Katie Newmark
April 20, 2004

The Effect of Charter Schools on Academic Achievement in North Carolina

Over the past fifteen years, school choice has moved to the forefront of education reform. Policymakers have considered both school vouchers and charter schools; while only a small number of governments have implemented voucher programs, 41 states and the District of Columbia have allowed charter schools. Minnesota passed the nation's first charter school law in 1991, and in March 2004, Washington became the latest state to adopt a charter school law.[1] As of January 2004, 2,996 charter schools nationwide served over 684,000 students (CER(a)).

Charter schools blend characteristics of both public and private schools. Like public schools, they are open to all students, paid for with tax dollars, and accountable to the public. Like private schools, students choose to attend, rather than being assigned, and the schools have considerable autonomy from government regulations. Charter school advocates argue that freedom from regulation will allow the schools to innovate better ways to meet the needs of their students. They therefore predict that charter school students will improve their achievement test scores. This claim has played a prominent role in discussions of charter school laws, but its validity is still uncertain.

Methodological problems have made it very difficult for researchers to determine if charter schools improve the achievement levels of their students. Ideally, they would like a randomized, controlled experiment. A diverse group of students would be randomly assigned to charter schools. The only difference between public school students and charter school students would be the random school assignment, and thus higher test scores in charter schools could be attributed solely to the practices of the charter schools. But charter school attendance is not

[1] For more information about school choice programs in the states, see Kafer (2003).

random. Families deliberately choose to leave the public schools to attend charter schools, and researchers must find a way to cope with the problem of selection bias.

The problem is that certain characteristics may make a family more likely to leave the public schools or choose a particular charter school. If those characteristics also affect test scores, it muddles the interpretation of charter school test scores: are those characteristics or the school itself driving the test score results? Observable characteristics such as income and race can be statistically controlled for, but many unobservable, unquantifiable characteristics also influence the choice to attend charter schools. For example, if students with motivated, involved parents are more likely to go to charter schools and those students are also more likely to have high test scores, then the test scores of charter school students will be high—but not necessarily because of the effect from the school itself. The type of students who have selected the school is also affecting the test score results.

A 2001 study evaluating the first three years of charter school operations in North Carolina highlights the need to control for selection bias (Noblit and Corbett 2001). Using just descriptive statistics, with no statistical controls, the authors find that for the reading and math end-of-grade tests in grades 3 through 8, charter school students were less likely than public school students to score at or above grade level. Individual students who attended the same charter school for all three years had higher prior achievement scores, by about one or two scale score points, than students who attended public schools for all three years, but tended to fall to one or two points *below* public school students by the end of the three years. Noblit and Corbett point out the problem of selection bias that affects these uncontrolled results, but do not try to control for it, even though their other findings strongly suggest that charter school students are systematically different from public school students. They report that charter school students

were more likely to be male and black and tended to have parents with higher levels of education. While noting the wide variance among the reasons charter schools gave as their primary reasons for existence, they also generalize charter schools as "establishing niches" for students whose needs are not being met in the public schools. In particular, they observe that in some cases, the school district and the charter schools are "beginning to informally stake out a relationship in which the charter schools use their small size to advantage in working with students who need extra help." This characterization suggests that charter schools are targeting themselves to certain populations, which enhances the selection bias problem because those certain populations will be more likely to choose charter schools.

In their analysis of charter school achievement across the country, Greene, Forster, and Winters (2003) attempt to minimize selection bias by comparing "untargeted" charter schools to nearby public schools. Because some charter schools target specific student populations, such as academically gifted students, at-risk students, or juvenile delinquents, they argue that "comparing targeted charter schools to regular public schools is like comparing apples and zebras." Restricting their study to untargeted charter schools allows Greene, Forster, and Winters to compare more similar student populations. They further control for student characteristics by looking at year-to-year test score changes to evaluate the value added by the schools and by comparing charter schools to the nearest untargeted public schools, using "geography to control for demography." For untargeted charter schools in North Carolina, which comprise 84.9 percent of the state's charter schools, they find the schools have a positive but statistically insignificant effect on both math and reading test scores.

Bifulco and Ladd (2004) exploit an extensive panel data set to track the performance of individual students for as many as six years as they move between traditional public schools and charter schools. Following the strategy of Hanushek, Kain, and Rivkin (2002), they compare the test score gains of charter school students to the gains made by those same students when they attend traditional public schools. Analyzing gains, as Greene, Forster, and Winters did, allows them to look at the value added by charter schools, and observing the same students at several different times allows them to control for unobservable individual characteristics that remain constant over time. With these strong controls for selection bias, they find that the effect of charter school attendance is negative and significant. Compared to the gains they make in traditional public schools, charter school students make gains nearly 0.10 standard deviations smaller in reading and 0.16 standard deviations smaller in math.

