NBER WORKING PAPER SERIES

SCHOOL CHOICE, COMPETITION, AND AGGREGATE SCHOOL QUALITY

Michael Gilraine
Uros Petronijevic
John D. Singleton

Working Paper 31328
http://www.nber.org/papers/w31328

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2023, Revised March 2025

We are grateful to Peter Arcidiacono, Natalie Bau, Sarah Cohodes, Angela Crema, Michael Dinerstein, Christopher Neilson, Miguel Urquiola, Iris Vrioni, and participants at ASSAs, Columbia University, Duke University, NBER Education, Purdue University, Simon Fraser University, SITE, University of Pennsylvania, Washington University in St. Louis, University of Wisconsin-Madison, and WEA for helpful comments. Ande Shen provided excellent research assistance. All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

School Choice, Competition, and Aggregate School Quality
Michael Gilraine, Uros Petronijevic, and John D. Singleton
NBER Working Paper No. 31328
June 2023, Revised March 2025
JEL No. H75, I21, I28

ABSTRACT

This paper develops and estimates an empirical model that allows for evaluating the impact of charter school choice on education quality in the aggregate. We estimate the model using student-level data from North Carolina. Using the model to counterfactually reverse North Carolina's removal of its statewide charter school cap, we show that the policy raised the average public school's value-added by around 0.01 (on the student test score distribution). This competitive response drives the aggregate learning gain, worth about $2,000 per student. We further use model simulations to quantify the expected human capital returns to the marginal charter school, which show that policies that steer entrants to disadvantaged and underserved neighborhoods stand to amplify the impact of charter school expansion.

Michael Gilraine
Economics Department
Simon Fraser University
West Mall Centre 3602
8888 University Drive
Burnaby, BC V5A 1S6
Canada
and NBER
gilraine@sfu.ca

Uros Petronijevic
York University
Department of Economics
Vari Hall
4700 Keele Street
Toronto, ON M3J 1P3
CANADA
upetroni@yorku.ca

John D. Singleton
University of Rochester
Department of Economics
Harkness Hall
Box 270156
Rochester, NY 14627
and NBER
john.singleton@rochester.edu

1 Introduction

Education policy debates often center on identifying the most effective reforms for improving student outcomes, and a range of policies—–such as school finance reforms, accountability measures, and school choice initiatives—–compete for limited public resources. In the U.S., school choice primarily takes the form of publicly funded and authorized, yet privately operated, charter schools, which have become a focal point in the broader discussion on how to improve educational quality.

Even though estimates of the returns to certain education policies, such as teacher quality improvements, are increasingly available, there remains little consensus on whether charter school expansions have historically improved outcomes for the average student. This is in part because, absent plausible natural experiments that generate exogenous variation in charter school policy across markets, the literature tackles pieces of the larger puzzle separately: one strand examines whether students who attend charter schools benefit, while another focuses on spillovers affecting those remaining in traditional public schools.[1] Moreover, although factors such as the types of charter schools authorized, their locations, and their consequence interactions with public schools are likely to influence the return to charter expansion, the existing literature lacks ex ante policy analyses that can guide policy design.

In this paper, we develop and estimate an equilibrium model that links the presence and characteristics of charter schools with the level and distribution of student outcomes. The model incorporates heterogeneity across households and differences in school quality that undergird attendance effects while also capturing how public schools respond to charter school competition. This integrated framework allows for conducting counterfactual analyses, which we use to evaluate the aggregate impact of charter school expansion on students ex post as well to quantify how policy can influence that impact by influencing the types, qualities, and locations of charter schools.

We estimate the model using rich, geocoded student-level data from North Carolina, which provides detailed information on student enrollment patterns and test scores.[2] We focus the analysis on elementary grades in the three largest Commuting Zones (CZs), where around 8% of students attended a charter school during the 2015-16 school year. The test

[1] Surveys of the charter school literature include Epple et al. (2016) and Cohodes and Parham (2021).

[2] The data are provided by the North Carolina Education Research Data Center (NCERDC).

score panel allows us to measure school quality using school-level estimates of value-added, while the information on individual students' demographics and residences enables flexibly estimating demand controlling for time-varying school-level unobservables.[3] Another key source of variation comes from North Carolina's decision to lift its statewide cap on charter schools in 2011. We use this policy change in two ways: First, to develop empirical moments that discipline how public school quality in equilibrium depends on (non-random) competitive exposure to charter schools; Second, as a case study for our policy analysis evaluating the aggregate impact of expanded charter school choice.

We begin by documenting several facts about the removal of the charter school cap that motivate the case study. First, while our estimates of value-added show that public schools and charter schools that opened prior to the cap being lifted are of a similar quality on average, we find that the average post-cap charter entrant is appreciably lower quality. This raises the question whether students who attend a charter school because of the expansion benefit, which will depend on the nature of selection. Second, using a difference-in-differences framework that compares the changes in value-added of public schools located nearby a newly-opened charter school with those located farther away, we find that public schools more competitively-exposed by the cap lifting increase their quality on average. We also document important heterogeneity: public schools' value-added only increases following the nearby entry of a charter school offering of a traditional curriculum.[4] This finding suggests that curricula horizontally differentiate charter schools, a possibility our model tests for, and raises the questions whether and how policy should consider this heterogeneity.

We estimate the model in two steps. We use student-level enrollment choices and the school quality estimates to first estimate demand. To then estimate the supply-side link between demand incentives and public schools' qualities, we leverage the differential exposure of public schools to competition following the lifting of the charter school cap, described above, as instruments for their own-quality (semi-)elasticity of demand. The demand estimates indicate that the average student would be willing to travel around 0.1 miles for a 10 percentile

[3]As discussed later, we employ an empirical Bayes estimator for the value-added estimates to address issues of statistical noise.

[4]This variable is manually-coded using information gleaned from charter schools' applications to the State Board of Education. A traditional curriculum stands in contrast with those that offer, as examples, project-based or experiential learning. Though the design differs in several ways, this result parallels the findings in Gilraine et al. (2021). Below, we discuss the prior paper in expanded detail.

point improvement in school quality.[5] But there is wide heterogeneity, with quality-sensitive students having the weakest preferences for charter schooling (all else equal). The estimates also allow us to characterize how curricula horizontally differentiate charter schools in terms of elasticities of substitution. We find that the the average charter school would lose fewer than 5% of its students were all public schools to raise their quality by 0.05σ (on the student test score distribution), but the average traditional charter school would lose around twice as many students as the average non-traditional charter.

Estimates in hand, we return to the case study to estimate the aggregate impact of charter school expansion on students. We do this by counterfactually removing the charter schools that opened after the 2011 cap lifting from markets and solving for school choices and public schools' qualities in that equilibrium to compare the predictions for student outcomes with the data. We find that the share of students attending a charter school would be 2.3 points lower absent the cap lifting. In that equilibrium, the average public school would supply about 0.01σ lower school quality (on the student test score distribution) and the average student's test scores would be about 0.005σ lower than in the data. The results further indicate that economically disadvantaged Black and Hispanic students benefited relatively more from the policy change.

Is the estimated impact on the average student's test scores from raising the charter school cap meaningfully large? We benchmark this by way of comparison with other education policies: the effect is roughly equivalent to one-fifth the estimated gain from replacing the bottom 5% of all teachers (Gilraine et al., 2020) and is around 50% larger than the impact of a \$1,000 per student increase in capital funding (Biasi et al., 2024). The estimated return to earnings from the first policy (Chetty et al., 2014b) imputes an increase of nearly \$2,000 in lifetime earnings for the average student arising from the policy-induced 2 point increase in charter school share.

These findings raise the question of how policy influences the impact of charter school expansion on students. We approach this by using the model as a framework for studying the problem facing charter school authorities: whether to approve or deny the marginal application to open a charter school. Charter school authorization frameworks commonly

[5] A point of comparison is with Campos (2024), who finds that families of LA high schoolers would travel 0.44 more miles to attend a school 10 percentile points better in school quality. Our smaller estimate is consistent with travel costs being much more salient at the elementary school level.

prioritize assessing the applicant's quality, which is difficult to predict ex-ante.[6] Those frameworks also ignore the potential for competitive spillovers and typically do not directly weigh in a charter school's curriculum type and intended location.[7] We do this by using simulations to build a matrix of ex-post student outcomes, wherein each cell represents the combination of where an additional charter school in the market locates (defined by Census tract); its type (traditional or non-traditional); and its quality. We juxtapose several possible locations types. The matrix is then used to approximate the expected return to the marginal charter school given that type and location are known, but quality is not.

The expected returns to the marginal charter school that we estimate contain several insights for policy. First, the expected return is always positive in the locations we consider, implying that approving the marginal applicant maximizes test scores. We show this occurs due to competitive effects: in fact, the expected effect of the marginal charter school on the students who will attend it is actually negative.[8] The findings also indicate that curricular information is useful—an authorizer would need better information about non-traditional charter schools to be indifferent to type—and that policies that successfully steer charter schools to disadvantaged and underserved neighborhoods stand to meaningfully amplify the returns to charter school expansion. We find that the expected impact of the marginal traditional charter school in a low income location is twice as large as if placed in a high income location. The numbers suggest that the policy effect would have been 8% larger (in expectation) had the cap lifting policy included just one additional traditional charter school in a low income location.

Contribution to the literature: Our paper primarily connects to a large literature that examines whether, and under what condition, policies that expand school choice improve educational outcomes for primary and secondary school students.

The literature on charter schools in the U.S. broadly subdivides into studies that ask whether students who attend charter schools benefit and those that test for competitive

[6] A second priority is typically that there is likely to be demand for the charter school.

[7] However, the presence of other charter schools already and likelihood the new charter will fiscally strain local public schools are factors explicitly weighed negatively by some authorizers.

[8] This parallels what we find about the charter cap lifting—sorting to the new charter schools attenuates the aggregate test score impact on average. In other words, the average student who remains in public school after the charter expansion benefits more than the average student.

spillovers on students in public schools.[9] While lottery-based evidence demonstrates that certain oversubscribed charter schools are highly-effective (Chabrier et al., 2016), test score value-added and matching estimates suggest that many charters are not better than the average public school (CREDO, 2009; Singleton, 2019).[10] Adherence to so-called "No Excuses" practices—a label describing programs that emphasize high expectations, comportment, and core math and reading skills (Carter, 2000; Thernstrom and Thernstrom, 2004)—predicts charter school effectiveness (Angrist et al., 2013; Dobbie and Fryer, 2013). The evidence on charter school competition likewise suggests that impacts depend on charter schools' characteristics. Gilraine et al. (2021), who use data from the same setting as the current paper, find evidence supporting competitive responses by public schools—except those exposed to "horizontally differentiated" charter schools.[11] Tobin (2024) meanwhile finds negative effects of charter school expansion and presents evidence this is driven by competition with for-profit charter schools.[12] Earlier evidence that does not consider such heterogeneity is mixed (e.g. Bettinger 2005; Sass 2006; Zimmer and Buddin 2009; Imberman 2011).

We contribute to the literature by developing an empirical model that places attendance effects and competitive impacts in a common framework for policy analysis. The model shares features in common with Ferreyra and Kosenok (2018a) and Walters (2018), who estimate demand models to assess impacts of charter school expansion on enrollees' welfare and outcomes, respectively. A primary feature of our model is that equilibrium public school quality responds endogenously to competition from charter schools, allowing us to evaluate impacts on the learning outcomes of all students.[13] A second feature of our model is that

[9]A distinct concern is that charter schools may instead have negative spillovers on public schools through the channel of peer composition. These effects could arise from peer effects, where charter students are positively selected, or from increases in the cost of education and could offset public school productivity gains. Our model to come does not allow for either. Our prior work (Gilraine et al., 2021) rules out these possibilities in this context by leveraging the year gap between charter approval and opening and showing that the treated public schools responded once they knew the charter would open nearby, but before the charter had actually opened and could affect peers or resources.

[10]Evidence from lotteries includes Hoxby and Murarka (2009); Abdulkadiroglu et al. (2011); Angrist et al. (2016b); Dobbie and Fryer Jr. (2015).

[11]Having a non-differentiated (i.e. traditional) curriculum also correlates highly with adherence to "No Excuses" practices (Gilraine et al., 2021). Slungaard Mumma (2022) also reports that spillover effects of academically-focused charters are more positive.

[12]The paper argues this is due to competition on non-academic amenities, which for-profit charter schools tend to boast more of. The paper focuses on public middle schools, which have more flexibility to adjust non-academic offerings than do elementary schools (this paper's focus).

[13]Mehta (2017) also models competition between charter and public schools in North Carolina, though for an earlier period. Beyond several differences in modeling framework, the present paper is distinct in a)

demand for a charter school depends on its curriculum, which may be traditional in its focus on core academic skills or not. This allows for quantifying the degree to which curricular differentiation softens quality competition (MacLeod and Urquiola, 2013) and for quantifying its aggregate importance.[14]

We estimate the aggregate impact of charter school expansion by simulating student outcomes when charter schools are counterfactually removed from markets. Chen and Harris (2023) instead approaches that question using national data on charter enrollments and student outcomes. Though focused on the same specific policy change as Gilraine et al. (2021) (the cap lifting in North Carolina in 2011), this paper's approach addresses a major limitation of that paper's estimated "policy effect": it applies only to students attending a public school prior to cap lifting, whose experience is unlikely to represent the average student's. A point our policy analysis illustrates is that expanded school choice can yield (expected) gains to the average student via the competitive channel, even when school choice alternatives are not on average better than traditional schooling options.

Finally, the insights we develop regarding how policy can enhance the impact of charter school expansion by influencing the type and location of charter schools are new to the literature, which contains limited evidence to inform charter school accountability and authorization frameworks.[15]

The remainder of this paper is organized as follows. Section 2 discusses the data used in our analysis, highlighting the institutional features and key data patterns. We then present the model in Section 3 and outline our estimation approach in Section 4. We present the estimated model parameters in Section 5 and, in Section 6, use the model to counterfactually evaluate the effects of several policies of interest on student outcomes. Section 7 offers concluding remarks.

---

connecting the model to school value-added; and b) using policy variation related to the cap lifting in North Carolina for identification.

[14]Along related lines, Epple et al. (2021) present a model where charter schools endogenously choose educational practices, which includes their curriculum, but do not consider the implications for public school productivity. The equilibrium effects of horizontal competition, though in terms of student-school match quality, in Pakistan are examined by Bau (2022).

[15]Such frameworks commonly prioritize achievement impacts, but test scores are only available for decisions about which charters to *renew*, not which applications to approve (Bross and Harris, 2016). The most relevant work shows that replicated "child" schools of effective charter schools are as effective (Cohodes et al., 2021), suggesting that an operator's track record can be a reliable indicator of quality.

2 Data Description and Key Data Patterns

In this section, we describe the setting and data. At a high level, three features of the data and institutional environment are of particular importance to our approach. First, the detailed student-level data allow us to construct choice sets for each student based on their residence and estimate the quality of each school. Second, charter school applications allow us to determine whether each charter school offers a traditional or non-traditional curriculum. Third, the removal of the state-wide cap on charter schools provides variation in the competitive pressure faced by public schools, which is used to inform the model estimates. We now discuss each of these features in turn.

2.1 Data

Our detailed, student-level administrative records are provided by the North Carolina Education Research Data Center (2008-2017). These data include information on all North Carolina public school students (charter and traditional public) for the 2007-08 through 2016-17 school years. The data contain test scores for each student in mathematics and reading on standardized end-of-grade exams in grades three through five. We standardize these test scores at the student level to have a mean of zero and a variance of one for each grade-year to ensure comparability of test scores across grades. We use these test scores to construct measures of school quality (see Section 2.2.1).

In addition to test scores, the student data contain information regarding each student's grade, socioeconomic status, and ethnicity. We also obtain information regarding students' residential locations in each school-year from the NCERDC, a crucial input into our demand model given the importance of distance in determining school choice. For confidentiality reasons, student location in the NCERDC data is reported at the Census block group level. We therefore define each student's location as the centroid of the block group in which he or she resides.[16] We link these student-level data to the universe of public and charter schools in the state. School locations are available from the Common Core of Data, allowing us to compute distances between students' residences and all schools in their education market.

Our empirical model also allows student demand to be a function of the type of curriculum

[16]The median area of a Census block group in North Carolina is 2.2 square miles.

a charter school offers. We classify charter school curricula following the methodology in Gilraine et al. (2021) using charter school applications. In particular, we use the information contained in the applications to manually classify each charter school as either following a "traditional" or "non-traditional" curriculum.[17] To do so, we classify charter schools that emphasize project-based or experiential learning (including Montessori) in their application as following a non-traditional curriculum. Charters are otherwise classified as following a traditional curriculum, which usually entails a focus on core math and reading skills. Importantly, we classify all charter schools in this way, including both those who opened prior to the charter school cap being lifted and those who opened after the removal of the statewide cap.[18]

Defining Markets: We focus our analysis on elementary grade-level school markets. As charter schools have no defined attendance zones, students can attend these schools even if they live outside the geographic school district the charter is located in. We therefore define education markets in our data based on Commuting Zones ('CZs'), which are aggregations of counties based on commuting patterns in the 1990 Census. These Commuting Zones are designed to span the area in which people live and work and therefore provide a natural way to partition markets. We focus our analysis on the three largest commuting zones in North Carolina: Charlotte, the 'Research Triangle' (i.e., Raleigh-Durham-Chapel Hill), and Greensboro-High Point. These CZs cover 60 percent of students in North Carolina. Furthermore, as charter schools tend to locate in more urban areas and these CZs include the five largest cities in the state, our data consist of 70 of the 114 elementary charters in the state (as of 2015-16).[19]

[17] The recent charter school applications are available online at `https://www.dpi.nc.gov/students-families/alternative-choices/charter-schools/applications/submitted-apps`. Charter school applications pre-2012 are available by request from the North Carolina Department of Public Instruction.

[18] An alternative way to categorize charter schools is by whether they contract with for-profit organizations for management services (e.g. Singleton 2017). We focus on curricula to connect with the literature on school practices (Angrist et al. 2013; Epple et al. 2021) and because we model parental demand over school characteristics. While parents can observe and assess a charter school's curriculum, it is not clear that profit status matters independently. This rationale is consistent with Tobin (2024), which uses for-profit status to proxy for a charter school offering more non-academic amenities.

[19] Similarly, two thirds of the charter elementary schools that opened following the cap lifting in 2012 located in one our sample commuting zones. These markets are quite geographically distinct: while there are two rural border counties between the Triangle and Greensboro CZs, students residing there are a very small minority of the sample.

Data Summary: Table 1 reports summary statistics, with column (1) doing so for the entire state while column (2) restricts the sample to our CZs of interest.[20] Overall, our restricted sample looks similar to the state overall in terms of demographics. Our three CZs are, however, more urban portions of the state and so students reside closer to schooling options and charter share is higher.

Columns (3) and (4) show summary statistics for our CZs of interest pre-charter cap lifting (2011-12) and post-charter cap (2015-16). Crucially, we observe a large increase in the number of charters post-cap which leads to a corresponding increase in charter share (from 4 to 8 percent) and a decline in distance to the nearest charter option (from 7.8 to 6.7 miles). Column (3) also indicates that pre-charter cap there was very limited enrollment in charters that follow a non-traditional curriculum. The post-cap charter share increase, however, was disproportionately concentrated in these non-traditional charters as their share increased by 138% (versus a 82% increase in charter share).

