Educational Evaluation and Policy Analysis
December 2015, Vol. 37, No. 4, pp. 419–436
DOI: 10.3102/0162373714558292

Do Charter Schools Improve Student Achievement?

Melissa A. Clark
Mathematica Policy Research, Princeton
Philip M. Gleason
Mathematica Policy Research, Geneva
Christina Clark Tuttle
Mathematica Policy Research, Washington, DC
Marsha K. Silverberg
Institute of Education Sciences, U.S. Department of Education

This article presents findings from a lottery-based study of the impacts of a broad set of 33 charter middle schools across 13 states on student achievement. To estimate charter school impacts, we compare test score outcomes of students admitted to these schools through the randomized admissions lotteries with outcomes of applicants who were not admitted. We find that impacts varied considerably across schools and students, with more positive impacts for more disadvantaged schools and students and more negative impacts for the more advantaged. On average across the schools in the study, the impacts of charter middle schools on student achievement were negative but not statistically significant, regardless of whether we examined the impact of the offer of admission or actual attendance at these schools.

Keywords: *educational policy, program evaluation, methodology, econometric analysis, experimental design*

Charter schools are a central component of current efforts to reform the public education system in the United States. These schools are publicly financed, but free of many of the regulations that govern traditional public schools, such as those involving staffing, curriculum, and budget decisions. As of the 2013–2014 school year, about 6,400 charter schools served about 2.5 million students in 40 states and the District of Columbia.[1] These numbers reflect rapid growth in the charter school sector in recent years; for example, there were just 2,800 charter schools serving 0.7 million students as of 2003.

Despite the policy emphasis on charter schools and the growth in their numbers, rigorous evidence of their effectiveness on a broad scale is limited. Previous research includes observational or non-experimental analyses across several school districts or states (see, for example, Ballou, Teasley, & Zeidner, 2008; Bifulco & Ladd, 2006; Booker, Gilpatric, Gronberg, & Jansen, 2007; Center for Research on Educational Outcomes [CREDO], 2013; Hanushek, Kain, Rivkin, & Branch, 2007; Sass, 2006; Zimmer, Gill, Booker, Lavertu, & Witte, 2012) and lottery-based studies that each focused on a single large urban area (Abdulkadiroglu, Angrist, Dynarski, Kane, & Pathak, 2011; Dobbie & Fryer, 2011, 2013; Hoxby, Murarka, & Kang, 2009; Hoxby & Rockoff, 2005) or a single state (Angrist, Parthak, & Walters, 2012). Most of the non-experimental studies have typically found impacts that were insignificant or negative, whereas the lottery-based studies have found impacts that were large and positive.[2]

Two key differences might explain why the lottery-based studies have tended to show more positive charter school impacts than the non-experimental studies. First, the studies differ in their designs, and the non-experimental studies' internal validity may be compromised if students attending charter schools in a given year differ from those who do not in ways that are not fully captured by the observed variables included in the model. In contrast, the lottery-based studies potentially provide strong internal validity by comparing lottery applicants who were randomly admitted to charter schools with those who were not. Second, the non-experimental studies have tended to cover broad sets of charter schools, whereas the lottery-based studies have focused primarily on urban charter schools.[3] If urban charter schools have more positive impacts than non-urban charter schools, this would explain the pattern of more positive findings in lottery-based studies noted above, even if the internal validity of both non-experimental and lottery-based studies is good.

Relatively little evidence from lottery-based studies exists on the impacts of non-urban charter schools. Through their focus on single large urban areas (Boston, New York, and Chicago) or a single state dominated by a large urban area (Massachusetts), the findings of prior lottery-based studies are most relevant to charter schools in large urban areas. Nationally, fewer than half of all charter middle schools are located in large, urban areas (Gleason, Clark, Tuttle, & Dwoyer, 2010).

This article presents findings from an evaluation of 33 charter middle schools in 13 states.[4] Using the results of these schools' randomized admissions lotteries as an instrumental variable for charter school attendance, the study provides internally valid evidence of charter school impacts on student achievement from a geographically diverse sample. The study team's careful in-person monitoring of the charter school admissions lotteries helped to ensure that the lottery procedures, and students' resulting admission status, were truly random, and that the study results are thus as rigorous as possible.

Consistent with many previous studies that have focused on broad sets of charter schools, we find no evidence that, on average, attending charter schools had a positive impact on student achievement. Instead, we found that attending a charter school had impacts on math and reading that were negative but not statistically significant. A potentially more policy-relevant finding, however, was that estimated impacts varied substantially across charter schools, and this variation was correlated with observable school characteristics. The average impact on achievement of attending charter schools serving lower achieving or more disadvantaged students was large and positive, whereas the average impact of charter schools serving higher achieving or more advantaged students was large and negative.

Background on Charter Schools

Charter schools are public schools that are established on the basis of a contract, or charter, that a private board holds with a charter authorizer over some pre-determined number of years. As part of the contract, charter schools are released from many state and district regulations that govern traditional public schools, such as those involving staffing, curriculum, and budget decisions. In exchange for this flexibility, charter schools are expected to be held accountable for the quality of student outcomes and may be closed by their authorizer if they fail to meet expectations. Most charter schools are open enrollment schools—any student within the district or state in which the school is located may attend the school if space is available. Proponents argue that the schools' autonomy allows them to innovate, test new ideas, and bring competitive pressures to improve traditional public school systems. Critics are concerned that these schools draw students and resources away from traditional public schools and that inadequate oversight will lead to many low-quality charter schools.

The first charter school opened its doors in Minnesota in 1992. The number of states permitting charter schools grew rapidly during the 1990s, as did the number of charter schools and students enrolled. The growth in the number of charter schools and the number of students they enrolled continued to increase in the 2000s, despite the fact that only four new states passed authorizing legislation between 1999 and 2003 and none did so between 2004 and 2009. Charter schools have continued to grow since 2009, perhaps due in part to recent federal initiatives such as the Race to the Top Fund established under the American Recovery and Reinvestment Act of 2009.

Study Design

The study's design relies on the admissions lotteries held by oversubscribed charter schools—schools that had a larger number of applicants than they had spaces available. Lottery winners offered admission to study schools form the treatment group for the study, whereas the lottery losers not admitted to these schools form the control group. The randomized lotteries ensure that the only systematic difference between the treatment and control groups is whether they were admitted to a study charter school—on average, there should be no differences in the characteristics, motivation, or expectations of the students or their parents. Because lottery admission is highly correlated with charter school attendance but otherwise uncorrelated with students' academic outcomes, using lottery admission as an instrumental variable for charter school attendance can provide unbiased estimates of the causal effects of attending a charter school.

The Sample of Charter Middle Schools

Schools were recruited for the study sample over a 2-year period from any state with eligible charter schools. To be eligible for the study, a charter school had to meet three criteria. First, it had to be a middle school—with an entry grade between Grades 4 and 7. Second, it had to have been operating as a charter school for at least 2 years at the time it was recruited. This minimized the chances that participating schools would still be under development and thus undergoing substantial change during the study period. Third, it had to be sufficiently oversubscribed—that is, to have more applicants than could be offered admission to the school—so that it could accommodate the study's experimental design.

