 to 1,455 subjects. The graduation data are based on examination of these 1,455 subjects.

[15]Unfortunately, the "misses" are not an indication of nonvoting/ registration. Rather, the misses indicate that there was not enough data on which to search. (The misses are made up of individuals who were not originally in the STAR database and thus have no recorded birth date.) These individuals also have no information on high school graduation. Because we have no data aside from original assignment (which is unrelated to missingness), we exclude them from all subsequent analysis, reducing the number of subjects under consideration to 1,455.

[16]The results are similar if we were to use years of education as opposed to high school graduation, modeling the system of equations

TABLE 4 STAR Experiment: Effects of Random Assignment on Graduation Rates and Voter Turnout

	All Subjects		Subjects with Known Graduation Status	
	Control	**Treatment**	**Control**	**Treatment**
Graduated from high school	–	–	85.0%	90.1%
Voted in 2002 or 2004 elections (broad match criteria)*	41.5%	43.8%	43.8%	48.4%
Voted in 2002 or 2004 elections (restrictive match criteria)**	38.8%	40.8%	42.2%	46.8%
N	1,026	429	559	252

Source: Voting records obtained from Polimetrix.
*The broad turnout measure is coded zero for subjects whose closest match to the voter file did not vote in 2002 and 2004.
**The restricted turnout measure is coded zero for any subject whose closest match to the voter file did not vote in 2002 and 2004 or who was, based on an imputation algorithm, thought to have a low probability of being a registered voter.

TABLE 5 Bivariate Probit Regression Results

	Perry Study	**IHAD Study**	**STAR* (Graduation Subset)**	**All Studies Pooled****
Dependent Variable = High School Graduation	–0.140	0.295	1.037	
Constant	(0.158)	(0.204)	(0.063)	
(SE)	0.529	0.513	0.246	
Assignment to Treatment Cohort	(0.228)	(0.383)	(0.118)	
(SE)				
Dependent Variable = Voter Turnout	–1.602	–1.013	–1.850	
Constant	(0.553)	(0.752)	(0.778)	
(Bootstrapped SE)	1.084	1.051	1.948	1.401
High School Graduate	(0.888)	(1.070)	(0.880)	(0.540)
(Bootstrapped SE)				
1–Tailed p–value for the effects of high school graduation on voter turnout	0.111	0.163	0.013	.005
Log pseudo–likelihood	–133.180	–73.473	–862.237	
Number of Observations	123	58	811	

*STAR estimates are based on the restrictive measure of voter turnout described in Table 3.
**Pooled results are derived using a precision-weighted average of the bivariate probit coefficients for each study. A precision-weighted average is the average of the three coefficients, where each is weighted by the inverse of its estimated sampling variance. Sampling variances for the voter turnout equation were estimated using bootstrapping, which gives a more conservative assessment of sampling variability.

positive and large, although it achieves statistical significance only in the STAR study. High school graduation increases turnout by 1.1 probits in the Perry dataset, 1.1 probits in the IHAD dataset, and 1.9 probits in the STAR dataset. To put these figures into perspective, bear in mind that one probit is the equivalent of moving one standard deviation in a normal distribution; e.g., moving from a 50% probability of voting to an 84% probability. Not surprisingly given the small sample sizes involved, each estimate is associated with a fair amount of sampling variability. Since none of the estimates is significantly different from one another,[17] it makes sense to pool the results together.

In order to calculate the pooled estimate, we follow standard meta-analytic practice and computed a precision-weighted average, that is, an average in which each estimate is weighted by the inverse of its squared standard error.[18] The resulting estimate is 1.40 with a standard

using a linear first-stage and a probit second stage. Years of education, however, is measured in only two of our three datasets.

[17]The standard error of the difference in coefficients is the square root of the sum of the squared standard errors. The largest z-ratio is therefore (1.948–1.084)/1.250 = 0.69, p > .25.

[18]An alternative approach to a precision-weighted average is to combine the three studies into a single dataset, adding fixed effects

error of 0.54, which implies a z-ratio of 2.59 (p = .005). Translated into percentage point terms, this probit estimate indicates that a high school dropout with a 15.6% chance of voting would have a 65.2% chance of turnout if randomly induced to graduate from high school. A 90% confidence interval surrounding the estimate of 1.4 ranges from .51 to 2.29. Even the low end of this interval still suggests an important effect: again translating into percentage terms, a high school dropout with a 15.6% chance of voting would have a 30.9% chance of voting if randomly induced to graduate from high school.