In light of Greene, Forster, and Winters's conclusion that North Carolina charter schools have no effect on test scores and Bifulco and Ladd's findings that the schools have a negative effect on achievement, it certainly seems strange that the number of charter school students in North Carolina has increased steadily in every year of their existence, from 0.37 percent of the total public school student population in 1997-98 to 1.36 percent in 2001-02 (Table 1). Why would parents want to uproot their children from traditional public schools to send them to charter schools that would, at best, have no effect on their achievement? Assuming that people make rational, fully-informed decisions, one possible explanation is that they judge schools by more than achievement. If, for example, charter schools are safer or allow greater student-teacher interaction than traditional public schools and families value these outcomes, they may prefer charter schools even if the schools do not affect achievement.

Another possible explanation is that these empirical analyses, despite their careful efforts, still have not adequately controlled for selection bias. In other words, charter schools are in fact improving the achievement of their students, but in the research to date, selection bias has masked that positive effect.

The purpose of this paper is to investigate the second explanation. Greene, Forster, and Winters develop an innovative methodology that recognizes the considerable heterogeneity of charter school missions, but their attempt to use "geography to control for demography" seems inadequate. Because charter schools enroll students from anywhere in a district, or even other districts, they do not draw from the same pool of students as the nearby traditional public schools that enroll only students in the surrounding neighborhoods. The student bodies of charter schools are thus likely to be demographically different from those of nearby traditional public schools, and the controls used by Greene, Forster, and Winters are inadequate.

Bifulco and Ladd, on the other hand, have extremely rich data that allows them to use a value-added fixed-effects model. Short of a randomized, controlled experiment, this type of model may be the best way to control for selection bias. Unfortunately, the extraordinary data requirements of this method make it difficult to imitate in other states or to replicate in North Carolina.[2] A less data-intensive method of controlling for selection bias would therefore be useful. If such a method showed that charter schools do not improve achievement, it would support Bifulco and Ladd's results. If it showed that charter schools do improve achievement, it would suggest that at least one of the two methods still gives biased results. Further research would be necessary to sort out these issues of bias.

[2] The data used by Bifulco and Ladd was collected by the North Carolina Education Research Data Center, and the procedures for obtaining access to the data are available at the following website: http://www.pubpol.duke.edu/centers/child/NC_Education_Research_Data_Center/procedures.html. In particular, student access is limited to PhD and Masters degree students.

The literature evaluating the effect of private schools on achievement suggests just such a method for controlling selection bias that has the advantage of being less data intensive than Bifulco and Ladd's methods and also has not yet been tried for North Carolina charter schools: using instrumental variables in a two-stage least squares (TSLS) regression.

When examining the effect of private schools on achievement, researchers want to use private school attendance as a regressor to explain the dependent variable achievement. The selection bias problem indicates that private school attendance may be correlated with unobservable characteristics that are also correlated with achievement. These unobservable characteristics cannot be included as regressors, so they are included in the error term. Private school attendance is thus correlated with the error term in violation of one of the basic OLS assumptions, which makes it an endogenous regressor that biases the resulting coefficient estimates. As explained earlier, charter schools are also choice schools subject to selection bias, and researchers studying the effect of charter schools on achievement thus face the exact same econometric difficulties. To confront these difficulties, researchers studying private school achievement have used instrumental variables in TSLS regressions (see, for example, Sander and Krautman 1995, Evans and Schwab 1995, Neal 1997, Neal 1998). A similar approach may be useful for analysis of charter schools.

The key idea of TSLS is to think of the endogenous regressor as having two components: one determined by endogenous factors, and the other by exogenous factors. The presence of the first component causes regressions using that regressor to have biased coefficient estimates. To eliminate that bias, TSLS uses instrumental variables to remove the problematic endogenous component of the endogenous regressor. An instrumental variable is one that is correlated with the endogenous regressor but not with the dependent variable. In the first-stage regression, the

endogenous variable is regressed on an appropriate instrument, which results in fitted values of the endogenous variable that contain only the exogenous portion. Thus purged of any endogeneity, these fitted values can then be used in the place of the original endogenous regressor in the second-stage regression and the resulting coefficient estimates will be unbiased (Stock and Watson 2003).

To date, only one paper has used TSLS to analyze the effect of charter schools on achievement. Bettinger (1999) examines whether competition from Michigan's charter schools affects achievement at nearby traditional public schools. Because charter school location may be determined endogenously, he develops an instrument based on a feature of Michigan's charter school law that allows state universities to approve charter schools. He reasons that traditional public schools closer to universities would be more likely to have charter schools located nearby, but proximity to a university would not have a direct effect on traditional public school achievement. Using this instrument, he finds that charter schools have little effect on traditional public school achievement. Unfortunately, because most states, including North Carolina, do not have similar provisions in their charter school laws (CER 2004(b)), there are limited possibilities for extending this instrumental variables strategy to other states.