The first cohort of schools were those holding admissions lotteries for the 2005–2006 school year, and the second cohort were those holding lotteries for the 2006–2007 school year. Using national databases, we identified 492 charter middle schools that had been open at least 2 years at the time they were recruited and were thus potentially eligible for the study. Among 146 schools deemed to be potentially oversubscribed after an initial survey screen, 49 did not agree to participate in the study and the remainder were more carefully assessed for eligibility. Although 77 schools both agreed to participate and initially appeared eligible for the study, ultimately 36 charter schools in 32 sites remained sufficiently oversubscribed through the study period (i.e., they had at least 10 students who participated in the lottery but remained too far down on the waiting list to be offered admission) and participated in the study in at least 1 of the 2 study years.[5] Of these, 33 schools in 29 sites had sufficient outcome test score data to be included in the analysis of student achievement, as discussed further below.

Table 1 compares the charter middle schools in the study sample with all other charter middle schools nationwide, based on a survey we administered to all charter middle school principals.[6] The charter schools included in the study sample were statistically similar to the national sample of charter middle schools along several dimensions. While most prior lottery-based studies of charter schools focused exclusively on those in large urban areas, as noted above, the proportion of charter schools in large urban areas in this study (36%) closely matched that of charter middle schools not included in the study (41%). The study sample also resembled the national sample in terms of student enrollment, student–teacher ratio, length of school day and year, teacher experience and certification, and revenue per pupil.

However, study schools differed in some respects from non-study charter middle schools (in addition to the fact that all study schools had more applicants than available seats and held admissions lotteries, which was not true of most non-study schools). Study schools had a higher percentage of White students, on average, and a lower percentage of Black students, than the non-study charter middle schools. They also had a lower percentage of students eligible for free or reduced price lunches and a higher percentage of

TABLE 1
Characteristics of Study Charter Schools and Non-Study Charter Middle Schools

	Charter middle schools in study	All other charter middle schools	Difference	*p* value
Located in large urban area (%)	36	41	−4	.602
Enrollment (*M*s)				
Total enrollment	387	298	90	.080
Enrollment per grade	111	88	23	.259
Student–teacher ratio	14.6	16.7	−2.1	.150
Time in school (*M*s)				
School day length in hours	7.3	7.0	0.3	.117
School year length in days	182.4	181.4	1.0	.968
Staff				
Experience of principal (*M* number of years as principal)	6.1	5.7	0.5	.562
% of schools at which two thirds of teachers have 5+ years experience	50	34	16	.060
Midpoint of teacher salary (US$) range at school (*M*)	48,168	44,280	3,888	.022*
% of teachers at school with full state certification (*M*)	77	78	−2	.924
Characteristics of students at school (*M*s)				
% Hispanic	26	25	1	.825
% White	53	38	15	.012*
% Black	16	29	−13	.024*
Average daily attendance rate (%)	95	92	4	.067
% of students receiving free or reduced price lunches	44	62	−18	.003**
% of students with learning disability and/or IEP	12	12	0	.705
% of students classified as LEP	3	9	−6	.069
Academic achievement of students at school (*M*s)				
% of seventh graders meeting state proficiency in math	66	51	15	.001**
% of seventh graders meeting state proficiency in reading	75	57	19	<.001**
Autonomy index (*M*)	4.6	5.2	−0.6	.083
Charter school characteristics				
Age of school (*M*)	7.0	5.9	1.2	.015*
Authorized by local school district (%)	56	44	12	.214
Operated by CMO (%)	11	20	−9	.384
Total US$ revenues per student, including private funding (*M*)	8,030	8,710	−679	.402
Accountability index (*M*)	2.59	2.45	0.14	.296
Sample size: Characteristics based on principal survey or common core of data	36	434		
Sample size: Characteristics based on principal survey alone	35	299		
Sample size: School test scores	36	380		

Note. The source of the information provided in this table includes a survey administered in fall 2006 or fall 2007 to the principals of all charter middle schools nationally, the Common Core of Data from the National Center for Education Statistics, and the School Data Direct database maintained by the State Education Data Center of the Council of Chief State School Officers. IEP = individualized education plan; LEP = limited English proficient; CMO = Charter Management Organization.
*Difference is statistically significant at the .05 level (two-tailed test). **Difference is statistically significant at the .01 level (two-tailed test).

seventh graders meeting their state proficiency standards in both reading and math. Finally, study schools had been in operation longer than non-study schools, on average (7.0 vs. 5.9 years).

These differences highlight the notion that the school sample was not nationally representative, and that impacts for the selected sample may differ from those of all charter middle schools nationwide. Nonetheless, the study included a broad set of charter middle schools across 13 states in both urban and non-urban areas, and included schools serving both highly disadvantaged populations and more advantaged populations. This variation among the charter middle schools in the study sample allows us to examine whether charter schools in different settings or serving different types of students also differ in their impacts on student achievement.

Charter Schools' Admissions Lotteries

Participating charter schools typically held their admissions lotteries in the winter or spring prior to the school year for which students were applying for admission. As the schools' admissions lotteries were central to our ability to produce valid estimates of the schools' impacts, the study team carefully monitored and documented the lottery process. A member of the study team attended each lottery to ensure that the mechanism for selecting lottery winners and determining the order of the waiting list was truly random. After documenting the lottery outcomes, we confirmed with the school that our record of the lottery results matched the record of the school and, if there were discrepancies, we worked to resolve them. We also documented any special features of the lottery, including exemptions, stratification, or special rules for siblings who applied at the same time. Finally, we documented whether sample members applied to more than one charter school participating in the study.

The information we obtained on schools' lotteries enabled us to create sampling weights to ensure that the control group of lottery losers formed an appropriate counterfactual for the treatment group of lottery winners in the analysis. The sample weights—based on each student's probability of "winning" the lottery and being offered admission to the charter school—ensured that the subtleties of the schools' lottery processes were reflected in the analysis. For example, students who applied to different charter schools would likely have differing probabilities of admission based on the number of applicants and the number that each school admitted through the lottery. Furthermore, a student who applied to more than one charter school in our sample (a "dual applicant") would have a higher probability of admission to a charter school than a student who applied to a single charter school, all else equal. Within the lottery at a given school, different groups of students might have different probabilities of admission if the school uses lottery stratification or special rules for siblings who apply together.[7]

Once we determined the probability of admission for each student, we used it to calculate the sampling weight for each student. For those who ultimately were selected into the treatment group, the weight was calculated as the inverse of the probability of admission. For students assigned to the control group, the weight was calculated as the inverse of the probability of *not* being admitted to a charter school (i.e., as the inverse of one minus the probability of admission). Finally, we normalized the weights, so that the weighted sample size would equal the actual sample size in each site. In the end, the construction of the sample weights ensured that both the weighted sample of treatment group students and control group students were representative of the students who applied to the study schools and participated in the schools' lotteries.[8] Gleason et al. (2010) provides more details on how these weights were calculated.

An alternative approach for dealing with different probabilities of admission would have been to include lottery fixed effects (or "risk sets") in the regression model used to estimate impacts. Sample members would be classified based on their probabilities of admission, and fixed effects would be included for each set of students with the same probability. For example, students who applied to more than one charter school in the study would be grouped with other students who applied to the exact same set of charter schools and a fixed effect would be included for this risk set. This approach has been used in several previous lottery-based studies of charter schools (Abdulkadiroglu et al., 2011; Angrist et al., 2012; Dobbie & Fryer, 2011). We conduct sensitivity tests comparing our results with the results based on this risk-set approach.

