

Competition in Public School Districts: Charter School Entry, Student Sorting, and School Input Determination

Nirav Mehta¹

University of Western Ontario

Abstract

This paper develops and estimates an equilibrium model of charter school entry, school input choices, and student school choices. The structural model renders a comprehensive and internally consistent picture of treatment effects when there may be general equilibrium effects of school competition. Simulations indicate that the mean effect of charter schools on attendant students is positive and varies widely across locations. The mean spillover effect on public school students is small but positive. Lifting caps on charter schools would more than double entry but reduce gains for attendant students.

¹nirav.mehta@uwo.ca The data were provided by North Carolina Education Research Data Center. I thank Kenneth Wolpin, Hanming Fang, and Elena Krasnokutskaya for their guidance, and the Editor for many useful suggestions. I have greatly benefited from discussions with Andrew Clausen, Tim Conley, Chao Fu, Eleanor Harvill, Lance Lochner, Rachel Margolis, Salvador Navarro, Seth Richards-Shubik, Shalini Roy, David Russo, Panos Stavrinos, Todd Stinebrickner, Petra Todd, and the participants in the Penn Empirical Micro lunch group and seminar. This research was supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305C050041-05 to the University of Pennsylvania and the SSHRC Insight Development Grant Program. The opinions expressed are those of the author and do not represent views of the U.S. Department of Education or SSHRC.

1 INTRODUCTION

The provision of school choice is often proposed as a way to improve educational outcomes for students in poorly performing public schools. Charter schools are at the center of the recent debate concerning education policy reforms, such as President Obama’s *Race to the Top*, which rewards states that lift legislative caps on the number of charter schools; such caps are present in most states with charter schools (White, 2009).² Policymakers would like to know how student achievement has been affected by existing charter schools and how it would be affected if they expanded the role of charter schools in public education. Advocates argue that charter schools improve the performance of students attending charters (“direct effect”) and students attending competing public schools (“spillover effect”). Critics of charter schools argue that they “cream-skin”—that is, the better outcomes at charter schools represent student selection, not test score gains—and that charter schools negatively affect students attending competing public schools.

The need to understand how charter schools affect student achievement has motivated a large body of empirical work. Some of this work uses lottery designs, which estimate direct treatment effects using oversubscribed schools (Hoxby and Rockoff, 2005; Angrist et al., 2012); other work estimates value-added models of test score growth using panel data on students who switch between public and charter schools (Bettinger, 2005; Bifulco and Ladd, 2006; Sass, 2006; Hanushek et al., 2007). Bifulco and Ladd (2006), Sass (2006), Chakrabarti (2008), and Imberman (2011) use a variety of methods to estimate spillover effects of school choice. Estimates of direct and spillover effects are widely mixed across studies, which is consistent with Gleason et al. (2010), who find substantial heterogeneity in charter school impacts on attendant students.

Prior research highlights the heterogeneity of charter school effects on student achievement but cannot provide a comprehensive evaluation of how charter school policy affects student achievement for several reasons. First, policymakers interested in the effect of lifting caps on the number of charter schools need a way to extrapolate findings from studies of existing charter schools to new charter schools serving different populations of students. Second, prior studies do not model why charter schools open in certain places but not in others (Hanushek et al., 2007). Understanding where new charter schools would open and

²Though charter schools are publicly funded and, therefore, technically a type of public school, for brevity I typically refer to them as “charter schools” and to traditional public schools as “public schools.” Like public schools, charter schools cannot selectively admit students. They typically have considerably more autonomy than public schools regarding personnel decisions, curricula, school hours, and pedagogical methods, and often have lower per-pupil resources due to a lack of separate capital funding streams. All students have access to a public school but not all students have access to a charter school because charter schools enter certain areas and not others.

their effects in these areas is crucial to the debate about lifting charter school caps. Third, increasingly popular lottery-based designs cannot quantify the effect of all existing charter schools on student achievement because they do not provide a way to extrapolate results from oversubscribed charter schools to those that are not oversubscribed. The potential for bias could be large if oversubscribed charter schools are also those that households believe will deliver stronger benefits. Fourth, although several authors have quantified spillover effects, none have provided a coherent framework that uses estimates of bias from spillovers to adjust estimated direct effects of school choice (e.g., Cullen et al., 2006). This could be important if a charter school improved outcomes for all students but students remaining at the public school benefited more than did those attending the charter, in which case a lottery-based design would find a negative direct effect.

This paper develops and estimates an equilibrium model of competition between charter and public schools to confront these issues. It does so by modeling three key components of the school choice debate: student school choices, charter and public school inputs, and charter school entry decisions. The model incorporates selection on student ability in two ways: charter school entry decisions take into account market-level ability distributions and, within markets, student sorting is a function of heterogeneous student ability and inputs at both charter and public schools. Modeling student school choices as a function of both student and school characteristics allows for generalization of estimates based on existing charter schools to charter schools that might enter in new markets were caps lifted, even in the presence of student sorting on unobserved ability. The model of school input choices allows charter schools to have heterogeneous treatment effects across markets through variation in input provision and also predicts what inputs for public schools would have been in the absence of charters, which is necessary to properly quantify spillover effects. The equilibrium framework provides an internally consistent method to quantify both spillover effects and the bias introduced when one ignores equilibrium responses by public schools, unifying the two strands of the literature where authors either estimate the direct effect of charter schools or use different students and schools to estimate spillover effects on public school students. By modeling charter school entry it is possible to quantify how many more charter schools would open and in which markets they would open were caps lifted, which is important if the effects of charter schools are heterogeneous across markets. This paper also estimates the extent to which peer ability affects student achievement at public and charter schools, meaning that the model allows for spillovers through changes in both student ability and public school input choices.

The model is estimated using maximum likelihood on administrative data from the North Carolina public school system. The data contain the universe of schools and students in the

North Carolina public school system from 1998 to 2001 and include variables that enable estimation of the model's demand and supply sides, such as public and charter school locations, charter school entry decisions, and detailed per-pupil school resources, which enter the model as a per-pupil capital index. The school attended, typical weekly hours of homework reported done, and standardized test scores are recorded for each student in each year. The data also contain student locations, which enter the model through a distance cost of attending a school, shifting the probability a student will attend a charter school. Student-level distance data are aggregated into market-level distance distributions, which affect ability sorting and provide variation to identify peer effects.

Average weekly hours of homework done by students at a school constitutes the measure of effort, the endogenous school input to test score production. Effort has an intuitive appeal as an input choice and average time spent doing homework is a natural candidate for the effort input for several reasons. First, learning course material takes time so the use of time spent doing homework is firmly grounded in economic theory. Second, time-use data are commonly used in the economics of education, as in Stinebrickner and Stinebrickner (2008) and Ferreyra and Liang (2012), who use time spent on studying and homework, respectively, as measures of student effort. These studies focus on how much students work, which relates to the notion of effort used in this paper because schools may face a cost of getting students to work some desired amount and students may face a cost of working for the amount of time desired by their school. Third, Stinebrickner and Stinebrickner (2004) show that time-use reports of own study time vary considerably even within a semester. The existence of noise, which may in part be due to transient shocks in reported homework, supports averaging student homework levels at the school-year level to measure how much they work.

The estimated model fits the data well and provides a natural framework to answer several questions crucial to informing charter school policy. The estimated mean direct effect of charter schools on attendant students is 11% of a standard deviation (sd) in test scores and there is a positive, though small, mean spillover effect of 2% sd on public school students in duopoly markets. These effects are driven by charter schools' more effective technology, the relative unresponsiveness of public school effort to charter schools, and a lack of strong mean sorting on student ability. Having observed their superior technology, one might consider a policy forcing charter schools to serve all students in markets they enter. However, this policy would reduce the number of charter schools and, because charters drastically reduce effort when forced to serve many more students, would also dramatically reduce average gains in student achievement. There is substantial between-market heterogeneity in the mean direct effect; in particular, the direct effect of a charter school is positively related to demand for it, which should caution policymakers seeking to generalize results from lottery-based studies

based on oversubscribed charter schools. Finally, more than twice as many charter schools would enter if caps were lifted, though the larger average market size and lower average per-pupil capital in new entry markets greatly diminishes their estimated direct effects.

This paper complements the literature modeling competition between public and private schools (Epple and Romano, 1998; Nechyba, 2000; Caucutt, 2002; Ferreyra, 2007). While this literature improves our understanding of student school choices, peer effects, and the effects of private school vouchers, public schools are assumed to be monolithic and do not make input choices, precluding effort-based spillover effects of school choice.³ In a related literature, Epple and Sieg (1999), Sieg et al. (2004), and Bayer et al. (2004, 2007) study equilibrium sorting across school districts and municipalities to understand housing markets and local public good provision (see Kuminoff et al., 2013 for a recent review). Such sorting could generate variation in public school capital levels—this paper takes such sorting, and therefore, public school resources, as given, instead focusing on entry and competition. Other work studies sorting for different types of schools, such as Epple et al. (2006), Fu (2014), and Kapor (2015), which estimate equilibrium models of the higher education market.

This paper is closest to Ferreyra and Kosenok (2015), which develops and estimates a model of charter school location and input decisions in Washington DC. As with the literature modeling competition between public and private schools, in their model public schools do not choose inputs. Instead, public schools follow reduced-form policy rules, precluding their framework from producing endogenous spillover effects operating through changes in input choices. Also relevant is Walters (2014), which uses lotteries to estimate a flexible model of household demand for charter schools, using these estimates to quantify treatment effects for students attending new charter schools. In contrast to the current paper, Walters models application behavior but does not model school input choices or charter school location decisions and does not allow for peer effects in the production technology.

There are also many papers studying either parts of this problem or related problems. Altonji et al. (2015) develop an econometric framework to quantify the extent to which private school voucher programs cream-skim from public schools, and find a small, but negative, effect on the performance of students remaining at the public school. Glomm et al. (2005) describe determinants of charter school location decisions. McMillan (2004) develops a theoretical model where a public school may reduce its provision of costly effort in response to competition from a private school. The current paper naturally complements this work by quantifying spillovers stemming from charter school competition. Hastings et al. (2009), Agarwal and Somaini (2015), and Calsamiglia et al. (2016) use households' stated preferences over schools to estimate household demand. The latter two papers study a context in which

³See Neilson (2013) for a more recent example.

households submit school rankings to a centralized school assignment mechanism; related to these papers is De Haan et al. (2015), which compares household survey data with rankings households submitted to a centralized mechanism. The current paper instead uses revealed preferences to estimate demand, while making progress towards understanding charter school entry and school input determination.

2 MODEL

2.1 Overview An economy in the model is one market and time period. Each economy contains one traditional public school, one potential charter school entrant, and a continuum of households, each with one student. The potential charter school entrant makes an entry decision, after which the public school and charter school—if it has entered—choose effort inputs to test score production.⁴ Households, which differ by student ability and location within the market, choose the school that maximizes their utility. All households are assigned to the public school if the charter school did not enter. In equilibrium, the charter school makes the optimal entry decision based on its expected equilibrium value of entry, neither school wishes to change its effort level, and no household wishes to switch schools.

Denote the public school tps (for “traditional public school”) and the potential charter school entrant ch . Schools and students are indexed by s and i , respectively. Market and time subscripts are suppressed in this section to simplify exposition. Variables in bold denote the pair of variables for both schools, e.g., $\mathbf{k} = (k_{tps}, k_{ch})$.