2.2 School Quality, Policy Backdrop, and Quasi-Experimental Variation

Central to our analysis is a measure of school quality. In this section, we first briefly describe how we estimate the academic quality of each school, followed by a description of several key patterns related to the removal of North Carolina's statewide cap on charter schools.

2.2.1 Estimating School Quality

We estimate school quality using mathematics[21] test score value-added.[22] In the model of student achievement we take to the data, education inputs (including school quality) are

[20]Table 1 reflects the sample restriction to students in grades K through 2 in the demand estimation, which we discuss later.

[21]Results using English test score value-added are the same sign and statistically significant – see Figure A.3 – but roughly one-third the magnitude, in line with the broad education literature whereby English test score responses are more muted.

[22]Principally known for its application to measuring teacher quality (e.g. Kane and Staiger 2008; Chetty et al. 2014b), test score value-added estimates rest on a selection-on-observabales assumption (where the conditioning set importantly includes students' lagged performance) but have been validated using random assignment (Deming et al., 2014; Angrist et al., 2017) and linked to students' long-run success (Kirkebøen, 2022). Lottery-based estimates of charter school impacts make weaker assumptions, but have the distinct disadvantage that they are only available for oversubscribed charter schools (whose lottery information can also be linked to student outcomes) and so cannot produce a full picture of the distribution of school quality.

Table 1: Summary Statistics

	All K-2 North Carolina Students[1] (1)	Demand Estimation Sample[2] (2)	Split by School-Year: Demand Estimation Sample (2011-12) (3)	Demand Estimation Sample (2015-16) (4)
Student Demographics				
% White	50.2	48.8	51.2	46.2
% Black	23.7	24.0	23.0	24.9
% Hispanic	17.8	19.1	18.3	19.9
% Asian	3.1	4.2	3.6	4.9
% Economically Disadv.	52.2	47.6	50.7	44.3
% Charlotte CZ	26.2	45.3	46.4	44.0
% Research Triangle CZ	21.2	36.6	36.6	36.7
% Greensboro-High Point CZ	10.5	18.1	17.0	19.3
School Attendance Summaries				
% Attend Assigned Public[3]	70.1	57.5	57.7	57.2
% Attend Charter School	5.39	6.30	4.45	8.11
% Attend Non-Traditional Charter	1.85	1.70	1.00	2.38
Distance to School Attended[4]	2.29	2.09	2.09	2.10
Distances to School Options				
Distance to Nearest Public (miles)[4]	1.70	1.48	1.49	1.46
Distance to Nearest Charter (miles)[4]	8.73	7.22	7.75	6.66
Observations (student-year)	588,765	337,855	165,959	171,896
# of public schools	1,237	606	597	595
# of charters	118	73	44	70

Notes:
[1] Data coverage: Kindergarten through second grade students for the 2011-12 and 2015-16 school years.
[2] Same as for column (1), but restricted to the three largest commuting zones in North Carolina: Charlotte, the Research Triangle, and Greensboro-High Point.
[3] Only reported when student residence is observed and can be assigned to a school attendance zone. The sample for these summary statistics is 360,050 and 277,865 student-year observations in columns (1) and (2), respectively. Note that while the link between student residence and school attendance zone can be missing for some rural areas it is near-universal in urban areas.
[4] Only reported when student residence is observed. The sample for these summary statistics is 484,300 and 280,740 student-year observations in columns (1) and (2), respectively.

additive in their effects and the achievement of a student i at school s in year t is written as:

$$y_{ist} = X'_{ist}\beta + q_{st} + \epsilon_{ist} \tag{1}$$

where y_{ist} is the student's mathematics test score, X_{ist} is a large vector of observable individual and school-level student characteristics, and ϵ_{ist} is a random test score shock which is assumed to be iid normal with variance σ^2_ϵ. The estimated value of q_{st} represents the contribution of school s in year t to test scores that is unexplained by student characteristics, or simply its value-added.

Crucial to the validity of value-added models is that the control vector, X_{ist}, is sufficiently rich so that potential test scores are independent of a student's school choice, conditional on observables. While not perfectly unbiased, school value-added measures tend to feature limited bias when lagged test scores are included in the control vector (Deming, 2014; Angrist et al., 2016a, 2017). We therefore control for flexible functions of lagged test scores, along with student-level and school-grade demographics.[23] We estimate school value-added using all fourth and fifth grade students in North Carolina from 2008-09 through 2015-16. Table A.1 provides summary statistics for the sample used to estimate school VA. Following the literature, the school's value-added in each year, q_{st} is estimated using empirical Bayes to minimize mean squared error.[24] Appendix B tests for forecast bias in our setting using a school-switcher quasi-experiment. We find that our forecast coefficients are close (but not equal) to 1, indicating that there is some bias, but it is limited.[25] Our view is therefore that using school VA as a proxy for school quality is unlikely to materially affect our results,

[23]Specifically, X_{ist} includes: (i) lagged test scores using a cubic polynomial in prior-year scores in math and English, interacted with grade dummies, (ii) demographics, including: economically disadvantaged status, ethnicity (six ethnic groups), gender, limited English status, gifted status, and disability status. We also include the following school-grade level controls: (iii) cubics in school-grade means of prior-year test scores in math and English (defined based on those with non-missing prior scores) interacted with grade dummies, (iv) cubics in school-grade means of all the demographic covariates, (v) school-grade size, and (vi) grade-by-year dummies. Note that our VA measure controls for peer influences, although we obtain similar results if we omit peer influences as the correlation between the two VA measures exceeds 0.9.

[24]Formally, $q_{st} = y_{st}\frac{\sigma^2_s}{\sigma^2_s + \sigma^2_\epsilon / n_{st}}$ where $y_{st} \equiv \frac{\sum_i^{n_{st}}(y_{ist} - X'_{ist}\beta)}{n_{st}}$ is the fixed effect of school s in year t in equation (1), n_{st} is the number of students in school s at year t, and σ^2_s and σ^2_ϵ are the variances of y_{st} and ϵ_{ist} (which we estimate via maximum likelihood estimation and plug-in).

[25]Our estimates of forecast bias are near-identical to Angrist et al. (2017) who use school choice lotteries to test the validity of school VA models and conclude that "The test results [...] suggest conventional VAM estimates are biased. At the same time, OLS VAM estimates tend to predict lottery effects on average, with estimated forecast coefficients close to 1. OLS estimates would therefore seem to be useful even if imperfect."

in line with a large structural education literature using school VA as a proxy for school quality (Neilson, 2017; Ferreyra and Kosenok, 2018b; Singleton, 2019; Allende, 2019; Bau, 2022; Crema, 2024; Corradini, 2024; Campos and Kearns, 2024).

2.2.2 Lifting of Charter School Cap, New Charter Entrants, and Effects on Public Schools

In this section, we examine several data patterns that motivate the policy analysis of North Carolina's removal of the cap on the number of charter schools to come. Initiated by its receipt of a federal Race to the Top grant, on June 6, 2011, North Carolina lifted its 100-school cap on the number of charter schools allowed to operate—a limit that was in place since charter schools first emerged in the state in the 1996-97 academic year. Rapid growth of charter schools followed, with there being just shy of 100 schools prior to the cap's lifting and 176 charter schools in operation by the 2016-17 academic year. (Figure A.1 displays the number of charters in North Carolina and our three markets of interest from 2008-09 through 2015-16.)

In Figure 1(a), we show the distribution of school value-added in 2015-16 across public and charter schools in North Carolina, subdividing the charter schools among pre-existing charters that opened before the charter cap was lifted and those that opened after. We also distinguish between enrolllment-weighted (the dark lines) and unweighted (the light lines) distributions. Value-added for pre-existing charters is, on average, slightly below that of public schools in North Carolina (by 0.05), but the distributions are relatively similar. The newly-opened charters, however, have significantly lower VA than the public schools (by 0.14). This difference is not a simple artifact of the new charter schools being young and there being learning-by-doing: Figure A.2 compares the newly-opened charter school value-added distribution in 2015-16 to 2018-19 (last year of data available) and shows that the two distributions are similar, with the mean value-added for these newly-opened charters only increasing by 0.007 over the three additional years. Figure 1(b) then contrasts the VA distribution for charters by curriculum type: The value-added for charters that follow a traditional curriculum is significantly higher than those that follow a non-traditional curriculum (by 0.10). These descriptive statistics are suggestive that the lifting of the charter school cap may actually make students who take up the new charter options *worse off* (in terms of school quality) due to the lower quality in these schools.

Comparing the enrollment-weighted VA distributions to their unweighted counterparts highlights that higher quality schools enroll more students as one would expect when students value school quality. Interestingly, the only school type where higher quality schools do not attract substantially more students is non-traditional charter schools, which could be consistent with the students choosing these schools having weak preferences for school quality.

School Summary Statistics: We next report school summary statistics in Table 2. The average charter school has a similar racial composition of students to the average traditional public school but it is smaller in size and has a much smaller fraction of economically disadvantaged students. In terms of location, compared to traditional public schools, charter schools tend to locate in more racially diverse, lower income, and dense census tracts, where higher fractions of the population have four-year college degrees. These differences are especially pronounced when comparing the locations of non-traditional charter schools to the average traditional public school.

Figure 2 visualizes the locations of charter schools by type. Specifically, it displays the population distribution through the 'Research Triangle' Commuting Zone and then overlays the location of charter schools that opened after the cap lifting, differentiating these charters by curriculum. We see that most charters locate in the densest region of the Commuting Zone formed by the 'triangle' of Raleigh-Cary, Durham, and Chapel Hill. The urban preference of charters—especially non-traditional charters—drives these aforementioned differences in location characteristics between charters and traditional public schools.

The last row in Table 2 shows that traditional charter schools are near public schools with similar value-added to their own. In contrast, non-traditional charter schools are located near higher value-added public schools compared to their value-added. We also look at the quality of nearby schools. Here, we see that the value-added of the school nearest to the non-traditional charter fell by 0.021σ from 2011-12 to 2015-16. In contrast, the school nearest to the traditional charter has a large increase in value-added of 0.031σ from 2011-12 to 2015-16. Although suggestive, these shifts are indicative that nearby schools responded to the post-cap entry of traditional charter schools by raising their quality, while those nearby non-traditional charter schools did not. This evidence, however, is merely suggestive and so

Figure 1: Distribution of Public and Charter School Value-Added

(a) Value-Added for Public Schools and Pre- and Post-Cap Charters

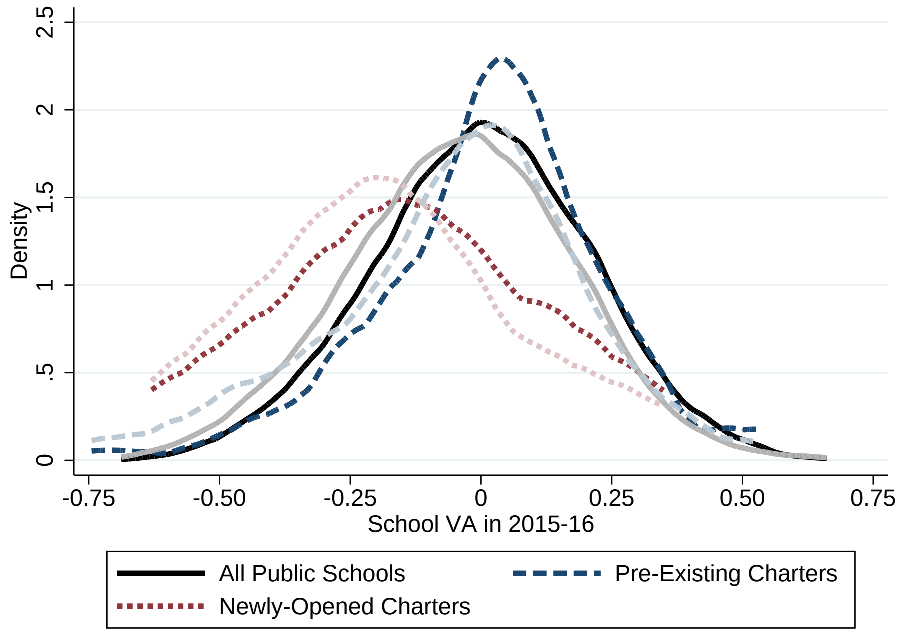

(b) Charter School Value-Added by Curriculum

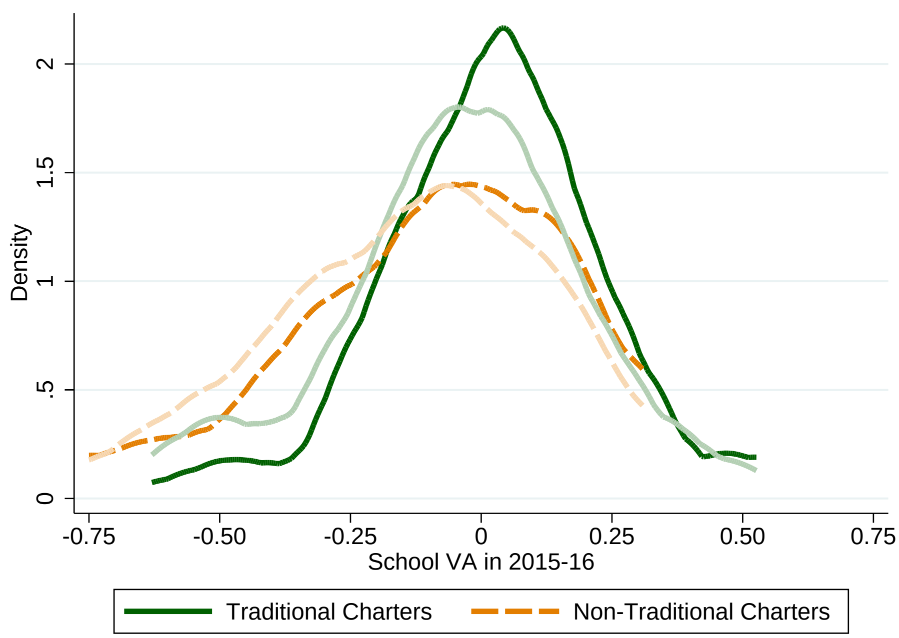

Notes: This figure shows both enrollment-weighted and unweighted distributions of school value-added for the 2015-16 school year. Enrollment-weighted distributions appear in dark colors and the unweighted distributions are shown using associated lighter shades. Figure 1(a) displays the value-added distributions separately for public and charter schools. The charter school VA distribution is further subdivided into 'pre-existing charters' which opened prior to the charter cap being lifted (i.e, pre-2012-13) and 'newly opened' charters that opened after the charter cap was lifted (i.e., 2012-13 or later). Figure 1(b) then presents the value-added distributions separately for charter schools that follow a traditional and non-traditional curriculum.

Table 2: Summary Statistics for Schools in 2015-16

	Public Schools (1)	Charter Schools (2)	Traditional Charter Schools (3)	Non-Traditional Charter Schools (4)
School Characteristics				
% White	47.9	50.1	52.1	46.5
% Black	25.1	29.5	30.8	27.2
% Hispanic	18.8	9.6	9.1	10.4
% Asian	4.0	5.0	3.3	8.0
% Economically Disadv.	46.7	28.6	29.3	27.3
School Size (K-2 only)	261.2	196.0	215.3	161.2
Value-Added	0.006	-0.028	-0.003	-0.072
Location Characteristics (Census Tract)				
% White	71.12	61.87	64.90	56.41
% Black	23.02	30.50	28.4	34.12
% Hispanic	9.35	10.75	10.36	11.47
% Asian	3.71	4.62	3.69	6.28
% Population in Labor Force	64.67	67.45	67.35	67.64
% Population with 4-yr College Degree	30.90	38.01	37.25	39.38
Density (Population/Square Mile Area)	1,259.07	2,087.85	1,673.74	2,833.26
Median Household Income ($ 2017)	60,046.44	57,890.03	60,847.09	52,567.32
Total K to 8 School Enrollment	783.23	748.24	785.89	680.48
Distances to Nearby Schools				
Distance to Nearest Public School (miles)	2.56	1.46	1.44	1.51
Distance to Nearest Newly-Opened Traditional Charter School (miles)	13.65	10.45	10.22	10.86
Distance to Nearest Newly-Opened Non-Traditional Charter School (miles)	18.88	13.49	15.17	10.45
Value-Added of Nearest Public School				
Value-Added of Nearest Public in 2011-12	0.024	0.019	0.025	0.007
Value-Added of Nearest Public in 2015-16	0.000	0.031	0.056	-0.014
Own Value-Added (2015-16) Minus Nearest Public Value-Added (2011-12)	-0.019	-0.040	-0.022	-0.070
# of schools	595	70	45	25

Notes: Data cover all schools in the three largest commuting zones in North Carolina: Charlotte, the Research Triangle, and Greensboro-High Point.

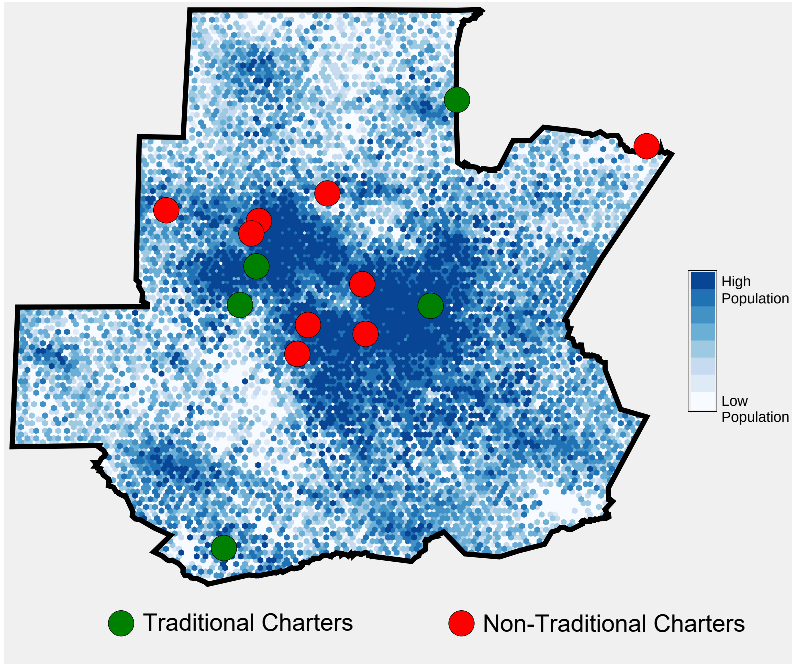

Figure 2: Population Density in the 'Research Triangle'

Notes: This figure shows the how the population is distributed across the 'Research Triangle' (Raleigh-Durham-Chapel Hill) Commuting Zone according to the 2010 Census. Specifically, we assign Census blocks to over 20,000 equidistant hexagons that we have placed throughout the commuting zone and then display the total population of all Census blocks located in that hexagon. The location of charter schools that opened post-cap are overlaid, differentiated by curriculum.

we now leverage quasi-experimental variation to show this differential response to charter entry by curriculum type.