After the lotteries were conducted, and lottery winners offered admission, the study charter schools continued to admit applicants from the randomly ordered waiting list as space became available. Students who were admitted in the lottery or were offered admission in proper order from the waiting list (whether or not they opted to attend) were included in the study's treatment group, whereas those who participated in the lottery but were not offered admission were included in the control group.

Student Sample

The full student sample included those who applied to a charter middle school in the study, who participated in the school's admissions lottery (i.e., were not exempted from the lottery), and for whom we obtained parental consent. Charter schools often exempt some students from their lotteries, admitting them outside of the lottery process. All charter schools in this study exempted students with siblings already attending the school, and some schools exempted other groups of applicants, such as children of school staff or siblings of alumni. Overall, study schools filled 32% of their open seats with applicants who were exempt from the lottery. Parental consent was obtained prior to the admissions lottery, ensuring that there would be no systematic relationship between the likelihood of consent and the outcome of the lottery (i.e., students' treatment status). Thus, the consent rate was similar for students offered admission to the school through the lottery (62%) and those who participated in the lottery but were not offered admission (61%).

These two sample restrictions—the fact that schools admitted some students outside of the lottery process and that only the 61% to 62% of students whose parents gave consent are included—imply that the study sample is not representative of all students who attended the study's charter schools, which limits the study's external validity. In particular, students exempted from the lottery—largely consisting of those with some prior connection to the school—and those for whom parental consent was not obtained are not represented in the study sample.

The full student sample included 2,904 students—1,744 in the treatment group and 1,160 in the control group—from two study cohorts that were each followed over a 2-year follow-up period. For the main analysis, we further restricted the sample to a set of 2,330 students (1,400 treatment and 930 control) for whom we could most reliably estimate charter school impacts, by imposing two additional restrictions. First, as discussed above, we excluded charter school sites at which we were able to obtain school records data on student outcomes for fewer than 5 students in either the treatment or control group. This occurred either because the charter school held lotteries with either few lottery winners or few lottery losers or because the district and state in which the school was located declined to provide schools records data.

Second, we included only sample members for whom we obtained baseline data on student achievement. We did this to minimize differences in the availability of outcome data for treatment and control group students, as these differences could bias the impact estimates. Students without baseline achievement data were less likely to have attended a public school in the baseline year. Thus, they were also less likely to have attended a public school and have achievement data in the follow-up year if they lost the lottery and were not admitted to the charter school (but would likely attend the public charter school if they did win the lottery).[9] This restriction led us to drop 538 students from the analysis sample and is consistent with analyses of charter school impacts reported in most of the other lottery-based studies of charter schools.[10] However, results are not sensitive to the restriction, as described further below.

Table 2 displays baseline characteristics of treatment and control group students in the main analysis sample. As expected given that the admissions lotteries were random, treatment and control group students exhibited few statistically significant differences in baseline characteristics. Of the 33 characteristics in Table 2, there were statistically significant differences between the treatment and control groups for only two. Treatment group students had higher pre-baseline mathematics scores (scores from 2 years before the treatment group enrolled in the study schools) than control group students. However,

TABLE 2

Baseline Characteristics of Treatment and Control Group Students in Main Analysis Sample (Proportions Unless Otherwise Indicated)

	M, treatment group	*M*, control group	Difference in *M*	*SE* of difference
Achievement (z-score units)				
Baseline reading score	0.42	0.43	−.01	.440
Pre-baseline reading score	0.47	0.38	.09	.066
Baseline math score	0.45	0.45	.00	.049
Pre-baseline math score	0.47	0.32	.15*	.069
Disciplinary measures				
Number of days absent in baseline school year	6.07	5.62	.46	.295
Student suspended in baseline school year	0.04	0.03	.01	.009
Demographic characteristics				
White, non-Hispanic[a]	0.57	0.55	.02	.022
Black, non-Hispanic[a]	0.10	0.09	.00	.014
Other race, non-Hispanic[a]	0.07	0.08	−.01	.013
Hispanic	0.27	0.28	−.02	.017
Male	0.46	0.48	−.01	.025
Young for grade	0.005	0.008	−.003	.004
Old for grade	0.090	0.089	.001	.016
Has IEP	0.18	0.16	.02	.027
Limited English proficiency/ELL	0.10	0.08	.02	.013
Family characteristics (proportions)				
Income to poverty ratio 0%–100%	0.13	0.12	.01	.016
Income to poverty ratio 100%–200%	0.21	0.19	.02	.021
Income to poverty ratio 200%–300%	0.18	0.16	.02	.020
Income to poverty ratio >300%	0.49	0.54	−.05*	. 023
Two-parent family	0.78	0.79	−.01	.021
Not two-parent family, but more than one adult	0.05	0.04	.01	.011
English main language spoken at home	0.89	0.90	−.01	.013
Mother’s education: High school or less	0.23	0.24	−.01	.021
Mother’s education: Some college	0.35	0.35	.00	.025
Mother’s education: College	0.42	0.42	.00	.024
Born in the United States	0.92	0.92	.00	.015
Family received TANF or food stamps in past 12 months	0.05	0.05	.00	.011
Free or reduced price lunch—Eligible	0.34	0.35	.00	.022
School enrollment (proportions)				
Enrolled in charter school at baseline	0.05	0.06	−.01	.013
Enrolled in private school at baseline	0.00	0.01	.00	.004
Enrolled in public school at baseline	0.94	0.93	.01	.014
Homeschooled at baseline	0.01	0.00	.01	.004
Baseline school type unknown	0.00	0.00	.00	.002
Number of students[b]	1,400	930		
Number of sites	29	29		

Note. Sample includes students in main analysis sample (students with non-missing baseline test score data in the sites included in the main impact analyses). Means are estimated at the site level and averaged across sites, giving equal weight to each site. Estimates are weighted to account for differential probabilities of assignment to the treatment and control groups in each site. IEP = Individualized Education Plan; ELL = English language learner; TANF = Temporary Assistance for Needy Families.

[a]Race categories are mutually exclusive and may not equal 100% due to rounding.

[b]Sample size differs for some of the individual baseline characteristics due to differential rates of missing data for different characteristics.

*Difference significantly different from 0 at .05 level (two-tailed test).

treatment and control group students had identical mean mathematics scores in the baseline year. Treatment group students were also less likely (47% vs. 52%) to have family incomes above 30% of the poverty line. Two statistically significant differences are approximately what we would expect due to chance when examining differences in 33 characteristics with a 5% critical value. This suggests that the treatment and control groups in the main analysis sample were well balanced according to baseline characteristics, providing a strong foundation for the impact evaluation. Comparisons of the baseline characteristics of treatment and control group students among the full sample, including those without baseline test scores, showed that these two groups were also well balanced with respect to baseline characteristics (Online Appendix Table 1 available at http://epa.sagepub.com/supplemental), as did comparisons of the characteristics of treatment and control group students among the main analysis sample with valid Year 2 test score data (Online Appendix Table 2 available at http://epa.sagepub.com/supplemental), which was the sample that contributed to the main impact estimates on Year 2 reading and math.