In this paper, I will analyze the effect of charter schools on the achievement of their students by using an instrumental variables strategy. The next section explains North Carolina's charter school law. The following section describes the data used in this study. Then I discuss the characteristics of charter schools and their students. The empirical section provides a model of student achievement, describes my instruments, and presents the results from estimating the model using TSLS. The final section summarizes this study and offers some directions for further research.

North Carolina's Charter School Law

Passed by the General Assembly in 1996, North Carolina's Charter Schools Act (North Carolina General Statute §115C-238.29) outlines six objectives for charter schools. The first, to "improve student learning," receives the most attention in charter school research, but other benefits to students, as well as to parents and teachers, are also emphasized. For instance, the fourth listed objective is to "create new professional opportunities for teachers, including the opportunities to be responsible for the learning program at the school site."[3]

Any person, group of people, or nonprofit corporation can apply for a charter. Applications are made to the State Board of Education (SBE), local boards of education, or the board of trustees of a school in the University of North Carolina system. Local boards and universities can give only preliminary approval; the SBE has the final say, so most applications are made directly to the SBE (McNiff and Hassel 215). Charters are initially granted for a period no longer than 5 years and subsequently renewed for periods no longer than 5 years. No more than five new charters per local school district per year may be granted, and there can never be more than 100 charter schools in the whole state (North Carolina Department of Public Instruction(a)).

Charter schools are controlled by boards of directors that decide, with few restrictions, "all matters related to the operation of the school, including budgeting, curriculum and operational procedures" (NCDPI(a)). A school's funding is determined on a per-pupil basis, and 100 percent of state and local funding follows students from traditional public schools to charter schools. In addition, a school may be eligible for federal funds for start-up or for programs such

[3] In addition, the law expresses particular concern for helping "students who are identified as at risk of academic failure or academically gifted," a goal that only increases the problem of selection bias.

as Title I. The board hires and fires teachers and other employees (i.e., the teachers are not employees of the local school district) and sets their salaries, bound only by the requirement that 75 percent of elementary teachers and 50 percent of secondary teachers must be certified. All interested students, regardless of race, gender, or disability, must be admitted, and after one year of operation, the racial composition of the charter school must "reasonably reflect" the racial balance of district. If the school is over-enrolled, students will be chosen by lottery. A charter school year is required to be a minimum of 180 days, but the length of the school day is determined by charter school.

Data

The data for this research comes primarily from three sources: the North Carolina Department of Public Instruction (NCDPI), the National Center for Education Statistics' Common Core of Data (CCD), and the 2000 Census.

NCDPI provides the achievement data from its ABCs of Public Education program. Designed in 1995 to provide school accountability and promote the teaching of basic skills, the ABCs program requires students in grades 3 through 12 to take standardized state end-of-grade and end-of-course tests. All schools, including charter schools, are evaluated by both student performance and growth in student achievement, and the results are widely publicized in school report cards. High-performing schools and schools with exceptional growth are recognized and rewarded by the state, while low-performing schools receive help from state assistance teams. These public features of the program make ABC test scores relevant to parents, who can easily use the standardized scores to objectively judge the quality of their children's schools. Although

parents may value other features of schools, the ready availability of ABC test scores make them important and relevant measures of charter school achievement. (NCDPI(b))

To evaluate school achievement, I use NCDPI's summary measure, the performance composite, for 2000. A school's performance composite is calculated by dividing the total number of tests on which its students scored at or above grade level by the total number of tests taken by its students. This number is thus a measure of the percentage of students at the school who are performing at or above grade level.[4] The performance composite is a good choice for an achievement measure in this study because the state uses it to evaluate the growth and performance of schools and because it is easily available and easily understood by researchers and parents alike.

All other school-level data is from the 2000 CCD. The sample of schools included in this study is restricted to all elementary schools, where an elementary school is defined as a school having a lowest grade no higher than third grade and a highest grade no lower than fourth grade. Although not a traditional definition of elementary schools, this broader definition is useful because approximately one-half of all charter schools with traditional elementary grades also have middle or high school grades (Noblit and Corbett). Limiting the elementary school definition to schools with only kindergarten through sixth grade would greatly reduce the number of charter schools in the sample and not provide an accurate picture of the competition between charter and traditional public schools for students of that age. Schools with only grades below third grade are excluded because the do not have any achievement data (ABCs testing

[4] Comparisons of performance composites across schools are complicated because the composition of the total number of tests taken can differ. For instance, only fourth, seventh, and tenth graders take writing tests, so the performance composites of schools without one or more of those grades represent a different spectrum of tests than schools with those grades. In addition, a school that has a greater percentage of fifth graders than another school will have the fifth grade reading and writing tests form a larger component of its performance composite. To the extent that fifth graders differ from other students, that school's performance composite will be different for a reason unrelated to school quality.

starts in third grade). State schools with specialized missions, such as the Eastern North Carolina School for the Deaf,[5] are also excluded because they are typically ungraded and because their administration and student body composition are so drastically different from those of district schools.