2.2 Students There is a continuum of students of measure μ . A student $i \in I$ has ability a_i , where $a_i \sim F(a_i)$ with density $f(a_i)$. Student i maximizes utility by choosing a school $s_i \in S_i$, their school choice set. If there is a charter school in student i ’s market then $S_i = \{tps, ch\}$, otherwise $S_i = \{tps\}$. Student (ex-post) choice-specific utility is $u_{is} = y_{is} - c_{is} + \eta_{is}$, where y_{is} is i ’s test score at school s , c_{is} is i ’s non-pecuniary cost of attending school s , and $\eta_{is} \sim N(0, \sigma_\eta^2)$ is a choice-specific preference shock.⁵

The test score depends on student ability, peer quality—defined as the mean ability of students at the school— \bar{a}_s , school effort e_s and capital k_s , and a productivity shock ν_{is}^y ,

⁴Charter school “entry” corresponds to either new entry or continued operation within a market. This distinction is irrelevant in the current section but matters for estimation and simulation.

⁵Charter and public schools cannot charge tuition.

which is realized after students choose schools:⁶

$$(1) \quad y_{is} = a_i \omega_s \bar{a}_s^{\theta_s} (\alpha_s e_s^{\beta_s} + (1 - \alpha_s) k_s^{\beta_s})^{\tau_s / \beta_s} + \nu_{is}^y = Ey_s(a_i, \bar{a}_s, e_s, k_s) + \nu_{is}^y,$$

where ω_s is Hicks-neutral total factor productivity (TFP) and $\nu_{is}^y \sim i.i.d. N(0, \sigma_{\nu^y}^2)$. Each production function has a constant elasticity of substitution (CES) between capital and effort. Similar to Hoxby (2002), who models charter and public schools as having different production functions, the CES parameters are allowed to differ between charter and public schools. The motivation for allowing the CES parameters to differ is that charter schools are typically organized and run differently from traditional public schools (Dobbie and Fryer Jr., 2013); well-documented differences include hiring different types of teachers (Hoxby, 2002; Podgursky, 2006) and having teachers who employ innovative instructional strategies (Caruthers, 2012). This paper models the potential for such innovation by allowing—but not forcing—charter and public schools to have different production functions.

Student i 's cost of attending school s depends on school effort, distance from the student to the school, r_{is} , a fixed cost of attending the charter school, c_{ch} , and an ability-specific cost of attending the charter school $c_{ch,a}$, according to $c_{is} = c_e e_s + c_r r_{is} + (c_{ch} + a_i c_{ch,a}) 1\{s = ch\}$.⁷ The choice to allow for non-pecuniary costs of attending charter schools is based on the observation that charters often have lower levels of amenities, such as gymnasiums or playgrounds. Valuation of these amenities may depend on student ability, hence the ability-specific charter school cost $c_{ch,a}$.

The optimal school choice policy for student i in a duopoly market is

$$(2) \quad s_i^* = \arg \max_{s \in S_i} \{E_{\nu_i^y} [u_{is}]\} = \arg \max_{s \in S_i} \{Ey_s(a_i, \bar{a}_s, e_s, k_s) - c_{is} + \eta_{is}\},$$

where $\nu_i^y = (\nu_{i,tps}^y, \nu_{i,ch}^y)$ is the pair of test score productivity shocks for i .

⁶Some authors in this literature include student covariates, such as race, in the test score outcome equation (Angrist et al., 2012) or student utility (Ferreira and Kosenok, 2015). This paper effectively includes such covariates in student-specific ability. Note that mean peer covariates are consequently subsumed by the mean ability of students at the school. The assumption that, within a location, households care more about peer ability than race has empirical support. Fischer (2003) finds that racial sorting into location is twice as strong as income-based sorting; income likely affects ability through the provision of costly inputs. Therefore, there may be more within-location variation in ability than in race.

⁷None of the parameters c_e, c_r, c_{ch} , or $c_{ch,a}$ are assumed to be “costs” in the estimation of the model in the sense that their signs are unrestricted. For example, effort may be costly for households because spending time doing homework (or spending time getting one’s kids to do their homework) imbues effort with an opportunity cost. On the other hand, if effort produces components of human capital that are valued by households but not reflected in standardized achievement tests, then households may have negative effort costs because effort serves as a stand-in for an unmeasured outcome desired by households.

2.3 Schools Each market is endowed with a public school and a potential charter school entrant. Each school is endowed with a capital level and location within the market, all of which are known to the potential charter school entrant before it makes its entry decision, denoted $z \in \{no\ entry, entry\}$.

The school's objective is a weighted average of the average test score of students at the school \bar{y}_s , school size μ_s , and cost function c_s :

$$(3) \quad v_s = \delta_{ys} \bar{y}_s + \mu_s - c_s = \delta_{ys} \bar{a}_s E y_s(1, \bar{a}_s, e_s, k_s) + \mu_s - c_s.$$

Student responses to school choices determine both \bar{y}_s (through mean ability \bar{a}_s) and μ_s (details in Appendix A.1). If the average test score did not directly enter the school's objective function ($\delta_{ys} = 0$), the model would predict that monopoly public schools would exert no effort because they would draw all students in their market without incurring any effort cost. The inclusion of school size μ_s helps capture administrators' incentives; e.g., a principal might be fired for poor management if very few students attend their school.

The school's cost function allows for interactions between effort and capital, school size, and mean student ability:

$$(4) \quad c_s = \psi_{es1} e_s + \psi_{es2} e_s \mu_s + \psi_{es3} e_s k_s + \psi_{es4} e_s \bar{a}_s - (\psi_{\text{mkt.size}} \mu + \psi_{\text{fr.Black}} \mu_{\text{Black}}) 1\{s = ch\},$$

where $\psi_{\text{mkt.size}}$ is the charter school's valuation of market size μ , $\psi_{\text{fr.Black}}$ is the charter school's valuation of being in a market with a share of Black students μ_{Black} . The inclusion of an interaction term between effort and mean ability ψ_{es4} allows the school's cost of getting students to spend the desired amount of time doing homework to depend on mean student ability. For example, more work may need to be assigned to have high-ability students spend a certain amount of time doing homework. Though schools are risk-neutral, a potential preference for larger markets $\psi_{\text{mkt.size}}$ allows charter schools to target larger markets to make sure they have enough students. The last term $\psi_{\text{fr.Black}}$ allows the model to capture the fact that charter schools may want to serve Black students, which may explicitly be part of their mission statements (Bifulco and Ladd, 2007).

If there is a charter school in the market each school has a policy

$$(5) \quad e_s^* = \arg \max_{e_s} E_{\nu^e} [v_s(\mathbf{e})] = \arg \max_{e_s} E_{\nu^e} [v_s(e_s, e_{-s})],$$

i.e., each school chooses its own effort e_s to maximize its expected objective, given the action of the other school e_{-s} and distribution of effort productivity shocks $\nu^e = (\nu_{tps}^e, \nu_{ch}^e)$, which

is realized after schools choose effort. Each shock ν_s^e determines observed effort according to

$$(6) \quad e_s^o = e_s \nu_s^e,$$

where $\ln \nu_s^e \sim i.i.d. N(0, \sigma_{\nu^e}^2)$.

Students in monopoly markets have no school choice so the average ability for students attending the monopoly public school is the market average \bar{a} and the measure of students attending is the market size μ , returning the monopolist objective $v_{tps}^{mono} = \delta_{y,tps} \bar{y}_{tps} + \mu - c_{tps}$. The monopolist public school has a policy

$$(7) \quad e_{tps}^{mono*} = \arg \max_{e_{tps}} E_{\nu_{tps}^e} [v_{tps}^{mono}(e_{tps})].$$

The charter school's optimal entry decision is

$$(8) \quad z^* = entry \Leftrightarrow E_{\nu^e} [v_{ch}(\mathbf{e}^*)] \geq \underline{v},$$

where \mathbf{e}^* is the pair of equilibrium effort levels chosen by both schools and \underline{v} is a random variable known to the charter school, which denotes an exogenous fixed cost of entry and/or operating.

2.4 Equilibrium The solution concept is subgame perfect Nash equilibrium, so the model is solved by backwards-induction.

A *Subgame Perfect Nash Equilibrium* for the period game is an entry decision and student and school decisions such that

- a: the charter school enters ($z^* = entry$) if and only if the entry cost shock is less than its expected payoff in the entry subgame (eq. (8)),
- b: if the charter school enters, the ensuing subgame equilibrium is the *Entry Subgame Equilibrium*, and
- c: if the charter school does not enter, the ensuing subgame equilibrium is the *Monopoly Subgame Equilibrium*.

An *Entry Subgame Equilibrium* consists of charter and public school effort choices, realized school effort, household school choices, and mean school abilities such that

- b.i*: each school's chosen effort e_s^* maximizes the expected objective of the school, given the other school's chosen effort (eq. (5)),
- b.ii*: realized school effort for each school e_s^o is chosen effort multiplied by the effort productivity shock (eq. (6)),
- b.iii*: student school choices s_i^* maximize utility given realized school effort (eq. (2)), and
- b.iv*: the pair of mean abilities at both schools $\bar{\mathbf{a}} = (\bar{a}_{tps}, \bar{a}_{ch})$ is consistent with realized school effort and student optimization (eq. (26) in Appendix A.1).

A *Monopoly Subgame Equilibrium* consists of a chosen public school effort, realized public school effort, household school choices, and mean public school ability such that

- c.i*: the public school's chosen effort e_{tps}^{mono*} solves the monopolist public school's problem (eq. (7)),
- c.ii*: realized public school effort $e_{tps}^{mono^o}$ is the monopolist's chosen effort multiplied by its effort productivity shock (eq. (6)),
- c.iii*: all students attend the public school ($s_i^* = tps, \forall i \in I$), and
- c.iv*: mean ability at the public school is market mean ability ($\bar{a}_{tps}^{mono} = \bar{a}$).

I solve for the equilibrium value of entry by first deriving student school choices as a function of school effort, their own ability, and their distance to charter and public schools. I then find a Nash equilibrium in chosen charter and public school effort by iterating best responses, which requires computation of a fixed point in mean student abilities at charter and public schools within every iteration. Equilibrium effort levels are then hit with productivity shocks, after which they enter household school choice problems. The model is solved numerically (details in Appendix A.1).⁸ I have proved existence of an equilibrium and though I have not proven there is a unique equilibrium, an extensive search has never found more than one equilibrium effort pair in the entry subgame (see Appendices B.1-B.2).

The model generates rich interactions between charter and public schools, which capture key parts of the debate on school choice and are key to evaluating charter school policy. Charter school entry decisions depend on market characteristics, such as student ability and locations and school capital. The model inherently produces heterogeneous treatment effects of charter schools so it can be used to extrapolate findings to markets charter schools have not yet entered. The model can generate direct and spillover effects of either sign. If the charter school is much more productive than the public school, high-ability students may attend it with very high probability, which will be smaller the further it is from students. Depending on market characteristics and parameter values, a public school may not find

⁸The model is solved using an approximate version of the schools' problems where schools do not integrate over ν^e , reducing computing time by around a factor of a million. I have used numerical integration to verify that this does not substantially change effort choices.

it advantageous to try to retain high-ability students but may instead decrease equilibrium effort from the monopoly level.

3 DATA

Model parameters are estimated using administrative data from North Carolina. The data are taken from the universe of public and charter schools and were provided by the North Carolina Education Research Data Center (NCERDC). The data contain variables necessary to estimate student-level test score production functions using school-level inputs, and include detailed panels on teachers, students, and charter and traditional public schools in the North Carolina public school system. For teachers, the data contain years of experience and the school in which they work. For students, the data contain demographic characteristics, which school they attend, grade in school, standardized reading and math test scores for students in grades 3-8 and grade 10, self-reported weekly hours of homework done, and student household locations. School-level data are used to calculate computers per pupil and district per-pupil revenues.⁹

Because solving the model requires knowing school capital levels for public and charter schools in all markets, including those without charter schools, I develop an algorithm linking per-pupil expenditures at public schools to capital at both public and charter schools. This algorithm essentially scales down per-pupil expenditures, an intuitive measure of school capital, at public schools to a level consistent with the lower levels observed at charter schools. Solving the model also requires an algorithm computing distance distributions for all markets. The average and median student in most markets is closer to the public school but there is also substantial variation in distance distributions across markets. In particular, there are markets where the majority of students are closer to the charter school. Further details about the capital and distance data can be found in Appendix B.3.

