

Mathematica Policy Research Q&A Charter School Study

What are the major findings of this study?

- This is the first charter school study to date to use a lottery-based experimental design across students with varying backgrounds in multiple districts in multiple states.
- We found that, on average, charter schools had no significant impacts on student achievement in math and reading.
- We found that charter school impacts varied widely across schools.
- Impacts were most positive among schools in large, urban areas and among those serving the most disadvantaged students.

Definitions of key terms:

- **Experimental versus Nonexperimental:** Experimental studies use a random procedure to assign students to the school or program of interest. This ensures that, on average, students in the school or program have the same characteristics as the comparison group of students who were not randomly selected. The lottery-based experimental design used in the Charter School Study represents the gold standard in charter school research design.
- **Oversubscription:** Charter schools that receive applications from more students than they can admit are “oversubscribed.” These schools hold lotteries to randomly select students who will be offered admission to the school.
- **Treatment and Control:** The term “treatment group” refers to students who receive the intervention of interest (in this study, the treatment was admission to a charter school through an admission lottery). The “control group” refers to students who do not receive the intervention of interest (admission to a charter school after applying through a lottery).
- **Lottery Winners and Losers:** Students who applied to attend charter schools through a lottery and were offered admission are defined as “lottery winners.” Those who applied through a lottery and were not offered admission are defined as “lottery losers.”
- **Statistical Significance:** For an outcome to be “statistically significant,” there must be less than a 5 percent chance that the observed impact or difference was created by random variation in the study’s sample of students rather than by a true effect of the charter school.
- **Traditional Public Schools (TPS):** Public schools that are not charter schools.

What are the key points related to this study's methods and findings?

- In order to address the debate about how charter schools influence student achievement and other outcomes, we examined 36 charter middle schools that used admission lotteries in 15 states across the country and estimated the impacts of these schools on student achievement and other student outcomes.
- The study is the first of its kind to use an experimental design to study charter schools in multiple states and types of communities. It also measures a broader range of outcomes than other experimental studies.
- To ensure the validity of the experimental design, we observed and monitored the schools' lotteries to verify that the selection of students was truly random, and we documented both the results of the lottery and schools' subsequent admissions processes. We then followed both the lottery winners and lottery losers over the next two years and measured their outcomes.
- The study focused only on charter middle schools that were popular enough to hold lotteries. On average, the charter schools in the study served a student population that was more advantaged than other charter middle schools nationally.
- Results varied widely across charter schools. Those in large, urban areas and those serving disadvantaged students were successful in boosting students' math achievement. Those serving a more advantaged student population had negative impacts on student achievement in both reading and math.

Were there any additional findings?

- On average, the study charter schools had a positive effect on lower-income and lower-achieving students, and a negative effect on higher-income and higher-achieving students.
- On average, the study charter schools did not affect most of the other outcomes examined, including attendance, homework completion, student behavior, and parents' involvement in their child's education.
- Charter schools had a positive effect on students' and parents' satisfaction with their schools. Lottery winners (and their parents) perceived their schools to be more challenging academically and to have a stricter disciplinary environment. Parents of lottery winners also reported that their children were more well adjusted than did the parents of lottery losers.
- Charter schools' impacts on student achievement were related to certain school characteristics and practices. We found that charter schools with smaller enrollments and those more likely to use ability grouping in math or English classes had more positive (or less negative) impacts on achievement.
- There was weaker evidence that charter schools with longer school days and/or years as well as those with higher per student revenues from public and private sources had more positive (or less negative) effects.

- By contrast, we found no evidence that aspects of the policy environment in which charter schools operate, such as their operational autonomy, were related to their impacts on achievement.

Additional Questions:

What are the most important study findings and what do they tell us about the direction we need to go in?

The most important finding from the study is the general pattern of impacts: overall, charter schools had no significant effect on achievement, but impacts were positive at some charter schools (those in large urban areas and those serving disadvantaged students) and negative at other charter schools (those serving more advantaged students).

With respect to the research, a major priority is to learn more about what makes these schools more or less successful. This study suggests that simply increasing the number of charter schools may not change overall student achievement. But if we can figure out what makes the best charter schools successful and promote those school attributes, student achievement can be improved.

How does this study add to our knowledge of other research on charter schools? What's the contribution of this research?

