

Findings from the City of Big Shoulders

by CAROLINE M. HOXBY and JONAH E. ROCKOFF

- home
 - unabridged
 - past issues
 - search
 - about
 - single issues
 - subscriptions
- read the unabridged article

The number of charter schools has grown very rapidly in the United States, from essentially none in 1990 to more than 3,400 today. Supporters believe that the flexibility granted these new public schools allows them to be more innovative and responsive to student needs than traditional public schools are. And the fact that no student attends a charter school unless his parents want to keep him there means that families can “vote with their feet.” When a parent leaves a charter, so does the funding associated with his child. Thus a charter school cannot survive without satisfied parents. But charter schools do not just answer to parents; they must also persuade an authorizer to recharter them every few years, and they must participate in statewide testing and accountability. Will this concoction of flexibility, answering to parents, and accountability to the government raise school quality? Bluntly put, do students in charter schools learn more than their counterparts in traditional public schools? More than they would have learned had they stayed put?



Younger Students Learn More in Charter Schools

A Lottery-Based Evaluation of Charter Schools

Getting a reliable answer to these questions is vital to the current policy debate, but researchers who try to answer them face considerable obstacles. First and foremost, most charter schools are new and small. They just don't yet have enough results for researchers to draw conclusions. Second, although all charter schools share the features mentioned above, they are otherwise a diverse lot. Many set up shop in urban areas, serve minority and low-income students, and rely on a strategy and curriculum associated with an education management organization. However, some charter schools serve very rural, mostly white students. Some are run as start-ups by parent or community groups that do not associate themselves with a particular strategy or curriculum. Even within the world of education management organizations, approaches to learning can differ substantially. In short, an assessment of some charter schools is useful for learning about similar charter schools, but we should not expect it to inform us about all charter schools.

Even when researchers can evaluate charter schools that are large enough to contribute useful results to a study, old enough to have a track record, and representative of a substantial share of all charter schools, they face a daunting analytical challenge: finding students in the regular public schools who are truly comparable to the charter school students. Students who apply to attend charter schools are a self-selected group, and simply comparing them with all other students in local public schools is likely to be misleading. We do not even know whether to expect self-selection to work for or against charter schools. On the one hand, parents who try out charter schools may be especially motivated. On the other hand, parents whose children are doing well may avoid being “guinea pigs” in relatively untried schools.

In our study, we overcome this challenge by exploiting a feature common to most charter schools: the lottery that schools use to admit students when they have more applicants than spaces. Such lotteries present an opportunity, rare in education, to conduct randomized experiments of the type more commonly used in medical research.

We use this lottery-based approach to evaluate three schools managed by the Chicago Charter School Foundation (CCSF). Our treatment group (those who, in medicine, would receive the pill) comprises charter school applicants who drew a lottery number that earned them a place at one of the charter schools (lotteried in). Our control group (those who would receive the placebo) comprises the applicants who were lotteried out. All told, the study focuses on 2,448 students who are divided between the lotteried-in and lotteried-out groups. It's important to realize that all of the students in the study applied to charter schools, so self-selection is the same for all of them. All that distinguishes the groups is their randomly drawn lottery numbers, so we can be confident that the

groups are comparable not only in observable ways (like race and income), but also in less tangible ways, such as motivation to succeed. Currently, we can compare the progress of both groups for up to four years following their application. We are continuing the study and will report further results as they become available.

Our results to date, which indicate clear positive effects of attending a charter school on the math and reading test scores of students who enter charter schools in kindergarten through 5th grade, represent the most credible evidence yet available on how charter schools affect student achievement. They are also uniquely informative for policymaking. In the long run, as charter schools become more established, almost all of their students will have entered in the early grades. Policymakers should therefore assign greater weight to studies that focus on such students than they do to studies that, because they lack experimental data, must focus on atypical students who enter charter schools when they are older.

The Chicago Charter School Foundation

Chicago is home to almost all the charter school students (8,817 of 9,980) in Illinois. Charter schools in Illinois are free to establish their own missions and curricula, but they participate in the state accountability system and must abide by personnel restrictions similar to those of regular public schools. In Chicago, charter schools receive a per-student fee equal to only 75 percent of the average per-pupil operating spending in traditional public schools. For the 2003–04 school year, it was \$5,279.

The Chicago Charter School Foundation is a charitable organization that has been operating since 1997. It oversees five primary schools, one high school, and one K–12 school. Together, its schools enroll more than half of Chicago’s charter school students. Seats in the charter schools are in demand. In the spring of 2004, CCSF schools had 2.4 applicants for every student they could admit. Most CCSF schools are run by nonprofit education management organizations, but one is run by a for-profit organization.