The test score in the model is the average of the reading and math test scores and is normalized to have a mean of 3 and standard deviation of 1 for each grade to make it comparable across grades. The average test score is 3 to ensure that all markets have ability distributions with positive means; otherwise, the model would predict that effort for any school in such a market would be zero if effort was costly for schools.¹⁰

Effort is computed using self-reported data on the hours of homework students say they typically do per week, which are then averaged to create a school-wide effort variable per year (details in Appendix B.3). It may be useful to think of the school as having the ability

⁹Charter schools are considered to be their own school districts in North Carolina.

¹⁰The estimation sample has a mean test score of 3.05 and standard deviation of 0.95, as some observations are lost while making sample restrictions.

to enforce students to do a certain amount of homework per week, but enforcement requires effort from the teachers or principal and is therefore costly. That is, school effort costs broadly capture how hard it is for schools to get their students to work a certain amount.¹¹ As discussed previously, the ability-effort interaction coefficient in school cost functions ψ_{es4} allows this cost to depend on student ability.

There is empirical support for this treatment of effort. First, Stinebrickner and Stinebrickner (2004) show that student self-reports of homework time can be very noisy, which means that averaging student homework levels may help capture effort differences between schools. It is important to note that effort levels do not simply capture between-student variation potentially caused by sorting: almost half of within-student variation in students' own homework levels in the data is explained by changes in the school attended. Another concern is that time spent doing homework may be a problematic measure of effort if higher "effort" were actually a signal of lower ability or if more time spent on homework were a signal of lower inputs elsewhere in school. In either case, increases in effort should be negatively related to academic achievement because it acts as a signal for a lower level of some other input. Regressions of student test scores on prior scores and school capital and effort levels return positive partial correlation coefficients for effort at both charter and public schools, mitigating this potential concern.

Both capital and effort are assumed to be the same for all students at the same school. The assumption is innocuous for capital, because most school capital is applied fairly evenly to students within schools.¹² Even if it were not, data on within-school capital expenditures are not available. In contrast, effort choices for individual students are observed. Assuming that there is only one effort level per school per year means I can avoid solving for students' individual effort choices and has the additional benefit of averaging out the measurement error likely contained in individual effort reports.

3.1 Markets Each student's school choice set must be defined to solve and estimate the model. A market in the model is similar to the catchment area for a traditional public school. In reality, markets partition North Carolina, though these markets do not necessarily correspond to public school attendance zones. To make the model solution independent across markets, charter schools may compete with at most one public school, and each public

¹¹De Fraja et al. (2010) include homework assigned as a measure of school effort in their study of the relationship between child, parent, and school effort, while Ferreyra and Liang (2012) use time spent doing homework as a measure of student effort. Fruehwirth (2014) also considers student-chosen effort. The model here can be thought of as a school choosing an assigned amount of homework, targeting how much time students at the school will take to complete the homework, which is the input to the production function.

¹²Special education and gifted and talented student programs are notable exceptions. Charter schools tend to have much smaller fractions of both types of students.

school may compete with at most one charter school. Each charter school is designated to be in the market of the public school closest in distance, within the same public school district. The assumption that each public school competes with one charter school is supported by the data. In the four instances when the same public school is the closest public school to more than one charter school, I designate the charter school closest to the public school as its competitor and exclude the other charter schools from the sample. Private schools are not part of school choice sets.¹³

There are two main advantages to structuring markets this way. First, limiting the extent of competition allows me to model policy-relevant interactions between charter and public schools. Solving this model is infeasible without some restriction on competition between charter and public schools: If charter schools could draw students from multiple public schools all public schools would then indirectly compete with each other, clearly an intractable scenario. Other models of school choice focus on distinct metropolitan areas that are much smaller than an entire state (e.g., Ferreyra and Kosenok, 2015). Because the current paper considers the entire state of North Carolina, it is natural to not allow every school to compete with every other school in the state.

Second, a geographic rule provides a simple mechanism to create each student’s school choice set and assigns potential charter school entrants to competing public schools even in their absence, which is essential to model entry (Appendix B.7). Such a rule is necessary because charter schools do not have geographic cut-offs for attendance, meaning one cannot always surmise which public school a student would have attended had he not chosen the charter school, nor which charter school could have been attended by a student observed in a public school. For example, elementary “feeder” schools sending students to both public and charter middle schools cannot be used to define markets because they cannot define markets in areas without charter schools. Appendix B.6 contains further discussion of market definitions.

3.2 Estimation Sample The estimation sample consists of middle schools (grades 6-8) from 1998-2001. Elementary, middle, and high schools may have different test score production functions, meaning it makes sense to focus on one school level; the model is estimated using only middle schools for several reasons: middle school provides a natural decision-point for students as most students switch schools between grades 5 and 6, the data contain standardized test scores for all middle school grade levels but only a limited number of

¹³Private schools in North Carolina seem to target different students than charter schools. The share of Black students at charter schools in the estimation sample is 25%, while the share of Black students in private schools in the state was 7% in 2010. Source: Author’s calculations from NCES (2010), <http://nces.ed.gov/surveys/pss/privateschoolsearch/>.

elementary or high school grades, and magnet schools, which have more open enrollment schemes than traditional public schools, are less pervasive at the middle school level: about 80% of students attending magnet schools are not in the grades 6-8 in the years 1999-2001. The estimation sample spans 1998, the first year charter schools were allowed in North Carolina, through 2001, as the statewide cap of 100 charter schools appeared close to binding in 2002.¹⁴ Estimating model parameters on data before the cap was binding obviates modeling the interdependence of charter school entry decisions that would be induced by the cap, meaning the same model can also be used to simulate entry decisions in the absence of caps. I called the North Carolina Department of Public Instruction to inquire about potentially oversubscribed charter schools, so that I could explore the feasibility of comparing my results to lottery-based studies, but because such data were not available such schools are treated as non-oversubscribed.¹⁵

The estimation sample contains markets that were stable over the sample period, students whose school choices were consistent with their market assignments, students observed for at least two years, and a random sub-sample of students from markets where charter schools were never observed during the sample period.¹⁶ The estimation sample includes 78,294 public school student observations in markets without charter schools, 63,216 public school student observations in markets with charter schools, 1,984 total public-school-years, and 4,911 charter school student observations in 128 charter-school-years. Due to a potential concern that excluding students may affect parameter estimates, I re-estimated the model excluding markets where more than 5% of students were observed crossing market boundaries; estimates of parameters and treatment effects remain substantially unchanged (results in Appendix B.7).

3.3 Descriptive Statistics Table 1 shows moments of the marginal distributions of market-level means and standard deviations of test scores in 1997 (i.e., the year before any charter schools entered). The first row presents moments of the distribution of market-level means, and shows that there is considerable spread in mean prior test scores between markets; the 75th percentile highest mean test score is about 15% higher than the 25th percentile mean test score, and the highest is over three times the lowest. The second row presents moments of the distribution of market-level standard deviations and shows that not only do test scores

¹⁴Additionally, data from 1997 are used to identify market ability distributions; see Section 4 for details.

¹⁵In conversations with staff at about one third of the charter schools open at some point during 1998-2001, I learned that schools are often uncertain even about whether they had waiting lists.

¹⁶This last restriction was adopted to ease computation of the likelihood. I sample 100% of students in markets with at most 100 students, 20% of students in markets with 101-200 students, 15% of students in markets with 201-300 students, 10% of students in markets with 301-400 students, and 5% of students in markets with more than 400 students.

vary quite a bit within markets, but there is also considerable heterogeneity in within-market variation of prior test scores. Of course, prior test scores do not directly measure student ability; rather, they are a combination of public school inputs and productivity shocks. Nonetheless, these data suggest that there may be substantial heterogeneity between markets in the distribution of student ability.

Table 1: Distribution of Market-Level Means and Standard Deviations of 1997 Test Scores

	Minimum	25th %ile	Median	Mean	75th %ile	Maximum
Mean	1.322	2.790	2.988	2.985	3.195	3.972
St. dev.	0.152	0.815	0.859	0.859	0.912	1.162

Entry: Charter schools operate in 6% of markets.¹⁷ Table 2 shows that duopoly markets tend to be larger (first column), markets in which charter schools would have higher capital (second column), and markets with lower prior mean test scores (third column).

Table 2: Fraction of Markets with Charter Schools by Market Characteristics

	Market Size μ	Charter Capital k_{ch}	Mean 1997 Test Score \bar{y}_{1997}
Above-median	0.067	0.095	0.050
Below-median	0.045	0.017	0.062

School inputs: On average, charter schools have about three-quarters of the per-pupil capital levels of public schools (0.43 for charter schools, versus 0.54 and 0.56 for monopolist and competitor public schools, respectively). Table 3 describes how mean school effort varies by market characteristics. The first column presents effort for charter schools, the second column presents effort for competing public schools (i.e., those in duopoly markets), and the third column presents effort for monopolist public schools; the first row presents the mean over all markets, the second and third rows present mean effort by market size, and the fourth and fifth rows present mean effort by capital level. Mean effort is higher in markets with charter schools (2.66 and 2.69 hours, respectively, for charter and public schools in duopoly markets, versus 2.43 hours for monopoly public schools). In contrast to their public school competitors, charter schools exert lower effort in larger markets. However, both charter and competing public schools exert more effort in high-capital markets; while this is also true for monopolist public schools, the difference is much smaller.

Student outcomes: The first column of Table 4 shows that students in charter schools have the highest average test scores, followed by students in public schools in duopoly markets, followed by students in public schools in markets without charters. The second column of

¹⁷Statistics pool data over all years in the estimation sample unless otherwise specified.

Table 3: Mean Effort (Hours of Homework) by School Type and Market Characteristics

	Duopoly Markets		Monopoly Markets
	Charter	Public	Public
Total	2.66	2.69	2.43
Above-median μ	2.51	2.85	2.49
Below-median μ	2.89	2.46	2.37
Above-median k_s	2.74	2.78	2.51
Below-median k_s	2.20	2.21	2.35

Table 4 shows that, on average, charter schools comprise only 9.2% of students in duopoly markets, i.e., charter schools are much smaller than their public school counterparts.

Together, the above facts highlight the need for a framework that can accommodate both within- and between-market heterogeneity. The model allows students to sort into schools based on their unobserved ability and, given the substantial between-market heterogeneity, we would naturally expect ability sorting patterns to differ across markets. Although ability is unobserved, eq. (1) shows that prior test scores may provide initial evidence about potential ability sorting. This evidence is easiest to interpret when using data from 1998—the last year before charter schools were allowed to enter, and, therefore, the last year before any ability-based sorting into schools. Therefore, I used a random coefficient model to decompose prior (i.e., from 1997) test scores of students in markets in which charter schools entered in 1998, based on market and whether those students attended the charter school that year, i.e., $s_{i,1998,m} = ch$:

$$(9) \quad y_{i,1997,m} = \lambda_0 + \lambda_m + \gamma_0 1\{s_{i,1998,m} = ch\} + \gamma_m 1\{s_{i,1998,m} = ch\} + \iota_{i,1997,m},$$

where the λ_0 and λ_m capture overall and market-specific mean test scores, respectively, γ_0 represents average sorting on prior test score into charter schools, and γ_m are market-specific deviations from this average sorting, which have a mean of zero and standard deviation σ_{γ_m} .¹⁸ Table 5 shows that there is (weak) evidence of negative average prior-test-score sorting into charter schools (i.e., $\gamma_0 < 0$, with a p-value of 0.584). It also shows that there is substantial between-market variation in this sorting (i.e., $\sigma_{\gamma_m} > 0$, where we reject the hypothesis that $\sigma_{\gamma_m} = 0$ with a p-value of 0.017); indeed, the standard deviation of market-specific charter school coefficients is more than four times bigger than the average charter-school coefficient across all markets.