Naturally, this estimate is subject to the usual caveats about local average treatment effects estimated from special populations. We do not know the effect of education among those who could not be induced to graduate as a function of the three interventions, nor can we be certain that the three studies' subjects are as responsive to educational attainment as the general population in the United States or other times or places. It may well be the case that the low SES students who make up the bulk of these experimental samples are especially responsive to exogenous interventions that lead to higher graduation rates. Nevertheless, one can still be impressed by the sheer magnitude of the effect we observe in all three studies and the fact that the pattern of results looks similar across three demographically distinct sets of subjects.

Conclusion

Prior to undertaking this project, the authors expressed skepticism about the causal claim that education increases voter turnout (Green 2005; Sondheimer 2006). The data presented here have led to a reversal of this assessment. Although each of the three studies is small and idiosyncratic, the pattern of evidence is convincing. Under the null hypothesis of no causal effect, the joint probability of obtaining three bivariate probit coefficients as large as the ones we observe is small. More work remains to be done, as the downstream consequences of recent randomized interventions such as school choice (Hastings, Kane, and Staiger 2005) remain to be investigated. The experimental evidence at hand, however, indicates that educational attainment profoundly affects voter turnout.

It is tempting but mistaken to dismiss these findings as merely confirming what scholars have long known about education's effects. As Gerber, Green, and Kaplan (2004) point out, observational findings confront two sources of uncertainty: sampling error, which diminishes as the number of studies and observations mounts, and specification error, which persists so long as the biases associated with observational research designs remain unclear. The accumulation of hundreds of observational studies has reduced sampling uncertainty to a quantity very close to zero. Yet model uncertainty has remained considerable, leading scholars to wonder whether what if any part of the observed correlation is causal. The experimental and quasi-experimental studies reported here represent an important turning point in a literature that has for decades found itself mired in uncertainty about whether to attach a causal interpretation to the correlation between education and political participation. The results suggest that the true causal relationship is profound and warrants renewed theoretical discussion on schooling and democratic citizenship and empirical investigation into possible causal linkages.

By what mechanism does education induce higher rates of voting? The extensive literature on education and turnout has proposed several possible mediating fac-tors.[19] One hypothesis is that education imparts the skills to negotiate bureaucratic hurdles associated with voting (Wolfinger and Rosenstone 1980). This argument is sometimes adduced to explain why education is less influentialin countries such as Britain, where voter registration is automatic (Milligan, Moretti, and Oreopoulos 2004), but it runs into a number of empirical problems. One finds a steep education gradient in voting even in states such as North Dakota, where there is no voter registration, or in states where voters can register on Election Day (Middleton 2007). Conversely, one finds no diminution in the relationship between education and turnout over time, despite the easing of registration requirements. Our experimental datasets provide tentative support for the bureaucratic competence hypothesis. Registration was measured in both the Perry and IHAD studies, and in both cases the assigned treatment group registered at higher rates than the assigned control group (22.9% vs. 17.0% for Perry; 42.1% vs. 41.0% for IHAD), although the results fall well short of statistical significance.

for study to the bivariate probit specification in equations (1–2). This approach makes the stronger assumption that all three interventions have the same effect on graduation rates and that the covariance between the disturbance terms is constant across studies. This specification produces an estimated education effect of 1.184 (SE = .762), p = .06.

[19]This discussion presupposes that the educational interventions really did influence education attainment. A skeptic might argue that these experiments produced apparent treatment effects because, by chance, the treatment groups comprised more intelligent and able children. Contrary to this hypothesis, the experimental groups turned out to be quite similar in late adolescence or adulthood in terms of intelligence and scholastic aptitude tests.

A second possibility is that education increases one's general interest in and knowledge of politics (Delli Carpini and Keeter 1996; Hyman, Wright, and Reed 1975). Further analysis of the IHAD data provides evidence of a causal relationship between education and political interest. Subjects in the IHAD treatment group scored approximately 1.3 points higher on average than the control group on a 0–7 point index measuring political interest. This interest in politics dovetails with findings concerning the independent relationship between education and internal political efficacy, an individual's belief in her own ability to participate in politics (Sondheimer 2006). Perhaps, then, schooling cultivates a sense of citizenship and efficacy that promotes political involvement, a proposition that demands much more attention from experimental researchers.