The Impact of Charter Openings on Nearby Public School Quality and Enrollment: We motivate our upcoming analysis by examining the impacts of charter school openings on nearby public schools. We do this using a combination of spatial variation in the cross-section and policy variation over time. This analysis thus has features in common with our prior work (Gilraine et al., 2021), but focuses on school-level impacts on public schools' enrollment and value-added (quality); these moments will inform the estimates of the structural model. Appendix C provides full details.

We compare changes in school enrollment and quality for public schools located near the newly-opened charter schools (treatment) to those farther away (control) following the removal of North Carolina's statewide charter school cap. We focus on charters opening in the first two years after the charter cap was lifted. This restriction provides cleaner pre-post comparisons since public schools knew where these first two cohorts of charters would open at the 2012-13 school year and so competition-induced responses would start in 2012-13.[26] In our three CZs of interest, this leaves us with 15 elementary charter openings to exploit. We then take all schools within 30 miles of a newly-opened charter and define 'treated' schools as those within 20 miles of the newly-opened charter and all those further away as 'control' schools.[27] This assumes that 'control' schools are unaffected, but among 'treated' schools we allow distance to scale the *intensity* of the treatment. We regress each outcome (either school enrollment or value-added) on the treatment indicator interacted with a post-cap dummy; the interaction of treatment, the post-cap dummy, and treatment distance; school-by-event fixed effects, and district-year fixed effects. Finally, the treatment variables are interacted with whether the newly-opened charter follows a traditional or non-traditional curriculum so that we can see whether effects differ based on the curriculum of newly-opened charters.[28]

Panel A of Figure 3 plots the estimated coefficients from an event-study specification that interacts the treatment indicator with year dummies (rather than a simple post-charter cap dummy).[29] It does so separately for when the newly-opened charters follow a traditional curriculum (Figure 3(a)) and a non-traditional curriculum (Figure 3(b)). The figures reveal that there is no evidence of significant differential trends prior to the cap lifting between control schools and schools that were treated by the entry of a traditional or non-traditional charter school. Once the charter cap lifts, however, we see a substantive increase in value-added in the treated schools compared to the control schools when the nearby charter follows a traditional curriculum. The point estimates plotted correspond to the impacts when a charter school opens next door (i.e. distance zero) to the incumbent public school. No such value-added increase is observed when the nearby charter follows a non-traditional curriculum.[30]

[26] This comes from the fact that after the charter cap was lifted there was a 'fast track' application where the waiting period was waived. As charters usually had to wait one year to open, this meant charters applying to open in either the 2012-13 or 2013-14 school year had their publicly-available applications submitted and approved by 2012-13 and so nearby schools would know of these entry plans and be able to respond in 2012-13. See Gilraine et al. (2021) for a more detailed timeline on the application and approval dates.

[27] Figure A.4 – alongside Gilraine et al. (2021) – shows robustness to the choice of treatment distance.

[28] For inference, we two-way cluster our standard errors at each charter opening event and at the school level since our data are stacked and may feature multiple observation per public school. As we have few clusters (i.e., 15 charter opening events), all inference is conducted via the wild cluster bootstrap (Cameron et al., 2011; MacKinnon and Webb, 2018).

[29] Figure A.3 replicates Figure 3 but replacing mathematics VA with English VA: Results are the same sign and statistically significant but roughly one-third the magnitude.

[30] These results stand in contrast with Tobin (2024), who similarly examines the 2011 cap lifting in North Carolina and finds negative effects of charter competition. However, Tobin (2024) focuses instead on middle schools, which have more flexibility to adjust non-academic offerings. The finding that negative effects are

Panel B shows the results for school enrollment: a similar pattern is evident whereby schools near a newly-opened charter following the traditional curriculum lose enrollment, while those near a newly-opened charter using a non-traditional curriculum are unaffected.

Table A.2 reports the results of the difference-in-differences regressions, which indicate that, relative to the control schools, treated schools near a newly-opened traditional curriculum charter see a 0.03 increase in value-added and a nine-student decline in enrollment. In line with the visual evidence, minimal effects are seen for schools located near non-traditional curriculum charters. Note that Figure A.4 shows our difference-in-differences point estimates for various definitions of treated, making clear that we obtain similar results when we shrink the treatment radius all the way to 5 miles (below that we have too few schools for meaningful results).[31] These event-study results show that charter competition causes nearby schools to lose enrollment which forces them to compete by raising quality, at least when the newly-opened charter follows a traditional curriculum.

The lack of enrollment response among public schools nearby charters that follow a non-traditional curriculum raises the natural question of where students attending these charters come from. We resolve this dissonance by investigating how nearby private school enrollment responds in Figure A.5. We find enrollment declines among private schools nearby a newly-opened non-traditional charter. No such enrollment declines are observed for private schools nearby charters that follow a traditional curriculum. Since private schools also have the leeway to offer a non-traditional curriculum, these results support the hypothesis that charter schools offering traditional curriculum compete with public schools while those offering non-traditional curriculum are competing with private school or home-schooling options.

concentrated where for-profit charters aligns with ours of no significant test score effects when public schools compete with non-traditional charter schools. At the same time, Tobin (2024) compares districts where new charters opened to those where they did not, while we compare public schools based on their proximity to new charters while controlling for district fixed effects. It is therefore possible that, notwithstanding the differences across middle and elementary schools' abilities to offer non-academic programming, Tobin (2024)'s (across-district) negative effect and our (within-district) positive effect can both be true.

[31] Recall that because we interact treatment with distance in each specification, the plotted point estimates always correspond to the impact of a charter school opening next door (i.e. distance zero) to a public school relative to the distance between a new charter and public school being greater than the treatment radius.

Panel A: Value-Added

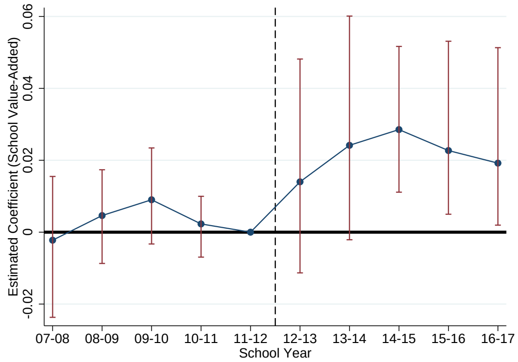

(a) Charter has Traditional Curriculum

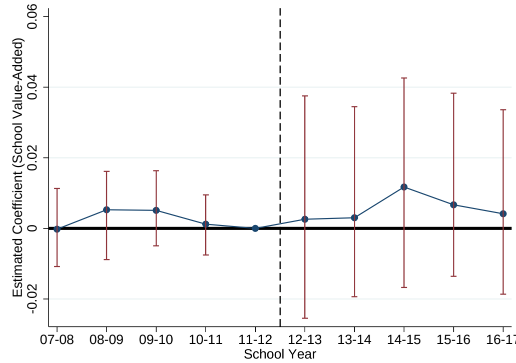

(b) Charter has Non-Traditional Curriculum

Panel B: Enrollment

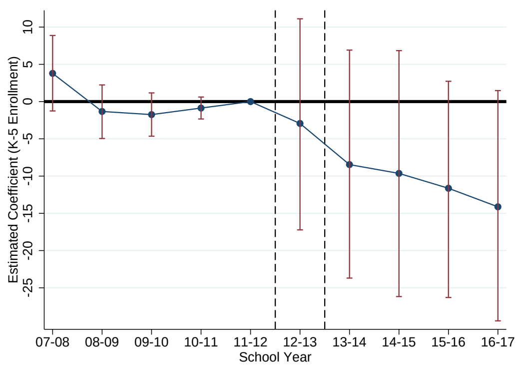

(c) Charter has Traditional Curriculum

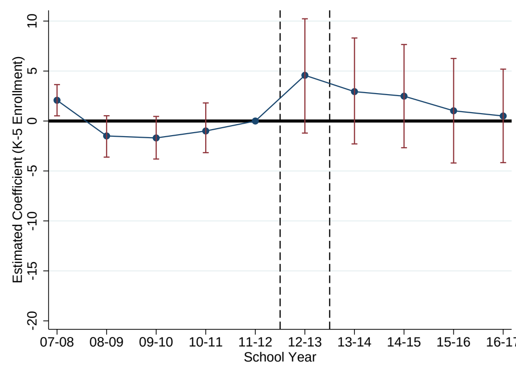

(d) Charter has Non-Traditional Curriculum

Figure 3: Event Studies: Newly-Opened Charter on Nearby School Value-Added and Enrollment

Notes: This figure shows the estimated value-added (Panel A) and enrollment (Panel B) difference between schools ‘treated’ by a newly-opened charter relative to ‘control’ schools by year. Data are restricted to schools in the Charlotte, the ‘Research Triangle’ (i.e., Raleigh-Durham-Chapel Hill), and Greensboro-High Point commuting zones. Treated schools are defined as schools located within 20 miles of a newly-opened charter that opened in 2012-13 or 2013-14. Control schools are defined as schools located between 20 and 30 miles of a charter schools that opened in 2012-13 or 2013-14. Results are subdivided by whether the newly-opened charter follows a traditional curriculum or not. Note that 2012-13 is considered the first ‘treated’ year because although the charters themselves opened in either the 2012-13 or 2013-14 school year, public schools would have known by the start of 2012-13 whether or not a charter was opening nearby or would open nearby in 2013-14. The dashed vertical line therefore separates the ‘pre-years’ from the ‘post-years.’ For the enrollment figures a second vertical line is added as the first ‘post’ year only had half the charters able to enroll students (as they had announced their opening, but not yet opened). The horizontal line represents a point estimate of zero. The dashed ‘whiskers’ represent 90 percent confidence intervals with inference conducted via wild clustered bootstrap (Cameron et al., 2011) two-way clustering by newly-opened charter and public school.

3 Empirical Model

Having described the data and key variation that will inform the estimates and evaluation to come, we now turn to our empirical model of student demand for schooling options and school supply. We build the model to leverage the detailed data described above as well as the policy-driven variation in charter school supply and public school quality.

3.1 Demand

On the demand side of the model, students choose from among the schools – public and charter – in their choice set to maximize utility. The indirect utility to student i from attending school j is given by:

$$\begin{aligned} u_{ij} =& \beta_i^V q_j + \beta_i^C Charter_j + \beta_i^H NonTrad_j \\ &+ \gamma_i \log(d_{ij}+1) + \gamma_i^C Charter_j \times \log(d_{ij}+1) + \kappa_i Assigned_{ij} \\ &+ \xi_j + \epsilon_{ij} \end{aligned} \tag{2}$$

where $Charter_j$ and $NonTrad_j$ are indicators for whether j is a charter school and, if a charter, whether the curriculum offered is non-traditional. q_j represents school j's quality (as measured by estimated value-added to student learning); β_i^V is thus i's "marginal utility" of value-added.[32] d_{ij} is the student's residence's distance (in miles) from school j's location, while $Assigned_{ij} = 1$ if j is their assigned local public school.[33] We allow the travel cost to differ by whether or not the school is a charter school. Finally, ξ_j is a structural error that represents (an index of) unobserved school qualities or amenities that is valued in common by students.

It is useful to re-write equation (2) as:

$$u_{ij} = \delta_j + \mu_{ij}(\theta)$$

where δ is the vector of "mean" utilities (which absorb the ξ's), while the ϵ's are idiosyncratic T1EV choice shocks. In contrast, $\mu_{ij}(\theta)$ captures systematic heterogeneity in

[32]While we use the term "preference parameters" throughout the paper for convenience, it should be understood that these valuations of school characteristics likely represent a combination of both tastes and information about schools.

[33]This information is obtained from the NCES School Attendance Boundary Survey from the 2010-11 and 2015-16 school years. Rather than assuming students cannot possibly attend public schools outside their attendance zone, the inclusion of this variable in the indirect utility estimates the "cost," which is held constant in the policy counterfactuals to come, of doing so.

preferences, governed by the parameters in θ:

$$\begin{aligned}\mu_{ij}(\theta) =& \tilde{\beta}_i^V q_j + \tilde{\beta}_i^C Charter_j + \tilde{\beta}_i^H NonTrad_j \\ &+ \gamma_i \log(d_{ij}+1) + \gamma_i^C Charter_j \times \log(d_{ij}+1) + \kappa_i Assigned_{ij}\end{aligned}$$

We allow for observed and unobserved heterogeneity in the demand model:

$$\begin{pmatrix} \tilde{\beta}_i^V \\ \tilde{\beta}_i^C \\ \tilde{\beta}_i^H \end{pmatrix} = \begin{pmatrix} \tilde{\beta}^V \\ \tilde{\beta}^C \\ \tilde{\beta}^H \end{pmatrix} W_i + \begin{pmatrix} \upsilon_i^V \\ \upsilon_i^C \\ \upsilon_i^H \end{pmatrix}$$

where W_i is a vector of household characteristics and $\upsilon_i \sim N(0, \Sigma)$ is a vector of random coefficients on tastes for value-added, charter schools, and non-traditional instruction. Preferences over distance and assigned public school only depend on observed characteristics. We include indicators for economic disadvantage and underrepresented minority in W_i.

3.2 Supply: Public School Value-Added

On the supply side, we model the decisions of public schools over how to set educational quality, making explicit how these decisions depend on the prevailing student demand in their local areas as well as the supply of traditional and non-traditional charter schools.

In the empirical model, public schools choose value-added (taking other schools' choices as given) in order to maximize a "rent-seeking" utility function (McMillan, 2004). For public school j at time t, this is given by:

$$U_j = F_j(D_j(\mathbf{q})) - C_j(\mathbf{q})$$

where $C_j(\mathbf{q}) = mc_j(q_j)D_j(\mathbf{q}) + fc_j$. $mc_j(q_j)$ is their marginal cost per pupil and depends on their value-added choice, while fc_j represents fixed costs. $F_j()$ is a function representing how public schools' value total student enrollment, $D_j(\mathbf{q})$. Enrollment is derived from the demand model and depends on all schools' quality choices, $\mathbf{q}$.

In the case of pure profit maximization, note that $F_j'() = p_j > 0$ where p_j is public school j's "price" – an object set by state funding formulas. While public schools are not profit-maximizing entities, in North Carolina almost the entirety of per pupil revenues from state and local sources follows students when they switch from public schools to charter schools, meaning that $F'_j() \neq 0$. Because public schools stand to lose funding when enrollment falls, on the margin, their incentive is to retain enrollment in the face of charter school competition.[34] Under these funding models, approximating public school objectives with the rent-seeking representation above has a long-standing history in the literature (Hoxby, 2002), and the implied incentives serve (at least implicitly) as the primary impetus for the vast literature exploring the competitive effects on public school quality stemming from both charter school penetration and private school voucher programs (e.g, Sass, 2006; Bifulco and Ladd, 2006; Booker et al., 2008; Imberman, 2011; Winters, 2012; Figlio and Hart, 2014; Cordes, 2018; Figlio et al., 2020).

The first-order condition of the maximization problem implies:

$$(\tau_j - mc_j(q_j^*))\sigma_j(\mathbf{q}^*) = mc'_j(q_j^*)$$

where $\tau_j = F'_j()$. In this expression, $\sigma_j(\mathbf{q}) = \frac{1}{D_j(\mathbf{q})}\frac{\partial D_j(\mathbf{q})}{\partial q_j}$ is public school j's own-value-added semi-elasticity of demand. Note that this object depends on the demand parameters and has no closed-form representation. The system of these equations for all schools in each market defines a Nash equilibrium in qualities.

We assume that $mc_j(q_j) = \pi_j + \kappa \exp q_j$. As in Neilson (2017), we can simplify the first-order condition to yield an important and intuitive expression for school j's equilibrium choice of value-added. This is given by:

$$q_j^* = \underbrace{\log\frac{\tau_j - \pi_j}{\kappa}}_{\text{Perfect comp. VA}} \underbrace{- \log[1 + \frac{1}{\sigma_j(\mathbf{q}^*)}]}_{\text{VA “markdown”}} \tag{3}$$

The expression consists of two parts: the level of quality that would be supplied under perfect competition (which depends on the public school utility function and cost parameters) and a value-added "markdown." This latter term embeds public schools' incentive to supply higher

[34] Note that this is not true in other states, such as Massachusetts and New York, where state aid is targeted at public school districts facing enrollment losses from charter schools. In North Carolina, the main exception to public-sourced revenues following students is capital appropriations, which public school districts do not have to share with charter schools on an equal per pupil basis.

value-added when competitive pressure is higher, as captured by the own-value-added semi-elasticity. Conversely, the "markdown" to value-added will be large for public schools that effectively operate as local monopolists. In this setup, the markdown expression differs from the profit-maximization case in that a public school's equilibrium quality may be shifted up or down by non-profit incentives and constraints.[35]

4 Estimation

The empirical model is estimated in several steps. As described in Section 2.2.1, we start by estimating school value-added offline.[36] We next estimate the heterogeneous demand parameters and recover mean utilities. This step is described below. We then leverage the spatial difference-in-differences based around charter exposure following North Carolina's cap removal to estimate $\bar{\beta}^V$ and public schools' quality policy function.

4.1 Estimating Demand

The demand model generates expressions for choice probabilities that can be mapped to the student-level choices via maximum likelihood. The probability that student i chooses school j in their choice set is given by:

$$p_{ijt} = \int \frac{\exp \delta_j + \mu_{ij}(\theta)}{\sum_{k \in C_{it}} \exp \delta_k + \mu_{ik}(\theta)} f(\tilde{\upsilon}_i) d\tilde{\upsilon}_i \tag{4}$$

We restrict choice sets (C_{it}) to public schools within 7 miles and charter schools within 30 miles.[37] θ represents the vector of heterogeneous demand parameters optimized over. The estimation procedure recovers the vector of mean utilities δ using the BLP contraction

[35] While equation (3) allows for incentives to supply quality if they scale with size, the setup does not allow for direct preferences over quality. The expression also rules out quality adjustment costs. We test these restrictions against the data later.

[36] To reduce noise, we use all school years 2008-09 through 2011-12 to estimate pre-cap school VA and all school years 2012-13 to 2015-16 to estimate post-cap school VA.

[37] For students with fewer than 10 nearby public schools, however, we expand their choice set to include the closest 10 irrespective of distance. Note that because there are many public schools, choice sets grow very quickly with the distance radius, raising computational burden. The 7 mile radius is not very restrictive, however: public school students' median travel distance is a little over 2 miles. We drop students whose school of attendance falls outside their choice set (about five percent of our sample). We suspect at least half of these cases are driven by coding errors as the distance between a student's residence and school of attendance is improbable (e.g., student attending a school over 100 miles from their home)

mapping to match predicted and observed shares and uses simulation to form the choice probabilities. We use quadrature to integrate the random coefficients and, for those students whose residence location is not known, we integrate out over demographic-specific densities estimated from the residential data.[38] We specify the random coefficient structure with a standard normal and estimate the preferences over value-added, charter schools, and non-traditional curricula associated with the unobserved type.[39]

We estimate the demand model on Kindergarten through 2nd grade students in six markets: 2011-12 (pre-removal of the charter cap) and 2015-16 (post-removal) for each of the three major Commuting Zones in North Carolina. The estimation sample includes 337,855 student-year observations.