The study contributes in a major way to the research on charter schools, because it is the first to apply an experimental design across a broad range of charter schools—36 schools in 15 states. We implemented the experimental design carefully, including observing the schools' lotteries and monitoring their subsequent admissions processes, so we can be confident that we have addressed a key concern about previous nonexperimental studies of charter schools—that they do not control for all the differences between students who choose to apply to charter schools versus those who remain in traditional public schools.

Moreover, the study allows us to make sense of previous findings that appeared at odds with one another. The most rigorous prior studies—those based on lotteries in large urban areas—had found positive impacts of a number of big city charter schools (in Chicago, New York City, and Boston). The more geographically diverse studies that examined a broader range of charter schools typically failed to find any positive impacts, and some found negative effects. Based on our results, we hypothesize that a key reason for this difference was the difference in school locations and types of students served by the schools in the different studies.

In summary, we consider this study to be the definitive investigation of charter schools' average impacts nationwide. Future research should prioritize identifying the factors that lead to better or worse charter school performance and examining the impacts of charter schools on longer-term outcomes such as high school graduation and college entry.

Does this study measure the impact of switching schools instead of the impact of attending a charter school?

No, this study does not measure the impact of switching schools. As described in Chapter III of the report, nearly all of the students in the study sample—including both those admitted to a charter school and those not admitted to a charter school—switched schools during the study’s two-year follow-up period. Students in the treatment group were a bit more likely than students in the control group to switch schools during the first year, but were less likely (by an equal amount) to switch schools during the second year.

How does this study compare to Mathematica’s recent report on KIPP schools? Are there KIPP schools in this study?

The initial report from Mathematica’s multiyear KIPP evaluation is consistent with both this study’s findings and some of the literature that suggests that urban charter schools serving disadvantaged youth have a positive impact on student achievement, because KIPP schools tend to be located in urban areas and serve a largely disadvantaged population.

The KIPP report was restricted to charter schools in the KIPP network, leading that study to focus on a single charter school model. The national sample of charter schools in the Charter School Study, in contrast, included a myriad of charter school approaches.

The initial report from the KIPP evaluation used a different methodology than the Charter School Study sponsored by the U.S. Department of Education. The Charter School Study used the gold standard of a lottery-based experimental design. The KIPP report relied on a matched comparison group design—a nonexperimental approach—although future reports in the KIPP evaluation will incorporate a lottery-based design.

As a condition of schools’ participation in this study, we cannot identify schools, and that includes confirming whether or not there were KIPP schools in our sample. The Charter School Study did include a variety of types of schools.

Were there significant differences in charter school impacts for different subgroups of students?

We examined whether the impacts of charter schools varied for subgroups of students based on the following characteristics:

- Income level
- Race/ethnicity (white, non-Hispanic versus nonwhite or Hispanic)
- Gender
- Baseline math/reading achievement

We found that impacts varied significantly by income level, as measured by eligibility for free or reduced-price meals. Impacts were positive for lower-income students and negative for higher-income students. Impacts were also related to students’ achievement level at the time they enter the charter schools—impacts were more positive for lower-achieving students and more negative for higher-achieving students.

We found no significant differences in impacts among subgroups defined by race/ethnicity or gender.

Did the study examine impacts for special education students or English language learners?

We did not examine special education students or English language learners, as there were too few of these students in our sample to produce reliable estimates.

Does the study suggest that specific teacher practices and policies may affect achievement?

This study did not examine specific teacher-level practices and policies in detail.

Why did the study not find more positive impacts?

These findings are consistent with a number of other charter school studies that have examined a broad range of charter schools across school districts and, in some cases, states. A common theme in nearly all studies of charter schools is that there is great variation across these schools—thus, it is not all that surprising that some of these schools would perform better than traditional public schools (TPS), while many others would perform the same as or worse than TPS.

Why did the study show negative effects for more advantaged and higher-achieving students and positive effects for more disadvantaged and lower-achieving students?

The question of why our study indicates that charter schools serving more advantaged students have negative effects while those serving a more disadvantaged set of students have positive effects is very important to examine. We believe that researchers and policymakers should focus more attention on this question in the future.

In our study, the primary goal was to estimate *what* effect a set of charter schools had, rather than *why* some schools were more successful than others. So we can't definitively determine why schools serving more advantaged students had more negative impacts. However, we did conduct some exploratory analyses in the study that are relevant to this question.