¹⁸Note that all 1997 test scores pertain to public school students because no charters had yet opened; the *tps* index is omitted to simplify exposition, as it is redundant.

To help understand this heterogeneity in sorting patterns, I now describe how sorting into charter schools based on prior achievement varies with respect to relative school inputs, effort and capital. This investigation is motivated by the model’s implication that, *ceteris paribus*, ability sorting into a particular school should be stronger the higher its effort and capital are relative to its competitor.¹⁹ Specifically, for each market with a charter school in 1998, I compute the mean prior test score for students attending the charter school in 1998 (i.e., $\bar{y}_{\{i:s_i,1998,m=ch\},1997,m}$), and public school in 1998 (i.e., $\bar{y}_{\{i:s_i,1998,m=tps\},1997,m}$). I then examine how the ratio of the mean 1997 test score for students attending the charter school in 1998 over that mean for students attending the public school in 1998, $\bar{y}_{\{i:s_i,1998,m=ch\},1997,m}/\bar{y}_{\{i:s_i,1998,m=tps\},1997,m}$, differs by whether the charter school in that market has an above-median level of effort and capital inputs, relative to its public school competitor. The first column of Table 6 shows that markets in which charter schools have above-median relative effort exhibit much stronger sorting of high-prior-score students into charter schools; the second column shows that this is also true for capital, though to a lesser extent. That is, Table 6 shows that variation in the relative inputs of charter and public schools can help explain patterns of heterogeneous sorting; relative input levels explain more than 50% of the variation in sorting into charters based on prior scores.

In summary, on average, charter schools have input levels that are no higher than those at public schools yet they also have higher test scores. This suggests that charter and public schools have different production functions. At the same time, there seems to be substantial variation in between-market distributions of prior scores and, hence, potentially student abilities and in the pattern of (potentially, ability-based) sorting between charter and public schools between markets.

Table 4: Sample Means of Test Scores and Market Share

Duopoly Markets	Test Scores \bar{y}_s	Market Share $\overline{\mu_s/\mu}$
Charter	3.137	0.092
Public	3.075	0.908
Monopoly Markets		
Public	3.028	1.000

¹⁹This may not be true if households had very high effort costs. However, as is shown in Section 5, effort “costs” are actually estimated to be negative.

Table 5: Decomposition of 1997 Test Score

Intercepts	Estimate	Std. Error
Mean (λ_0)	3.300	0.104
Std. Dev (σ_{λ_m})	0.401	0.077
Random Coefficients		
Mean (γ_0)	-0.130	0.238
Std. Dev (σ_{γ_m})	0.560	0.239

Notes: 4530 observations

See eq. (9) for parameter definitions.

Table 6: Sorting on 1997 Test Scores by Ratio of Input Levels

	Effort (e_{ch}/e_{tps})	Capital (k_{ch}/k_{tps})
Above-median ratio	1.12	0.98
Below-median ratio	0.71	0.87

4 ESTIMATION

The model is estimated using maximum likelihood. The game is played in every market $m \in 1, \dots, M$ and period $t \in 1, \dots, T$. In every game there is a new public school and potential charter school entrant, each endowed with per-pupil capital levels and locations, and a new measure of students endowed with abilities and locations. Observed outcomes for each market and period include charter school entry/operating decisions, school effort levels, student school choices, and student test scores. A market is linked across periods through its time-invariant ability distribution and whether a charter entered in the previous period in the market, which determines the entry/operating cost shock distribution.

It is necessary to recover unobserved market-level ability distributions to solve and estimate the model. Ability is assumed to be normally distributed within markets, meaning it suffices to recover its mean and variance.²⁰ If students can choose between schools in a market, sorting on unobserved ability may complicate recovery of market ability distributions. By using test score distributions from 1997, the year before charter school authorization, there is no such selection problem because all students attended public schools in their respective markets that period. The recovered ability distributions can then be treated as observed when integrating over student ability in school maximization problems and student likelihood statements. Because they are functions of the public school test score production function parameters, market ability distributions must be recovered jointly with the estimation of the model. Effectively, student achievement depends on prior test scores via the recovered ability distributions.

Using the production function for public schools (1), the mean test score for market m

²⁰Market ability distributions are non-parametrically identified given the public school test score production function and test score productivity shock distribution. Details are available upon request.

in 1997 is

$$\begin{aligned}
\bar{y}_{tps,97,m} &= \int a E y_{tps}(1, \bar{a}_m, e_{tps,97,m}^o, k_{tps,97,m}) f_m(a) da + \underbrace{\int f_{\nu^y}(\nu_{tps}^y) d\nu_{tps}^y}_0 \\
(10) \quad &= E y_{tps}(1, \bar{a}_m, e_{tps,97,m}^o, k_{tps,97,m}) \int a f_m(a) da = \bar{a}_m^{1+\theta_{tps}} E y_{tps}(1, 1, e_{tps,97,m}^o, k_{tps,97,m})
\end{aligned}$$

where $(e_{tps,97,m}^o, k_{tps,97,m})$ are observed inputs, f_m is the density of ability in market m , and \bar{a}_m is mean ability. The normalization of public school TFP $\omega_{tps} = 1$ affects the mean of recovered ability distributions; the charter school's TFP ω_{ch} can be thought of as productivity relative to that of a public school. The second moment can be used to recover the variance of ability for each market using the recovered \bar{a}_m and analogous reasoning.²¹

Let Φ and ϕ denote the standard normal CDF and density, respectively. The potential charter school entrant enters with probability

$$(11) \quad \Pr\{z_{tm}^o = entry | z_{t-1,m}^o\} = \Phi\left(\frac{E_{\nu^e} [v_{ch}(\mathbf{e}_{tm}^*)] - \mu_{\underline{v}, z_{t-1,m}^o}}{\sigma_{\underline{v}, z_{t-1,m}^o}}\right),$$

where \mathbf{e}_{tm}^* is equilibrium effort and $\mu_{\underline{v}, z_{t-1,m}^o}$ and $\sigma_{\underline{v}, z_{t-1,m}^o}$ depend on last period's charter school entry decision $z_{t-1,m}^o$.²² The likelihood of observed effort at school s is the density of the difference between effort predicted by the model e_{stm}^* , which depends on the entry decision, and observed effort:

$$(12) \quad L\{e_{stm}^o | z_{tm}^o\} = \frac{1}{\sigma_{\nu^e}} \phi\left(\frac{\ln e_{stm}^o - \ln e_{stm}^*}{\sigma_{\nu^e}}\right).$$

The probability student i attends the charter school ($s_{itm}^o = ch$) is a function of ability

²¹The variance of the ability distribution cannot be separated from that of the test score shock, so σ_{ν^y} is set to 0.40. For market-level ability distributions to be feasible, σ_{ν^y} was chosen in order to not create any markets with a negative variance of student ability or degenerate ability distribution (the smallest market-level standard deviation of prior test score is 0.405). The chosen value of the standard deviation of the test score shock has the added benefit of targeting a key value in the literature – the residual in Cunha et al. (2010) accounts for about 34% of the variation in outcomes, while $\sigma_{\nu^y} = 0.40$ implies that the “analogous” residual accounts for about 22% of the variation of outcomes here. A much larger value would render the ability distribution in at least one market to be either degenerate or infeasible. Note that both the shock variance as well as the assumption of equal variances for shocks at public and charter schools may play less of a role than they typically would. This is because households maximize ex-ante expected utility and the test score shock enters in an additively separable manner, meaning the variance of the test score shock does not affect their optimal decisions. That is, though the level of σ_{ν^y} affects recovered ability distributions it does not affect household decisions given ability.

²²Recall again that “entry” means either new entry or continued operation. Note that $\mu_{\underline{v}, z_{t-1,m}^o}$ should not be confused with μ and μ_s in the model, which denote measures of students. There was no entry before the first period of the model.

a_i , mean ability at both schools $\bar{\mathbf{a}}_{itm}$, observed school inputs $(\mathbf{e}_{itm}^o, \mathbf{k}_{itm})$, the pair of distances from the student to both schools \mathbf{r}_{itm} , and own ability a_i :

$$(13) \quad \Pr\{s_{itm}^o = ch | z_{itm}^o = entry, a_i\} = \Phi\left(\frac{a_i \Delta g(\bar{\mathbf{a}}_{itm}, \mathbf{e}_{itm}^o, \mathbf{k}_{itm}) - \Delta c_{itm}}{\sigma_{\Delta\epsilon}}\right),$$

where eq. (20) defines Δg and Δc_{itm} . Because mean schools abilities are unobserved, it is necessary to compute the fixed point of mean abilities at charter and public schools $\bar{\mathbf{a}}_{itm}$, given observed effort and capital (details in Appendix A.1). The likelihood of observed test score y_{istm}^o is

$$(14) \quad L\{y_{istm}^o | z_{itm}^o, s_{itm}^o = s, a_i\} = \frac{1}{\sigma_{\nu^y}} \phi\left(\frac{y_{istm}^o - E y_s(a_i, \bar{\mathbf{a}}_{istm}, \mathbf{e}_{istm}^o, \mathbf{k}_{istm})}{\sigma_{\nu^y}}\right).$$

Entry cost shocks \underline{v}_{itm} , effort productivity shocks ν_{istm}^e , test score productivity shocks ν_{istm}^y , and preference shocks η_{istm} are assumed to be independent.

The total likelihood is in Appendix A.2. It combines the above school- and student-level likelihood statements, integrating the latter according to recovered market-specific ability distributions. Asymptotic standard errors are computed using the sum of the outer product of the observation-level scores. Given a market's ability distribution and prior entry status, student and school observations are independent over time. For asymptotic analysis, let the number of markets M go to infinity, holding constant the average number of households within each market and number of time periods T .²³

4.1 Identification of Test Score Production Functions Though the model is estimated using maximum likelihood, it is useful to discuss how variation in distances helps to identify the test score production functions in the presence of potential sorting on both unobserved student ability and the unobserved mean ability of students' peers. This section provides an overview of the argument; Appendix B.4 develops the below argument in more detail.

Between-market variation in distance distributions and within-market variation in distances between students and schools both contribute to identification of the test score production function. Consider a market in which the charter school is located very far from students, relative to the public school. The public school in this market would effectively

²³Data are assumed to be missing randomly. There are some charter school observations where homework was not reported for any students, so those schools did not contribute to the effort likelihood and affected students did not contribute to the likelihood. About 5% of observations in charter and public schools are missing test score data, so these students also do not contribute to the test score likelihood. I integrate the likelihood over market distance distributions for students missing location data. The assumption that these addresses are missing at random is justified by the fact that an indicator for whether the address is missing is not significantly associated with a student's test score when controlling for student ethnicity.

serve all students, implying that variation in mean achievement in this market over time would be due to changes in capital and effort, not sorting on student ability.²⁴ This identifies the portion of the public school’s production function that does not depend on ability, $Ey_{tps}(1, 1, e_{tps}, k_{tps})$. Analogous reasoning, in markets where the charter school was closer to students, would identify $Ey_{ch}(1, 1, e_{ch}, k_{ch})$.