County-level demographic data, including racial composition, average household income, and adult education attainment, comes from the 2000 Census. County data is useful because in North Carolina, for the most part, each county is a school district. Each charter school is technically its own district, but each charter school is inside only one county and both NCDPI and the National Center for Education Statistics classify charter schools by their host counties. It therefore seems reasonable to associate each charter school with the demographics of a host traditional public school district. A more difficult exception to the equivalence of counties and school districts is that 12 counties have both a county school district and one or more city school districts.[6] For the purposes of this study, city districts will be merged with their surrounding county districts, as is done in the CCD.

Characteristics of Charter Schools and Their Students

The racial composition of charter school students is significantly different from that of traditional public schools. The average charter elementary school is nearly half black, while the average traditional public elementary school is about 32 percent black (Table 2). The average percentages of Asian and Hispanic students are small for both charter and traditional public

[5] The other schools in this category are the Governor Morehead School (for the blind), the North Carolina School for the Deaf-Morganton, the Central North Carolina School for the Deaf, and the Juvenile Evaluation Center (for students in the juvenile justice system).

[6] The 12 county districts with city school districts are: Buncombe County and Asheville city; Cabarrus County and Kannapolis city; Catawba County, Hickory city, and Newton-Conover city; Cleveland County, King's Mountain city, and Shelby city; Columbus County and Whiteville city; Davidson County and Thomasville city; Halifax County, Roanoke Rapids city, and Weldon city; Iredell-Statesville and Mooresville city; Orange County and Chapel Hill-Carrboro city; Surry County, Elkin city, and Mt. Airy city; and Sampson County and Clinton city.

schools, but charter schools have fewer students of both races. The percentage of white students is also smaller at the average charter school, but not significantly so. To the extent that students of different races differ in socioeconomic status, academic performance, and unobservable characteristics, these racial composition statistics suggest a selection bias problem.

Parent and community characteristics may also be important for determining student achievement, and here too, the average charter school is significantly different from the average traditional public school. To measure parent income and educational attainment, I use the average characteristics of the adults in the zip code of the schools.[7] Since charter schools enroll students from a wider geographic range than just the nearby neighborhoods, zip code level data is an imperfect measure of charter school parents. (Greene, Forster, and Winters (2003) face a similar problem when they attempt to use "geography to control for demography.") Furthermore, not all adults in a zip code are parents of elementary school students; it is possible that parents of school-age children are systematically different from other adults, and thus a zip code average would be inaccurate. Lacking school-level information about parents, however, zip code level data is a reasonable proxy. This proxy data reveals that, on average, the median income and poverty rate of charter school parents are similar to those of traditional public school parents. On the other hand, the average charter school parents are significantly more highly educated than the average traditional public school parents.

A substantial difference between the communities in which charter schools locate and the ones in which traditional public schools locate is that charter schools are located more frequently in urban areas: the percentage of the population in a charter school's zip code that lives in an urban area averages about 67 percent, while for traditional public schools the percentage

[7] Data for the percentage of students receiving free and reduced price lunch, a traditional measure of parental income levels, is not available for individual charter schools.

averages about 53 percent. The difference is statistically significant. Urban areas tend to have greater crime and poverty rates, for example, and these characteristics may indirectly affect students' achievement scores by affecting their home lives. Charter schools also tend to be located in wealthier school districts, although the difference is not significant for any measure except percentage of students receiving free or reduced-price lunch. In addition, charter schools have higher average per pupil expenditures, which may be partly explained by being located in wealthier school districts. If schools with more money can provide better instruction, then the differences in district wealth suggest yet another potential source of bias.

In sum, the average charter school has more black students than the average traditional public school, as well as more highly-educated parents, and is located in a more urbanized zip code in a wealthier school district that spends more per student. Any attempt to determine whether charter schools improve the achievement of their students must control for these factors to isolate the effect of the charter school itself from the effects of student demographics or community characteristics.

The descriptive statistics suggest some basic ways in which the charter school itself may influence achievement (Table 3). In addition to a charter school's administrative policies, hiring practices, and curriculum choices, the size of the school can be important. Compared to the average traditional public elementary school, the average charter elementary school enrolls less than half as many students but over a larger range of grades. Over half of charter elementary schools include middle and/or high school grades, while less than ten percent of traditional public elementary schools do (Table 3). Although charter schools do not have a significantly smaller average student-teacher ratio, their overall smaller size may still provide a more nurturing atmosphere where teachers and administrators can get to know their students well.