Data

To measure the effects of charter schools on student achievement, the evaluation relied on test score data from state assessments. These data were obtained from schools, districts, or states for the baseline year and the preceding "pre-baseline year" as well as for the 2 follow-up years. Among members of our analysis sample, in Year 1, we obtained valid math scores for 94% of the treatment group and 89% of the control group, and valid reading scores for 95% of the treatment group and 89% of the control group. In Year 2, we obtained valid math scores for 90% of the treatment group and 84% of the control group, and valid reading scores for 91% of the treatment group and 84% of the control group.

Because sample members were spread across 13 states, each of which administered a different assessment, we converted all scores to *z* scores, defined as the student's raw score on the state assessment minus the mean score on the test among all students in the state who took the test, divided by the standard deviation of the scores for that same group, by grade level.[11] Thus, students' *z* scores reflect their performance on the state assessment relative to the typical student in that state and grade.

Additional covariates for the impact analysis were obtained from a baseline survey completed by parents when their children applied to a study charter school. The survey collected demographic and socioeconomic information from parents at the time of application, as well as their reasons for applying to the participating charter school and information on other schools to which they were applying. The overall response rate on the baseline survey among analysis sample members was 91%—92% among the treatment group and 90% among the control group.

Analytic Methods

Estimating the Impact of Charter School Attendance on Student Achievement—Main Approach

The goal of the analysis presented in this article is to estimate the impact of attending a charter school on student achievement in reading and math. Our main approach to this analysis is to estimate the impact of attending a charter school in each site in the study and then to average these impacts across sites. Because the decision to attend a charter school may be correlated with student achievement, we use results from the admissions lotteries as an instrumental variable for whether the student attends a charter school in the year following the lottery. While most (78%) of the students admitted to a charter school via lottery attended this school in the year following the lottery and a few attended a non-study charter school, 19% of lottery winners did not attend a charter school. Among lottery losers, 6% attended a study charter school and 9% attended another nearby charter school.

To estimate site-level impacts, we used the following regression model:

$$y_{ij} = \alpha_j + \beta' \boldsymbol{X_{ij}} + \delta_j C_{ij} + \varepsilon_{ij}, \qquad (1)$$

where y_{ij} is the test score outcome of interest for student *i* in site *j*[12]; α_j is a site-specific intercept; $\boldsymbol{X_{ij}}$ is a vector of characteristics of student *i* in site *j*, including an indicator for whether the student was in Cohort 1 or 2 of the sample; C_{ij} is a binary variable equal to one if student i in site j attended a charter school in the school year immediately following the lottery (whether or not that school was in our lottery sample), ε_{ij} is a random error term that reflects the influence of unobserved factors on the outcome; and $\boldsymbol{\beta}$ and δ_j are parameters or vectors of parameters to be estimated. The coefficient δ_j represents the impact of attending a charter school for students in site j. As noted above, observations were weighted to account for unequal selection probabilities in the charter school lotteries. Covariates included baseline and pre-baseline reading and math test scores; student absences and suspensions in the baseline year; demographic characteristics such as race/ethnicity, gender, and age; binary indicators for whether the student has an individualized education plan and limited English proficiency; household characteristics such as income, family structure, mother's education, home language, and public assistance receipt; and type of school attended at baseline—the full set of covariates is listed in Online Appendix Table 3 available at http://epa.sagepub.com/supplemental. Missing values of covariates were imputed as the mean value of the covariate by site and sample cohort, and missing value dummy variables were also included in the model.

Because the decision to attend a charter school may be correlated with student achievement in ways that are unobservable in our data, ordinary least squares estimates of Equation 1 may be biased. Thus, we used lottery admission as an instrumental variable for whether the student attended a charter school in the year following the admissions lottery. The first-stage equation for this instrumental variables estimation is a linear probability model which takes the following form:

$$C_{ij} = \lambda_j + \boldsymbol{\zeta}'\boldsymbol{X_{ij}} + \pi_j T_{ij} + \eta_{ij}, \tag{2}$$

where C_{ij} and $\boldsymbol{X_{ij}}$ are as defined above, λ_j is a site-specific intercept, T_{ij} is a binary variable equal to 1 if student i received an admissions offer via lottery from a study charter school in site j, η_{ij} is a random error term that reflects the influence of unobserved factors on the decision to attend a charter school, and $\boldsymbol{\zeta}$ and π_j are parameters or vectors of parameters to be estimated. The estimated coefficient on treatment status in site j, π_j, represents the impact of being admitted to a study charter school via lottery on the student's decision to attend a charter school. Because lottery admissions offers were randomly determined, instrumental variables estimates of δ_j in Equation 1 represent the causal effect on achievement of attending a charter school for a student in site j.[13]

To obtain an overall estimate of the average impact of attending a charter school on student achievement, we averaged the site-specific impact estimates $\hat{\delta}$ over the J sites included in the estimation, taking an equally weighted average as follows:

$$\hat{\delta} = \frac{1}{J}\sum_{j=1}^{J}\hat{\delta}_j. \tag{3}$$

By equally weighting the estimated impacts from each site, we allowed each impact to have an equal influence on the overall impact estimate, thereby providing unbiased estimates of attending a charter school in the average site in the study. However, we also tested the sensitivity of our results to our approach for calculating the average impact by first according more weight to more precisely estimated site-level impacts and then by using a risk-set approach that weights sites according to the number of students in the sample in each site, as described below.

Sub-Group Estimates

In addition to estimating overall effects of study charter school admission for the full study sample, we estimated the impact of attending a charter school for several population sub-groups. To estimate these impacts, we used the following regression model:

$$y_{ij} = a_j + \boldsymbol{B}'\boldsymbol{X_{ij}} + d_j C_{ij} + f_j S_{ij} + g_j C_{ij} S_{ij} + e_{ij}, \tag{4}$$

where S is an indicator for whether the student is in sub-group S, and all other variables are as defined as in Equation 1. Following the same approach as for our main model, we used lottery admission interacted with sub-group as an instrumental variable for whether the students in each sub-group attended a charter school in the year following the admissions lottery. The estimated

TABLE 3
Charter School Impacts on Student Achievement

	Main model			Sensitivity tests of 2SLS results			
	First stage impact of winning lottery on attending a charter school in Year 1	Reduced form impact of winning lottery of student achievement (ITT impact estimate)	2SLS impact of attending a charter school on student achievement	No covariates	Include students with missing baseline scores, non-response weights	Weight site-level impact estimates by their inverse variance in computing average	Risk-set model
Reading							
Year 1 reading score	.698**	−.037	−.058	−.047	−.025	−.037	−.028
SE	.019	.032	.051	.082	.054	.042	.042
Number of students	2,141	2,141	2,141	2,141	2,371	2,141	2,141
Year 2 reading scores	.699**	−.059	−.070	−.040	−.106	−.093*	−.095*
SE	.019	.032	.050	.076	.055	.043	.048
Number of students	2,032	2,032	2,032	2,032	2,261	2,032	2,032
Math							
Year 1 math scores	.697**	−.047	−.076	−.022	−.115	−.073*	−.066
SE	.019	.029	.047	.080	.064	.037	.040
Number of students	2,120	2,120	2,120	2,120	2,351	2,120	2,120
Year 2 math scores	.699**	−.058	−.065	−.011	−.081	−.096*	−.095
SE	.019	.039	.067	.091	.073	.049	.053
Number of students	2,027	2,027	2,027	2,027	2,252	2,027	2,027