The model also allows unobserved ability to play a role in test score production, through the parameter θ_s . Manski (1993) taxonomizes peer effects as *endogenous* (determined by the mean achievement of attendant students), *correlated* (determined through a common shock to students in the same school), and *contextual* (operating directly through the mean ability of attendant students), and argues that a “reflection problem” makes it difficult to separate these channels when using commonly invoked linear-in-means models of peer effects. I follow other structural work studying competition between schools (such as Epple and Romano, 1998; Nechyba, 2000; Caucutt, 2002; Ferreyra, 2007) by assuming that neither endogenous social interactions nor correlated shocks affect test scores.²⁵ Conditioning on the estimated $Ey_s(1, 1, e_s, k_s)$, variation in achievement in the above “no ability sorting” markets, combined with the fact that peer quality is a known (up to θ_s) function of student ability in such markets, pins down θ_s . Though I allow for a relationship between distributions of student ability and distance *between* markets, ability and distance are assumed to be independently distributed *within* markets, as in Card (1993). This means that, though not strictly necessary, within-market variation in net distance from the charter school also helps estimate the importance of peer effects.

The assumption that ability and distance distributions are independent within markets could be violated if charter schools were able to target students within markets but is justified for several reasons. First, the model already allows charter schools to target students to a large degree, in the sense that markets have different ability and distance distributions and charter schools make entry decisions. Charter schools likely have less control over their specific locations within the relatively small areas inside markets because they typically locate in densely populated areas that have already been developed, so their optimal choice within a market (from a purely spatial perspective) is unlikely to be available. Because they do not receive capital funding streams, charter schools often use existing buildings to avoid expensive capital expenditures incurred when building a new school, meaning that even if their (spatially) optimal location within a market were available, it would unlikely trump a low-rent option close-by within the market. Cullen et al. (2005) check the plausibility of this

²⁴Recall that market-level ability distributions are assumed to be constant.

²⁵Even if there were a reflection problem, the production technologies are non-linear and, therefore, in principle both endogenous and correlated effects could also be identified.

assumption in a similar environment by checking whether observed student characteristics are related to distance, and do not find evidence of a linear relationship. I perform a similar check and also find little evidence that ability and distance are related within markets.²⁶

5 ESTIMATION RESULTS

5.1 Parameters Table 7 presents the estimated parameters for the model. The first nine rows are the test score production function parameters.²⁷ Both schools have effort shares in test score production that are higher than capital shares. Capital is not estimated to be productive at public schools—that is, the effort share at public schools is not statistically distinguishable from one ($\alpha_{tps} = 0.966$, with a standard error of 0.116). This is consistent with Hanushek (2003). Charter schools do not seem to be significantly different from being homogeneous of degree one ($\tau_{ch} = 0.953$, with a standard error of 0.052). In contrast, we can reject that public schools are homogeneous of degree one ($\tau_{tps} = 0.060$, with a standard error of 0.009). The elasticity of substitution between capital and effort is 0.60 for charter schools and 0.77 for public schools; inputs are more substitutable at charter schools than at public schools but neither school’s technology is close to Cobb-Douglas, which has an elasticity of substitution of one. Contextual peer effects at charter schools are much less important than those at public schools ($\theta_{ch} = 0.026 < \theta_{tps} = 0.995$). The finding that peer quality is an important determinant of achievement at public schools is consistent with Ferreyra (2007) in the realistic scenario in which student ability is increasing in household income; comparable estimates do not exist for charter schools. Although variance in mean ability across charter schools does little to affect student test scores, charter schools have an estimated TFP that is much higher than that at public schools ($\omega_{ch} = 1.74 > \omega_{tps} = 1$). Taken together, production functions substantially differ between charter and public schools.²⁸

Household preference parameters are denominated in standard deviations of test scores. Students prefer to attend public schools, *ceteris paribus* ($c_{ch} = 15.28$), and the interaction between ability and disutility from attending charter schools is indistinguishable from zero ($c_{ch,a} = -0.36$, with a standard error of 0.21). The disutility from attending charter schools may, in part, capture enrollment caps which, in reality, limit charter school sizes.²⁹ As

²⁶A linear regression of net distance of a student from both charter and public schools on and market and time indicator variables had an R^2 of 0.7715 when the student test score was excluded and 0.7716 when the student test score was included.

²⁷Recall that all public schools share the same test score production technology and all charter schools share (a different) test score production technology.

²⁸A likelihood ratio test rejects a restricted model in which charter and public schools share a common technology (except for TFP ω_s) with a p-value smaller than 10^{-10} .

²⁹Though this parameter may seem large, it should not create significant cause for concern about how any potential bias of estimates of household preference parameters would impact the estimated distribution of

Table 7: Parameter Estimates

Parameter	Estimate	Standard error	Description
Test score production functions			
ω_{ch}	1.744	0.114	TFP for charter school relative to public
α_{ch}	0.627	0.049	effort share, charter school
β_{ch}	-0.675	0.114	substitution parameter, charter school
α_{tps}	0.966	0.116	effort share, public school
β_{tps}	-0.295	0.991	substitution parameter, public school
τ_{ch}	0.953	0.052	return to scale, charter school
τ_{tps}	0.060	0.009	return to scale, public school
θ_{ch}	0.026	0.029	peer effects, charter school
θ_{tps}	0.995	0.012	peer effects, public school
Student cost			
c_{ch}	15.280	3.662	student charter school cost
$c_{ch,a}$	-0.361	0.211	student charter school cost interact ability
c_e	-3.869	0.968	student effort cost
c_r	0.726	0.179	student distance cost
School preference shock			
σ_η	7.985	1.971	st. dev. student school preference shock
School valuation of test scores			
$\delta_{y,ch}$	20.960	3.748	value of average test score, charter school
$\delta_{y,tps}$	19.634	3.483	value of average test score, public school
School effort cost functions			
$\psi_{e,ch,1}$	0.000	7.432	disutility of effort, charter school
$\psi_{e,tps,1}$	1.930	0.398	disutility of effort, public school
$\psi_{e,ch,2}$	14.137	3.559	effort, school size interaction, charter school
$\psi_{e,tps,2}$	-0.071	0.021	effort, school size interaction, public school
$\psi_{e,ch,3}$	-15.508	5.003	effort, capital interaction, charter school
$\psi_{e,tps,3}$	-0.917	0.265	effort, capital interaction, public school
$\psi_{e,ch,4}$	7.423	1.980	effort, mean ability interaction, charter school
$\psi_{e,tps,4}$	0.000	0.109	effort, mean ability interaction, public school
Entry cost			
$\mu_{\underline{v},no\ entry}$	96.444	29.416	mean entry cost distribution, no entry last period
$\sigma_{\underline{v},no\ entry}$	20.169	10.529	st. dev. entry cost distribution, no entry last period
$\mu_{\underline{v},entry}$	34.621	18.483	mean entry cost distribution, entry last period
$\sigma_{\underline{v},entry}$	14.053	17.145	st. dev. entry cost distribution, entry last period
$\psi_{mkt.size}$	8.580	6.491	weight for market size
$\psi_{fr.Black}$	23.479	14.077	weight for fraction of Black students in market
Effort productivity shocks			
σ_{ν_e}	0.185	0.003	st. dev. effort productivity shock

mentioned previously, data on charter school capacities and applications, which are required to separate this cost from capacity constraints, were not available. The disutility of effort is negative ($c_e = -3.87$), which means households prefer attending the school where they have to work more, even after taking into account increased test scores. The per-kilometer distance cost is about three-quarters of a standard deviation ($c_r = 0.73$) of test scores.

Schools are estimated to face positive costs of exerting effort. It would be hard to rationalize interior effort level choices otherwise, given that the estimates point towards positive marginal products of effort at both public and charter schools, and interior effort data cannot be explained by household effort costs, which are negative (i.e., households would like more effort, *ceteris paribus*). Charter schools have much larger size-effort cost interactions than public schools ($\psi_{e,ch,2} = 14.14 > \psi_{e,tps,2} = -0.07$). The cost of exerting effort is mitigated by capital at both school types ($\psi_{e,ch,3} = -15.51, \psi_{e,tps,3} = -0.92$). Finally, charter schools suffer a larger effort cost for educating higher ability students ($\psi_{e,ch,4} = 7.42$). Higher per-pupil capital levels may make it easier for the school to create, assign, and grade homework because there are more computers per student or if there are smaller class sizes. On the other hand, designing curricula for high ability students may be more demanding.

Finally, the mean of the entry cost shock distribution is lower when there was a charter school in the market in the previous period ($\mu_{v,entry} = 34.62 < \mu_{v,no\ entry} = 94.44$). Charter schools are also more likely to enter larger markets ($\gamma_{mkt.size} = 8.58$) and in markets the higher the share of the market is Black students ($\gamma_{fr.Black} = 23.48$), though these parameters are somewhat imprecisely estimated.

5.2 Model Fit The model captures charter school entry, charter and public school effort, student choice, and student test score patterns for North Carolina. The model does a good job of predicting the fraction of markets with charter schools; on average, 6% of markets have charters, the same as in the data. Table 8 shows the fraction of markets with charter schools, by market characteristics. It shows that the model also captures the facts that charter schools are more likely to enter larger markets (first and second columns), markets where they would receive higher per-pupil capital (third and fourth columns), and markets with lower mean 1997 test scores (fifth and sixth columns).

treatment effects because the primary manner in which the model shifts demand for charter schools is through the mean effect on random utility, c_{ch} , not the term measuring the interaction between student ability and the non-pecuniary cost of attending charter schools, $c_{ch,a}$. This means that estimates of treatment effects, and the parts of school objectives pertaining to student achievement, would be relatively unaffected by charter school size caps. The concern about biased household preference parameters would be more salient if the estimated interaction between student ability and the cost of attending a charter school were large, in which case the estimated distribution of treatment effects might be biased. Appendix B.8 examines sensitivity of treatment effects to the non-pecuniary cost c_{ch} .

Table 8: Fit: Fraction of Markets with Charter Schools by Market Characteristics

	Market Size μ		Charter Capital k_{ch}		Mean 1997 Test Score \bar{y}_{1997}	
	Observed	Predicted	Observed	Predicted	Observed	Predicted
Above-median	0.067	0.080	0.095	0.080	0.050	0.056
Below-median	0.045	0.042	0.017	0.041	0.062	0.065

The model can fit school effort patterns. The top row of Table 9 shows that the model captures the fact that charter and public schools in duopoly markets exert more effort than monopolist public schools, though it slightly under-predicts charter school effort. Higher effort levels in duopoly markets could be driven by both charter schools’ tendency to enter markets with higher per-pupil capital levels and their effect on public school effort. The rest of Table 9 shows that the model also captures the relationships between effort, market size, and capital. As opposed to both duopoly and monopoly public schools, charter schools have much lower predicted effort levels in larger markets (second versus third row). Charter schools in markets with above-median capital (fourth row) exert more effort than those in markets with below-median capital (fifth row). The same is true for public schools in both duopoly and monopoly markets. Two channels underlie these patterns. First, capital directly augments test score production, increasing effort choices through capital-effort complementarity—especially at charter schools. Second, capital reduces effort costs at both school types via the negative interaction between capital and effort in school effort cost functions. The latter channel helps explain why public schools exert higher effort in markets where they have higher per-pupil capital, despite the small direct (i.e., via the production technology) effect of capital on the achievement of public school students.

The first two columns of Table 10 show that the model captures the ranking average of test scores in the estimation sample: students at charter schools have the highest average test scores, followed by students attending public schools in markets charter schools have entered, followed by students in public schools in markets without charters. The third and fourth columns of Table 10 show model fit for the fraction of students choosing charter schools, and show that the model also fits the pattern that charter schools are smaller than competing public schools.