Our current report relies on CCSF’s oldest schools, all of which have been in operation since the late 1990s and have produced enough results to be evaluated: Longwood (K–12), with 1,200 students in Washington Heights; Bucktown (K–8), with 600 students in Logan Square; and Prairie (K–8), with 350 students in Roseland. Longwood is run by Edison Schools, which is for-profit, while Bucktown and Prairie are operated by the nonprofit American Quality Schools. Although these education management organizations differ somewhat, their strategies are fairly typical of organizations geared toward urban, disadvantaged children. They feature a structured school day and curriculum, combined with a family-oriented approach designed to get parents involved.

The charter schools we study are all located in neighborhoods where the population is disproportionately minority and poor, but the schools are not alike. Longwood is in a very black neighborhood, and 99 percent of its students are black. Bucktown and Prairie are in neighborhoods that are mixed ethnically, but they draw students who are disproportionately likely to be Hispanic and in need of bilingual education (see Figures 1a through 1c).



Note: Nearby public schools are schools within a three-mile radius of the respective charter school, with the exception of Prairie which draws from an Hispanic area with a smaller radius.

SOURCES: Consortium on Chicago School Research data and National Center for Education Statistics data

The charter school students are about as likely to be eligible for special education and for the free or reduced-price lunch program as are students in the regular Chicago public schools. It's very important to use the regular public schools' classifications of students into lunch program, special education, and bilingual education. Otherwise, the classifications could reflect differences in how often the charter schools place students in these programs rather than their students' traits.

The effects of attending a charter school reported in any study can only safely be extrapolated to students and schools like those included in the study. The students in our study are urban, dominated by racial and ethnic minorities, and largely disadvantaged. All of the students in our study applied to a charter school, so our results pertain to students who want to attend charter schools. Of course, these are precisely the students in whom policymakers are interested. No one suggests that students who do not want to attend charter schools should be forced to enroll in them, so learning whether they would have done better or worse in such schools is irrelevant.

The CCSF Lotteries

The charter school lotteries we study are pretty standard. A separate lottery is held for each school and grade. For example, if Bucktown has 60 kindergarten places available for 120 applicants and five 2nd-grade places available for 25 applicants, there would be a kindergarten lottery and a 2nd-grade lottery. After a charter school's first year of operating a particular grade, it is normal for the most places to be available in kindergarten. In each lottery, applications are assigned a random number and ordered according to it. Using this ordering, the places available in each grade in each school are filled. (If a student is lotteried in, then his or her siblings are also automatically granted a place if they apply in a subsequent year and in a grade for which there is space available.)

In this article, we focus on students who participated in the lotteries held in spring 2000, 2001, and 2002. The Consortium on Chicago School Research generously agreed to match as many of these students as was possible to the Chicago Public Schools' student database using their names, dates of birth, and the school and grade they reported attending when they applied. These data provide us with information on achievement, as measured by the Iowa Tests of Basic Skills (ITBS), before students applied and, even more crucially, with post-application achievement data for students who remained in Chicago's regular public schools.

All students who enrolled in a charter school were matched to a Chicago Public Schools record, as

were 73 percent of the charter school applicants who applied but did not enroll. We ultimately limit our analysis to the 2,448 of these students who applied from a Chicago public school or applied to kindergarten (and thus were not in any school when they applied). We do this because the correct comparison for a student who applies from a private school is a lotteried-out student who would not appear in the Chicago Public Schools database. Our results should therefore be interpreted as the effect of attending a CCSF charter school on students who would otherwise be attending a regular public school, not the effect on students who would otherwise be attending a private school.

Enrollment in Regular Public Schools

One oft-stated concern about charter schools is that they will draw away the highest-achieving students from the regular public schools around them. We can address this issue by comparing the prior test scores of charter school applicants in our data with the test scores of students in regular public schools in their neighborhoods (within three miles). This exercise assumes that students would attend local schools if charter schools did not exist. If this is basically correct, the comparisons give a sense of how a charter school's existence affects regular public schools around it.

Longwood's applicants, before applying to the charter school, had similar reading scores but lower math scores (5 percentile points lower) than other students in neighboring regular public schools. Bucktown's applicants had similar reading scores but lower math scores (7 percentile points lower) compared with students in neighboring regular public schools. Applicants to Prairie score about the same in math as students in the neighboring regular public schools, but their reading scores are higher (4 percentile points higher). As we mentioned above, however, Prairie draws from a neighborhood with a smaller radius than the one we allow, and its students' earlier scores are typical of that smaller neighborhood. In short, the charter schools draw students who are, on average, somewhat lower achieving than public school students in the neighborhoods where the schools are located (see Figure 1d). While the differences in incoming achievement are not dramatic, they certainly do not support the theory that charter schools drain regular public schools of their best, most-advantaged students. Remember that the above differences in earlier achievement do not affect our results because they are between applicants and nonapplicants. For our control group, we use lotteried-out applicants, not nonapplicants.