Note. Estimates are based on regression models that, in addition to charter school admission, include the covariates listed in Online Appendix Table 3 available at http://epa.sagepub.com/supplemental. In all but the risk-set model, impacts are estimated at the site level and averaged across sites. Test scores were standardized across states by converting to z scores (raw scores minus the state mean score for that subject and grade, divided by the standard deviation of scores for that subject and grade), and impact estimates represent charter schools' effects on student scores expressed in terms of statewide standard deviations of scores for the student's grade. Reported sample sizes indicate the number of students or sites with non-missing data for at least one of the four test score outcomes. Sample sizes vary for individual outcomes. 2SLS = two-stage least squares; ITT = intent to treat. *Difference between lottery winners and losers is statistically significant at the .05 level (two-tailed test). **Difference between lottery winners and losers is statistically significant at the .01 level (two-tailed test).

coefficient on charter school attendance, d_j, provides an estimate of the impact of attending a charter for students not in sub-group S in site j, and the estimated coefficient on charter school attendance interacted with sub-group, g_j, represents the difference in impacts between students in sub-group S and students not in sub-group S in site j. Summing d_j and g_j thus provides an estimate of the impact for students in sub-group S in site j. We then averaged the impact estimates for each sub-group across all sites to obtain an overall impact estimate for that sub-group (following the same approach used to average impact estimates for the full sample in Equation 3).

Sensitivity Analyses

To assess the sensitivity of our main estimates to the specific estimation method used, we also estimated impacts using several alternative approaches, described below.

Exclusion of Covariates. Our main model controlled for baseline student test scores and other baseline student characteristics. Controlling for baseline characteristics improves the precision of the impact estimates. However, as noted by Freedman (2008), theory suggests that inclusion of baseline covariates may bias impact estimates, although in practice this bias tends to be small (Schochet, 2010). To assess the sensitivity of our models to inclusion of baseline covariates, we estimated models that did not include any covariates other than site fixed effects and site treatment status interactions.

Inclusion of Students With Missing Baseline Test Scores. As described above, to minimize the possibility of bias attributable to differential rates of missing test score outcome data between the treatment and control groups, we limited the sample to students with valid baseline test score data. Such students were more likely to have non-missing follow-up test scores regardless of admission to a study charter school. As an alternative to accounting for missing outcome data, we estimated impacts by using data from all sample members, regardless of whether they had valid baseline test scores, and adjusted for differential rates of missing outcome data by using non-response weights.[14]

Giving Greater Weight to More Precisely Estimated Site-Level Impacts. To obtain our main impact estimates, we computed an equally weighted average of the site-level impact estimates (Equation 2). Thus, sites with estimated impacts based on relatively small samples received equal weight as sites with impacts based on large samples. To test the sensitivity of our results to this approach for weighting site-level impact estimates, we also estimated average impacts by weighting each site-level impact estimate by its inverse variance.

Estimating Impacts Using a Risk-Set Approach. Our main model estimated site-level impacts and averaged across sites, using sample weights to adjust for differential probabilities of admission. However, another common approach in the literature is the use of a risk-set model in which the outcome variable is regressed on an indicator for charter school attendance and lottery-specific risk sets—a set of fixed effects indicating the lottery or lotteries to which each student applied (Abdulkadiroglu et al., 2011; Angrist et al., 2012; Dobbie & Fryer, 2011).[15] Estimates from the risk-set approach have a slightly different interpretation than estimates from our main model. The risk-set model implicitly weights each lottery-level impact estimate according to the number of sample members who participated in that lottery, while under our main approach we give equal weight to each study site. Thus, the risk-set estimates reflect the impact of charter school attendance on the average student in the study, while estimates from our main model reflect the impact of charter schools in the average site in the study.

The Impact of Attending a Charter School on Student Achievement

On average, estimated impacts of attending a charter school on student achievement in the 2 years following the admissions lotteries were negative but not statistically significant in our main model (Table 3). Estimates from our main two-stage least squares model indicate that attending a charter school in the year following the admissions lottery lowered student reading achievement by .058 standard deviations and student math achievement by .076 standard deviations in the school year immediately following the lottery, and lowered achievement by similar levels (.070 standard deviations in reading and .065 standard deviations in math) in the second school year after the lottery. Results are similar in the sensitivity tests we conducted. Estimates for all two-stage least squares specifications are negative and of similar magnitude, ranging from −.011 to −.115 standard deviations, and are generally not statistically significant, except in the model that weights site-level impact estimates by the inverse variance of the estimate in each site in computing the average or, in the case of Year 2 reading scores, the risk-set model.

As an alternative measure of student achievement, we also examined impacts on the proportion of students achieving various proficiency levels on their state assessments in reading and math in Years 1 and 2. These results, shown in Online Appendix Table 4 available at http://epa.sagepub.com/supplemental, indicate that charter school attendance had little effect on students' proficiency rates. Two-stage least squares estimates of the impact of charter school attendance on the percentage of students scoring in the most advanced proficiency category in their state were all negative, ranging from −2 percentage points in Year 1 reading and Year 2 math to −5 percentage points in Year 1 math, but none were statistically significant. Impacts on the percent of students scoring in their state's "proficient" category or higher or in their state's "partially proficient" category or higher were small and not statistically significant, and varied in sign across outcome variables.

We also estimated impacts for sub-groups of students (Table 4). There are no statistically significant differences (or clear pattern of differences) across sub-groups defined by students' race (White

TABLE 4

2SLS Impact Estimates for Sub-Groups of Students

	2SLS impact estimate	*SE*	2SLS impact estimate	*SE*	Difference in impacts	*SE*
	Non-White and/or Hispanic		White, non-Hispanic		Difference between sub-groups	
Reading achievement						
Year 1	−.032	.059	.012	.043	.044	.072
Year 2	−.090	.068	−.091*	.044	−.001	.080
Math achievement						
Year 1	.009	.050	−.112*	.041	−.121	.064
Year 2	−.042	.083	−.115*	.053	−073	.098
Number of students	1,026		1,099			
	Female		Male		Difference between sub-groups	
Reading achievement						
Year 1	.001	.043	−.026	.050	−.027	.066
Year 2	−.077	.041	.027	.053	.104	.067
Math achievement						
Year 1	−.054	.037	−.008	.041	.046	.055
Year 2	−.092	.048	.023	.058	.115	.074
Number of students	1,127		1,003			
	Certified for free or reduced price lunch		Not certified for free or reduced price lunch		Difference between sub-groups	
Reading achievement						
Year 1	−.068	.058	−.025	.043	.043	.072
Year 2	.050	.060	−.124**	.042	−.174*	.074
Math achievement						
Year 1	.063	.051	−.135**	.036	−.198**	.062
Year 2	.179**	.057	−.135*	.054	−.315**	.078
Number of students	770		1,333			
	Baseline reading achievement below median		Baseline reading achievement above median		Difference between sub-groups	
Reading achievement						
Year 1	−.093*	.046	−.026	.046	.067	.066
Year 2	−.017	.049	−.148**	.049	−.131	.069
Math achievement						
Year 1	−.017	.036	−.044	.045	−.027	.057
Year 2	.050	.050	−.148**	.056	−.198**	.073
Number of students	1,089		1,025			
	Baseline math achievement below median		Baseline math achievement above median		Difference between sub-groups	
Reading achievement						
Year 1	−.003	.042	−.084	.046	−.081	.063
Year 2	.013	.042	−.110*	.048	−.123	.064
Math achievement						
Year 1	−.041	.037	−.046	.045	−.004	.058
Year 2	.055	.051	−.116*	.054	−.171*	.074
Number of students	1,078		1,051			