One possibility is that the charter schools in urban areas serving disadvantaged students are doing something different from non-urban charter schools, and that these differences make them more successful. We examined school characteristics including enrollment, hours of operation, student-teacher ratio, teacher experience, and use of ability grouping. While there are some differences between urban and non-urban charter schools with respect to these characteristics, even after we accounted for these characteristics we found that schools serving more disadvantaged students were more successful (in the sense of having more positive/less negative impacts). So we did not find evidence that the variation in impacts across schools can be explained by specific strategies the schools were using.

It is also possible that variation in charter school impacts is created more by the quality of nearby TPS. In other words, the charter schools serving advantaged students may be competing against

TPS that serve these students fairly well already, while the TPS competing with urban charter schools may be weaker schools with fewer resources.

In the absence of a rigorous measure of the effectiveness of the nearby TPS for their particular populations of students, it is not possible to determine whether the TPS in urban areas are “worse” than those in the non-urban areas, but this remains a viable hypothesis that should be explored in future research.

Does any particular type of charter school (e.g., KIPP) seem to have better success?

We did not specifically study any particular school model. We did examine the conditions under which charter schools were most likely to have positive effects, and found that those in large urban areas and those serving disadvantaged students had the most positive effects.

Eleven percent of the charter schools in our study were operated by charter management organizations (CMOs). However, we found no evidence that the impacts of these schools were higher or lower than impacts of schools not operated by CMOs.

Is it fair for this study to generalize across all charter schools?

It’s important to note that, while a broad range of schools participated in the study (36 schools across rural, suburban, and urban areas in 15 states), the participating schools are not nationally representative of charter middle schools.

Based on these study findings, what policy recommendations would be appropriate for the federal government and states? Do you think changes to charter school standards are called for? Should states and districts use the findings as a signal to stop supporting certain types of charter schools?

It is important to debate how this study should influence current and future charter school policies. We found that, on average, the charter schools in our study did not positively or negatively affect student achievement, although perhaps more relevant for policymakers is the wide variation in performance across the schools in the study. Charter schools as a whole are not a magic bullet for improving student performance, but some charter schools can be effective in raising student achievement. For instance, we found that impacts were most positive among schools in large, urban areas and among those serving the most disadvantaged students. Future research that attempts to discern *why* some charter schools are more effective than others can help policymakers further refine charter school policies and regulations to more positively influence student achievement.

Why did the study only include middle schools? What about the middle school experience leads you to believe that an analysis of charter middle schools would be illustrative of charter schools in general?

The study focused on middle schools because we wanted to assess the effects of charter schools on student achievement as measured by performance on state tests. Using state tests to measure achievement had two main benefits:

1. Nearly all public school students from grades 3 through 8 take these tests every year, and test results are generally available from states or districts.
2. State tests are the achievement measure most often used by policymakers to assess school quality. They are a policy-relevant measure of achievement.

However, since the tests are only available in most states for grades 3 through 8, in order to conduct a study in which we monitored the lottery process and followed students for the next two years, we needed to study middle schools.

While we have no particular reason for believing that the effects of charter middle schools differ from those of charter elementary or high schools—the research evidence is mixed—the results of our study are not statistically representative of those schools.

Did the study design include any analysis of why parents would have their children leave a TPS for a charter school? Were there variations in performance based on different motivations?

One advantage of the experimental design we used in this study is that it guarantees that, at the time of the lottery, the motivation of parents of the students who ended up in charter schools and parents of those who applied but did not get in was equivalent, on average. It is not the case that there were differences in the motivation levels of students (or their parents) who were admitted to charter school and our control group of students who attended TPS. For example, the proportion of students who applied to more than one charter school at the time of the lottery was the same (about 20 percent) among both groups of students.

We did not examine whether the impacts of charter schools were different for students whose baseline levels of motivation appeared to be different.

Since charter schools receive, on average, only 78 percent of the typical per pupil allocation of a traditional school, does the study indicate that charter schools provide a more cost-effective education than traditional schools?

The schools in our study may have been funded at higher or lower levels than the typical charter schools. While the study did examine average per-pupil revenues, these figures may have been influenced by differences in the student population attending each school. We did not measure costs at the TPS that sample members attended. Thus, we could not conduct a detailed analysis of revenues of the charter schools in our study relative to the TPS that students attended if they couldn't attend the charter school. Nor did we conduct a cost-benefit analysis of the charter schools in the study.