ALL PHOTOS / TONY ARMOUR

Attrition and Noncompliance

Some charter-school applicants do not comply with the treatment that the lottery "assigns" them. A small share of lotteried-in students do not actually enroll in a charter school. Instead, they enroll in a private school, a public school in another district, or—most often—continue in a regular Chicago public school. Also, some lotteried-out students do not continue to attend the Chicago public school from which they applied. They switch to a private school, a public school in another district, or even a different charter school.

Students who remain somewhere in the Chicago public school system (including charter schools) appear in our database, making them "observed noncompliers." Accounting for observed noncompliance in a randomized experiment is a fairly simple matter; we can adjust the estimated effect of attending a charter school to reflect the fact that some lotteried-in students did not attend.

The remaining noncompliers, however, are not observed because they disappear from the database when, for instance, they move to a suburban school district or switch to a private school. Unobserved noncompliers ("attriters") are a problem in a randomized study if the characteristics of students who attrite among the lotteried in are different from the characteristics of students who attrite among the lotteried out. Fortunately, this problem does not arise in our study: the patterns of attrition are very similar among lotteried-in and lotteried-out students. Most important, the lotteried-out students who attrite are neither higher nor lower scoring than the lotteried-in students who attrite. Thus our results should not be affected by the fact that we are not able to track every student through every postlottery school year.

Large Lotteries Work Best

CCSF followed careful procedures to ensure that their lotteries were truly random. But do randomized lotteries automatically generate treatment and control groups that are comparable? For the group of applicants as a whole, they apparently did. Looking at earlier test scores and demographic characteristics, we find no statistically significant differences between the lotteried-in and lotteried-out groups. And the fact that the groups are so similar in their outward traits suggests they are also similar in unobservable traits like motivation.

Yet in addition to checking whether the lotteried-in and lotteried-out students are comparable as whole groups, we also need to check that subgroups of students, sorted by the grade to which they applied, are comparable. That is, we need to check the comparability of lotteried-in and lotteried-out students who entered as, say, 2nd graders. The reason we need to check these subgroups is that separate lotteries were run for each grade of entry. To see this point, let's consider a concrete example. Suppose that a school held a lottery among 100 applicants for 50 kindergarten places and held a lottery among 20 applicants for two 6th-grade places. For the same reason that flipping a coin 100 times would probably result in about 50 percent heads, randomization would probably ensure comparability between the 50 lotteried-in and 50 lotteried-out kindergarteners. But we can be much less confident about the 6th-grade lottery, even if we know it is random, because there are so few applicants and places. Randomization could easily produce two lotteried-in students who just happen to be quite different from the 18 lotteried-out students.

In fact, the example is close to the truth. We find that randomization does ensure comparable groups in grades at which quite a few students are admitted. However, randomization is insufficient to ensure comparable groups in grades of entry that are rarely used. Since most students start in charter schools in early grades (kindergarten and 1st grade alone account for about 50 percent of new students), there are comparable groups for students who enter in kindergarten through grade 5. The 6th, 7th, and 8th grades account for, respectively, only 8, 5, and 4 percent of all CCSF admittees, and higher grades account for even tinier percentages. Thus it should be no surprise that the lotteried-in and lotteried-out groups are not comparable for grades of entry like 6 through 12.

In short, we confidently estimate the effect of attending a charter school for students who enter kindergarten through grade 5. We cannot use the lottery-based method with any confidence to estimate the effect of attending a charter school on students who enter in atypical grades, like grades 6 through 12. This is a limitation, because we might be intellectually curious about how charter schools affect the rare student who enters as, say, a 12th grader. However, it is a limitation that is largely irrelevant to policymakers. Most charter school students will, by definition, enter in grades that are typical grades of entry.

Charter Schools and Student Achievement

Because our evaluation is based on data from a randomized assignment, our analytic strategy is relatively simple. In essence, we simply compare the achievement of lotteried-in and lotteried-out applicants through the spring of 2004, or up to four years following their initial application. The results we report are adjusted to reflect the fact that not all lotteried-in students enrolled in charter schools. They therefore represent the effect of actually attending a charter school, not simply of drawing a lottery number low enough to gain admission. To refine the comparison, we account for the slight differences in the observable traits, including earlier test scores, that emerged by chance between lotteried-in and lotteried-out applicants. This refinement makes little difference in practice because randomization ensured that the groups were comparable.