Having more grades within one school allows these relationships to last longer and also eliminates the disruptions associated with changing schools.[8] Once the endogenous variation in charter school attendance is removed, charter schools may prove to have a positive effect on achievement because of environmental differences such as these.

Empirical Models and Results

A raw comparison of charter school achievement with traditional public school achievement shows that charter schools have lower scores. The average traditional public school's performance composite is 74.6, while the average charter school has a performance composite of only 55.7 (Table 4). Although the average charter and traditional public school are located in districts with similar achievement scores, most charter schools score far below their districts. Moreover, they score below the majority of traditional public schools in their district (Figure 1). Perhaps because the sample of charter schools is small, or perhaps because charter schools target different populations, the variance among achievement scores is greater for charter schools. The range of charter school scores is nearly 15 points higher than the range for traditional public schools, and the standard deviation for charter schools is almost twice that of traditional public schools.

Of course, as discussed previously, charter school students differ from traditional public school students in observable characteristics, and likely in unobservable characteristics as well. To evaluate the effect of charter schools on achievement, independent of observable student characteristics, we must control for those characteristics. In other words, the question of interest is: if two students were identical except one of the students attended a charter school, would the

[8] Hanushek, Kain, and Rivkin (2001) find a small short-term negative effect on achievement caused by switching schools.

charter school student have higher achievement scores? To investigate this question, consider the following model of achievement for student *i* at time *t*:

$$\log(achievement_{it}) = \alpha\ charterschool_{it} + \beta X_{1it} + \gamma X_{2it} + \varepsilon_{it}, \quad (1)$$

where $charterschool_{it}$ is a dummy variable indicating whether the student attends a charter school, X_{1it} is a vector of student characteristics including race and sex, X_{2it} is a vector of parental characteristics including income and educational attainment, and ε_{it} is the error term. Without individual-level data, however, I cannot estimate this model. Instead, I will use the following model with data aggregated by school district *d*, where each district has I_d students:

$$\sum_{i=1}^{I_d} \log(achievement_{it}) / I_d = \sum_{i=1}^{I_d} (\alpha\ charterschool_{it}) / I_d + \sum_{i=1}^{I_d} (\beta X_{1it}) / I_d + \sum_{i=1}^{I_d} (\gamma X_{2it}) / I_d + \sum_{i=1}^{I_d} (\varepsilon_{it}) / I_d \quad (2)$$

Assuming that α, β, and γ are equal for each individual—a common assumption in aggregation models that amounts to saying that certain characteristics have the same effect on different individuals (Intriligator, Bodkin, and Hsiao 1996)—then equation 2 simplifies to:

$$\overline{\log(achievement_{dt})} = \alpha\ \overline{charterschool}_{dt} + \beta \bar{X}_{1dt} + \gamma \bar{X}_{2dt} + \bar{\varepsilon}_{dt}, \quad (3)$$

where $\overline{\log(achievement_{dt})}$ is the log of average student achievement score in the district, $\overline{charterschool}_{dt}$ is the percentage of elementary school students in the district who attend charter schools, $\bar{X}_{1dt}$ is a vector of district student characteristics including the percentage of students of different races and the percentage of male students, X_{2dt} is a vector of district parental characteristics including average income and average educational attainment, and $\bar{\varepsilon}_{dt}$ is the error term. Note that the estimated coefficients that can be obtained from equation 3 are equivalent to

the estimated coefficients that can be obtained from equation 1; aggregated county-level data can be used to find individual-level effects. In particular, $\hat{\alpha}$ is an estimate of the effect of charter school attendance on a student's achievement, after controlling for observable characteristics.

Aggregating the data introduces two potential sources of bias. One source is the achievement measure. Without individual-level achievement data, I constructed $\overline{achievement_{dt}}$ as an average, weighted by enrollment, of the performance composite scores of each county's elementary schools. But because different students may take different numbers of tests, my measure of $\overline{achievement_{dt}}$ is not quite the same as the average of $achievement_{it}$ for all students in the district.[9] As long as this discrepancy between measures is not correlated with the right-hand-side variables, however, it will not bias the results. A more problematic source of bias comes from the parental characteristics measures. Without data on the parents of individual students, I use income and educational attainment data for the entire county. As explained earlier for zip code-level data, this county-level data is an imperfect proxy for parental characteristics. The error in these measures will bias the estimated coefficients towards zero.