Do your findings suggest that it would have been better for students not to transfer to a charter school?

We would not draw that conclusion from the study's findings. Overall, we found that the average charter school in the study did neither better nor worse than nearby TPS in boosting student achievement. Moreover, there was a lot of variation in these impacts both across students and across schools, with some schools having positive effects, some having negative effects, and some having no effects on student achievement. In addition, the study did not examine longer-

term outcomes that could be influenced by charter school attendance, such as high school achievement or college enrollment rates.

Finally, we found that being admitted to one of these charter schools led to an increase in students' and parents' satisfaction with their school. So the students in our sample who were admitted to charter schools appeared to be generally satisfied with their choice to apply.

In your national study across multiple geographic areas, was there one standardized test that you used across all sample schools/students? If not, how did you factor in the level of difficulty of the assessment in your comparison?

To measure student achievement, we used students' performance on their state assessment, in part because we believed it to be the most policy-relevant measure available. These tests do differ across states, and we studied charter schools in 15 different states. To address these differences in tests, we used the following two-part strategy:

1. Since we estimated impacts separately for each charter school in the study, we could ensure that all students in our sample at any given charter school site—including those admitted to the school and those not admitted—took the same test. Thus, in each case, our comparison of charter school students with TPS students involves a comparison among a set of students who all took the same test.
2. It is the case, however, that students in different sites may have taken different tests. To address this issue, we measured each student's performance on his or her state test relative to the average student in that state. We did this by calculating what researchers call a z-score. If the value of this score is positive, it means that the student in our sample did better than the average student in that state; if negative, it means that the sample member did worse than the average student in the state. By using this relative measure, we were able to combine the results across charter schools in the study located in different states.

Can an experimental design really be used to evaluate charter schools, given all the differences in state regulations, operators, models, and other characteristics?

We believe that an experimental design is the best way to evaluate charter schools, since it provides the highest quality evidence about the true effects of being admitted to a charter school. It is true that there are many differences between charter schools with respect to their characteristics, the policy environment they face, and the general context in which they operate. All this variation, however, does not imply that we should turn away from experimental designs. Instead, these differences suggest that it is important to not only examine the average effects of charter schools, but also look for differences in effects among charter schools that have very different characteristics. We have attempted to explore these issues in our study.

Was there any difference in the performance of charter schools by the ethnicity, race, or socioeconomic status of students? For example, did charter schools with large, urban, African American populations do better than charter schools with large, urban, white or Hispanic populations?

We examined whether the impacts of charter schools in our study varied by students' socioeconomic status, race/ethnicity, gender, and baseline achievement levels. We found that the effects of charter schools differed significantly for students with different socioeconomic status backgrounds. In particular, impacts were positive in math for low-income students eligible for federal free or reduced-price school meals, but negative in math and reading for higher-income students not eligible for these benefits. We also found that impacts were significantly correlated with the achievement level of students at the time they entered the charter schools—impacts were more negative for higher-achieving students and more positive for lower-achieving students. Impacts were not significantly related to a student's gender.

We examined whether charter school impacts for white, non-Hispanic students differed from those for other students (including both black and Hispanic students). We did not find any significant differences in charter school impacts for these white versus non-white students. Our samples of black and Hispanic students were too small to reliably examine them separately.

What experience does Mathematica have in conducting studies of this type that involve complex education issues?

Mathematica has a long history of conducting large and complex studies of education interventions for the U.S. Department of Education and other clients. Many of our previous studies, like this one, have involved large-scale experimental research designs. For example, Mathematica's Evaluation of Teacher Preparation Models used an experimental design across 73 schools to examine how different teacher preparation methods contributed to the academic achievement of students. Similarly, Mathematica's Impacts of Comprehensive Teacher Induction Study used an experimental design that included 394 schools to test the impacts of teacher induction programs on the retention and instructional practices of novice elementary school teachers and on the achievement of their students.

Mathematica has also conducted—and continues to conduct—numerous studies of charter schools. In addition to this study, for example, we are currently conducting an evaluation of KIPP schools for the KIPP Foundation, as well as a nationwide study of charter management organizations (CMOs) and their schools for the New Schools Venture Fund, the Bill & Melinda Gates Foundation, and the Walton Family Foundation.