Third, it is possible that increased educational attainment expands one's social network and thus likelihood of participating in community and political endeavors. Educated individuals are more likely to have politically involved people in their network of friends and coworkers and are also more likely to receive attention from political campaigns (Rolfe 2004). Thus, education may set in motion not only changes in outlook but also changes in the way that one is engaged by one's social and political environment. Again, surveys of the Perry and IHAD subjects provide some suggestive evidence in favor of this hypothesis, as treatment subjects were more likely to report belonging to community organizations (34.6% vs. 33.9% for Perry; 52.6% vs. 31.6% for IHAD).

These three explanations by no means exhaust the list of potential causal pathways connecting educational attainment to voting. Other possibilities include voters' increased sense of efficacy (Abramson and Aldrich 1982) or strength of partisan attachment (Shaffer 1981) or simply education-induced affluence and accompanying economic interests (Doherty, Gerber, and Green 2006; Sears and Citrin 1982). In sum, the list of potential mediators is long, and the fact that these mediators may be correlated with unobserved causes of voting and with one another makes the empirical task of explaining the relationship between education and turnout particularly daunting (Judd and Kenny 1981).

The challenge for researchers seeking to assess systematically the mediating role of each of the complementary explanations is to devise experiments (or uncover natural experiments) in which the intervention targets a single pathway. For example, if education increases voting by expanding the range of one's social network, the content of education is largely irrelevant; an intervention that increases educational attainment through an enhanced science or performing arts curricula would be expected to increase turnout. Conversely, if education increases turnout by imparting civic knowledge and political interest, the random assignment of civics coursework should increase political participation. And if education transmits its influence by making people more affluent, randomly assigned transfer programs alone should increase turnout.

One aim of this article is to set in motion a line of research that begins by using experiments to identify the effects of education and then employs ever more refined experimental designs in order to answer questions about causal pathways. Another aim is to use experiments in order to determine which, if any, macrolevel trends stand in need of explanation. Based on our experimental results, it appears that the anomaly that Brody (1978) identified is real. If schooling indeed generates higher rates of electoral participation, why have turnout rates fallen since the nineteenth century, and why have they not increased since the 1960s? A full treatment of this question clearly goes beyond the scope of this article, but the leading hypotheses fall into two interrelated categories, environmental factors and changes in voters' attributes. Environmental factors include institutional changes and trends in the way that campaigns interact with voters. The institutional changes by and large have encouraged participation by eliminating poll taxes, easing registration requirements, and extending the length of time that voters may cast ballots. On the other hand, the decline in the number and vitality of civic organizations coupled with the political parties' increasing reliance on mass media and other impersonal campaign tactics have arguably diminished voters' level of civic engagement and voter turnout (Putnam 2000). Although countries outside the United States tend not to have its web of nonparty political organizations, there is growing evidence that implicates increasingly centralized and media-centered campaigns in the apparent trend toward lower turnout observed comparatively (Smith 2006).

In light of the findings presented here, hypotheses about turnout-depressing trends at the macropolitical level warrant renewed attention. Efforts to link aggregate trends in voter turnout to corresponding trends in voters' sense of efficacy or partisan attachment (Abramson and Aldrich 1982; Cassel and Luskin 1988) have stalled in the wake of methodological disputes about the microlevel causal relationship between these variables and voter turnout. That dispute is not unlike the methodological deadlock that has long beset the analysis of education and voter turnout, and the way forward is arguably the same. In order to ascertain whether variables such as political efficacy or organizational involvement influence turnout, we must obtain experimental or near-experimental data in which these factors are changed exogenously through some type of intervention. And in order to ascertain whether education influences efficacy or involvement, we must continue to investigate the consequences of exogenous changes to educational attainment.

References

Abramson, Paul R., and John H. Aldrich. 1982. "The Decline of Electoral Participation in America." *American Political Science Review* 76(3): 502–21.

Achilles, Charles M. 1999. *Let's Put Kids First, Finally: Getting Class Size Right*. Thousand Oaks, CA: Corwin Press.

Angrist, Joshua D. 2001. "Estimation of Limited Dependent Variable Models with Dummy Endogenous Regressors: Simple Strategies for Empirical Practice." *Journal of Business and Economic Statistics* 19(1): 2–16.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91(434): 444–55.

Angrist, Joshua D., and Jorn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Aron, Laudan Y., and Burt S. Barnow. 1994. "Evaluation of New York City's Class of 1992 'I Have a Dream' Program." New York: New York City "I Have a Dream" Foundation.