Table 5 presents the results from estimating equation 3 with OLS. The estimated effect of charter school attendance is positive and statistically significant. If a student switches from a traditional public school to a charter school, his achievement score will increase by an estimated 22 percent.[10] The other coefficients, except percentage male and percentage Hispanic, have the expected signs, but only the coefficient on the percentage of white students, is statistically

[9] For instance, assume that School A and School B are the only two schools in a district and each school has only two students. At School A, student 1 scores at or above grade level on 2 out of 3 tests (i.e., his performance composite is 0.67), student 2 scores at or above grade level on 1 out of 2 tests. At School B, student 1 scores at or above grade level on 2 out of 2 tests, student 2 scores at or above grade level on 1 out of 2 tests. The performance composite for School A is (2+1)/(3+2), or 0.6, and the performance composite for School B is (2+1)/(2+2), or 0.75. The constructed aggregate achievement measure for the district is (0.5)*(0.6) + (0.5)*(0.75), or 0.675. However, the average of all student performance composites is (0.67 +0.5+1.0+0.5)/4, or 0.667.

[10] This estimate is an upper bound on the effect since switching schools causes a short-term decrease in test scores (Hanushek, Kain, and Rivkin 2001).

significant. With the narrowing of the sample size that results from aggregation, the large variances for most of the coefficient estimates is to be expected.

The OLS estimates are only a first approximation to answering the question of how charter school attendance affects test scores. Because the students choosing to attend charter schools may differ from other students in unobservable ways, charter school attendance is endogenous and its estimated coefficient is contaminated by selection bias. To correct for the selection bias, I use an instrumental variables strategy.

My choice of instruments recalls Noblit and Corbett's characterization of North Carolina charter schools as filling "niches" not adequately served by the traditional public schools. If it is more difficult for a family to find a traditional public school that matches their preferences, they will be more likely to choose a charter school that matches their particular niche. Indeed, charter school advocates cite the benefits of improved family-school matches as one of the primary advantages of charter schools. Since traditional public schools are administered at the district level, there is not much room for diversity among the schools. If most of the people in the district have similar school preferences, this situation will be fine; but if families in the district have varied preferences, they will be less likely to find traditional public schools that match their needs and more likely to send their children to charter schools. A measure of the diversity of school preferences in a district is thus a good candidate for an instrument variable if it is also uncorrelated with student achievement.

School preferences are correlated with family characteristics such as race, income, and parental education levels. [citations] Therefore, district-level Herfindahl indexes of these characteristics can quantify diversity of school preferences. Herfindahl indexes, typically used in industry concentration studies, measure whether market power in concentrated in just a few firms

or in many firms. Calculated using the formula $H = \sum_i s_i^2$, where s_i is the market share of firm i and $\sum_i s_i = 1$, a Herfindahl index value ranges from zero to one; the closer H is to zero, the more diversity there is in the market. For Herfindahl indexes of family characteristics, s_i represents the percent of the district population with a particular characteristic. Using 2000 Census data, I calculate district-level Herfindahl indexes for race, income, and adult educational attainment.[11] In addition, I calculate district-level Herfindahl indexes for traditional public school ABC performance composite scores.[12] Districts with a greater diversity of school achievement scores have students with a greater variety of needs, and families will be more likely to be interested in charter schools targeted towards at-risk or high-achieving students. As measures of school preference diversity, these Herfindahl indexes are likely correlated with district charter school attendance. Using similar logic, Glomm, Harris, and Lo (2003) find that Herfindahl indexes for race and adult education levels are negatively correlated with charter school location. That is, districts with more diverse populations are more likely to attract charter schools.

Furthermore, these Herfindahl indexes are also unlikely to be correlated with student achievement. Race, income, and adult education levels certainly influence achievement, but a Herfindahl index is a nonlinear combination of those factors. More importantly, a Herfindahl index considers only the distribution of these characteristics, not whether they are high or low. A district made up almost entirely of high-income families will have the same index value as a

[11] For these indexes, I use market shares in the classifications used by the Census. For race, the categories are: white; black; Asian; American Indian and Alaskan native; Native Hawaiian or other Pacific Islander; Hispanic; more than one race; and other. For household income, the categories are: less than $10,000; $10,000 - $14,999; $15,000 - $19,999; $20,000 - $24,999; $25,000 - $29,999; $30,000 - $34,999; $35,000 - $39,999; $40,000 - $44,999; $45,000 - $49,999; $50,000 - $59,999; $60,000 - $74,999; $75,000 - $99,999; $100,000 - $124,999; $125,000 - $149,999; $150,000 - $199,999; and $200,000 or more. For educational attainment, the categories are: less than a high school diploma; high school graduate; some college, less than 1 year; some college, 1 or more years, no degree; associate degree; bachelor's degree; master's degree; professional school degree; and doctoral degree.