Note. 2SLS = two-stage least squares.
*Difference between lottery winners and losers is statistically significant at the .05 level (two-tailed test). **Difference between lottery winners and losers is statistically significant at the .01 level (two-tailed test).

non-Hispanic vs. non-White or Hispanic) or gender. However, estimated impacts in Year 2 math achievement were positive for more disadvantaged students as measured by certification for free or reduced price lunch, and large and negative for more advantaged students (for both subjects in Year 1 as well as math in Year 2). These differences between the two groups of students were statistically significant at the 5% level.[16] These same patterns persisted for Year 2 math scores for sub-groups defined by students' baseline achievement in reading or math (defined by whether the student scored above or below the sample median on the respective test). In other words, impacts were significantly more positive (or less negative) for students who were lower achieving at baseline than for those who were higher achieving.

Exploring Variation in Charter School Impacts

While the overall average impacts of the study charter schools were negative and not statistically significant, impacts varied across charter schools in the study. Figure 1 presents the distribution of estimated impacts on Year 1 reading and mathematics scores across study charter school sites, arranged by magnitude of impact. The figure shows substantial variation in the impacts. Impacts on Year 1 reading z scores ranged from −.96 to +.76, with a standard deviation of .33. One estimated impact was statistically significant and negative and one statistically significant and positive, with the remainder not significantly different from zero at the 5% level. Impacts on Year 1 mathematics z scores ranged from −.82 to +.54, with a standard deviation of .31. Four of the site-level estimated impacts were statistically significant and negative and two were statistically significant and positive, with the remainder not significant. F tests reject the null hypothesis that impacts are equal across sites, with a p value of .009 for Year 1 reading and .000 for Year 1 math.

To further investigate the circumstances under which charter schools are more or less effective relative to nearby public schools, we estimated impacts for several sub-groups of charter schools in the sample (Table 5). Consistent with the findings for the student sub-group analysis, these results showed that charter schools serving a high proportion of students certified for free or reduced price lunch had a positive and significant impact on Year 2 math achievement, whereas charter schools serving a low proportion of these students have negative and significant impacts on math and reading achievement. Differences across these sub-groups were statistically significant for math achievement in Years 1 and 2. Similarly, schools serving high proportions of students with low baseline achievement have more positive impacts than those serving a lower proportion of these students—these differences were statistically significant for reading and math in Year 2.

Impacts for charter schools in urban areas were similar to those in non-urban areas for reading achievement, but impacts in math were different for urban and non-urban schools. Charter schools in urban areas had positive and significant impacts in math achievement (in Year 2), whereas non-urban charter schools had negative and significant impacts. This pattern of more positive impacts in charter schools serving more disadvantaged students and in urban areas and more negative impacts in charter schools serving more advantaged students and non-urban areas is consistent with findings from previous lottery-based studies (Angrist et al., 2012; Dobbie & Fryer, 2011; Hoxby et al., 2009).

Because impacts are measured relative to the counterfactual in each site—the schools charter school students would have attended had they not attended a charter school—variation in impacts could reflect variation in the effectiveness of the charter schools as well as variation in the effectiveness of the non-charter schools available to students in that area. The more positive impacts among charter schools serving more disadvantaged students could indicate that these schools are more effective than charter schools serving more advantaged students, but they could also indicate that the non-charter schools available to less advantaged students are weaker than those available to more advantaged students. We are not able to rigorously disentangle these two possibilities in our data, but regardless of the reasons, our results show a clear pattern of charter school benefits for more disadvantaged students and a clear pattern of negative impacts for more advantaged students.

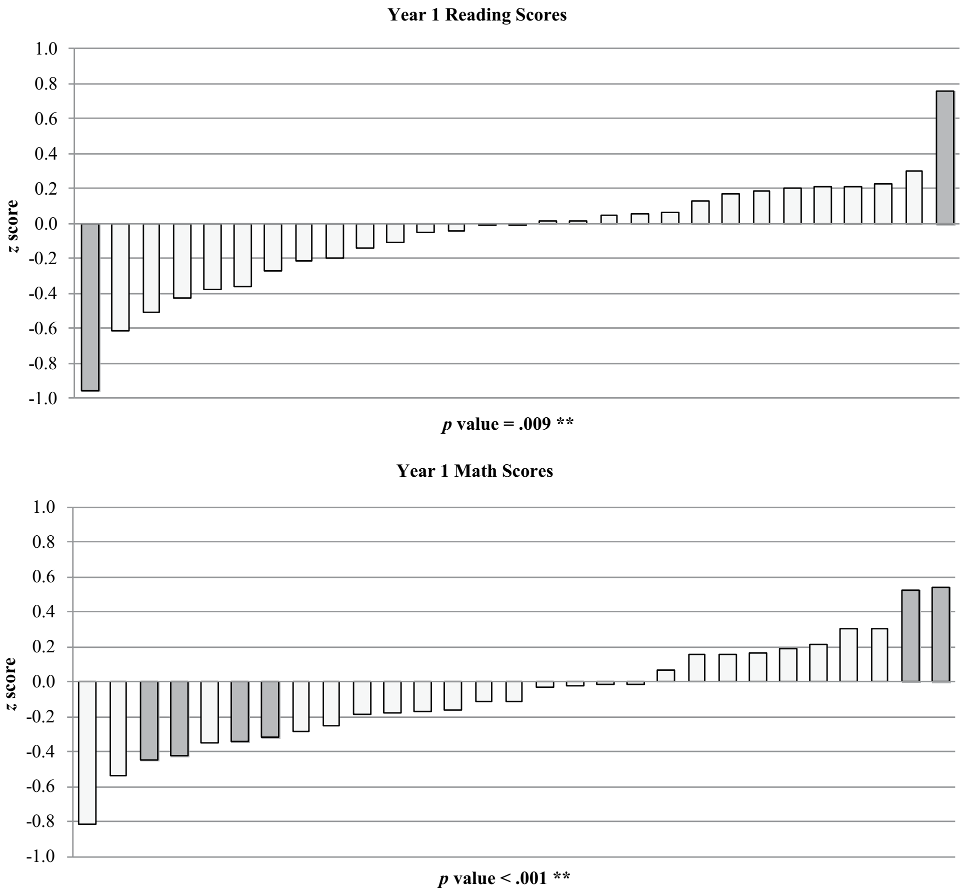

FIGURE 1. *Distribution of site-level two-stage least squares impact estimates.*
Note. Shaded bars are statistically significant impacts at the .05 level (two-tailed test). p values are from F tests of the null hypothesis that impacts are equal across sites.
*Variation in impacts is statistically significant at the .05 level (two-tailed test). **Variation in impacts is statistically significant at the .01 level (two-tailed test).