Barnett, W. Steven. 1985. "Benefit-Cost Analysis of the Perry Preschool Program and Its Policy Implications." *Educational Evaluation and Policy Analysis* 7(4): 333–42.

Bernstein, Robert, Anita Chadha, and Robert Montjoy. 2001. "Overreporting Voting: Why It Happens and Why It Matters." *Public Opinion Quarterly* 65(1): 22–44.

Boruch, Robert, Dorothy De Moya, and Brooke Snyder. 2002. "The Importance of Randomized Field Trials in Education and Related Areas." In *Evidence Matters: Randomized Trials in Education Research*, ed. F. Mosteller and R. Boruch. Washington, DC: Brookings Institution Press, 50–79.

Brady, Henry E. 1999. "Political Participation." In *Measures of Political Attitudes*, ed. J. P. Robinson, P. R. Shaver, and L. S. Wrightsman. New York: Academic Press, 737–801.

Brody, Richard A. 1978. "The Puzzle of Participation in America." In *The New American Political System*, ed. A. King. Washington, DC: American Enterprise Institute for Public Policy Research, 287–324.

Card, David. 1999. "The Causal Effect of Education on Earnings." In *Handbook of Labor Economics*, ed. O. Ashenfelter and D. Card. Amsterdam, Neth.: Elsevier Science, 1801–63.

Cassel, Carol A., and Robert C. Luskin. 1988. "Simple Explanations of Turnout Decline." *American Political Science Review* 82(4): 1321–30.

Converse, Philip E. 1972. "Change in the American Electorate." In *The Human Meaning of Social Change*, ed. A. Campbell and P. E. Converse. New York: Russell Sage, 263–337.

Dee, Thomas S. 2004. "Are There Civic Returns to Education?" *Journal of Public Economics* 88: 1697–1720.

Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.

Dewey, John. 1916. *Democracy and Education: An Introduction to the Philosophy of Education*. New York: Free Press.

Doherty, Daniel J., Alan S. Gerber, and Donald P. Green. 2006. "Personal Income and Attitudes toward Redistribution: A Study of Lottery Winners." *Political Psychology* 27(3): 441–58.

Evertson, Carolyn M., and Catherine H. Randolph. 1989. "Teaching Practices and Class Size: A New Look at an Old Issue." *Peabody Journal of Education* 67(1): 85–105.

Finn, Jeremy D., and Charles M. Achilles. 1999. "Tennessee's Class Size Study: Findings, Implications, Misconceptions." *Educational Evaluation and Policy Analysis* 21(2): 97– 109.

Finn, Jeremy D., Susan B. Gerber, and Jayne Boyd-Zaharias. 2005. "Small Classes in the Early Grades, Academic Achievement, and Graduating from High School." *Journal of Educational Psychology* 97(2): 214–23.

Franklin, Mark N. 2004. *Voter Turnout and the Dynamics of Electoral Competition in Established Democracies since 1945*. New York: Cambridge University Press.

Gerber, Alan S., Donald P. Green, and Edward H. Kaplan. 2004. "The Illusion of Learning from Observational Research." In *Problems and Methods in the Study of Politics*, ed. I. Shapiro, R. Smith, and T. Massoud. New York: Cambridge University Press, 251–72.

Green, Donald P. 2005. "On Evidence-Based Political Science." *Daedalus* 134(3): 96–100.

Green, Donald P., and Alan S. Gerber. 2002. "The Downstream Benefits of Experimentation." *Political Analysis* 10(4): 394–402.

Gutmann, Amy. 1987. *Democratic Education*. Princeton, NJ: Princeton University Press.

Hanushek, Eric A. 1999a. "The Evidence on Class Size." In *Earning and Learning: How Schools Matter*, ed. S. E. Mayer and P. E. Peterson. Washington, DC: Brookings Institution Press, 131–68.

Hanushek, Eric A. 1999b. "Some Findings from an Independent Investigation of the Tennessee STAR Experiment and from Other Investigations of Class Size Effects." *Educational Evaluation and Policy Analysis* 21(2): 143–63.

Hastings, Justine S., Thomas J. Kane, and Douglas Staiger. 2005. "Parental Preferences and School Competition: Evidence from a Public School Choice Program." Yale Economic Applications and Policy Discussion Paper No. 10.

Herrnstein, Richard J., and Charles Murray. 1994. *The Bell Curve: Intelligence and Class Structure in American Life*. New York: Simon and Schuster.