[12] The categories for this index are ten-point intervals.

district made up almost entirely of low-income families. It is difficult to see how the district-wide distribution of race, income, educational attainment, or test scores would have an effect on achievement, other than through its effect on charter school attendance. For that reason, these Herfindahl indexes may be useful instrumental variables.

The two-stage least squares (TSLS) regressions will first estimate the following equation:

$$\overline{charterschool}_{dt} = \delta Z_{dt} + \lambda \bar{X}_{1dt} + \mu \bar{X}_{2dt} + \eta_{dt}, \quad (4)$$

where Z_{dt} is a vector of instruments and η_{dt} is the error term. The fitted values of $\overline{charterschool}_{dt}$ from this first-stage regression will then be used to estimate the following second-stage equation:

$$\overline{\log(achievement_{dt})} = \alpha\, \overline{charterschool}_{dt} + \beta \bar{X}_{1dt} + \gamma \bar{X}_{2dt} + \bar{\varepsilon}_{dt}. \quad (5)$$

If the instruments are powerful enough, the coefficients estimated by equation 5 will be free of selection bias.

Table 6 presents the TSLS results. For comparison, the OLS results from Table 5 are also included. To test the exogeneity of the instruments, I included the instruments as regressors in the original OLS regression to determine their effect on student achievement. None of the instruments have coefficients significantly different from zero in these regressions; therefore, the instruments are exogenous. Models 1 through 4 use each instrument individually. All of the instruments have some explanatory power for charter school attendance, as seen in the higher R^2 values for the TSLS first-stage regressions than for the OLS "first-stage" regression without any instruments. All of the first-stage R^2 values are still negative, however. In addition, each of the instruments except the achievement index has the expected sign, but none of them has a statistically significant effect on charter school attendance. Model 5, in which the three

instruments that had the expected signs are combined, does not improve matters. Thus, although the instruments are exogenous, their effect on charter school attendance is so weak that the estimates from the second stage are likely to be unreliable. Indeed, the second-stage estimates of the effect of charter school attendance vary greatly with the model. In some cases, the TSLS estimate is greater than the OLS estimate, and in other cases, it is lower. In no case, however, is it significantly different from zero. Leaving aside the issue of weak instruments, these results indicate that selection bias inflates the effect of charter schools on achievement; students choosing to attend charter schools have unobservable characteristics that improve their test scores. Once those characteristics are controlled for, the charter school effect disappears.

The preceding results may be distorted by one outlying data point. In Pamlico County, there are only two elementary schools. One is a charter school and thus the percentage of students in the county attending charter schools is 43.1 percent. In contrast, the county with the next highest percentage of charter school students is Nash County with only 8.65 percent. Because Pamlico County is anomalous, I repeated the above regressions excluding Pamlico County.

The results of this second set of regressions are reported in Table 7. The OLS estimate for the charter school effect is again positive but no longer statistically significant. The model of charter school attendance without any instruments has a positive R^2 value, and adding the educational attainment Herfindahl index or the income Herfindahl index increases the R^2 value. On the other hand, the race and achievement indexes have little explanatory power, as seen in their lower R^2 values. All of the instruments are exogenous and have the expected sign. Additionally, the first-stage coefficient for the educational attainment index is statistically significant. None of the second-stage coefficients for charter school attendance are significant,

however. Model 2, in which the instrument is the race index, gives particularly strange results. Models 5 through 7 try various combinations of the instruments, excluding the race index. Although these models have relatively good explanatory power for charter school attendance, they also do not yield significant second-stage estimates. Furthermore, as was the case in the regressions that included Pamlico County, no consistent pattern regarding the direction of selection bias emerges. Both with and without controlling for unobservable characteristics, charter schools have no statistically significant effect on the achievement of their students.

Discussion

Because students must choose to attend charter schools, there is the potential for selection bias when addressing the question of how charter schools affect the achievement of their students. Two previous studies have attempted to examine this question for North Carolina charter schools while controlling for selection bias. Greene, Forster, and Winters (2003) compare untargeted charter schools to nearby traditional public schools, reasoning that these schools will have similar student bodies, and Bifulco and Ladd (2004) use rich student-level panel data to measure each student's performance against his own past performance. In this paper, I attempt a third strategy: two-stage least squares.

A direction comparison of the charter and traditional public schools indicates the need to control for selection bias. The average charter school has more black students than the average traditional public school, as well as more highly-educated parents, and is located in a more urbanized zip code in a wealthier school district that spends more per student. To the extent that these factors influence achievements, there will be selection bias.

Without controlling for student characteristics, charter schools have lower achievement scores than traditional public schools. The important question to ask then is whether these lower schools are a result of the self-selected charter school student population or of the schools themselves. Controlling only for observable student characteristics, charter schools have a statistically significant positive effect on student achievement scores, but it is not enough to control only for observable characteristics.