Conclusion

This article presents results from a large-scale, lottery-based study of charter schools, covering 33 charter middle schools in 13 states. We found that, on average, the charter schools in the study had negative but non-significant impacts on student achievement in reading and math. Impacts generally did not vary across sub-groups defined by students' race or gender. However, impacts were insignificant or positive for more disadvantaged students and negative and significant for more advantaged students, and this same pattern persisted across groups defined by baseline test scores. There was also considerable variation in impacts across schools, with some charter schools in the study positively affecting student achievement and others negatively affecting achievement. Charter schools in urban areas or serving more disadvantaged populations had more positive (or less negative) impacts than those in non-urban areas or serving more advantaged populations. These patterns could indicate that charter schools serving more disadvantaged students are more effective than those that serve more advantaged students but could also indicate that the counterfactual schools available to more advantaged students are weaker than those available to more disadvantaged students, leading to greater opportunity for positive charter school impacts among more disadvantaged populations.

It is important to keep in mind that charter schools were not randomly selected for the study,

TABLE 5
2SLS Impact Estimates for Sub-Groups of Sites

	2SLS impact estimate	*SE*	2SLS impact estimate	*SE*	Difference in impacts	*SE*
	High percent eligible for free or reduced price school meals		Low percent eligible for free or reduced price school meals		Difference between sub-groups	
Reading achievement						
Year 1	−.065	.045	−.022	.046	.044	.064
Year 2	.005	.044	−.109*	.044	−.114	.061
Math achievement						
Year 1	.040	.043	−.116**	.038	−.155**	.057
Year 2	.190**	.055	−.238**	.053	−.428**	.076
Number of students	1,141		1,006			
	Average baseline reading and math achievement in site below median		Average baseline reading and math achievement in site above median		Difference between sub-groups	
Reading achievement						
Year 1	.015	.048	−.086	.045	−.101	.067
Year 2	.062	.045	−.155**	.044	−.217**	.062
Math achievement						
Year 1	.018	.043	−.098**	.038	−.116*	.057
Year 2	.177**	.057	−.228**	.051	−.405**	.076
Number of students	1,017		1,133			
	Low percent White		High percent White		Difference between sub-groups	
Reading achievement						
Year 1	−.096*	.041	.018	.051	.115	.066
Year 2	−.093*	.040	−.021	.049	.072	.063
Math achievement						
Year 1	−.050	.038	−.042	.042	.008	.057
Year 2	.002	.049	−.098	.060	−.100	.076
Number of students	1,309		841			
	Urban		Not urban		Difference between sub-groups	
Reading achievement						
Year 1	−.042	.052	−.041	.041	.001	.066
Year 2	−.018	.052	−.076	.040	−.058	.066
Math achievement						
Year 1	.072	.055	−.099**	.033	−.171**	.065
Year 2	.164*	.072	−.140**	.046	−.304**	.086
Number of students	678		1,472			

Note. 2SLS = two-stage least squares.
*Difference between lottery winners and losers is statistically significant at the .05 level (two-tailed test). **Difference between lottery winners and losers is statistically significant at the .01 level (two-tailed test).

the number of study schools is small relative to the total population of charter middle schools, and the resulting sample is thus not nationally representative. The study included only oversubscribed charter schools that held admissions lotteries, and impacts for these schools could be systematically more positive than impacts of charter schools that are not oversubscribed. In addition, the study focuses on charter middle schools and does not produce evidence on the effects of charter schools at the elementary or high school levels. Similarly, our finding that the study charter schools in urban areas had more positive (or less negative) impacts than the study charter schools in non-urban areas does not imply that any charter school opened in an urban area will have positive impacts on student achievement—results only apply to the particular set of charter and non-charter schools in our study. Despite these limitations, our findings add significantly to the growing empirical evidence base on this important aspect of educational reform and management.

Authors' Note

This study was conducted for the Institute of Education Sciences of the U.S. Department of Education under Contract No. ED-01-C0039/0010.

Acknowledgment

We thank Hanley Chiang, Tom Cook, Brian Gill, Jay Greene, Laura Hamilton, Rick Hanushek, Paul Hill, Helen Ladd, Rebecca Maynard, Robert Meyer, Larry Orr, and Peter Schochet for their helpful comments and feedback.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. National Alliance for Public Charter Schools. "Details from the Dashboard." Available at http://www.publiccharters.org/wp-content/uploads/2014/02/New-and-Closed-Report-February-20141.pdf

2. Exceptions to the findings of negative or insignificant effects in the fixed-effects literature include Witte, Weimer, Shober, and Schlomer (2007), who found positive impacts in Wisconsin, and Ballou, Teasley, and Zeidner (2008), who found positive impacts for charter elementary schools but no statistically significant impacts for charter middle schools in Idaho.

3. In addition, lottery-based studies focus only on oversubscribed charter schools, popular schools that have more applicants than available seats. These oversubscribed charter schools may be more effective than non-oversubscribed charter schools, if applicants choose to apply there on the basis of an accurate assessment of school quality. Angrist, Parthak, and Walters (2012) presented some evidence that lottery schools in Massachusetts had more positive impacts than schools not holding a lottery.

4. The research presented here was part of an evaluation of charter schools conducted by Mathematica Policy Research for the U.S. Department of Education's Institute of Education Sciences (Gleason, Clark, Tuttle, & Dwoyer, 2010). Data used in the analyses are available in a restricted use file which researchers can request from the U.S. Department of Education's National Center for Education Statistics (NCES) through its Electronic Application System, available at http://nces.ed.gov/statprog/instruct.asp. In accordance with NCES publication policy, sample sizes from analyses presented in this article that were not previously reported in Gleason et al. (2010) are rounded to the nearest 10.

5. In general, each site corresponded to a single charter school. However, five pairs of participating charter schools had common applicants to their lotteries—we refer to these as "dual applicants." We treated these pairs of schools as single, combined sites in the analysis. (If a pair of schools had common applicants in one cohort but not the other, they were treated as a single site in the cohort in which they shared applicants and as individual sites in the other cohort.)

6. The "other" charter middle schools—those not participating in the study—include charter schools that did not receive enough applications to hold a lottery, that held a lottery but ended up offering admission to most or all of the lottery losers who ended up on a waiting list, and that refused to participate in the study. The response rate to the principal survey among these schools was 70%.

7. Lottery stratification occurs when schools first group students into categories or strata, according to some characteristic, and then conduct separate lotteries for each group. This usually results in a different probability of admission for each group of students. Schools sometimes use "sibling rules" to ensure that students from the same family who apply at the same time end up with the same lottery outcomes—that is, are either both admitted or both not admitted. For example, some schools will enter each sibling separately in the lottery but if one sibling is admitted through the lottery, the other is automatically offered admission to the school regardless of his or her lottery draw.

8. Without using sample weights to account for students' probability of admission, particular students may have an undue influence on the estimated treatment or control group means and on the estimated impacts. For example, as students who apply to more than one study school would have a higher probability of admission than those who apply to just a single school, all else equal, these "dual applicants" would have been more heavily represented in the treatment group than in the control group in an unweighted analysis.