Despite the difference in mean test scores, model simulations indicate that the mean ability of students at charter schools and public schools competing with charters is around the same—the average of the ratio of mean abilities of public over charter school students across markets is 0.996. However, there is also a wide range of sorting behaviors between markets, with the lowest ratio of mean abilities implying that public school students have only 60% of the peer quality that charter school students have and the highest value this

Table 9: Fit: Mean Effort (Hours of Homework) by School Type and Market Characteristics

	Duopoly Markets				Monopoly Markets	
	Charter		Public		Public	
	Observed	Predicted	Observed	Predicted	Observed	Predicted
Total	2.66	2.51	2.69	2.65	2.43	2.44
Above-median μ	2.51	2.37	2.85	2.76	2.49	2.52
Below-median μ	2.89	2.77	2.46	2.44	2.37	2.37
Above-median k_s	2.74	2.61	2.78	2.73	2.51	2.51
Below-median k_s	2.20	2.31	2.21	2.50	2.35	2.38

Table 10: Fit: Mean Test Scores and Market Share

	Test Scores \bar{y}_s		Market Share $\frac{\mu_s}{\mu}$	
	Observed	Predicted	Observed	Predicted
Duopoly Markets				
Charter	3.137	3.122	0.092	0.082
Public	3.075	3.062	0.908	0.918
Monopoly Markets				
Public	3.028	3.043	1.000	1.000

ratio obtains is 1.48, which corresponds to public school students having 48% higher peer quality than charter school students.

Variation in both effort and capital of charter and public schools explains over 40% of the variation in ability sorting patterns.³⁰ The model can account for such heterogeneous sorting via several channels. First, variation in relative school capital directly causes variation in ability sorting through the productive effect of capital at charter schools, via their technology. This direct effect is augmented by two indirect effects of capital that increase schools' chosen effort levels: via the negative interaction between school capital and effort in school effort cost functions and the effect of these higher effort levels in school technologies. Finally, subsequent productivity shocks to chosen effort create further between-market variation in relative effort levels. However, without some countervailing force, ability sorting would be too strong; the effects of variation in relative inputs is blunted by the non-pecuniary cost of attending charter schools and the fact that charter schools are typically located further from students than traditional public schools. Note that there is also between-market variation in relative distance of students from charter schools, which allows for sorting to be stronger or weaker between markets, *ceteris paribus*.

³⁰Details available upon request.

6 COUNTERFACTUAL SIMULATIONS

Charter school policy depends critically on how they affect the achievement of attendant students, students at competing public schools, and students in markets where charter schools would enter if caps on the total number of charters were lifted. To quantify these effects, I simulate the model 30 times, where I first simulate charter school entry decisions and school effort choices under duopoly and monopoly scenarios, and then simulate household school choices and test scores for each household in the market’s charter school, public school competitor, and monopolist public school.

6.1 Definitions of Treatment Effects The effect of being in a charter school (“direct effect”) for student i who lives in market m is the difference between the test score the student would have received at the charter school and that she would have received at the monopolist public school in her market:

$$\begin{aligned} \Delta_{im}^{direct} &= y_{i,ch,m} - y_{i,tps,m}^{mono} \\ (15) \quad &= Ey_{ch}(a_i, \bar{a}_{ch,m}, e_{ch,m}, k_{ch,m}) + \nu_{i,ch}^y - Ey_{tps}(a_i, \bar{a}_m, e_{tps,m}^{mono}, k_{tps,m}) - \nu_{i,tps,mono}^y, \end{aligned}$$

where $\bar{a}_{ch,m}$ is the mean ability of students attending the charter school, $e_{ch,m}$ is the effort level at the charter school, \bar{a}_m is the average ability at the monopolist public school (equal to mean ability in the market), $e_{tps,m}^{mono}$ is the effort level at the monopoly public school, and ν_{is}^y are ex-post test score productivity shocks.³¹ The effect of attending the public school competitor (“spillover effect”) for student i is the difference between the test score she would have received at the public school, when competing with the charter school, and that she would have received at the monopolist public school in that market: $\Delta_{im}^{spill} = y_{i,tps,m} - y_{i,tps,m}^{mono}$, where, analogous to $y_{i,ch,m}$, the term $y_{i,tps,m}$ includes the relevant public school inputs.

The model is a useful tool for program evaluation because it provides potential outcomes that can be used to generate both direct and spillover effects for all students, regardless of charter school presence or student school choice. Researchers typically focus on mean treatment effects, which are expected values of Δ^{direct} and Δ^{spill} for different choice-based sets of students. Consider the treatment of attending a charter school. The mean direct effect of treatment on the treated (direct TOT) is the mean effect of attending a charter

³¹The effort levels in all simulations are those chosen by schools and then hit with effort productivity shocks. In the model section there were no market or time subscripts because the analysis was done within one market and one time period. I drop the time subscript here to simplify exposition, but add a market subscript to make clear comparisons across markets.

school among students who would choose it. In market m , the mean direct TOT is

$$(16) \quad \bar{\Delta}_m^{direct,TOT} = \text{E} [\Delta_{im}^{direct} | s_{im} = ch] = \int \Delta_{im}^{direct} f_m(a_{im}; \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k} | s_{im} = ch) da_{im},$$

where $f_m(a_{im}; \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k} | s_{im} = ch)$ is the density of ability for students choosing the charter school in market m (see eq. (25)).³² The mean direct effect of treatment on the untreated (TOU) in m is the mean effect of attending a charter school among students who would choose the public school:

$$(17) \quad \bar{\Delta}_m^{direct,TOU} = \text{E} [\Delta_{im}^{direct} | s_{im} = tps] = \int \Delta_{im}^{direct} f_m(a_{im}; \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k} | s_{im} = tps) da_{im}.$$

The mean direct average treatment effect (ATE) in m averages over all students in the market $\bar{\Delta}_m^{direct,ATE} = \text{E} [\Delta_{im}^{direct}] = \int \Delta_{im}^{direct} f_m(a_{im}) da_{im}$, where f_m is the density of ability in market m . Market mean spillover effects are calculated analogously, substituting Δ_{im}^{spill} for Δ_{im}^{direct} and tps for ch for student school choices.

Market-level treatment effects are weighted by market size and entry status to aggregate treatment effects across markets; e.g., the mean direct TOT across all duopoly markets is

$$(18) \quad \bar{\Delta}^{direct,TOT,entry} = \sum_m 1\{z_m = entry\} \mu_{ch,m} \bar{\Delta}_m^{direct,TOT} / \sum_m 1\{z_m = entry\} \mu_{ch,m},$$

where $\mu_{ch,m}$ is the measure of students choosing the charter school in m . Other aggregated treatment effects are calculated analogously.

Researchers who exploit lotteries among applicants to an over-subscribed charter school compare test scores of applicants who were randomized into the charter school with those of applicants randomized into the competing public school. Denote the treatment effect estimated for student i in such a study as

$$(19) \quad \check{\Delta}_{im}^{direct} = y_{i,ch,m} - y_{i,tps,m} = \Delta_{im}^{direct} - \Delta_{im}^{spill},$$

i.e., $\check{\Delta}_{im}^{direct}$ is the difference between the direct and spillover effects for i . Intuitively, bigger changes in public school inputs caused by the charter school's presence increase biases in estimated treatment effects in lottery studies. For example, suppose a public school drastically changed its behavior in response to charter school entry, such that $\Delta_{im}^{spill} > \Delta_{im}^{direct} > 0$. In this case, a researcher using a lottery design would *incorrectly sign* the direct effect. Theory provides no a priori sign on the spillover effect, which means even the sign of this bias cannot

³²All simulations use households' distances from schools; I only suppress the dependence of f on household distance from schools to simplify exposition.

be determined without further structure.

6.2 Effects of Charter Schools on the Distribution of Test Scores Table 11 summarizes mean direct and spillover effects of charter schools on test scores for different subsets of students. It also reports the mean bias on the direct effect induced by ignoring spillover effects. All results are reported in percentages of a standard deviation (sd) of the average of math and reading test scores. The top panel of the table (“Duopoly markets”) reports results for markets in which charter schools are present and the bottom panel (“Monopoly markets”) reports what results would be in markets in which charter schools are not present. The row within each panel indicates which subset of households is being considered: households who would choose charters, households who would choose competing public schools, or all households in such markets (ATE). For example, the number associated with the first column ($\overline{\Delta}^{direct}$) and the row “Attend charter” in the top panel of the table is the mean direct effect of treatment on the treated in duopoly markets, $\overline{\Delta}^{direct, TOT, entry}$, i.e., the mean direct effect for students who chose the charter school in markets where charter schools are present.

Table 11: Mean Direct and Spillover Treatment Effects by School Choice in Duopoly and Monopoly Markets

Duopoly Markets	$\overline{\Delta}^{direct}$	$\check{\Delta}^{direct}$	$(\overline{\Delta}^{direct} - \check{\Delta}^{direct})$	$\overline{\Delta}^{spill}$
TOT (Attend charter)	0.108	0.104	0.005	0.005
TOU (Attend public)	-0.039	-0.054	0.016	0.016
ATE	-0.027	-0.042	0.015	0.015
Monopoly Markets: TOU				
Attend charter	0.018	0.0186	-0.001	-0.001
Attend public	-0.114	-0.124	0.010	0.010
ATE	-0.103	-0.111	0.009	0.009

Charter schools have positive effects on the test scores of attendant students and negligible spillover effects on students attending public schools. The first column shows that the mean direct effect of charter schools is highest for students who choose to attend charter schools in duopoly markets, about 11% sd. Most of the direct effect comes from charter school technology. The mean test score for students in duopoly markets who attend charters is 3.128; their mean score would be reduced to 3.026 when using charter school inputs and the public school technology, but would be reduced by a much smaller amount, to 3.013, when also using monopoly public school inputs.

The third column shows the mean bias induced in the estimation of direct effects by ignoring spillover effects (column 2) and the fourth column shows mean spillover effects; note that the third and fourth columns must be equal. The average bias introduced by ignoring spillover effects is small, on the order of 2% sd. This finding is consistent with other work studying the spillover effects of charter schools, such as Bifulco and Ladd (2006)—who study spillovers using similar data from North Carolina—and Sass (2006). Average treatment effects are much closer to the estimated effects for students attending public schools because the share of students attending charter schools is small (about 9% of duopoly markets), resulting in a negative direct ATE.

Mean treatment effects would be smaller in monopoly markets. The mean direct TOT of charter schools would be about 2% sd in monopoly markets and mean spillover effects are also slightly smaller. These differences are driven mostly due to lower levels of public and charter school inputs chosen in those markets. The average difference in effective inputs between charter schools and monopolist public schools, i.e., $\overline{Ey_{ch}(1, \bar{a}_{ch}, e_{ch}, k_{ch}) - Ey_{tps}(1, \bar{a}_m, e_{tps}^{mono}, k_{tps})}$, is 0.070 in duopoly markets, while the same difference is 0.016 in monopoly markets; this corresponds to over a four-fold difference.

Although estimated treatment effects are a combination of state-level charter school authorization laws and other institutional characteristics, market- and school-level characteristics, student characteristics, and identification strategy, one might want to compare the results presented here with those from lottery studies. Using a lottery design on an oversubscribed Massachusetts KIPP school, Angrist et al. (2012) find that applicants randomized into charters have 36% sd higher math test scores and 12% sd higher reading test scores than applicants randomized into competing public schools; their finding of a 24% sd increase in the average of reading and math test scores for $\overline{\Delta}^{direct}$ is much larger than this paper’s finding of 11% sd. Section 6.2.1 discusses this comparison in detail.