Higgins, Catherine, Kathryn Furano, Catharine Toso, and Alvia Y. Branch. 1991. "I Have a Dream in Washington, D.C. Initial Report." Philadelphia, PA: Public/Private Ventures.

Hogan, Robert E. 1999. "Campaign and Contextual Influences on Voter Participation in State Legislative Elections." *American Politics Research* 27: 403–33.

Hyman, Herbert. 1944. "Do They Tell the Truth?" *Public Opinion Quarterly* 8(4): 557–59.

Hyman, Herbert H., and Paul B. Sheatsley. 1956. "Attitudes toward Desegregation." *Scientific American* 195: 35– 39.

Hyman, Herbert H., and Paul B. Sheatsley. 1964. "Attitudes toward Desegregation." *Scientific American* 211: 16–23.

Hyman, Herbert H., Charles R. Wright, and John Shelton Reed. 1975. *The Enduring Effects of Education*. Chicago: University of Chicago Press.

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62(2): 467–75.

Judd, C. M., and D. A. Kenny. 1981. "Process Analysis: Estimating Mediation in Treatment Evaluations." *Evaluation Review* 5: 602–19.

Kahne, Joseph, and Kim Bailey. 1999. "The Role of Social Capital in Youth Development: The Case of 'I Have a Dream' Programs." *Educational Evaluation and Policy Analysis* 21(3): 321–43.

Kam, Cindy D., and Carl L. Palmer. 2008. "Reconsidering the Effects of Education on Political Participation." *Journal of Politics* 70(3): 612–31.

Katosh, John P., and Michael W. Traugott. 1981. "The Consequences of Validated and Self-Reported Voting Measures." *Public Opinion Quarterly* 45: 519–35.

Krueger, Alan B. 1999. "Experimental Estimates of Education Production Functions. *Quarterly Journal of Economics* 114(2): 497–532.

Krueger, Alan B., and Diane M. Whitmore. 2001. "The Effect of Attending a Small Class in the Early Grades on College-test Taking and Middle School Test Results: Evidence from Project STAR." *Economic Journal* 111(468): 1– 28.

Kuboyama, Emlei M. 2000. "East Palo Alto 'I Have a Dream' Program Evaluation Report." Stanford, CA: Stanford University School of Education.

Levinson, Meira. 1999. "Liberalism, Pluralism, and Political Education: Paradox or Paradigm?" *Oxford Review of Education* 25(1/2): 39–45.

Luskin, Robert C. 1990. "Explaining Political Sophistication." *Political Behavior* 12(4): 331–61.

McDonald, Michael P., and Samuel L. Popkin. 2001. "The Myth of the Vanishing Voter." *American Political Science Review* 95(4): 963–74.

McGrath, Robert E., and Judy Hayman. 1998. "The Paterson, New Jersey, I Have a Dream Program: Academic Performance and Outcomes." Teaneck, NJ: Fairleigh Dickinson University School of Psychology.

Merriam, Charles E., and Harold F. Gosnell. 1924. *Nonvot-ing, Causes and Methods of Control*. Chicago: University of Chicago Press.

Middleton, Joel. 2007. "Is the Relationship between Education and Turnout Weaker in States with Same-Day Registration?" Working paper, Yale University.

Milligan, Kevin, Enrico Moretti, and Philip Oreopoulos. 2004. "Does Education Improve Citizenship? Evidence from the United States and the United Kingdom." *Journal of Public Economics* 88: 1667–95.

Mosteller, Frederick. 1999. "How Does Class Size Relate to Achievement in Schools?" In *Earning and Learning: How Schools Matter*, ed. S. E. Mayer and P. E. Peterson. Washington, DC: Brookings Institution Press: 117–29.

Neuman, W. Russell. 1986. *The Paradox of Mass Politics: Knowledge and Opinion in the American Electorate*. Cambridge, MA: Harvard University Press.

Neyman, Jerzy Splawa. [1923] 1990. "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles." Section 9. *Statistical Science* 5(4): 465–80.

Nie, Norman H., Jane Junn, and Kenneth Stehlik-Barry. 1996. *Education and Democratic Citizenship in America*. Chicago: University of Chicago Press.

Nye, Barbara, Larry V. Hedges, and Spyros Konstantopoulos. 1999. "The Long-Term Effects of Small Classes: A Five Year Follow-Up of the Tennessee Class Size Experiment." *Educational Evaluation and Policy Analysis* 21(2): 127–42.