9. More than half (52%) of the students without baseline achievement data attended a private school or were homeschooled when they applied to a study charter school, compared with less than 1% of those with baseline achievement data. Among those who attended a private school or were homeschooled when they applied to the charter school, 90% of treatment group students attended a public school (typically the study charter school) during the first follow-up period, compared with only 34% of control group students.

10. In their lottery-based study of charter schools in Boston, for example, Abdulkadiroglu, Angrist, Dynarski, Kane, and Pathak (2011) used a similar sample restriction. Hoxby, Murarka, and Kang (2009) restricted the sample upon which their impact estimates were based to students with some test score availability, although they allowed this to be either in the baseline or follow-up period. The non-experimental fixed-effects studies of charter school impacts that compare test scores of students in charter schools with their test scores prior to their entry into a charter school also restrict the sample to those with valid achievement data during a baseline period (e.g., Bifulco & Ladd, 2006; Hanushek, Kain, Rivkin, & Branch, 2007; Sass, 2006; Zimmer, Gill, Booker, Lavertu, & Witte, 2012).

11. This approach for analyzing state assessment data in educational studies involving multiple states is one of the approaches recommended by a recent report on the use of state tests in education experiments (May, Perez-Johnson, Haimson, Sattar, & Gleason, 2009). It is also similar to the approach used by two other recent multistate studies of charter school impacts (Center for Research on Educational Outcomes [CREDO], 2013; Zimmer et al., 2012).

12. In most cases, a site refers to a single charter middle school participating in the study. In several cases, however, two study schools were close to one another and some students in the sample applied to both of these schools. In these cases, the site consists of the two study schools that share some student applicants.

13. We also estimated a reduced form version of Equations 1 and 2, in which the test score outcome was regressed on the variables in $\boldsymbol{X}$ and the lottery outcome (T). The estimate of the coefficient on T from this reduced form model is the "intent-to-treat" (ITT) estimate of a lottery offer of admission to a charter middle school on student achievement.

14. In particular, the non-response weights accounted for differences between the characteristics of sample members for whom we have outcome data versus those for whom we do not have outcome data.

15. These studies estimate the impact of years of charter school attendance on student achievement (using lottery admissions as an instrumental variable for years of attendance), whereas we estimate the impact of attending a charter school in the year immediately following the lottery. We opted to follow the latter approach because the former approach assumes that the effects of years of attendance are linear, an assumption that was not supported in our data.

16. This pattern of impacts is not simply a function of the particular charter schools attended by large numbers of disadvantaged students in the sample, as sub-group estimates were computed in each site and the overall estimate for each sub-group was computed as an equally weighted average of the site-level estimates. Thus, they suggest that on average, the charter schools in the study had more positive impacts for more disadvantaged students than for more advantaged students.

References

Abdulkadiroglu, A., Angrist, J. D., Dynarski, S. M., Kane, T. J., & Pathak, P. (2011). Accountability and flexibility in public schools: Evidence from Boston's charters and pilots. *Quarterly Journal of Economics*, *126*, 699–748.

Angrist, J. D., Parthak, A. P., & Walters, C. R. (2012, August). *Explaining charter school effectiveness* (NBER Working Paper Series No. 17332). Cambridge, MA: National Bureau of Economic Research.

Ballou, D., Teasley, B., & Zeidner, T. (2008). Charter schools in Idaho. In M. Berends, M. G. Springer, & H. J. Walberg (Eds.), *Charter school outcomes* (pp. 221–241). New York: Lawrence Erlbaum.

Bifulco, R., & Ladd, H. F. (2006). The impact of charter schools on student achievement: Evidence from North Carolina. *Education Finance and Policy*, *1*, 50–90.

Booker, T. K., Gilpatric, S., Gronberg, T., & Jansen, D. (2007). The impact of charter school attendance on student performance. *Journal of Public Economics*, *91*, 849–876.

Center for Research on Educational Outcomes. (2013). *National charter school study: 2013*. Stanford, CA: Center for Research on Educational Outcomes, Stanford University.

Dobbie, W., & Fryer, R. G., Jr. (2011). Are high-quality schools enough to increase achievement among the poor? Evidence from the Harlem children's zone. *American Economic Journal: Applied Economics*, *3*, 158–187.

Dobbie, W., & Fryer, R. G., Jr. (2013, July). *The medium-term impact of high-achieving charter schools on non-test score outcomes* (NBER Working Paper No. 19581). Cambridge, MA: National Bureau of Economic Research.

Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics*, *40*, 180–193.

Gleason, P., Clark, M., Tuttle, C. C., & Dwoyer, E. (2010, June). *The evaluation of charter school impacts: Final report* (NCEE 2010-4029). Washington, DC: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, U.S. Department of Education.

Hanushek, E. A., Kain, J. F., Rivkin, S. G., & Branch, G. F. (2007). Charter school quality and parental decision making with school choice. *Journal of Public Economics*, *91*, 823–848.

Hoxby, C. M., Murarka, S., & Kang, J. (2009, September). *How New York City's charter schools affect student achievement: August 2009 Report* (Second report in series). Cambridge, MA: New York City Charter Schools Evaluation Project.

Hoxby, C. M., & Rockoff, J. E. (2005). Findings from the city of big shoulders. *Education Next*, *5*(4), 52–58.

May, H., Perez-Johnson, I., Haimson, J., Sattar, S., & Gleason, P. (2009, November). *Using state tests in education experiments: A discussion of the issues* (Final report). Princeton, NJ: Mathematica Policy Research.

Sass, T. R. (2006). Charter schools and student achievement in Florida. *Education Finance and Policy*, *1*, 91–122.

Schochet, P. Z. (2010). Is regression adjustment supported by the Neyman model for causal inference? *Journal of Statistical Planning and Inference*, *140*, 246–259.

Witte, J. F., Weimer, D. L., Shober, A., & Schlomer, P. (2007). The performance of charter schools in Wisconsin. *Journal of Policy Analysis and Management*, *26*, 557–573.

Zimmer, R., Gill, B., Booker, K., Lavertu, S., & Witte, J. (2012). Examining charter school achievement effects across seven states. *Economics of Education Review*, *31*, 213–224.

Authors

MELISSA A. CLARK (mclark@mathematica-mpr.com) is a senior researcher at Mathematica Policy Research. She specializes in the design and implementation of rigorous impact evaluations, with a focus on education policy including highly selective alternative teacher certification programs such as Teach for America and principal professional development.

PHILIP M. GLEASON (pgleason@mathematica-mpr.com) is a senior fellow at Mathematica Policy Research. His research has focused on education policy, including charter schools and teacher effectiveness, and research on federal nutrition programs.

CHRISTINA CLARK TUTTLE (ctuttle@mathematica-mpr.com) is a senior researcher at Mathematica Policy Research. Her work focuses on the design and implementation of rigorous impact evaluations of educational interventions, particularly involving random assignment.

MARSHA K. SILVERBERG (marsha.silverberg@ed.gov), an economist, is a team leader at the U.S. Department of Education's Institute of Education Sciences (IES), National Center for Education Evaluation. She has more than two decades of experience designing and conducting rigorous evaluations of federal programs and federally supported strategies, particularly in the areas of college and career transitions and school choice.

Manuscript received April 17, 2014
Revision received July 13, 2014
Accepted September 29, 2014