We find that students in charter schools outperformed a comparable group of lotteried-out students who remained in regular Chicago public schools by 5 to 6 percentile points in math and about 5 percentile points in reading. (See Figure 2.) These are the key results of our analysis, and they translate into gains of 2.5 to 3 points for each year spent in the charter schools. The results are based on students who enter charter schools in kindergarten through grade 5, the grades of entry for which we can confidently estimate effects. To put the

The Elementary Effect of Charters (Figure 2)

Between 2000 and 2004, the lotteried-in CCSF school students in grades K-5 scored significantly higher on both the math and reading segments of the Iowa Tests of Basic Skills (ITBS) than the students who were lotteried out.

Effect of CCSF Schools on ITBS Percentile Score of students in Grades K-5

gains in perspective, it may help to know that 5 to 6 percentile points is just under half of the gap between the average disadvantaged, minority student in Chicago public schools and the average middle-income, nonminority student in a suburban district. If the students continued to make such gains for each year they spent in charter schools (a big “if”), then the gap between the charter school students and their suburban counterparts would close entirely after about five years of school. Right now, such projections are necessarily very speculative, but they help to give some sense of the magnitude of the charter-school effect.



* Significant at the 5% level
 ** Significant at the 10% level

SOURCE: Authors' calculations from Consortium on Chicago School Research data

The Virtues of Randomized Experiments

While the small number of students entering charter schools in midstream grades, like grades 6 through 12, precludes our estimating effects for them, the resulting focus is on the whole desirable. After a charter school is established, the vast majority of its students enter in the early elementary grades; for the most part, places in higher grades become available only when a student leaves.

In contrast, the rareness of late-grade entry poses serious problems for value-added analyses of charter schools, such as that by Robert Bifulco and Helen Ladd, who study North Carolina charter schools (see [“Results from the Tar Heel State,”](#) p. 60). Such studies, which compare the annual gains made by students in charter schools with the gains made by the same student while attending a traditional public school, draw only on the experiences of students who were tested for at least two years in the regular public schools before attending a charter school. Because they rely on state tests that are administered for the first time in the 3rd grade, almost all the students included entered charter schools in 5th grade or later. These students are most likely unrepresentative; after all, they are engaging in behavior that is rare.

The fact that 5th-grade entrants are rare is not accidental; it results from parents’ hesitancy to move children between schools. Logic would suggest that students who are moved midstream are more likely to be struggling socially or academically, and any such differences would cause results based on their experience to be misleading. It is dangerous to apply such results to more typical charter-school students, and it is wrong to portray them as representative in the absence of independent evidence that they are.

Our own data set can provide some indication of the magnitude of the problem. Fifth-grade entrants comprise only 13 percent of CCSF’s total admittees and only about 6 percent of the admittees in our analysis, which excludes applicants from private schools and does not include charter schools that are in their first year of operation. If we limit the analysis to the 5th-grade applicants for whom we can compute value-added estimates, the number of student-year observations included immediately falls by about 85 percent. If we use standard value-added methods to estimate the effects of attending a charter school for these students, the results do not match well with those of our lottery-based analysis. In short, studies that use value-added methods to evaluate charter schools are at best misleading. The students included are too atypical for the results to be interpreted in a straightforward way.

Conclusion

We have analyzed established charter schools in Chicago that are overseen by the Chicago Charter School Foundation. Our results demonstrate that, among students who enter in a typical grade, attending a charter school improves reading and math scores by an amount that is both statistically and substantively significant. We believe that these results can safely be extrapolated to similar schools that serve similar students. In particular, the results are most useful for understanding the effects of charter schools run by education-management organizations on student populations that comprise largely low-income and racial/ethnic minorities. We cannot confidently extrapolate the results to very different charter schools, students from very different backgrounds, or students who enter in atypical grades. Our results should be helpful for many policymakers who are concerned about urban students like those we study. However, we do not claim that the results are helpful for

all policymakers.

Research on charter schools, like the schools themselves, is fairly new. We are not aware of any other studies that use lotteries to isolate the effects of attending a charter school. Standard value-added analyses, which are often used to evaluate charter schools, rely entirely on an unusual group of students who switch from regular public schools to charter schools late in their elementary-school careers. Our analysis confirms that estimates of the effects of attending a charter school that rely on this peculiar group of students differ dramatically from estimates that are representative of students who apply to charter schools.

These differences probably stem from the tendency of parents to move children in the middle of elementary school only if they are already struggling. Thus we doubt that value-added analysis will ever produce results that have relevance beyond the peculiar set of students on which they depend. Evaluations of charter schools should rely on students who are typical of charter school applicants, not on students who are atypical. Randomization provides us with estimates that are inherently better than those based on value-added analysis.

Caroline M. Hoxby is professor of economics, Harvard University. Jonah E. Rockoff is assistant professor of economics and finance, Columbia Business School.

Published by the [Hoover Institution](#) © 2005 by the Board of Trustees of [Leland Stanford Junior University](#)

[home](#) | [past issues](#) | [search](#) | [about](#) | [subscriptions](#)