4.2 Identifying $\bar{\beta}^V$ Using Spatial DiDs

For public school j, we have two structural equations from the empirical model. For their mean utility (recovered in the above step), we have:

$$\delta_{jt} = \bar{\beta}^V q_{jt} + \xi_{jt} \tag{5}$$

Mean utility depends on the education quality, as measured by value-added, of the public school, q_{jt}, and the quality unobserved to the econometrician, ξ_{jt}. Although we do not explicitly model peer quality here, it is important to note that our model does not rule out parental preferences for peers. Indeed, ξ_{jt} may contain school-year-specific measures of peer quality, which are simply absorbed by the school demand residuals separately from preferences for test score value-added. We then have the quality policy function, equation (3), which depends on the demand parameters (including $\bar{\beta}^V$) and cost and objective function parameters.

It is commonplace to first decompose equation (5) using instruments for q_{jt} and then to estimate the policy function in a second step (conditional on the estimate of $\bar{\beta}^V$). Crucially, relevant and valid instruments for q_{jt} are needed to carry this out. In practice, many appli-

[38]We use 1,000 residence location draws. We estimate the residential densities from the 2011-12 (i.e. pre-charter cap removal) data.

[39]Note that this is an equivalent normalization to jointly estimating the variance on unobserved preference for one characteristic, e.g. value-added, and its correlation with preference for the other two characteristics.

cations typically rely on a combination of 1) market-level price indices; 2) product location space instruments (Berry et al. 1995); and 3) natural experiments, such as arising from policy changes.

In this paper, we instead identify and estimate $\bar{\beta}^V$ and the policy function in a single step that is based on the spatial difference-in-differences variation summarized earlier. Intuitively, this estimation approach asks: what value of $\bar{\beta}^V$ rationalizes the reduced-form effects of charter school exposure on public school value-added (given the utility function estimates)? The estimating equation is thus derived from the policy function, which links the competitive environment to a public school's choice of quality:

$$q_{jt} = -\log[1 + \frac{1}{\sigma_{jt}(\bar{\beta}^V)}] + \tau X_{jt} + \pi_j + \psi_{d(i)t} + \omega_{jt} \tag{6}$$

This equation re-writes the policy function, equation (3), such that the parameters other than $\bar{\beta}^V$ represent reduced-form objects that are directly estimated and held constant in the policy analyses (plus an error term ω_{jt}). These objects include linear functions of observed pre-determined cost shifters X_{jt}, the school fixed effects $\boldsymbol{\pi}$, and district-specific trends $\boldsymbol{\psi}$. From the demand model, the semi-elasticity as a function of $\bar{\beta}^V$ is given by:

$$\sigma_{jt}(\bar{\beta}^V) = \frac{1}{D_{jt}} \sum_i \int (\bar{\beta}^V + \tilde{\beta}_i^V) p_{ijt}(1 - p_{ijt}) f(\tilde{v}_i) d\tilde{v}_i \tag{7}$$

where $\tilde{\beta}_i^V$, p_{ijt} (given by equation (4)), and $f(\tilde{v}_i)$ are estimated alongside the other heterogeneous demand parameters in the previous estimation step.

For identification, note that $\sigma_{jt}(\bar{\beta}^V)$ is endogenous, but is shifted by exposure to charter school entry post-2012. We make the assumption, analogous with the spatial difference-in-differences estimates presented earlier, that charter schools do not choose location based on within-district *innovations* to ω. Importantly, this identification allows for the possibility that charters sort on innovations to ξ and accordingly avoids exclusion restrictions that require ruling out impacts on school choice through channels other than quality adjustments. The exposure variables then isolate exogenous variation in $\sigma_{jt}(\bar{\beta}^V)$, which can be used to estimate $\bar{\beta}^V$. This set of assumptions implies a nonlinear GMM estimator, which we detail next.

4.2.1 GMM

We estimate equation (6) using data from 2012 and 2016. To deal with incidental parameters, we first difference the equation and re-arrange:

$$\Delta q_j = -\log \frac{1 + \sigma_{j2016}(\bar{\beta}^V)}{\sigma_{j2016}(\bar{\beta}^V)} \frac{\sigma_{j2012}(\bar{\beta}^V)}{1 + \sigma_{j2012}(\bar{\beta}^V)} + \tau \Delta X_j + \psi_{d(i)} + \Delta\omega_j \tag{8}$$

The differencing cancels out the school fixed effects in equation (6) and the district-specific trends become district fixed effects. The parameters to be estimated are therefore $\bar{\beta}^V$ (which enters non-linearly), τ, and $\psi_{d(i)}$.

The moment condition is $E[\Delta\omega|Z] = 0$, where Z is a vector of instruments. These instruments include the right-hand side controls in equation (8)—X_{j2016}, X_{j2012}, and district fixed effects—as well as excluded instruments. Variables in X, discussed below, include a district cost index and treatment by an accountability program. The excluded instruments are the charter entry exposure variables. To use variation from multiple exposures, we create rows for each public-charter entrant (within 30 miles) pair. The baseline four excluded IVs are then: whether the public school is treated (i.e. the charter entrant is within 20 miles) and the entrant is $NonTrad$, whether treated and the entrant is $Trad$, and the treatment distance interacted with each treatment indicator.[40] In our preferred specification, we generate additional IVs by interacting these with the 2012 cost index and with the components of the school's overall quality semi-elasticity (equation (7)) for 2012 recovered from the first estimation step. These interactions leverage variation in constraints and incentives across public schools that influence their response to charter competition.[41]

4.2.2 The Costs of Providing School Quality and Accountability Pressure

We now briefly discuss the cost index and accountability program pressure we include in the X vector above.

[40] Our main results equally weight the rows, but we examine the sensitivity of the estimates to weighting the rows by the 1 over the total number of charter exposures.

[41] Specifically, we interact with $\frac{1}{D_{j2012}} \sum_i \int p_{ij2012}(1 - p_{ij2012}) f(\tilde{v}_i) d\tilde{v}_i$ and $\frac{1}{D_{j2012}} \sum_i \int \tilde{\beta}_i^V p_{ij2012}(1 - p_{ij2012}) f(\tilde{v}_i) d\tilde{v}_i$. The first variable is proportional to schools' semi-elasticty of demand among "mean" (i.e. non-URM, non-economically disadvantaged) students. The second is the difference in the semi-elasticity between "mean" and all other student types. Note we do not use these variables on their own as instruments, so as not to impose they are uncorrelated with cost shocks.

The Comparable Wage Index for Teachers A key determinant of the educational quality schools deliver is the cost of the labor involved in providing that quality, especially in the form of teacher salaries. We measure variation in these labor costs with the comparable wage index for teachers (CWIFT), an index made available by the National Center for Education Statistics and designed to identify geographic variation in (regression-adjusted) wages for college-educated workers outside of teaching, thereby serving as a proxy for the area-specific costs of hiring teachers.[42] We use data on the CWIFT at the school district level in 2012 and 2016.[43]

Turning Around North Carolina's Lowest-Achieving Schools Our second determinant of the change in school quality is a school's membership in North Carolina's Turning Around Lowest-Achieving Schools (TALAS) initiative. As part of its Race to the Top grant, North Carolina created the TALAS program to target schools and school districts for improvement plans based on inadequate proficiency or graduation rates. TALAS implementation started in the 2010-11 academic year and was fully implemented by 2011-12 (Henry and Guthrie, 2019). The program was multi-faceted, as treated schools experienced principal replacement, instructional reform, increased learning time, and financial incentives for teachers and principals when students realized adequate test score growth (Heissel and Ladd, 2018). We provide more detail about the TALAS program in Appendix D and, using a difference-in-differences framework, illustrate the TALAS-driven variation in school quality that informs estimates of equation (8).

5 Estimation Results

This section presents our estimation results. We first present parameter estimates and report elasticities of substitution, which speak to how heterogeneous preferences and curriculum heterogeneity differentiates charters.

[42] For a full description of the CWIFT, including a discussion of measurement and interpretation, see Cornman et al. (2019).

[43] Because we only have CWIFT data at the school district level and our main specification absorbs district trends, the cost index only enters our analysis through *interactions* with other variables included in equation (8). We thus allow the effects of changes in other determinants of school quality to depend on the costs of hiring teachers in the school district.

5.1 Estimates

Table 3: Demand Estimates

	Coef	SE
Log distance+1	-2.82	0.01
Log distance+1 × Econ disadv	0.11	0.01
Log distance+1 × URM	-0.07	0.01
Log distance+1 × Charter	1.70	0.03
Log distance+1 × Charter × Econ disadv	-0.80	0.04
Log distance+1 × Charter × URM	-0.45	0.05
Assigned	1.44	0.01
Assigned × CMS	0.78	0.01
Assigned × Greensboro	-0.10	0.01
Charter × Econ disadv	0.61	0.09
Charter × URM	0.69	0.11
Charter × unobs type	-1.84	0.20
VA × Econ disadv	-0.47	0.05
VA × URM	-0.95	0.05
VA × unobs type	1.95	0.23
NonTrad × Econ disadv	-0.77	0.05
NonTrad × URM	-0.03	0.04
NonTrad × unobs type	-1.22	0.25
Student-years	337,855	

Notes: Table reports estimated coefficients and corresponding standard errors for idiosyncratic component of utility underlying school demand (i.e. "heterogeneous parameters"). Not reported are estimates of interactions between an indicator for missing VA information and economic disadvatange/URM. The unobserved student type is drawn from a standard normal distribution. Standard errors calculated by taking square root of the diagonal of the inverse of the Hessian.

Table 3 presents demand estimates corresponding to the heterogeneous parameters in $\mu_{ij}(\theta)$ in equation (4). While not directly interpretable, the estimates are indicative of several important qualitative features of school demand. First, travel costs and school assignment are highly-salient. Students, especially non-disadvantaged majority students, are more willing to travel to charter schools, all else equal. Second, disadvantaged students have stronger preferences (relative to non-disadvantaged, white and Asian students) for charter schools overall (holding distance fixed) and weaker preferences for school value-added and for non-traditional charters schools. Students from underrepresented minority backgrounds also have stronger preferences for charter schooling (holding distance fixed) and weaker preferences for value-added but do not exhibit any differential preference for non-traditional charters. The estimates on students' unobserved type, which corresponds to a draw from a standard normal distribution, reveal how preferences along unobserved lines are correlated across school characteristics. The estimates indicate that observed and unobserved preferences for value-added and for charter schooling are strongly negatively correlated, with students who value value-added highly especially disliking non-traditional charter schools. In the next subsection, we examine the implications of these estimates for elasticities of substitution.

Table 4: Policy Function Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
$\bar{\beta}^V$	2.593***	2.453***	3.128***	3.111***	3.098***	2.919***
	(0.707)	(0.444)	(0.586)	(0.656)	(0.595)	(0.555)
λ					21.297	
					(54.705)	
η						-0.155
						(0.122)
District FE	Y	Y	Y	Y	Y	Y
Excluded IVs	Entry	Entry× Cost	Entry × Cost × Demand$_0$			
wt = 1/exposures	N	N	N	Y	N	N

Notes: Table reports results from estimating public schools' value-added policy function (equation (8)) via GMM. Each column includes N=2,249 public-charter school pair observations. The specifications in every column control for an indicator for TALAS and its interaction with the change in the district cost index. Column (2) interacts the four entry instruments with district costs in 2010; columns (3) through (6) add interactions with 2012 demand elasticity components (see footnote 41 for more detail). Column (4) weights observations by the inverse of the total number of charter exposures within 30 miles. Column (5) tests for quality adjustment costs, while column (6) tests the model by allowing for heterogeneity in response based on whether the entrant is a non-traditional charter school. Estimation uses a robust weight matrix. ***,** and * denote significance at the 1%, 5% and 10% levels, respectively.

Table 4 presents estimates of the value-added policy function, including $\bar{\beta}^V$, the "marginal utility" of value-added (for non-disadvantaged, non-URM students). Column (1) presents the estimate for $\bar{\beta}^V$ using only the entry exposure variables as excluded IVs. Column (2) interacts the entry variables with the district cost index in 2012 for additional IVs, while columns (3) through (6) include the full set of interactions between the entry variables, 2012 cost index, and 2012 elasticity components as excluded IVs.

Our preferred estimate of $\bar{\beta}^V$, which we use for the results that follow, is in column (3) of Table 4; column (4) examines robustness to how multiple charter exposures are weighted. To understand the economic meaning of the $\bar{\beta}^V$ estimate, Figure 4 plots the implied willingness-to-pay for an increase in school value-added in terms of travel distance to school. The figure shows that the average student would be willing to travel about 1.7 miles on average to experience a school quality improvement of 1 student test score distribution standard deviation.[44] The willingness-to-pay is heterogeneous with respect to unobserved sensitivity

[44] It is worth comparing our implied willingness-to-pay for school value-added to estimates from the prior

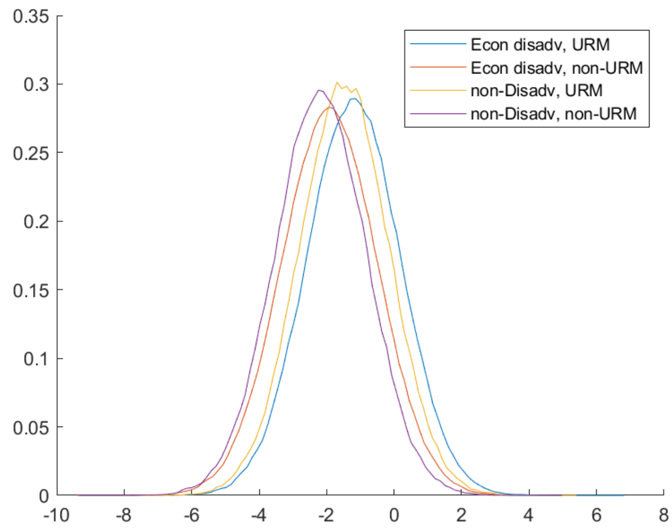

Figure 4: MWTP for Value-added
in terms of travel distance

Notes: This figure plots estimated MWTP distributions for a 1 student test score distribution standard deviation in school quality in terms of miles of travel distance (to public school) by student characteristics.

to quality and, as indicated by the utility function estimates in Table 3, is heterogeneous according to students' observed characteristics. This is shown in the figure by the rightward shifts for underrepresented minority and economically disadvantaged students

Columns (5) and (6) test restrictions placed on the data by our model of quality supply. In column (5), we allow for the presence of a quality adjustment cost to public schools costs of the form $\lambda \exp(q_{jt})$; the results cannot reject the hypothesis that the adjustment cost is zero. Column (6) tests the model restrictions by instead checking if the implied response to non-traditional charter competition differs. This is accomplished by interacting the semi-elasticity of demand variable with an indicator for whether the entrant is a non-traditional charter; η is the coefficient on that interaction. The result in column (6) shows that the estimate of $\bar{\beta}^V$ based on responses to traditional charter competition also successfully rationalizes the (muted) responses to non-traditional charters.

literature. Abdulkadiroğlu et al. (2020) finds that preferences for New York City high schools do not appear correlated with a school's value-added (once the school's peer composition is accounted for). In contrast, Campos (2024) shows that school quality indeed predicts school choices conditional on peer quality among LA high school students, finding that families would travel 0.44 more miles to attend a school 10 percentile points better in school quality. Our estimate of willingness-to-pay for school value-added is smaller and implies a willingness to travel only 0.1 miles, on average, for the same school quality improvement—consistent with travel costs being much more salient at the elementary school level.

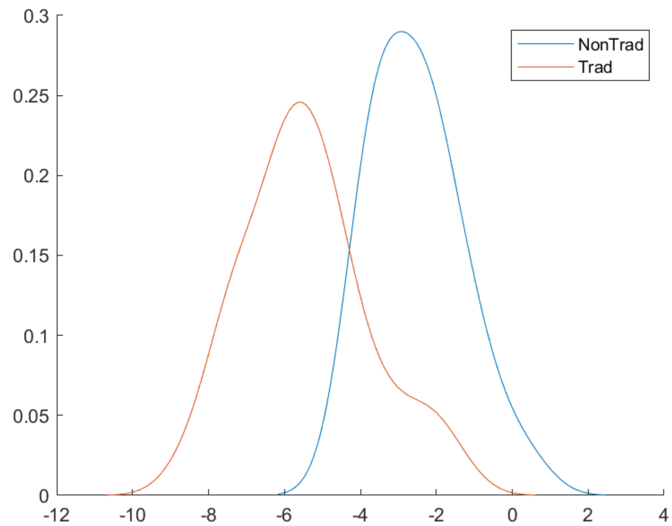

Figure 5: Elasticities of Substitution
Counterfactual: 0.05 SD increase in VA at ALL public schools

Notes: This figure plots distributions of the estimated percent of enrollment lost due to a 0.05 student test score distribution standard deviation increase in VA at all public schools by non-traditional ('NonTrad') charter schools and traditional ('Trad') charter schools.

5.2 Elasticities, Differentiation, and Competition

In this subsection, we explore what the model estimates mean for how public and charter schools compete.

Figure 5 shows the elasticity of demand for charter schools with respect to the quality of public schools. We counterfactually increase the quality of all public schools by 0.05σ (on the student distribution) and then compute the percentage change in enrollment. The figure plots densities of enrollment losses by type of charter school. The average charter schools loses fewer than 5% of its overall enrollment. This is indicative of the role of horizontal differentiation arising from heterogeneous preferences and travel costs. Moreover, the figure shows that demand for non-traditional charters is relatively more inelastic than demand for traditional charters; the average enrollment loss for non-traditional charter schools is around half the average enrollment loss for traditional charters. This finding indicates that, as hypothesized, curricular differentiation softens quality competition with public schools.

Figure 6 panel (a) plots public schools' own-value-added semi-elasticity of demand. These elasticities measure how elastic demand for a public school is with respect to its choice of value-added. These estimates can then be used to recover public schools' markdowns, which are reflective of their degree of market power. Figure 6 panel (b) displays their "perfect competition" quality levels alongside the distribution of value-added in the data. The figure indicates that, on average, quality is marked down by around 0.6σ (on the student distribution) from the level that would be supplied under perfect competition.

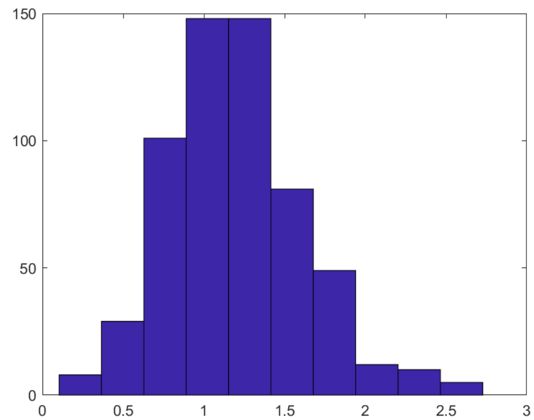

(a) Own-VA Semi-elasticity of Demand

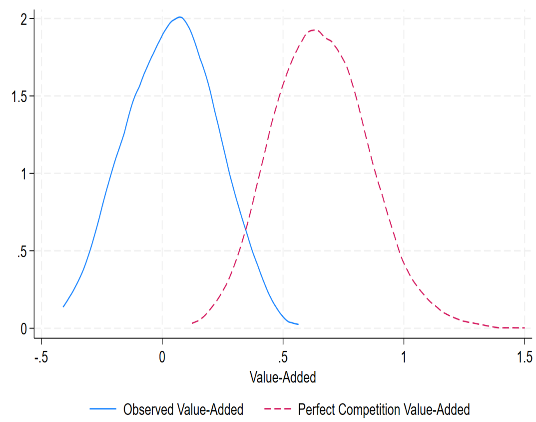

(b) Value-added Markdowns

Figure 6: Competition and Supply of School Quality

Notes: Figure 6a shows a histogram of public schools' estimated own-value-added semi-elasticity of demand (in 2016). Figure 6b plots the distribution of public schools' observed value-added in the data and the distribution of their estimated "perfect competition" level of value-added (in 2016).