To attempt to control for unobservable differences between charter and traditional public school students and eliminate selection bias, I use instruments that measure the diversity of school preferences in a district. Families in districts with a greater diversity of preferences will be more likely to exit to charter schools that can match their specific needs more exactly, yet the diversity of preferences will not affect student achievement. TSLS regressions using these instruments provide estimates of the charter school effect that vary widely in magnitude, but are not statistically different from zero. Excluding an outlying data point improves the fit of the first-stage regressions but still does not produce significant estimates of the charter school effect. This instrumental variables strategy thus does not provide a definitive answer to how selection bias influences the analysis of the effect of charter schools on achievement.

Instruments that are more highly correlated with charter school attendance would of course improve the TSLS results, but other improvements are possible, even without individual-level data. Attempting to distinguish between targeted and untargeted charter schools would provide an additional control for selection bias. A value-added model would also help control for selection bias, although such a model would be complicated by the fact that the student populations of schools change from year to year. In addition, the study would be strengthened by investigating alternative measures of charter school performance. Because families may be

choosing charter schools for reasons other than test scores, an analysis of the effects of charter schools on those outcomes, such as safety or teacher quality, would be useful.

Works Cited

Bettinger, Eric (1999). "The Effects of Charter Schools on Charter Students and Public Schools," Teachers College, Columbia University, National Center for the Study of Privatization in Education, Occasional Paper No. 4. Online: http://www.ncspe.org/publications_files/182_OP04.pdf

Bifulco, Robert and Helen F. Ladd (2004). "The Impacts of Charter Schools on Student Achievement: Evidence from North Carolina," unpublished.

The Center for Education Reform(a) (2004). "Charter School Facts." Online: http://www.edreform.com/index.cfm?fuseAction=stateStatChart&psectionid=14&cSectionID=44

The Center for Education Reform(b) (2004). "Charter Law." Online: http://www.edreform.com/index.cfm?fuseAction=cLaw

Evans, William and Robert Schwab (1995). "Finishing High School and Starting College: Do Catholic Schools Make a Difference?" *Quarterly Journal of Economics* 110: 947-974.

Glomm, Gerhard, Doug Harris, and Te-Fen Lo (2003). "Charter School Location," unpublished. Online: http://mypage.iu.edu/~gglomm/charter.pdf

Greene, Jay P., Greg Forster, and Marcus A. Winters (2003). "Apples to Apples: An Evaluation of Charter Schools Serving General Student Populations," Manhattan Institute, Education Working Paper 1. Online: www.manhattan-institute.org/ewp_01.pdf

Hanushek, Eric A., John F. Kain, and Steven G. Rivkin (2001). "Disruption versus Tiebout Improvement: The Costs and Benefits of Switching Schools," NBER Working Paper #8479.

Hanushek, Eric A., John F. Kain, and Steven G. Rivkin (2002). "The Impact of Charter Schools on Academic Achievement," unpublished. Online: http://edpro.stanford.edu/eah/papers/charters.aea.jan03.PDF

Intriligator, Michael D., Ronald G. Bodkin, and Cheng Hsiao (1996). *Econometric Models, Techniques, and Applications*, Upper Saddle River, NJ: Prentice Hall.

Kafer, Krista (2003). *School Choice 2003: How States Are Providing Greater Opportunity in Education*, The Heritage Foundation.

McNiff, Michelle Godard and Bryan C. Hassel (2002). "Charter Schools in North Carolina: Confronting the Challenges of Rapid Growth" in Vergari, Sandra ed., *The Charter School Landscape*, University of Pittsburgh Press, 208-229.

Neal, Derek (1997). "The Effects of Catholic Secondary Schooling on Educational Achievement," *Journal of Labor Economics* 15(1).

Neal, Derek (1998). "What Have We Learned about the Benefits of Private Schooling," *Federal Reserve Bank of New York Economic Policy Review*, 79-86.

Noblit, George W. and Dickson Corbett (2001). "North Carolina Charter School Evaluation Report." Raleigh: North Carolina State Board of Education. Online: http://www.ncpublicschools.org/accountability/evaluation/charter/evalreport.pdf

North Carolina Department of Public Instruction(a). "Frequently Asked Questions about Charter Schools." Online: http://www.ncpublicschools.org/charter_schools/facts.html

North Carolina Department of Public Instruction(b). "The ABCs Accountability Model," Online: http://www.ncpublicschools.org/abcs/

Sander, William and Anthony Krautmann (1995). "Catholic Schools, Dropout Rates, and Educational Attainment," *Economic Inquiry* 33: 217-233.

Stock, James H. and Mark W. Watson (2003). *Introduction to Econometrics*. Boston: Addison-Wesley.