Having observed the positive mean direct effects and negligible mean spillover effects of charter schools, policymakers may be tempted to increase the number of students in charter schools. Indeed, holding their entry and effort at their baseline levels, making charter schools monopolists would increase test scores by almost 30% sd—almost three times the mean direct effect. Therefore, I next explore how making charter schools monopolists, by assigning all students in a market and the public school’s capital to them after they have entered, would affect student achievement. On average, charter schools would reduce effort by more than an hour due to the large effort costs incurred when serving so many more students. Consequently, monopolist charters would have an average direct TOT of 2% sd, less than a fifth of the mean direct TOT in the duopoly baseline. Even restricting attention to markets where charter schools have positive ATEs in the baseline would result in much

smaller, or even negative, monopolist charter school ATEs. Charter schools would also enter fewer markets when forced to serve the entire market (5%, compared with 6% under the baseline) and would be much less likely to enter larger markets than they would under the duopoly scenario (2%, compared with 8% under the baseline). Altogether, this exercise shows that understanding charter school entry and input choices is crucial for policy.

6.2.1 Heterogeneity of Treatment Effects by Market Figure 1 shows how treatment effects vary across markets by plotting mean market-level direct and spillover TOT, where circle size represents market size. Mean market-level direct TOTs are quite heterogeneous; this heterogeneity relates to market size. The 75th percentile mean market-level direct TOT is 10% sd while the 25th percentile is *negative* at -21% sd. The finding of substantial between-market heterogeneity in the direct TOT is consistent with Gleason et al. (2010), who find substantial between-school variation in the mean direct effects of charter schools. Moreover, as in this paper, Gleason et al. (2010) find many charter schools with negative average direct effects. The school-size interaction in the charter school effort cost function plays a key role here: mean direct effects are much more likely to be positive in smaller markets (represented by smaller circles). There is much less variation in mean spillover TOTs, which is due to the relative ineffectiveness of public school inputs.

Figure 1: Market-Level Mean Direct TOT vs. Spillover TOT





Figure 2: Market-Level Mean Direct TOT by Fraction of Students Choosing Charter School



By design, lottery studies estimating a direct TOT using oversubscribed charter schools likely consider charter schools that are highly demanded by households. If households value

student achievement, such studies may only recover part of the distribution of charter school treatment effects. Figure 2 plots the mean direct TOT for each market against the fraction of students in that market who would choose to attend the charter school. In markets with below-median demand (i.e., those with a below-median fraction of students who would choose the charter school, were it to enter), the mean direct TOT is -19% sd in test scores, compared with slightly more than 11% sd for markets with above-median demand. When researchers study only over-subscribed, that is, highly demanded, charter schools they draw from schools on the right side of Figure 2, which shows the difficulty in generalizing findings from such studies.

6.3 Effect of Allowing Unlimited Charter School Entry One of the most contentious charter school policy debates is whether caps on the number of charter schools should be lifted. Because the model is estimated on data from years before the statewide cap came close to binding, it can be solved for the years 2002-2005 to quantify the effect of lifting charter school entry caps on the distribution of test scores.³³


The results suggest the cap on the total number of charter schools in North Carolina was binding. During 2002-2005, charter schools would operate in 15% of markets if caps were lifted, up from 6% for the period 1998-2001. Figure 3 shows how entry and other key variables would change. The overall effect of charter schools on attendant students is attenuated as charters enter larger markets with lower levels of capital (top two panels) and subsequently reduce effort provision (the third panel), which eventually creates a slight *negative* direct TOT (-0.4% sd). Spillover TOT impacts of charter schools in new entry markets are similar to those estimated for the first four years of charter school authorization (on average 1% sd). A much larger share of students would be in markets where charter schools operate (21% up from 6%). Overall, charter schools affect more students when caps are lifted but treatment effects are attenuated.

7 CONCLUSIONS

There is a long-standing and contentious debate about the effects of charter schools and, more generally, school choice, on student achievement because there are many moving parts that complicate analysis. The equilibrium model developed and estimated in this paper captures

³³Recall that this paper studies middle school grades. Although the 100-charter-school cap was nearly binding in 2001, precluding new entry in 2002 for charter schools serving any grade levels, there are only 43 schools in the estimation sample operating in 2001. Also note that, in the model, charter schools make entry decisions in each year. In this sense, the free-entry counterfactual does not restrict additional entry to markets in which there are not charter schools. Rather, it lets the model run for additional years, during which time the total number of entrants may exceed the cap.

Figure 3: Model Outcomes, 1998-2005



key mechanisms pertinent to the debate about charter school policy, such as where charter schools choose to locate and how many would enter were entry caps lifted (by modeling which markets they choose to enter), how charter school quality is determined (by endogenizing charter school inputs), the potential for spillover effects on competing public schools (by endogenizing public school inputs and allowing peer quality to enter test score production), and the potential for cream-skimming through student sorting (by modeling how students choose schools based on their unobserved ability). Student distances from schools provide a source of variation for student school choices, generating credible estimates of both charter and public school test score production technologies. The model fits patterns of charter school entry, public and charter school input provision, student school choices, and test scores quite well. The structure afforded by the model allows for a coherent definition and quantification of direct and spillover treatment effects in an equilibrium framework and simulation of the effects of changes to charter school policy, such as lifting caps on entry.

The approach developed here provides a more comprehensive picture of treatment effects than studies based on only a subset of charter schools, such as experimental or lottery-based designs, because there is substantial heterogeneity that such studies may not recover. Both the direct effect and spillover effects of charter schools are on average positive for treated

students, although the direct effect is much larger than the spillover effect. One striking result is the heterogeneity in the mean direct effects of charter schools on attendant students across markets. Spillover effects for students attending public schools are quite small, which suggests that lottery-based designs may have internal validity. However, the size of direct treatment effects strongly relates to demand for charter schools, casting doubt on the notion that lottery designs have external validity. Caps on the total number of charter schools are estimated to be binding and students attending charters in new entry markets would experience lower gains.

One limitation of this paper is that the creation of markets in which only one charter and one public school compete could bias estimates of treatment effects. I address this concern by first providing evidence that cross-market competition does not substantially affect charter school entry or direct or spillover effects. Moreover, I re-estimated the model, excluding markets where more than 5% of students were observed attending charter schools across market boundaries, and found that estimates of parameters and treatment effects remain largely unchanged. Finally, I use the estimated model to quantify the extent to which sample selection might affect the direct effect of charter schools on attendant students, finding evidence that the mean direct effect would be positive, though smaller, for excluded students and that the mean direct effect on all charter school students would change very little were excluded students included.

The framework developed in this paper takes an important first step towards gaining a deeper understanding of how schools compete in the public education sector and, while quite rich, could be extended in interesting ways. For example, the intertemporal dynamics in this paper are relatively simple and may warrant further investigation in future work. Additionally, incorporating data measuring other school inputs, such as teacher quality, could comprise a fruitful branch of future research.

A APPENDIX

A.1 Model Solution This section shows how to calculate school size and average test score, given an effort level pair $\mathbf{e} = (e_{ch}, e_{tps})$, which are necessary to solve for equilibrium of the

entry subgame. A student with ability a_i chooses a charter school if and only if

$$\begin{aligned}
& \mathbb{E}_{\nu^y} [y_{i,ch}] - c_{i,ch} + \eta_{i,ch} \geq \mathbb{E}_{\nu^y} [y_{i,tps}] - c_{i,tps} + \eta_{i,tps} \\
& \Leftrightarrow \\
& a_i \underbrace{(Ey_{ch}(1, \bar{a}_{ch}, e_{ch}, k_{ch}) - c_{ch,a} - Ey_{tps}(1, \bar{a}_{tps}, e_{tps}, k_{tps}))}_{\Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k})} + \underbrace{(\eta_{i,ch} - \eta_{i,tps})}_{\Delta \epsilon_i} \geq \underbrace{c_e(e_{ch} - e_{tps})}_{\Delta c_e} + \underbrace{c_r(r_{i,ch} - r_{i,tps})}_{\Delta c_{ri}} + \underbrace{c_{ch}}_{\Delta c_{ch}} \\
& \Leftrightarrow \\
(20) \quad & a_i \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) + \Delta \epsilon_i \geq \underbrace{\Delta c_e + \Delta c_{ri} + \Delta c_{ch}}_{\Delta c_i},
\end{aligned}$$

where $\Delta \epsilon_i \sim N(0, \sigma_{\Delta \epsilon}^2)$ and $\sigma_{\Delta \epsilon}^2 = 2\sigma_{\eta}^2$. Ability within the market is distributed according to $F(a_i) = N(\bar{a}, \sigma_a^2)$, making the left hand side of (20) the sum of two independent normally distributed random variables:

$$(21) \quad a_i \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) + \Delta \epsilon_i \sim N(\bar{a} \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}), \sigma_a^2 \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k})^2 + \sigma_{\Delta \epsilon}^2).$$

This provides an analytical expression for the share of students attending the charter school $\mu_{r,ch}$, given a distance difference Δc_{ri} :

$$(22) \quad \mu_{r,ch}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) = 1 - \Phi \left(\frac{\Delta c_e + \Delta c_{ri} + \Delta c_{ch} - \bar{a} \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k})}{\sqrt{\sigma_a^2 \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k})^2 + \sigma_{\Delta \epsilon}^2}} \right),$$

where Φ denotes the standard cumulative normal distribution.

There are $\rho \in 1, \dots, R$ separate distance pairs in the market, each with a measure μ_{ρ} . The total measure of students at the charter school is the sum of the shares of students of each distance, weighted by the measure of students in each distance bin:

$$(23) \quad \mu_{ch}(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) = \sum_{\rho=1}^R \mu_{\rho} \mu_{\rho,ch}(\Delta c_{\rho i}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}).$$

A student with ability a_i and relative charter distance cost Δc_{ri} will choose the charter if and only if $\Delta \epsilon_i \geq \Delta c_i - a_i \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k})$, which happens with probability $\Phi \left(\frac{a_i \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) - \Delta c_i}{\sigma_{\Delta \epsilon}} \right)$. By Bayes' Rule, the average ability of a student attending the charter school is

$$(24) \quad \bar{a}_{r,ch}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) = \int a_i f_r(a_i; \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k} | s_i = ch) da_i,$$

where the density of the ability of students at the charter school is

$$(25) \quad f_r(a_i; \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k} | s_i = ch) = \frac{\Phi\left(\frac{a_i \Delta g(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) - \Delta c_i}{\sigma_{\Delta \epsilon}}\right) f(a_i)}{\mu_{r, ch}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k})}.$$

The average ability of students attending the charter school is the weighted average of the average abilities of attendant students from each bin:

$$(26) \quad \begin{aligned} \bar{a}_{ch} &= \sum_{\rho=1}^R \mu_{\rho} \mu_{\rho, ch}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) \bar{a}_{\rho, ch}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) / \mu_{ch}(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) \\ \bar{a}_{tps} &= \sum_{\rho=1}^R \mu_{\rho} \mu_{\rho, tps}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) \bar{a}_{\rho, tps}(\Delta c_{ri}, \bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) / \mu_{tps}(\bar{\mathbf{a}}, \mathbf{e}, \mathbf{k}) \end{aligned}$$

where $\mu_{\rho, tps}$, μ_{tps} , $\bar{a}_{\rho, tps}$ and \bar{a}_{tps} are defined analogously for public schools. The pair of equations (26) define a fixed point for average ability at the charter and public schools given pairs of effort \mathbf{e} and capital \mathbf{k} .