Figure 7 visualizes the spatial distribution of these markdown changes from the pre- to the post-cap period in the 'Research Triangle.' Specifically, the figure displays a heat map of the change in public schools' value-added markdowns from 2011-12 to 2015-16 with the location of post-cap charter entries overlaid, differentiating between traditional and non-traditional charters. A clear visual pattern is apparent: There are sharp reductions in the value-added markdown (indicated by the lighter colors) wherever a traditional charter school opened. In contrast, a consistent pattern of markdown changes is not observed in areas where non-traditional charters entered. The fact that areas with a traditional charter entry (but not a non-traditional charter entry) experienced reduced value-added markdowns shows that our model is capturing the competitive responses we found in our reduced-form analysis.

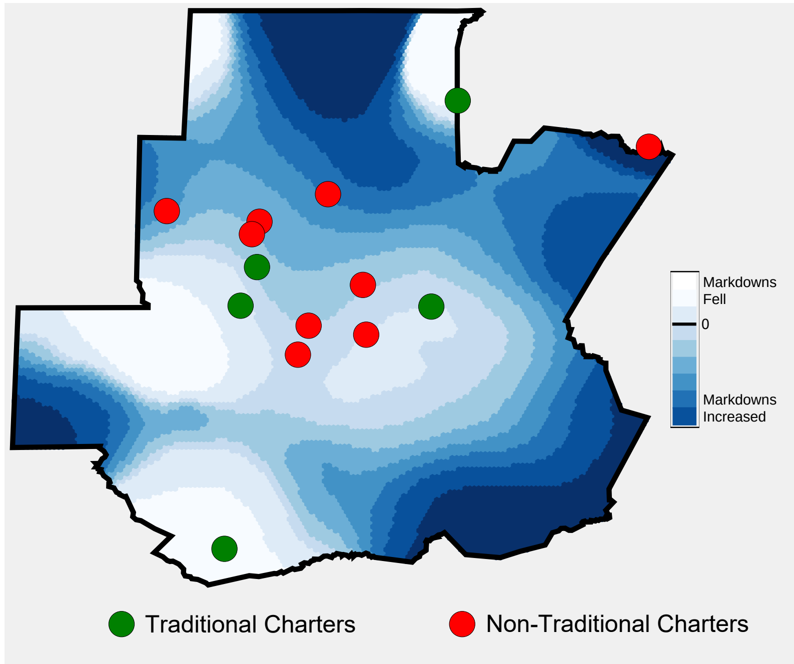

Figure 7: Changes in Value-Added Markdowns in the Research Triangle from 2011-12 to 2015-16

Notes: This figure displays a heat map of the change in value-added markdowns from 2011-12 to 2015-16 for the 'Research Triangle' (Raleigh-Durham-Chapel Hill) Commuting Zone; darker colors indicate that the value-added markdown increased over this time period while lighter colors indicate the value-added markdown fell. Specifically, we show the change in the distance weighted average markdown from 2011-12 to 2015-16 evaluated at over 20,000 equidistant grid points that we have placed throughout the commuting zone. The location of charter schools that opened post-cap are overlaid, differentiated by curriculum.

6 Counterfactual Policy Analyses

In this section, we present results of several policy analyses of interest. First, we consider the aggregate effects of North Carolina's removal of the charter school cap in 2011. We examine the effects across students and explore the supporting mechanisms.

We then turn to counterfactual simulations that quantify how policy can influence the aggregate effect of charter school expansion by influencing the types, qualities, and locations of charter schools.

6.1 Charter Cap Not Lifted

We first use the model to evaluate the aggregate effects on student learning of North Carolina's lifting of its statewide cap on charters. We examine this by solving for counterfactual school enrollments and value-added in an equilibrium where the 29 post-2012 charter school

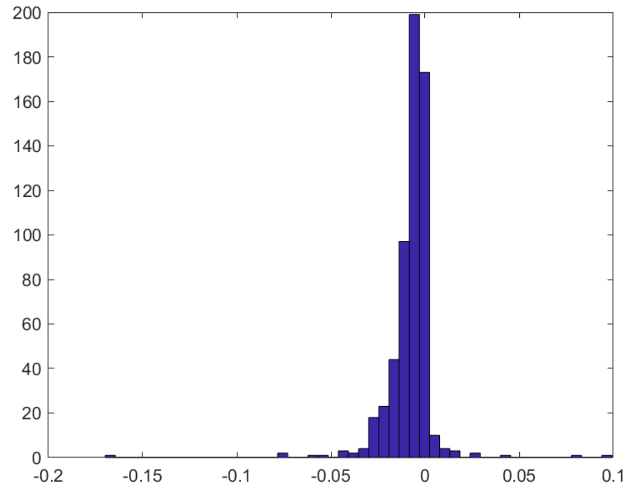

Figure 8: Difference in Value-added (school-level)
Counterfactual: No post-2012 charter entry

Notes: This figure plots a histogram of public schools' difference in VA between the data and the counterfactual in which the post-2012 entrants are removed.

entrants are removed from the three markets.[45]

Figure 8 summarizes the changes in school-level value-added between the data and the no entry equilibrium. The figure shows that the policy has a limited effect on the average public school, but a left tail reduces its quality meaningfully absent charter entry.[46] The average value-added reduction is 0.009σ (on the student distribution), as shown in Table 5. Table 5 also shows that charter school enrollment would have been about 2.3 points lower and the average student's test scores about 0.005σ lower without the cap lifting. The test score effect is smaller in magnitude than the average school's quality reduction because students re-optimize their enrollment choices. Table 5 also summarizes the test score impacts across student groups. Economically disadvantaged students who also belong to a an underrepresented minority benefit relatively more from the cap lifting, which causes a 0.007σ increase in their test scores on average.

Is the aggregate effect on student learning from raising the charter school cap—a 0.005σ increase in test scores on average—economically meaningful? We benchmark this effect against policies that raise teacher quality and increase infrastructure spending. First, the

[45] We use a contraction mapping to find new equilibria, using the qualities in the data as starting values. We also do not allow the new equilibrium values to exceed 1 or be less than -1.

[46] Several schools *increase* their quality somewhat because their residual demand becomes relatively more inelastic, as in McMillan (2004).

Table 5: Counterfactual: No post-2012 charter entry

Δ charter share	-0.023
Δ VA (school-level)	-0.009
Δ test scores (student-level)	
On average	-0.005
Non-disadv. & non-URM	-0.004
Econ. disadv. & non-UR	-0.005
Non-disadv. & URM	-0.004
Econ. disadv. & URM	-0.007
Due to competition	-0.008
Due to sorting	0.003

Notes: This table reports the changes to charter school share, school value-added, and student test scores that result from counterfactually removing post-2012 charter school entrants from our three markets.

test score effect we estimate of raising the charter cap (which increased the charter share by 2.3 points) is about one-fifth of the predicted impact of replacing the bottom 5% of teachers with average quality teachers (Gilraine et al., 2020). Chetty et al. (2014b) estimate the present value earnings impact of that policy to be approximately $250,000 per class, implying that the charter school cap lifting increases the average student's lifetime earnings by $1,931. An alternative benchmark is the impact of infrastructure spending. Biasi et al. (2024) find that a $1,000 per-pupil increase in capital spending over five years leads to a 0.05 standard deviation increase in district-level test scores. Their estimate converted to the student level implies an increase of 0.017 student-level standard deviations. Assuming students are exposed to the post-cap environment for five years to set the time horizon on an equal footing with that of the infrastructure spending policy, the cap lifting increases test scores by 1.5 times more over the same time period.

How does competition contribute to the gains from lifting the charter school gap? We assess this by decomposing the average test score effect into the change in school-level qualities $(q_j^1 - q_j^0)$ holding enrollments fixed at the data (p_{ij}^0) and the change in enrollments holding qualities fixed at the new equilibrium, i.e.:

$$\overbrace{\frac{1}{N}\sum_i\sum_j q_j^1 p_{ij}^1 - \frac{1}{N}\sum_i\sum_j q_j^0 p_{ij}^0}^{\text{Average test score impact}} = \underbrace{\frac{1}{N}\sum_i\sum_j p_{ij}^0 (q_j^1 - q_j^0)}_{\text{Competition}} + \overbrace{\frac{1}{N}\sum_i\sum_j q_j^1 (p_{ij}^1 - p_{ij}^0)}^{\text{Sorting}} \quad (9)$$

The first object on the right hand side captures competition by asking how test scores would have changed if students did not change where they attend school. The second object, which we term sorting, then measures how changes in where students go contribute to the aggregate test score effect. These calculations are provided in the bottom two rows of Table 5. The results reveal that it is public schools' competitive responses to charter school entry that is driving the aggregate gains from lifting the charter cap. In fact, the sorting channel actually reduces aggregate test scores, as students switch into charter schools when the cap is lifted that are lower quality than the counterfactual public school they would have attended.[47]

6.1.1 Effect of Removing All Charter Schools

While the previous counterfactual evaluates the effect of the 2011 removal of the statewide charter school cap by removing the entrants, the model structure also allows us to compute outcomes in the counterfactual scenario where all charter schools are removed from the markets. This policy analysis examines the effect of reducing the amount of charter school choice to zero.

Table 6 reports the results for school-level value-added and student test scores from removing all charter schools and reducing charter school enrollment to zero. The results indicate that the average public school's value added would be 0.025σ lower than in the data. For the average student, the reduction translates into a 0.02σ reduction in test scores. The average test score impact is very similar across student groups. Note that the effect size on test scores is several times larger than the effect of turning off just post-2012 charter school entry. One way of emphasizing this is that the estimated effect from shutting down all charter schools is indistinguishable from the effect of firing the bottom 5% of teachers

[47] This negative effect has an interesting parallel with recent findings from U.S. voucher programs (e.g. Abdulkadiroğlu et al. 2018) and implicitly values the utility gains from choice expansion, which includes valuations placed on other non-test outcomes, for these families.

Table 6: Counterfactual: No charter schools

	No post-2012 entry	No charters
Δ charter share	-0.023	-0.077
Δ VA (school-level)	-0.009	-0.025
Δ test scores (student-level)		
On average	-0.005	-0.021
Non-disadv. & non-URM	-0.004	-0.020
Econ. disadv. & non-UR	-0.005	-0.018
Non-disadv. & URM	-0.004	-0.022
Econ. disadv. & URM	-0.007	-0.024

Notes: This table compares differences in equilibrium outcomes from the data between the no entry counterfactual and a counterfactual in which all charter schools are closed.

(worth $9,000 in lifetime income per student per Chetty et al. 2014b).

6.2 Estimating the Expected Returns to the Marginal Charter School

In this section, we use the empirical model to estimate the expected return to adding one additional charter school in the market. We frame this analysis by appealing to the problem facing an authorizer who, with incomplete information, must decide whether or not to approve the marginal charter school application.[48] We thus compute these expectations over draws of the prospective entrant's ex-post quality, which is difficult for authorizers to reliably assess ex-ante, while conditioning on observable aspects of the school—its type and location. We do this by building a matrix of ex-post impacts: we compute the change (compared to the data) in the average student's test scores when we simulate adding one charter school, varying where it locates, its type (traditional or non-traditional), and its quality.

We simulate the aggregate across six different locations, defined by Census tracts.[49]

[48] In our setting and sample period, prospective charter schools submitted applications to the Charter Schools Advisory Board (now called the Charter School Review Board), typically at least 18 months before they planned to begin operations. The applications contained information about where the charter school intended to locate, intended grade levels, projected enrollment, leadership and governance, mission, instructional program, and statements of goals and educational focus. In the counterfactual analyses that follow, we treat prospective charter schools' curriucular focus and proposed location as the key pieces of information the authorizing board uses to guide its decision-making.

[49] To ease the computational burden of building the matrix, we include only the Triangle and CMS CZ. Within these markets, there are 903 unique tracts a charter school could be added to.

First, we examine outcomes when the marginal charter enters an "average" location, which we identify by locating the tract whose characteristics are most similar to the mean tract. The characteristics of the mean tract are reported in Appendix Table A.3 and show that its median household income and local public school quality (within 3 miles of the tract centroid) are about \$80,000 and 0.04σ, respectively. We also consider the outcomes when the charter enters the typical tract that charter schools enter—these locations tend to have somewhat lower local public school quality and a greater share of households who are black. Appendix Table A.3 shows, however, that there is wide variation across tracts that charters locate in.[50] This dispersion motivates us to juxtapose outcomes in high income and low income locations as well as high and low local public school quality locations. To do this, we find the tract in each market at the 84th (high) and 16th (low) percentile of that characteristic, respectively. Appendix Table A.3 reports the characteristics of these tracts: whereas low local public school quality locations also have low median income (\$59,000), local public school quality is actually somewhat higher than the mean in low median income tracts.

Within each location and charter school type, we simulate the outcomes when the entering charter school draws a low ($z = -1$), average ($z = 0$), or high ($z = 1$) quality "shock." Importantly, the realized value of the additional charter's quality will depend on this draw and its type. As the description of the value-added estimates shows, non-traditional charter schools have lower quality on average (by about 0.08σ) and our calculations will embed this difference in order to reflect the idea that the authorizer has equally limited information.[51] Across all cells of the matrix of outcomes, we assume that the charter school draws a median unobserved demand (ξ) for its type and restrict the school to enrolling 200 students at most.[52] When more students want to attend the new charter school than that, we allocate spots according to MWTP.[53]

The matrix of results is reported in Table 7 and contains several major findings: First, as

[50]For example, the standard deviation of median income across tracts with a charter school is \$24,000 and the standard deviation of local public VA is 0.18σ.

[51]Whereas holding constant quality itself would compute outcomes when a better draw of a non-traditional type were obtained.

[52]One way of thinking about holding the unobserved demand shock constant is that the authorizer has reliable information ex-ante that the charter school will be attractive to households. Note that this capacity is a bit bigger than the average post-2012 charter school entrant is in the data.

[53]Mechanically, this means that in the simulations we solve for a common "disutility" from the new charter school that ensures the constraint is satisfied; disutility is 0 when the school is not constrained.

Table 7: Aggregate Impact of Marginal Charter School by Location and Type

	(1)	(2)	(3)	(4)	(5)	(6)
Location Type:	Average	Avg Tract	Median Income		Local Public VA	
	Tract	With Charter	High	Low	High	Low
			Traditional			
VA Draw:						
$z=-1$	-0.06	-0.05	-0.07	-0.02	-0.08	-0.02
$z=0$	0.07	0.08	0.06	0.10	0.05	0.07
$z=1$	0.18	0.20	0.18	0.21	0.16	0.20
$\mathbb{E}$ impact	0.07	0.08	0.06	0.09	0.04	0.08
			Non-Traditional			
VA Draw:						
$z=-1$	-0.07	-0.07	-0.08	-0.06	-0.12	-0.03
$z=0$	0.03	0.04	0.02	0.07	0.02	0.02
$z=1$	0.13	0.14	0.13	0.18	0.13	0.07
$\mathbb{E}$ impact	0.03	0.04	0.03	0.07	0.01	0.02

Notes: This table reports the estimated impact on the average student's test scores of an additional charter entrant by location (columns), quality draw (rows), and by type as obtained by simulations. Simulated entrant's enrollment is capped at 200 students and entrant is assigned a median unobserved demand shock (given type) in all cells. Impacts are normalized by the cap lifting impact (0.0049σ). Expected impact approximated using Gauss-Hermite weights. Average Tract refers to Census tract most similar to average tract in the CZ; Avg Tract With Charter refers to the tract most similar to average among tracts containing at least one charter; "High" and "Low" refer to -1σ (16th percentile) and +1σ (84th percentile) tracts, respectively. See Appendix Table A.3 for descriptive characteristics of locations. Note that value of a VA draw depends on charter school type; e.g. a $z=-1$ non-traditional charter school has lower absolute quality than a $z=-1$ traditional charter school. Appendix Table A.4 reports decompositions of impacts into competitive and sorting effects.

the $z=-1$ rows show, low quality charter schools reduce outcomes for the average student irrespective of its type or where it is located. At the same time, high quality charters of either type have meaningfully positive impacts everywhere. For assisting interpretability, we scale the changes in the average student's test scores arising from the additional charter school by the charter school cap lifting's impact on the average student (0.0049σ). Thus, the table shows that just the addition of one $z=1$ traditional charter school to the kind of location charters tend to enter in (see column (2)) would have amplified the cap lifting's impact by 20%.

To approximate the expected returns to a marginal charter school, we apply Gauss-Hermite weights across the outcomes by quality draw.[54] This calculation shows that—across all locations considered and for either type—adding an additional charter school to the market will improve the average student's test scores. Put differently, and otherwise

[54] An important assumption that this embeds is that the quality shock distribution is independent of location.

ignoring costs, approving the marginal charter school applicant maximizes test scores. But the decompositions of the overall impacts in Table A.4 are revealing: the expected return to the marginal charter school applicant is instead negative almost everywhere *when its aggregate impact via competition is ignored.*[55]

Although always (net) positive, the expected returns reported in in Table 7 are heterogeneous and this heterogeneity is an additional source of policy insights. In particular, the results suggest that, while for some locations and some realized quality draws the difference in ex-post outcomes may not necessarily be very large, the expected return to the marginal traditional charter school is always greater than the expected return to the marginal non-traditional charter school. In fact, the expected return to a traditional charter can be as much as 4 times higher (in locations with low vs. high local public school VA – see columns (5) and (6)). This implies than an expected test score maximizing authorizer would need greater confidence in the quality of a non-traditional charter to be indifferent to type. This finding suggests that using curricula as an input to charter school authorization decisions on the margin will generate larger aggregate impacts.

The expected impacts across locations in Table 7 also speak to the value of leveraging information about the charter school's prospective location in authorizing frameworks. Specifically, we find that low income and low public school quality locations have significantly larger expected returns than their high counterparts. For example, the expected return to the marginal traditional charter school is twice as big—amounting to an 8% increase over the cap lifting's impact—in the low public school quality neighborhood. Given numerous charter schools locate in places similar to the high median income and high local public quality tracts (see Appendix Table A.3), this finding implies that policies that steer charters to disadvantaged and/or underserved neighborhoods will amplify the returns to charter school expansion.

[55]The lone exception is the low public school quality location, where its expected impact (ignoring competition) would be zero.

7 Conclusion

It has been over a quarter of a century since the first charter schools in the U.S. opened, but consensus regarding their aggregate effects on students learning remains elusive. A major reason for this is the bifurcated focus of prior empirical work on, on the one hand, treatment effects on students who attend charter schools or, on the other hand, competitive impacts on students who remain in public schools. Simply put, this paper asks: what do these prior findings add up to?