Pate-Bain, Helen, B. DeWayne Fulton, and Jayne Boyd-Zaharias. 1999. "Effects of Class-Size Reduction in the Early Grades (K-3) on High School Performance." HEROS Inc.

Powell, G. B. J. 1986. "American Voter Turnout in Comparative Perspective." *American Political Science Review* 80(1): 17–43.

Putnam, Robert D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.

Rolfe, Meredith. 2004. "Taking Social Structure Seriously: Education and Voter Turnout." Prepared for presentation at the annual meeting of the Midwest Political Science Association.

Rosenstone, Steven J., and John Mark Hansen. 1993. *Mobilization, Participation, and Democracy in America*. New York: Macmillan.

Rosenzweig, Mark R., and Kenneth I. Wolpin. 2000. "Natural 'Natural Experiments' in Economics." *Journal of Economic Literature* 38(4): 827–74.

Rubin, Donald B. 1978. "Bayesian Inference for Causal Effects: The Role of Randomization." *The Annals of Statistics* 6(1): 34–58.

Rubin, Donald B. 1990. "Neyman (1923) and Causal Inference in Experiments and Observational Studies." *Statistical Science* 5(4): 472–80.

Schuman, Howard, Charlotte Steeh, and Lawrence Bobo. 1985. *Racial Attitudes in America: Trends and Interpretations*. Cambridge, MA: Harvard University Press.

Schweinhart, Lawrence J., Helen V. Barnes, and David P. Weikart. 1993. *Significant Benefits: The High/Scope Perry Preschool Study Through Age 27, Monographs of the High/Scope Educational Research Foundation*. Ypsilanti, MI: High/Scope Press.

Sears, David O., and Jack Citrin. 1982. *Tax Revolt: Something for Nothing in California*. Cambridge, MA: Harvard University Press.

Shaffer, Stephen D. 1981. "A Multivariate Explanation of Decreasing Turnout in Presidential Elections, 1960–1976." *American Journal of Political Science* 25(1): 68–95.

Shoemaker, Margaret, and Molly Sims. 1997. "Delivering on a Promise: An Evaluation of the I Have a Dream Foundation Pasadena, California." Pasadena, CA: "I Have a Dream" Foundation, Pasadena.

Silver, Brian D., Barbara A. Anderson, and Paul R. Abramson. 1986. "Who Overreports Voting?" *American Political Science Review* 80(2): 613–24.

Smith, Jennifer K. 2006. "Campaign Professionalization and the Transformation of Party Organization in Britain and Germany." Unpublished manuscript, Yale University.

Sondheimer, Rachel M. 2006. "The Education-Participation Nexus: Rethinking Conventional Wisdom with

Randomized and Natural Experiments." Unpublished manuscript, Yale University.

Strusinski, Marianne. 1997. "Evaluation of the 'I Have a Dream Project'." Miami, FL: Dade County Public Schools Office of Educational Evaluation.

Tenn, Steven. 2007. "The Effect of Education on Voter Turnout." *Political Analysis* 15(4): 446–64.

Terza, Joseph V., W. David Bradford, and Clara E. Dismuke. 2008. "The Use of Linear Instrumental Variables Methods in Health Services Research and Health Economics: A Cautionary Note." *Health Services Research* 43(3): 1102–20.

Verba, Sidney, and Norman H. Nie. 1972. *Participation in America: Social Equality and Political Democracy*. New York: Harper and Row.

Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics*. Cambridge, MA: Harvard University Press.

Verba, Sidney, Kay Lehman Schlozman, Henry E. Brady, and Norman H. Nie. 1993. "Race, Ethnicity and Political Resources: Participation in the United States." *British Journal of Political Science* 23(4): 453–97.

Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who Votes?* New Haven, CT: Yale University Press.

Wooldridge, Jeffrey. 2002. *Econometric Analysis of Cross Section and Panel Data*. 1st ed. Cambridge, MA: MIT Press.

Word, Elizabeth R., John M. Johnston, Helen Pate-Bain, B. DeWayne Fulton, Jayne Boyd-Zaharias, Charles M. Achilles, Martha Nan Lintz, John Folger, and Carolyn Breda. 1990. "The State of Tennessee's Student/Teacher Achievement Ratio (STAR) Project: Technical Report 1985–1990." Tennessee State Department of Education.