To answer this question, this paper combines school-level value-added estimates and quasi-experimental identification of public schools' competitive responses to charter entry in an equilibrium model of U.S. elementary education markets. In the model, students choose schools, public schools choose quality (value-added), and charter schools choose entry and location. We estimate the model using geocoded student-level data from North Carolina, whose lifting of the statewide charter school cap in 2011 both provides important variation that we leverage in the estimation and the first policy that we evaluate using the estimated model. The combination of data, empirical model, and identification thus allow us to assess the aggregate return to charter school choice—e.g. whether (and by how much) the average student benefits—as well as to quantify how policy can amplify that return by influencing the types and locations of charter schools.

We report several major findings. The first is that lifting the charter school cap generated economically-meaningful aggregate human capital returns (nearly $2,00 per student in lifetime income on average). We further show that competition is the channel driving these overall gains, as students induced into choosing charter schools due to the expansion experience negative test score impacts. Counterfactual simulations indicate the expected returns to the marginal charter school are always positive (due to competitive effects), suggesting that approving new charters on the margin would increase test scores. The findings also highlight the importance of curricular information for decision-making, as traditional charter schools always have larger expected net impacts. This is in part because curriculum horizontally differentiates charter schools, as hypothesized in Gilraine et al. (2021). Finally, the results emphasize that directing both traditional and non-traditional charter schools to disadvantaged areas can significantly enhance their impact.

These findings are informative about the aggregate impacts on student learning of policies that expand school choice. With respect to charter school choice, they permit, for the first time, calculation of aggregate returns from charter school expansion and comparison with other large-scale U.S. education reforms in dollar terms. More generally, the results illustrate that expanded school choice can yield gains to the average student via the competitive channel, even if school choice alternatives are not on average better than traditional schooling options. In addition, the results highlight the role that strategic differentiation by schools, on dimensions other than location such as curriculum, can play in education markets. The findings also have broad implications for the design of school choice programs. U.S. states, for example, have taken different approaches to screening and authorizing charter schools. Our results speak directly to the policy choices facing charter school authorities.

References

Abdulkadiroglu, Atila, Joshua D. Angrist, Susan M. Dynarski, Thomas J. Kane, and Parag A. Pathak (2011), "Accountability and flexibility in public schools: Evidence from Boston's charters and pilots." *Quarterly Journal of Economics*, 126, 699–748.

Abdulkadiroğlu, Atila, Parag A. Pathak, and Christopher R. Walters (2018), "Free to choose: Can school choice reduce student achievement?" *American Economic Journal: Applied Economics*, 10, 175–206.

Abdulkadiroğlu, Atila, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters (2020), "Do parents value school effectiveness?" *American Economic Review*, 110, 1502–39, URL `https://www.aeaweb.org/articles?id=10.1257/aer.20172040`.

Allende, Claudia (2019), "Competition under social interactions and the design of education policies." URL `https://www.gsb.stanford.edu/faculty-research/working-papers/competition-under-social-interactions-design-education-policies`.

Angrist, Joshua, Peter Hull, Parag Pathak, and Christopher Walters (2016a), "Interpreting tests of school VAM validity." *American Economic Review*, 106, 388–92.

Angrist, Joshua D., Sarah R. Cohodes, Susan M. Dynarski, Parag A. Pathak, and Christopher R. Walters (2016b), "Stand and deliver: Effects of Boston's charter high schools on college preparation, entry, and choice." *Journal of Labor Economics*, 34, 275–318.

Angrist, Joshua D., Peter D. Hull, Parag A. Pathak, and Christopher R. Walters (2017), "Leveraging lotteries for school value-added: Testing and estimation." *Quarterly Journal of Economics*, 132, 871–919.

Angrist, Joshua D., Parag A. Pathak, and Christopher R. Walters (2013), "Explaining charter school effectiveness." *American Economic Journal: Applied Economics*, 5, 1–27.

Bau, Natalie (2022), "Estimating an equilibrium model of horizontal competition in education." *Journal of Political Economy*, 130, 1717–1764.

Bettinger, Eric P. (2005), "The effect of charter schools on charter students and public schools." *Economics of Education Review*, 24, 133–147.

Biasi, Barbara, Julien M Lafortune, and David Schönholzer (2024), "What works and for whom? effectiveness and efficiency of school capital investments across the u.s." Working Paper 32040, National Bureau of Economic Research, URL `http://www.nber.org/papers/w32040`.

Bifulco, Robert and Helen F. Ladd (2006), "The impacts of charter schools on student achievement: Evidence from North Carolina." *Education Finance and Policy*, 1, 50–90.

Booker, Kevin, Scott M. Gilpatric, Timothy Gronberg, and Dennis Jansen (2008), "The effect of charter schools on traditional public school students in Texas: Are children who stay behind left behind?" *Journal of Urban Economics*, 64, 123–145.

Bross, Whitney and Douglas N. Harris (2016), "How (and how well) do charter authorizers choose schools? evidence from the recovery school district in new orleans." Technical report, Education Research Alliance for New Orleans, URL https://educationresearchalliancenola.org/files/publications/Bross-Harris-How-Do-Charter-Authorizers-Choose-Schools.pdf.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2011), "Robust inference with multiway clustering." *Journal of Business & Economic Statistics*, 29, 238–249.

Campos, Christopher (2024), "Social interactions, information, and preferences for schools: Experimental evidence from los angeles." Working paper, University of Chicago Booth School of Busines, URL https://www.cqcampos.com/assets/pdf/zoc_experiment.pdf.

Campos, Christopher and Caitlin Kearns (2024), "The impact of public school choice: Evidence from los angeles's zones of choice*." *Quarterly Journal of Economics*, 139, 1051–1093, URL https://doi.org/10.1093/qje/qjad052.

Carter, Samuel Casey (2000), *No Excuses: Lessons from 21 High-Performing, High-Poverty Schools*. Heritage Foundation, Washington, DC., URL https://eric.ed.gov/?id=ED440170.

Chabrier, Julia, Sarah Cohodes, and Philip Oreopoulos (2016), "What can we learn from charter school lotteries?" *Journal of Economic Perspectives*, 30, 57–84.

Chen, Feng and Douglas N Harris (2023), "The market-level effects of charter schools on student outcomes: A national analysis of school districts." *Journal of Public Economics*, 228, 105015.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2014a), "Measuring the impacts of teachers I: Evaluating bias inTeacher Value-Added Estimates." *American Economic Review*, 104, 2593–2632.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2014b), "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review*, 104, 2633–2679.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2017), "Measuring the impacts of teachers: Reply." *American Economic Review*, 107, 1685–1717.

Cohodes, Sarah R and Katharine S. Parham (2021), "Charter schools' effectiveness, mechanisms, and competitive influence." URL http://www.nber.org/papers/w28477.

Cohodes, Sarah R, Elizabeth M Setren, and Christopher R Walters (2021), "Can successful schools replicate? scaling up boston's charter school sector." *American Economic Journal: Economic Policy*, 13, 138–167.

Cordes, Sarah A. (2018), "In pursuit of the common good: The spillover effects of charter schools on public school students in New York City." *Education Finance and Policy*, 13, 484–512.

Cornman, Stephen Q., Laura C. Nixon, Matthew J. Spence, Lori L. Taylor, and Douglas E. Geverdt (2019), "American community survey comparable wage index for teachers (acscwift) may 2019." Technical report, National Center for Education Sciences, URL https:

//nces.ed.gov/programs/edge/docs/EDGE_ACS_CWIFT_FILEDOC.pdf.

Corradini, Viola (2024), "Information and access in school choice systems: Evidence from new york city." Working paper, Columbia Business School, URL https://drive.google.com/file/d/1QdH31DkS1CD38Wi3c1yS0WJBzIgnz41r/view.

CREDO (2009), "Multiple choice: Charter school performance in 16 states." Technical report, Center for Research on Education Outcomes, Stanford University, URL https://credo.stanford.edu/reports/MULTIPLE_CHOICE_CREDO.pdf.

Crema, Angela (2024), "School competition, classroom formation, and academic quality." Technical report, Yale School of Management, URL https://drive.google.com/file/d/1kzQe863WwuQ5XIwE_ZSOOAlyQXxTR6vv/view.

Deming, David J. (2014), "Using school choice lotteries to test measures of school effectiveness." *American Economic Review*, 104, 406–11.

Deming, David J., Justine S. Hastings, Thomas J. Kane, and Douglas O. Staiger (2014), "School Choice, School Quality, and Postsecondary Attainment." *American Economic Review*, 104, 991–1013.

Dobbie, Will and Roland G. Fryer, Jr. (2013), "Getting Beneath the Veil of Effective Schools: Evidence From New York City." *American Economic Journal: Applied Economics*, 5, 28–60.

Dobbie, Will and Roland G. Fryer Jr. (2015), "The medium-term impacts of high-achieving charter schools." *Journal of Political Economy*, 123, 985–1037.

Epple, Dennis, Francisco Martinez-Mora, and Richard Romano (2021), "Charter school practices and student selection: An equilibrium analysis." Working Paper 29529, National Bureau of Economic Research, URL http://www.nber.org/papers/w29529.

Epple, Dennis, Richard Romano, and Ron Zimmer (2016), "Charter schools: A survey of research on their characteristics and effectiveness." In *Handbook of the Economics of Education* (Ludger Woessmann Eric A. Hanushek, Stephen Machin, ed.), volume 5, 139–208, Elsevier.

Ferreyra, Maria Marta and Grigory Kosenok (2018a), "Charter school entry and school choice: The case of Washington, D.C." *Journal of Public Economics*, 159, 160–182.

Ferreyra, Maria Marta and Grigory Kosenok (2018b), "Charter school entry and school choice: The case of Washington, DC." *Journal of Public Economics*, 159, 160–182.

Figlio, David and Cassandra M. D. Hart (2014), "Competitive effects of means-tested school vouchers." *American Economic Journal: Applied Economics*, 6, 133–156.

Figlio, David N., Cassandra M.D. Hart, and Krzysztof Karbownik (2020), "Effects of scaling up private school choice programs on public school students." Working Paper 26758, National Bureau of Economic Research, URL http://www.nber.org/papers/w26758.

Gilraine, Michael, Jiaying Gu, and Robert McMillan (2020), "A new method for estimating teacher value-added." Working Paper 27094, National Bureau of Economic Research, URL http://www.nber.org/papers/w27094.

Gilraine, Michael and Odhrain McCarthy (2024), "The effect of the prior teacher on value-added." *Review of Economics and Statistics*, 1–45.

Gilraine, Michael, Uros Petronijevic, and John D Singleton (2021), "Horizontal differentiation and the policy effect of charter schools." *American Economic Journal: Economic Policy*, 13, 239–76.

Heissel, Jennifer A. and Helen F. Ladd (2018), "School turnaround in North Carolina: A regression discontinuity analysis." *Economics of Education Review*, 62, 302–320.

Henry, Gary T., Shanyce L. Campbell, and Charles L. Thompson (2014), "Evaluation of district and school transformation school-level coaching and professional development activities." Technical report, Consortium for Educational Research and Evaluation – North Carolina.

Henry, Gary T. and Edward J. Guthrie (2019), "The effects of race to the top school turnaround in North Carolina." Working Paper 19-107, EdWorkingPapers, URL https://www.edworkingpapers.com/ai19-107.

Henry, Gary T., Edward J. Guthrie, and LaTricia W. Townsend (2015), "Outcomes and impacts of north carolina's initiative to turn around the lowest-achieving schools." Technical report, Consortium for Educational Research and EvaluationNorth Carolina.

Hoxby, Caroline M. (2002), "School Choice and School Productivity (or Could School Choice be a Tide that Lifts All Boats?)." Working Paper 8873, National Bureau of Economic Research, URL http://www.nber.org/papers/w8873.

Hoxby, Caroline M. and Sonali Murarka (2009), "Charter schools in New York City: Who enrolls and how they affect their students' achievement." Working Paper 14852, National Bureau of Economic Research, URL http://www.nber.org/papers/w14852.

Imberman, Scott A. (2011), "The effect of charter schools on achievement and behavior of public school students." *Journal of Public Economics*, 95, 850–863.

Kane, Thomas J. and Douglas O. Staiger (2008), "Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation." Working Paper 14607, National Bureau of Economic Research, URL http://www.nber.org/papers/w14607.

Kirkebøen, Lars J (2022), "School value-added and long-term student outcomes."

MacKinnon, James G. and Matthew D. Webb (2018), "The wild bootstrap for few (treated) clusters." *Econometrics Journal*, 21, 114–135.

MacLeod, W. Bentley and Miguel Urquiola (2013), "Competition and education productivity: Incentives writ large." In *Education Policy in Developing Countries*, 243–284.

McMillan, Robert (2004), "Competition, incentives, and public school productivity." *Journal of Public Economics*, 88, 1871–1892.

Mehta, Nirav (2017), "Competition in Public School Districts: Charter School Entry, Student Sorting, and School Input Determination." *International Economic Review*, 58, 1089–1116.

Neilson, Christopher (2017), "Targeted vouchers, competition among schools and the academic achievement of poor students." URL https://christopherneilson.github.io/work/documents/Neilson_SEPVouchers.pdf.

North Carolina Education Research Data Center (2008-2017), "Student, class and personnel files." URL http://childandfamilypolicy.duke.edu/research/hc-education-data-center/.

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

Singleton, John D. (2017), "Putting dollars before scholars? Evidence from for-profit charter schools in Florida." *Economics of Education Review*, 58, 43–54, URL http://www.sciencedirect.com/science/article/pii/S0272775716302527.

Singleton, John D. (2019), "Incentives and the supply of effective charter schools." *American Economic Review*, 109, 2568–2612.

Slungaard Mumma, Kirsten (2022), "The effect of charter school openings on traditional public schools in Massachusetts and North Carolina." *American Economic Journal: Economic Policy*, 14, 445–74.

Thernstrom, Abigail and Stephan Thernstrom (2004), *No Excuses: Closing the Racial Gap in Learning*. Simon and Schuster.

Tobin, Zachary (2024), "How do public schools respond to competition? evidence from a charter school expansion." *Economics of Education Review*, 99, 102519, URL https://www.sciencedirect.com/science/article/pii/S027277572400013X.

Walters, Christopher R. (2018), "The Demand for Effective Charter Schools." *Journal of Political Economy*, 126, 2179–2223.

Winters, Marcus A. (2012), "Measuring the effect of charter schools on public school student achievement in an urban environment: Evidence from New York City." *Economics of Education Review*, 31, 293–301.

Zimmer, Ron and Richard Buddin (2009), "Is charter school competition in California improving the performance of traditional public schools?" *Public Administration Review*, 69, 831–845.

Online Appendix

A Appendix: Supplemental Statistics

Figure A.1: Number of Charter Schools in North Carolina by Year

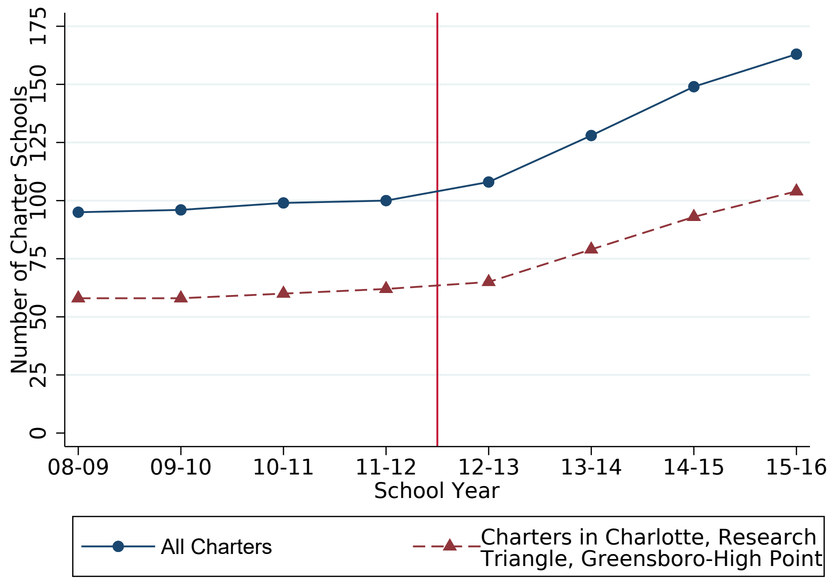

Notes: This figure displays the number of charter schools by year in North Carolina from 2008-09 to 2015-16, excluding two virtual charter schools that opened in 2015-16. The dashed line then displays the number of charter schools by year in the three commuting zones (Charlotte, Research Triangle, and Greensboro-High Point) that make up the "educational markets" that we analyze. The vertical line represents the lifting of the 100 school charter cap for the 2012-13 school year.

Figure A.2: Distribution of Newly-Opened Charter School Value-Added in 2015-16 and 2018-19

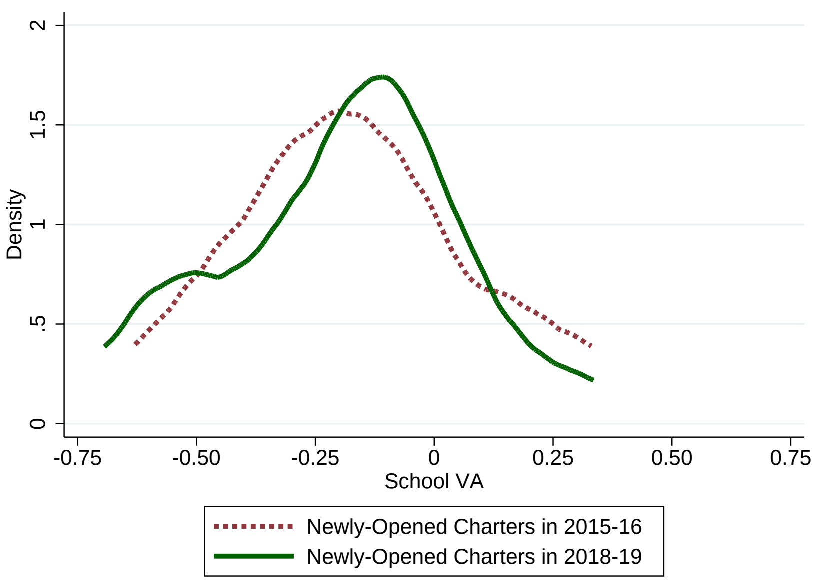

Notes: This figure investigates the potential for schools to improve over time by comparing the distribution of school value-added among newly-opened charters in the 2015-16 school year to their distribution in the 2018-19 school year. The 2015-16 value-added distribution is identical to the one in Figure 1(a).

Figure A.3: Event Studies: Newly-Opened Charter on Nearby half the School Value-Added

(a) New Charter has Traditional Curriculum

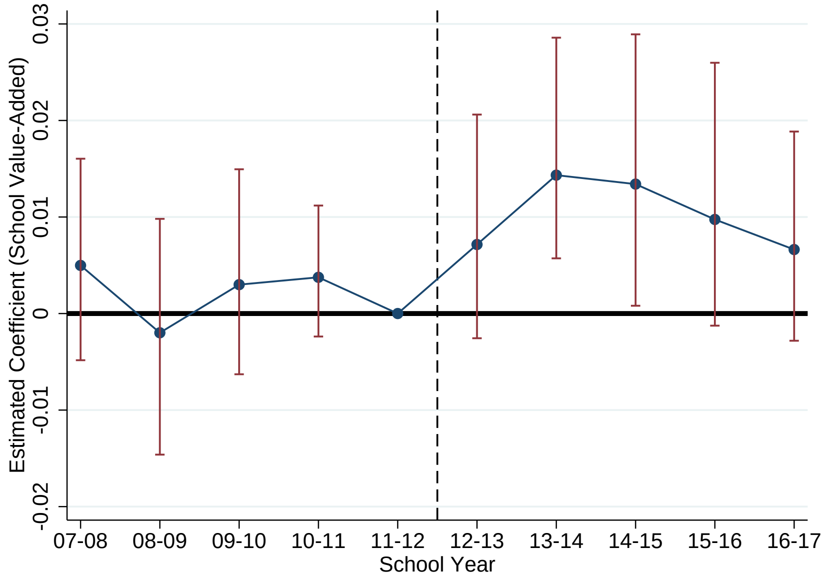

(b) New Charter has Non-Traditional Curriculum

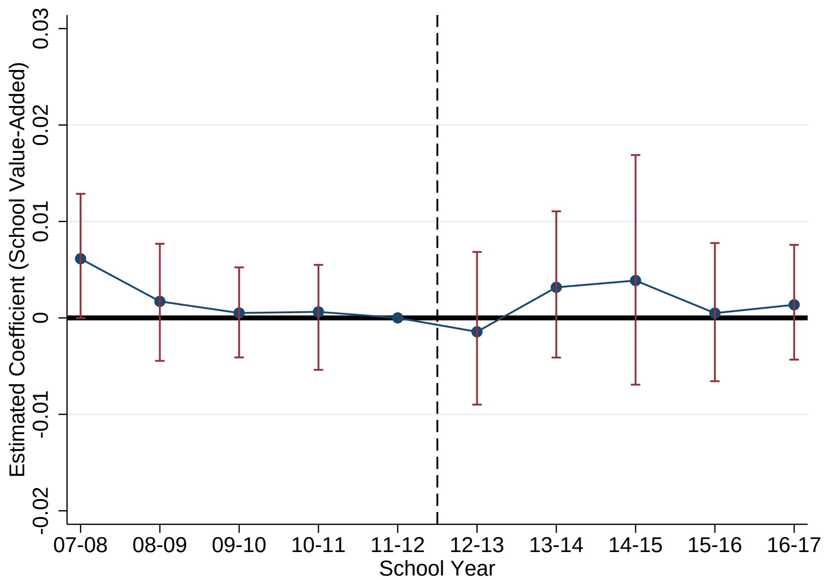

Notes: This figure shows the estimated English value-added difference between schools 'treated' by a newly-opened charter relative to 'control' schools by year. This figure is therefore identical to Figure 3, but uses value-added estimated using English test scores (rather than math test scores) as the outcome. Data are restricted to schools in the Charlotte, the 'Research Triangle' (i.e., Raleigh-Durham-Chapel Hill), and Greensboro-High Point commuting zones. Treated schools are defined as schools located within 20 miles of a newly-opened charter that opened in 2012-13 or 2013-14. Control schools are defined as schools located between 20 and 30 miles of a charter schools that opened in 2012-13 or 2013-14. Results are subdivided by whether the newly-opened charter follows a traditional curriculum (Figure A.3(a)) or not (Figure A.3(b)). Note that 2012-13 is considered the first 'treated' year because although the charters themselves opened in either the 2012-13 or 2013-14 school year, public schools would have known by the start of 2012-13 whether or not a charter was opening nearby or would open nearby in 2013-14. The dashed vertical line therefore separates the 'pre-years' from the 'post-years'. The horizontal line represents a point estimate of zero. The 'whiskers' represent 90 percent confidence intervals with inference conducted via wild clustered bootstrap (Cameron et al., 2011) two-way clustering by newly-opened charter and public school.

Figure A.4: Robustness: Varying 'Treatment' Definition for Reduced-Form Difference-in-Differences Regressions

(a) School Value-Added

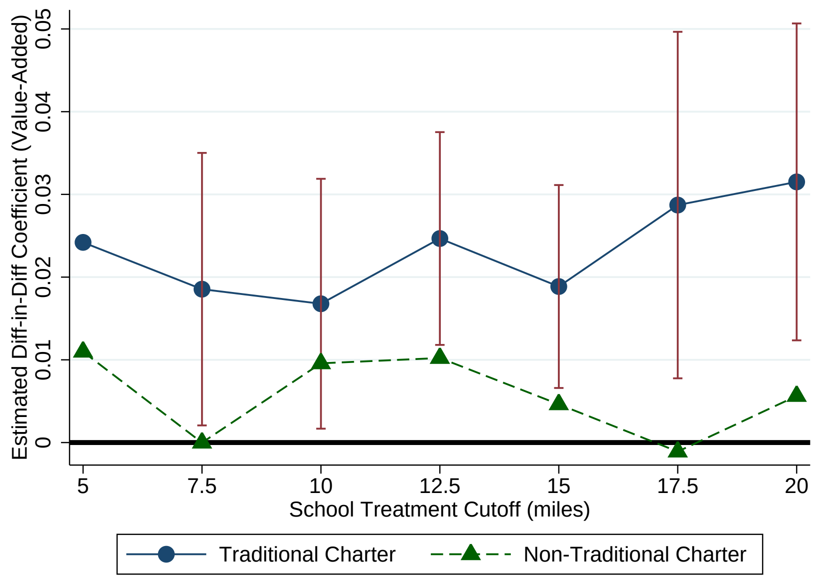

(b) School Enrollment

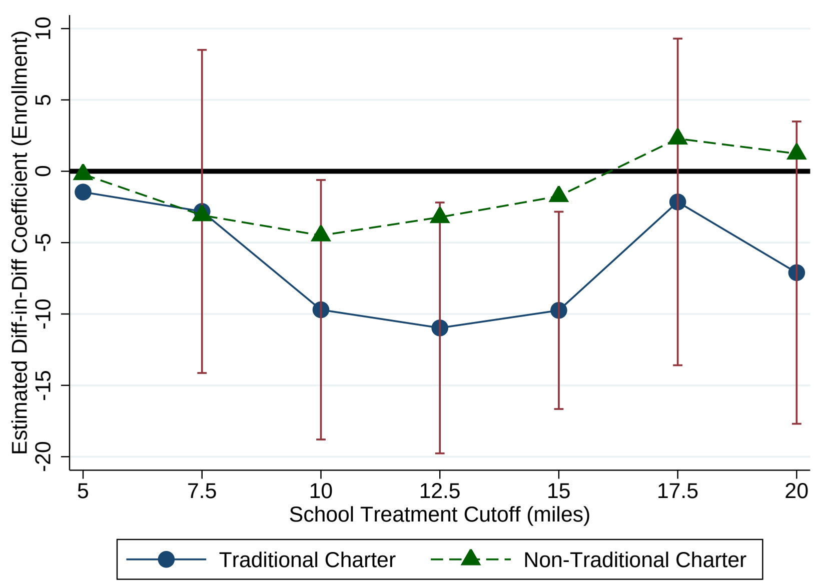

Notes: This figure shows robustness to Table A.2 in terms of the radius where schools are defined as 'treated.' In particular, we shrink the treatment radius from the 20 miles we used in Table A.2 in 2.5 miles increments all the way down to 5 miles. We then report the coefficients from this regression, both for traditional newly-opened charters (solid line) and non-traditional newly-opened charters (dashed line). Data are restricted to schools in the Charlotte, the 'Research Triangle' (i.e., Raleigh-Durham-Chapel Hill), and Greensboro-High Point commuting zones. Therefore the point estimates we report for the treatment radius of 20 miles are identical to those in columns (2) and (4) of Table A.2. The 'whiskers' represent 90 percent confidence intervals with inference conducted via wild clustered bootstrap (Cameron et al., 2011) two-way clustering by newly-opened charter and public school.

Figure A.5: Event Studies: Newly-Opened Charter on Nearby **Private** School Enrollment

(a) New Charter has Traditional Curriculum

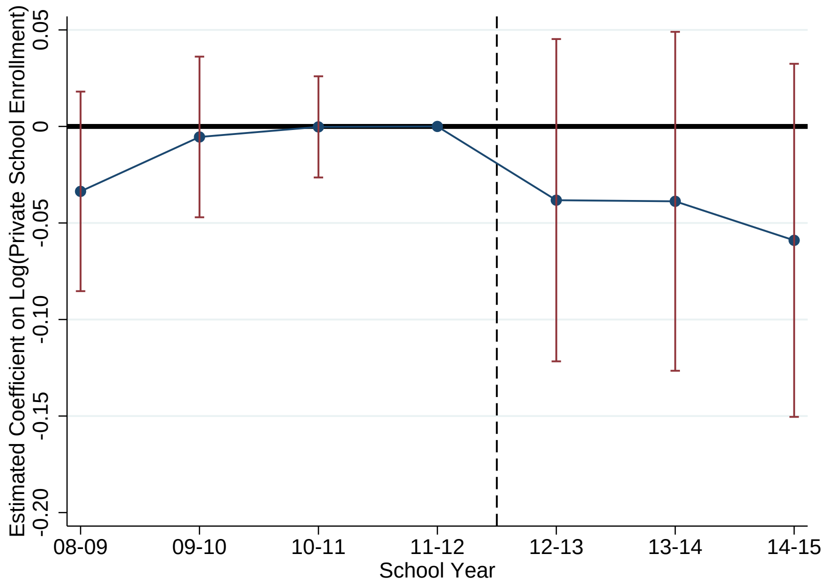

(b) New Charter has Non-Traditional Curriculum

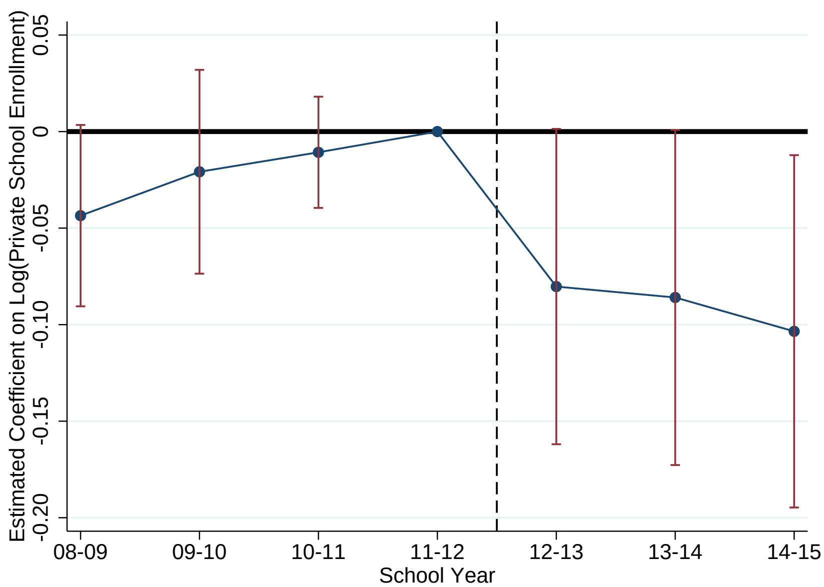

Notes: This figure shows the estimated school-level enrollment difference between **private** schools 'treated' by a newly-opened charter relative to 'control' private schools by year. Data are restricted to private schools that cover a K-2 grade and are located in Charlotte, the 'Research Triangle' (i.e., Raleigh-Durham-Chapel Hill), and Greensboro-High Point commuting zones. Treated private schools are defined as schools located within 20 miles of a newly-opened charter that opened in 2012-13 or 2013-14. Control private schools are defined as schools located between 20 and 30 miles of a charter schools that opened in 2012-13 or 2013-14. Results are subdivided by whether the newly-opened charter follows a traditional curriculum or not. Note that 2012-13 is considered the first 'treated' year because although the charters themselves opened in either the 2012-13 or 2013-14 school year, private schools would have known by the start of 2012-13 whether or not a charter was opening nearby or would open nearby in 2013-14. The dashed vertical line therefore separates the 'pre-years' from the 'post-years'. The horizontal line represents a point estimate of zero. The 'whiskers' represent 90 percent confidence intervals with inference conducted via wild clustered bootstrap (Cameron et al., 2011) two-way clustering by newly-opened charter and public school.

Table A.1: Summary Statistics: Value-Added Sample

	Full Sample[1] (1)	Value-Added Sample[2] (2)
Mean of Student Characteristics		
Mathematics Score (σ)	0.00	0.02
Reading Score (σ)	0.00	0.02
Lagged Mathematics Score (σ)	0.02	0.02
Lagged Reading Score (σ)	0.02	0.02
% White	52.0	52.4
% Black	25.4	25.3
% Hispanic	14.4	14.3
% Asian	2.9	2.7
% Economically Disadvantaged	52.0	52.2
% English Learners	6.1	5.5
% Gifted	15.1	15.9
% Students with Disability	13.2	12.9
# of Students	1,284,838	1,191,936
Observations (student-year)	2,238,703	2,084,317

[1] Data coverage: grades 4-5 from 2007-08 through 2016-17.

[2] The difference in sample sizes comparing columns (1) and (2) arises because we drop 154,386 million student-year observations that do not have contemporaneous or lagged math scores.

Table A.2: Difference-in-Differences Results: Public School Value-Added and Enrollment

Outcome:	Value-Added		Enrollment	
	All Schools	Charlotte, Triangle, Greensboro CZs	All Schools	Charlotte, Triangle, Greensboro CZs
	(1)	(2)	(3)	(4)
Panel A. Pooled				
All Newly-Opened Charters	0.021**	0.021*	-2.79	-3.32
	[0.002, 0.049]	[-0.001, 0.050]	[-10.38, 5.61]	[-10.95, 5.54]
Panel B. Heterogeneous				
Newly-Opened Charter follows	0.033**	0.034**	-7.32**	-9.18**
Traditional Curriculum	[0.006, 0.072]	[0.001, 0.074]	[-16.33, -0.58]	[-19.82, -1.13]
Newly-Opened Charter follows	0.003	0.004	3.81	3.96
Non-Traditional Curriculum	[-0.021, 0.038]	[-0.022, 0.045]	[-7.01, 11.12]	[-7.14, 11.84]
Test of Equality by Curriculum (p-value)	0.02	0.06	0.02	0.01
Observations (school-year)	28,358	24,003	28,364	24,007

Notes: This table shows difference-in-differences estimates from equation (13), whereby schools located within 20 miles of a newly-opened charter school are considered 'treated' while those located 20-30 miles from a newly-opened charter are considered 'control' and the effect is allow to differ by whether the newly-opened charter school follows a traditional or non-traditional curriculum. 'Test of Equality by Curriculum' reports the p-value of the hypothesis test that the point estimate for traditional curriculum charters is the same as the one for non-traditional curriculum charters. All inference is conducted via wild clustered bootstrap (Cameron et al., 2011) two-way clustering by newly-opened charter and public school. As confidence intervals are asymmetric under the wild clustered bootstrap, 95% confidence intervals are reported below the point estimate in square brackets. ***,** and * denote significance at the 1%, 5% and 10% levels, respectively.

Table A.3: Summaries of Census Tracts and Simulated Charter School Locations

	Average	Avg Tract With Charter		Median Income		Local Public VA	
	Tract	Mean	SD	High	Low	High	Low
Local public VA	0.04	0.02	0.18	0.04	0.05	0.21	-0.14
Density	10900	13018	9543	26804	12529	16163	15285
Median income	80956	77918	24105	111672	55452	116502	58991
N local publics	3.51	3.82	2.51	9.50	4.50	6.00	5.00
% Black	24.72	30.37	17.11	18.10	32.74	11.04	38.18
N tracts	903	56		2	2	2	2

Notes: Table reports mean characteristics across all Census tracts in the Triangle and Charlotte-Mecklenburg CZs (Average); across only tracts that contain a charter schools (Avg Tract With Charter; the 84th (High) and 16th (Low) percentile tracts in each market in terms of Median Income; and the 84th (High) and 16th (Low) percentile tracts in each market in terms of the VA of local public schools. N tracts is 2 in the four rightmost columns because the column includes 1 tract from each market. Note that averages and standard deviations are reported for Avg Tract With Charter. Characteristics correspond to total and averages within 3 miles of tract centroid.

Table A.4: Decomposition of Aggregate Impact of Marginal Charter School by Location and Type

		(1)	(2)	(3)	(4)	(5)	(6)
Location Type:		Average	Avg Tract	Median Income		Local Public VA	
		Tract	With Charter	High	Low	High	Low
				Traditional			
VA Draw:							
$z = -1$	Competition	0.12	0.13	0.11	0.13	0.11	0.07
	Sorting	-0.18	-0.18	-0.18	-0.15	-0.19	-0.09
$z = 0$	Competition	0.11	0.13	0.11	0.11	0.10	0.09
	Sorting	-0.04	-0.04	-0.04	-0.01	-0.05	-0.01
$z = 1$	Competition	0.11	0.12	0.10	0.10	0.10	0.11
	Sorting	0.08	0.08	0.08	0.11	0.07	0.10
$\mathbb{E}$ impact $\mid \neg$Competition		-0.05	-0.05	-0.05	-0.02	-0.06	0.00
				Non-Traditional			
VA Draw:							
$z = -1$	Competition	0.09	0.10	0.09	0.13	0.10	0.04
	Sorting	-0.16	-0.16	-0.17	-0.19	-0.22	-0.06
$z = 0$	Competition	0.09	0.10	0.09	0.12	0.10	0.04
	Sorting	-0.06	-0.06	-0.06	-0.05	-0.08	-0.02
$z = 1$	Competition	0.09	0.10	0.09	0.11	0.09	0.05
	Sorting	0.04	0.04	0.04	0.07	0.04	0.03
$\mathbb{E}$ impact $\mid \neg$Competition		-0.06	-0.06	-0.06	-0.05	-0.09	-0.02

Notes: This table reports decomposition into contributions of competition and sorting of estimated impact on the average student's test scores of an additional charter entrant by location (columns), quality draw (rows), and by type as obtained by simulations. Simulated entrant's enrollment is capped at 200 students and entrant is assigned median demand shock (given type) in all cells. Impacts are normalized by the cap lifting impact (0.0049σ). Expected impact ignoring competition approximated using Gauss-Hermite weights over Sorting cells only. Average Tract refers to Census tract most similar to average tract in the CZ; Avg Tract With Charter to tract most similar to average among those containing at least one charter; "High" and "Low" refer to -1σ (16th percentile) and +1σ (84th percentile), respectively. See Appendix Table A.3 for descriptive characteristics of locations. Note that value of a VA draw depends on charter school type; e.g. a $z = -1$ non-traditional charter school has lower quality than a $z = -1$ traditional charter school. Table 7 report overall impacts (sum of competition and sorting).

B Appendix: Validity of School Value-Added

School value-added (VA) is the contribution of a school to the achievement growth of their students, once all other determinants of student learning have been taken into account. We estimate school value-added using a standard model of student achievement in which education inputs (including school quality) are additively separable in their effects. The achievement of a student i at school s in year t is written as:

$$y_{ist} = X'_{ist}\beta + q_{st} + \epsilon_{ist} \tag{10}$$

where y_{ist} is the student's test score, X_{ist} is a large vector of observable individual and school-level student characteristics, and ϵ_{ist} is a random test score shock which is assumed to be iid normal with variance σ^2_ϵ. Our control vector X_{ist} includes: (i) lagged test scores using a cubic polynomial in prior-year scores in math and English, interacted with grade dummies, (ii) demographics, including: economically disadvantaged status, ethnicity (six ethnic groups), gender, limited English status, gifted status, and disability status. We also include the following school-grade level controls: (iii) cubics in school-grade means of prior-year test scores in math and English (defined based on those with non-missing prior scores) interacted with grade dummies, (iv) cubics in school-grade means of all the demographic covariates, (v) school-grade size, and (vi) grade-by-year dummies. The estimated value of q_{st} represents the contribution of school s in year t to test scores that is unexplained by student characteristics, or simply its value-added.

The validity of value-added models hinge on the control vector, X_{ist}, being sufficiently rich so that potential test scores are independent of a student's school choice, conditional on observables. The key control used in value-added models is previous student achievement (i.e., lagged test scores) which is taken as as a sufficient statistic for the unobserved history of inputs received by children. Two commonly-cited criticisms of the value-added approach is that lagged test scores fail to account for the sorting of students to schools and past inputs may decay at differential rates. The prior literature has tried to assess the validity of school value-added measures by leveraging random school assignment from school choice lotteries (Deming, 2014; Angrist et al., 2016a, 2017) and find that school value-added measures feature some bias, although it is limited and therefore school value-added is likely to be a useful proxy for school quality. For instance, Angrist et al. (2017) state that "conventional VAM estimates are biased. At the same time, OLS VAM estimates tend to predict lottery effects on average, with estimated forecast coefficients close to 1. OLS estimates would therefore seem to be useful even if imperfect."

We ensure in this appendix section that our school value-added added measures do not feature quantitatively significant bias which could materially affect our structural results. To do so, we assess the validity of our school value-added measure using two standard tests of forecast unbiasedness in the value-added literature: (i) out-of-sample forecasts, and (ii) student-switching quasi-experiments.

Out-of-Sample Forecasts: As a first check, we ensure that school value-added can appropriately forcecast student achievement gains. Intuitively, an OLS regression of (residualized) student test scores, y_{ist}, on estimated school value-added, $\hat{q}_{st}$, should yield a coefficient of one by construction. To ensure that this is the case, we follow Chetty et al. (2014a) and estimate measures of school value-added using equation (10) that leave out (jackknife) years $t-1$, t, and $t+1$. We leave out data from years $t-1$, t, and $t+1$ when estimating school value-added in year t as otherwise given that students are usually in school s both this year *and* either the prior or future year (given that our data consist of grade 4 and 5 students only) failing to leave these data out would introduce the same estimation errors on both the left- and right-hand side of the regression and produce biased school value-added estimates (Chetty et al., 2014a). We therefore regress:

$$\tilde{y}_{ist} = a + \lambda \hat{q}_{st}^{-\{t-1,t,t+1\}} + v_{ist} \tag{11}$$

where $\hat{q}_{st}^{-\{t-1,t,t+1\}}$ denotes the best (linear) forecast of school value-added for school s in period t, omitting data from years $t-1$, t, and $t+1$ in forming the prediction. Our outcome measure, $\tilde{y}_{ist}$, either represents student test score growth (i.e., $y_{ist} - y_{is,t-1}$) or residualized test scores.[56]

Columns (1) and (2) of Table B.1 report our estimate of out-of-sample forecast bias. Regardless of whether we use raw test score growth or residualized test scores as our outcome,

[56] We residualize student test scores using the same vector of observable characteristics X_{ist} in equation (10) and including school fixed effects so that we only estimate β using within-school variation.

we find a coefficient near 1, indicating that our estimated school value-added measure is able to forecast student test score growth out-of-sample. As this relationship could either be driven by the causal impact of schools (q_{st}) or differential student sorting based on unobservables (ϵ_{ist}), we turn to leveraging quasi-experimental variation coming from students switching schools.

Student Switcher Quasi-Experiment: We assess the validity of school value-added (VA) in our setting by leveraging a quasi-experiment that occurs often in administrative education data: students switching schools. This student switcher quasi-experiment follows a similar quasi-experiment in Chetty et al. (2014a) which leverages teacher switching, but adapted to leverage students switching schools. To do so, we regress test score residuals (or raw test score growth) on the change in the (appropriately jackknifed) school VA of students. Crucially, school VA is always measured as predicted VA in year t, so that there is no change in school VA among students who do not switch schools and so our regression is solely identified off students who switch schools.

Let $\hat{Q}_{ist}^{-\{t+1,t,t-1,t-2\}}$ be the jackknifed estimated value-added of school s in year t which omits years $t+1$, t, $t-1$, and $t-2$ from the value-added calculation.[57] Define $\Delta\hat{Q}_{ist} = \hat{Q}_{ist} - \hat{Q}_{is,t-1}$ as the change in school (jack-knife) VA in school s between years t and $t-1$ for student i. Importantly note that both $\hat{Q}_{ist}$ and $\hat{Q}_{is,t-1}$ omit years $t+1$, t, $t-1$, and $t-2$ and are therefore identical if the student attends the same school in both t and $t-1$.[58] Given this, the only variation in our regression will come from students who switch schools as $\Delta\hat{Q}_{ist}$ will be zero for school stayers. We then check for bias in our school VA measure using the following regression:

$$\Delta y_{ist} = a + \lambda\Delta\hat{Q}_{ist} + \epsilon_{ist} \tag{12}$$

where $\Delta y_{ist} = y_{ist} - y_{is,t-1}$ is the change in student i's test score between years t and $t-1$.

[57]We must omit years $t+1$, t, $t-1$, $t-2$ since we will regress student test score on school value-added. Since students are in the school this year and could be in the school last year or next year (as we use grade 4 and 5 data), if data from years $t+1$, t, and $t-1$ is used to construct value-added, then student i test score will enter both the left and right-hand sides of the regression leading to bias. $t-2$ must also be omitted as use $\Delta\hat{Q}_{is,t-1}$ in our first-difference equation. See Chetty et al. (2017) for a discussion of this issue.

[58]If one were to include drift in the school VA model, then one would need to instrument for $\Delta\hat{Q}_{ist}$ using appropriately jack-knifed VA forecasts for schools in a different period (say $t-2$) only as done in Gilraine and McCarthy (2024).

(Alternatively, we use student i's residualized test score[59] which should yield similar results given the lagged test score controls.)

Table B.1: Value-Added Validity Tests

Validity Test:	Out-of-Sample Forecast		School Switchers	
Outcome:	Test Score Growth (1)	Test Score Residuals (2)	Test Score Growth (3)	Test Score Residuals (4)
Coefficient (λ)	0.992	1.001	0.961	0.847
	(0.004)	(0.003)	(0.010)	(0.010)
P-Value Coef. Equals 1 ($\lambda = 1$)	0.061	0.646	0.000	0.000
Observations	1,854,643	1,854,643	1,851,275	1,851,275

Notes: Table reports results from our two forecast unbiasedness checks: (i) out-of-sample forecast, and (ii) student-switching quasi-experiment. Columns (1) and (2) conduct our out-of-sample forecast test described in equation (11) with raw test score growth and residualized test scores as the outcome, respectively. Columns (3) and (4) report results from our student-switching quasi-experiments described in equation (12) using raw test score growth and residualized test scores as the outcome, respectively. When using test score gains as the outcome we include our control vector, X_{ist}, excluding the student-level lagged test score controls. Standard errors clustered at the student-level are reported in brackets. The third row indicates the p-value of the hypothesis that the coefficient equals one. Data cover grades 4-5 from 2007-08 to 2017-18.

Columns (3) and (4) of Table B.1 report estimates of forecast bias leveraging the student switcher quasi-experiment. We find a forecast coefficient of 0.961 and 0.847 when using raw test score growth and residualized test scores as outcomes, respectively. These coefficients statistically differ from 1, indicating that school value-added estimates are forecast bias. That said, the estimates are near 1 indicating that our estimated school value-added measure is able to forecast student test score growth out-of-sample. We also would like to highlight the remarkably similar forecast estimates we find using the student switching quasi-experiment to those found by Angrist et al. (2017) who leverage school choice lotteries: Angrist et al. (2017) estimate a forecast coefficient of 0.950 and 0.864 for their test score growth and test score residual models, respectively. Those estimates are near-identical to ours, even though those authors exploit randomization via lotteries while we exploit quasi-experimental variation for students switching schools.

[59]We residualize student test scores using the same vector of observable characteristics X_{ist} in equation (10) and including school fixed effects so that we only estimate β using within-school variation.

C Appendix: Further Details on Charter Entry Event Studies

This Appendix sets out the estimating equations we use in Section 2.2.2 to capture the impact of nearby charter openings on public school quality and enrollment. To construct our data, we create a separate dataset for each of the 15 relevant charter school entries (see Section 2.2.2) that consist of all public schools within 30 miles of the newly-opened charter school. We then create dataset indicators and append our data together. We then run the following stacked difference-in-differences regression:

$$\begin{aligned} y_{scdt} = \delta_{sc} + \lambda_t + \gamma_d * t + \beta_{\textit{non-trad}} Post_t * treat_{sc} + \mu distance_{sc} * Post_t * treat_{sc} \\ + Trad_c \left(\lambda_t + \gamma_d * t + \beta_{trad} Post_t * treat_{sc} + \nu distance_{sc} * Post_t * treat_{sc} \right) + \epsilon_{scdt} \, , \end{aligned} \tag{13}$$

where y_{scdt} is school VA or K-2 enrollment in school s nearby newly-opened charter c in district d at time t, $Post_t$ is a post-charter cap indicator (e.g., equals 1 if 2012-13 or later), $treat_{sc}$ is an indicator equal to one if a school is within 20 miles of the newly-opened charter and zero within 20-30 miles, $Trad_c$ is an indicator equal to one if the newly-opened charter follows a traditional curriculum, $distance_{sc}$ is the distance to the newly-opened charter, $\gamma_d * t$ are district (linear) time trends, and δ_{sc} and λ_t are school-by-charter opening and year fixed effects. The parameter $\beta_{nontrad}$ captures the average change between treated and untreated schools when a non-traditional charter opens, while the sum $\beta_{\textit{non-trad}} + \beta_{trad}$ estimates the effect when traditional charters open. The results of equation (13) are reported in Table A.2.

To build Figure 3, we estimate an event-study version of equation (13) where we use event time indicators in lieu of the post indicator. Specifically, we regress:

$$\begin{aligned} y_{scdt} = \delta_{sc} + \lambda_t + \gamma_d * t + \sum_{\tau \neq -1} \beta^{\tau}_{\textit{non-trad}} (D^{\tau}_t * treat_{sc}) + \mu distance_{sc} * Post_t * treat_{sc} \\ + Trad_c \left(\lambda_t + \gamma_d * t + \sum_{\tau \neq -1} \beta^{\tau}_{trad} (D^{\tau}_t * treat_{sc}) + \nu distance_{sc} * Post_t * treat_{sc} \right) + \epsilon_{scdt} \, , \end{aligned} \tag{14}$$

where D^{τ}_t are indicators equal to one if year t is τ years after (or before, if negative) 2012-13 (i.e., the year of charter entry) and 0 otherwise. All other variables are defined in equation (13). The coefficients β^{τ}_{trad} and $\beta^{\tau}_{non\text{-}trad}$ (along with their confidence intervals) are then plotted in Figures 3(a) and 3(b), respectively.

Private School Enrollment: To investigate the impact of nearby charter entry on private school enrollment we gather enrollment data from all private schools in the state from 2008-09 to 2014-15.[60] (We lack private school test scores so we cannot investigate quality responses by the private schools.) Unfortunately, these data only report school-level enrollment and the grades taught by the school. Therefore, we cannot create a K-2 enrollment measure, although we do drop all private schools that do not teach grades K-2. We also focus on schools with an enrollment of 10 or more students leaving us with a sample of 335 private schools of which 189 are located in our three CZs of interest.

We conduct a similar exercise as above to look at the impact of nearby charter entry on private school enrollment. We once again construct our data by creating a separate dataset for each of the 15 relevant charter school entries (see Section 2.2.2) that consist of all private schools within 30 miles of the newly-opened charter school. We then create dataset indicators and append our data together. We then run the event-study regression described by equation (14) using (log) school-level private school enrollment as the outcome. The coefficients (along with their confidence intervals) are then plotted in Figure A.5.

[60] Data are available from `https://ncadmin.nc.gov/public/private-school-information/nc-directory-private-schools`.

D Appendix: Further Details About TALAS

In this Appendix, we provide a greater overview of the TALAS program in North Carolina. Schools were placed into TALAS in two ways. Individual schools statewide fell into the program if their 2009-10 proficiency rates or (in the case of high schools) graduation rates were below a given threshold. In addition, the District and School Transformation division of the North Carolina Department of Public Instruction identified 12 school districts with substandard (district-wide) aggregate performance composites in the 2009-10 year and placed *all* schools in those districts into a similar district-level turnaround program. A total of 118 schools qualified as TALAS schools due to substandard performance composites or graduation rates, while an additional 188 schools fell into the program via the district route, resulting in a total of 306 treated schools.[61]

Previous evaluations of North Carolina's TALAS program are sensitive to whether an average or local average treatment effect is estimated and, by extension, whether schools that entered TALAS under the district-level program are included as treated schools in the analysis. Using a difference-in-differences framework to identify an average treatment effect, Henry et al. (2014) and Henry et al. (2015) find positive effects of TALAS on school proficiency rates. Importantly, the difference-in-differences framework also includes in the analysis schools that were treated under the district-level program and indicates especially strong gains among the previously lowest-performing schools. In contrast, Heissel and Ladd (2018) and Henry and Guthrie (2019) both use a regression discontinuity design to estimate a local average treatment effect among schools who qualified for TALAS because their 2009-10 proficiency rate was below the set threshold, thereby discarding schools that were treated under the district-level program and restricting the comparison to schools with previous proficiency rates close to the cutoff. Both studies find no (or even a small negative) effect on student achievement among elementary and middle schools.

To maximize the available data variation, in this paper we use a difference-in-differences approach to estimate an average treatment effect, thereby allowing us to include in the analysis all TALAS schools, not just those with a 2009-10 proficiency rate near the TALAS

[61] For more information on how North Carolina implemented its Race to the Top turnaround programs, see `https://web.archive.org/web/20120919064916/http:/www.ncpublicschools.org/schooltransformation/overview/`.

threshold. Figure D.1 shows the effect of TALAS on year-by-year school value-added. We show raw trends in value-added over time for TALAS treated and untreated schools in Figure D.1(a). The value-added of TALAS schools jumps when the program is fully implemented in 2011-12 and then exhibits a steady decline in subsequent years, eventually reverting back to pre-reform levels when the funding for TALAS program expires after the 2014-15 school year. As expected, the value-added of untreated schools follows a more stable trend. Figure D.1(b) presents regression-adjusted estimates (along with 95-percent confidence intervals) of year-by-year differentials in value-added between TALAS treated and untreated schools. Relative to untreated schools, the value-added of treated schools increases by 0.1 units, or 43 percent of a standard deviation, in the first year of the program's full implementation; by the 2015-16 academic year, however, the effect of the program disappears.

In Figure D.2, we show that average TALAS treatment effects are remarkably similar across the two types of treated school (that is, those that entered the program via either the school- or district-level routes), lending credence to our approach of pooling all treated schools to improve precision (as we do in Figure D.1) in our estimation approach in equation (8).

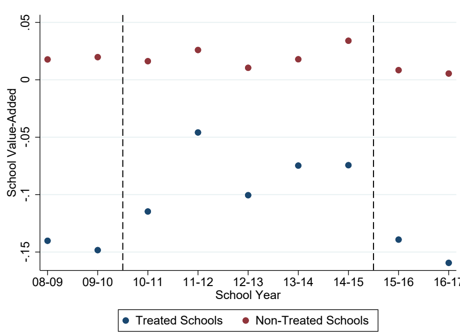

(a) Trends in School Value-Added

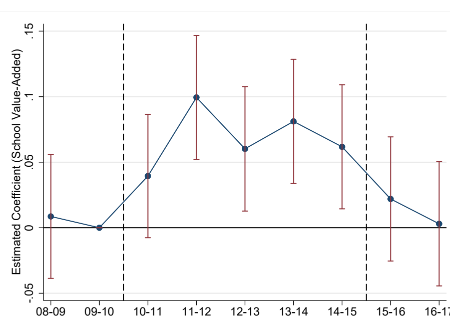

(b) Regression Adjusted Differential Changes

Figure D.1: TALAS-Driven Variation in School Value-Added

Notes: This figure shows trends in school value-added over time for TALAS and non-TALAS schools. Panel (a) shows mean school-level value-added in each year for both types of school. Panel (b) plots regression-adjusted estimates of the difference in mean value-added across TALAS and non-TALAS schools in each year. The blue circles in panel (b) represent the estimated coefficients on academic-year-TALAS-indicator variables from a regression of school-year value-added on these variables, year fixed effects, and school fixed effects. The red bars represent the 95-percent confidence intervals associated with the coefficient estimates, with standard errors clustered at the school level.

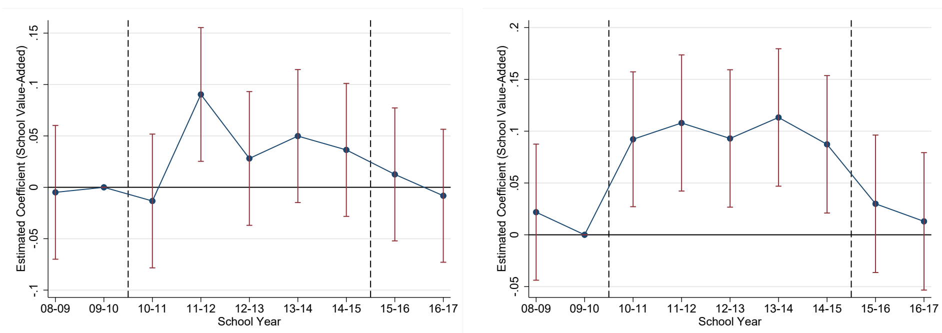

(a) Effects of District-Level Treatment (b) Effects of School-Level Treatment

Figure D.2: TALAS-Driven Variation in School Value-Added by Treatment Type

Notes: This figure shows regression-adjusted estimates of the difference in mean value-added across TALAS and non-TALAS schools in each year. In panel (a), we define TALAS schools as only those schools that entered the program because the aggregate performance composite of their district in 2009-10 was below a threshold; in panel (b), we define TALAS schools as only those schools that entered the program because their school-specific performance composite in 2009-10 was below a threshold. In both panels, the blue circles represent the estimated coefficients on academic-year-TALAS-indicator variables from a regression of school-year value-added on these variables, year fixed effects, and school fixed effects. The red bars represent the 95-percent confidence intervals associated with the coefficient estimates, with standard errors clustered at the school level.