After solving for \bar{a}_s , the average test score at school s is $\bar{y}_s = \bar{a}_s \omega_s \bar{a}_s^{\theta_s} (\alpha_s e_s^{\beta_s} + (1 - \alpha_s) k_s^{\beta_s})^{\tau_s / \beta_s} = \bar{a}_s E y_s(1, \bar{a}_s, e_s, k_s)$, which, along with μ_s , is substituted into school objectives when solving for optimal school effort.

A.2 Likelihood The likelihood function combines the previous probability and likelihood statements for markets and students, and integrates over the ability distribution in a market, given all the data X and parameters θ :

$$(27) \quad \begin{aligned} L(\theta | X) &= \left(\prod_{m \in M} \prod_{t \in 1, \dots, T} \Pr\{z_{tm}^o = entry | z_{t-1, m}^o\}^{1\{z_{tm}^o = entry\}} (1 - \Pr\{z_{tm}^o = entry | z_{t-1, m}^o\})^{1\{z_{tm}^o = no\ entry\}} \right. \\ &\quad \left. \left(\prod_{m \in M} \prod_{s \in S_{tm}} \prod_{t \in T} \left(L(e_{ch, tm}^o | z_{tm}^o = entry) L(e_{tps, tm}^o | z_{tm}^o = entry) \right)^{1\{z_{tm}^o = entry\}} \cdot L(e_{tps, tm}^o | z_{tm}^o = no\ entry)^{1\{z_{tm}^o = no\ entry\}} \right) \right. \\ &\quad \left. \left(\prod_{m \in M} \int_{a_i \in A_m} \left(\prod_{i \in I_{tm}} \prod_{t \in T} \left(\Pr\{s_{itm}^o = ch | z_{tm}^o = entry, a_i\} L(y_{i, ch, tm}^o | z_{tm}^o = entry, s_{itm}^o = ch, a_i) \right)^{1\{s_{itm}^o = ch\}} \right. \right. \right. \\ &\quad \left. \left. \left((1 - \Pr\{s_{itm}^o = ch | z_{tm}^o = entry, a_i\}) L(y_{i, tps, tm}^o | z_{tm}^o = entry, s_{itm}^o = tps, a_i) \right)^{1\{s_{itm}^o = tps\}} \right)^{1\{z_{tm}^o = entry\}} \right. \\ &\quad \left. \left. L(y_{i, tps, tm}^o | z_{tm}^o = no\ entry, s_{itm}^o = tps, a_i) \right)^{1\{z_{tm}^o = no\ entry\}} \right) dF_m(a_i). \end{aligned}$$

REFERENCES

- AGARWAL, N. AND P. SOMAINI, “Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism,” *NBER Working Paper No. 20775* (2015).
- ALTONJI, J. G., C.-I. HUANG AND C. R. TABER, “Estimating the Cream Skimming Effect of School Choice,” *Journal of Political Economy* 123 (2015), 266–324.

- ANGRIST, J. D., S. M. DYNARSKI, T. J. KANE, P. A. PATHAK AND C. R. WALTERS, “Who Benefits From KIPP?,” *Journal of Policy Analysis and Management* 31 (2012), 837–860.
- BAYER, P., F. FERREIRA AND R. McMILLAN, “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *Journal of Political Economy* 115 (2007), 588–638.
- BAYER, P., R. McMILLAN AND K. RUEBEN, “An Equilibrium Model of Sorting in an Urban Housing Market,” *NBER Working Paper No. 10865* (2004).
- BETTINGER, E., “The Effect of Charter Schools On Charter Students and Public Schools,” *Economics of Education Review* 24 (2005), 133–147.
- BIFULCO, R. AND H. LADD, “The Impacts of Charter Schools On Student Achievement: Evidence From North Carolina,” *Education Finance and Policy* 1 (2006), 50–90.
- BIFULCO, R. AND H. F. LADD, “School Choice, Racial Segregation, and Test-Score Gaps: Evidence From North Carolina’s Charter School Program,” *Journal of Policy Analysis and Management* 26 (2007), 31–56.
- CALSAMIGLIA, C., C. FU AND M. GÜELL, “Structural Estimation of a Model of School Choices: Boston Versus Its Alternatives,” *Working Paper* (2016).
- CARD, D., “Using Geographic Variation in College Proximity to Estimate the Return to Schooling,” *NBER Working Paper No. 4483* (1993).
- CARRUTHERS, C. K., “New Schools, New Students, New Teachers: Evaluating the Effectiveness of Charter Schools,” *Economics of Education Review* 31 (2012), 280–292.
- CAUCUTT, E. M., “Educational Vouchers When There Are Peer Group Effects – Size Matters,” *International Economic Review* 43 (2002), 195–222.
- CHAKRABARTI, R., “Can Increasing Private School Participation and Monetary Loss in a Voucher Program Affect Public School Performance? Evidence From Milwaukee,” *Journal of Public Economics* 92 (2008), 1371–1393.
- CULLEN, J. B., B. A. JACOB AND S. LEVITT, “The Effect of School Choice on Participants: Evidence From Randomized Lotteries,” *Econometrica* 74 (2006), 1191–1230.

- CULLEN, J. B., B. A. JACOB AND S. D. LEVITT, "The Impact of School Choice On Student Outcomes: An Analysis of the Chicago Public Schools," *Journal of Public Economics* 89 (2005), 729–760.
- CUNHA, F., J. HECKMAN AND S. SCHENNACH, "Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Econometrica* 78 (2010), 883–931.
- DE FRAJA, G., T. OLIVEIRA AND L. ZANCHI, "Must Try Harder: Evaluating the Role of Effort in Educational Attainment," *The Review of Economics and Statistics* 92 (2010), 577–597.
- DE HAAN, M., P. A. GAUTIER, H. OOSTERBEEK AND B. VAN DER KLAUW, "The Performance of School Assignment Mechanisms in Practice," *CEPR Discussion Paper No. 10656* (2015).
- DOBBIE, W. AND R. G. FRYER JR., "Getting Beneath the Veil of Effective Schools: Evidence From New York City," *American Economic Journal: Applied Economics* 5 (2013), 28–60.
- EPPLE, D. AND R. ROMANO, "Competition Between Private and Public Schools, Vouchers, and Peer-Group Effects," *American Economic Review* 88 (1998), 33–62.
- EPPLE, D., R. ROMANO AND H. SIEG, "Admission, Tuition, and Financial Aid Policies in the Market for Higher Education," *Econometrica* 74 (2006), 885–928.
- EPPLE, D. AND H. SIEG, "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy* 107 (1999), 645–681.
- FERREYRA, M. M., "Estimating the Effects of Private School Vouchers in Multidistrict Economies," *American Economic Review* 97 (2007), 789–817.
- FERREYRA, M. M. AND G. KOSEK, "Charter School Entry and School Choice: The Case of Washington, DC," *Working Paper* (2015).
- FERREYRA, M. M. AND P. J. LIANG, "Information Asymmetry and Equilibrium Monitoring in Education," *Journal of Public Economics* 96 (2012), 237–254.
- FISCHER, M. J., "The Relative Importance of Income and Race in Determining Residential Outcomes in US Urban Areas, 1970-2000," *Urban Affairs Review* 38 (2003), 669–696.
- FRUEHWIRTH, J. C., "Can Achievement Peer Effect Estimates Inform Policy? A View From Inside the Black Box," *Review of Economics and Statistics* 96 (2014), 514–523.

- FU, C., “Equilibrium Tuition, Applications, Admissions, and Enrollment in the College Market,” *Journal of Political Economy* 122 (2014), 225–281.
- GLEASON, P., M. CLARK, C. TUTTLE AND E. DWOYER, “The Evaluation of Charter School Impacts: Final Report. NCEE 2010-4029,” *National Center for Education Evaluation and Regional Assistance* (2010).
- GLOMM, G., D. HARRIS AND T.-F. LO, “Charter School Location,” *Economics of Education Review* 24 (2005), 451–457.
- HANUSHEK, E., “The Failure of Input-Based Schooling Policies,” *The Economic Journal* 113 (2003), F64–F98.
- HANUSHEK, E., J. KAIN, S. RIVKIN AND G. BRANCH, “Charter School Quality and Parental Decision Making with School Choice,” *Journal of Public Economics* 91 (2007), 823 – 848.
- HASTINGS, J., T. KANE AND D. STAIGER, “Heterogeneous Preferences and the Efficacy of Public School Choice,” *Working Paper* (2009).
- HENNINGSEN, A. AND G. HENNINGSEN, “Econometric Estimation of the ‘Constant Elasticity of Substitution’ Function in R: Package micEconCES,” (2011).
- HOXBY, C. M., “Would School Choice Change the Teaching Profession?,” *Journal of Human Resources* 37 (2002), 846–891.
- , “School Choice and School Productivity. Could School Choice Be a Tide That Lifts All Boats?,” in *The Economics of School Choice* (University of Chicago Press, 2003), 287–342.
- HOXBY, C. M. AND J. E. ROCKOFF, “The Impact of Charter Schools On Student Achievement,” *Working Paper* (2005).
- IMBERMAN, S. A., “The Effect of Charter Schools on Achievement and Behavior of Public School Students,” *Journal of Public Economics* 95 (2011), 850–863.
- KAPOR, A., “Distributional Effects of Race-Blind Affirmative Action,” *Working Paper* (2015).
- KUMINOFF, N. V., V. K. SMITH AND C. TIMMINS, “The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets,” *Journal of Economic Literature* 51 (2013), 1007–1062.

- MANSKI, C. F., “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies* 60 (1993), 531–542.
- MCMILLAN, R., “Competition, Incentives, and Public School Productivity,” *Journal of Public Economics* 88 (2004), 1871–1892.
- NCES, “PSS Private School Universe Survey Data for the 2009-2010 School Year,” (2010).
- NECHYBA, T. J., “Mobility, Targeting, and Private-School Vouchers,” *American Economic Review* 90 (2000), 130–146.
- NEILSON, C., “Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students,” *Working Paper* (2013).
- PODGURSKY, M., “Teams Versus Bureaucracies: Personnel Policy, Wage-Setting, and Teacher Quality in Traditional Public, Charter, and Private Schools.,” *Education Working Paper Archive* (2006).
- SASS, T. R., “Charter Schools and Student Achievement in Florida,” *Education Finance and Policy* 1 (2006), 91–122.
- SEIM, K., “An Empirical Model of Firm Entry with Endogenous Product-Type Choices,” *The RAND Journal of Economics* 37 (2006), 619–640.
- SIEG, H., V. K. SMITH, H. S. BANZHAF AND R. WALSH, “Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods,” *International Economic Review* 45 (2004), 1047–1077.
- STINEBRICKNER, R. AND T. R. STINEBRICKNER, “Time-Use and College Outcomes,” *Journal of Econometrics* 121 (2004), 243–269.
- , “The Causal Effect of Studying On Academic Performance,” *The BE Journal of Economic Analysis & Policy* 8 (2008), 1–55.
- WALTERS, C. R., “The Demand for Effective Charter Schools,” *Working Paper* (2014).
- WHITE, J., “States Open to Charters Start Fast in ‘Race to Top’, Press Release, US Department of Education,” <http://www2.ed.gov/news/pressreleases/2009/06/06082009a.html> (2009), accessed October 13, 2010.

B FOR ONLINE PUBLICATION

B.1 Existence of Equilibrium These proofs are for the case where there are no productivity shocks to chosen effort levels; the results also go through when there are shocks. To use Brouwer's Fixed Point Theorem, the pair of best-response functions must be a continuous self-map on a compact and convex set. I prove first there is a unique best response of one school to another, then that this best response function is continuous, and finally apply Brouwer's Fixed Point Theorem. As in the model, both school production functions satisfy Inada conditions and schools have convex effort costs.

Lemma 1. *The solution to the charter school's effort choice problem e_{ch}^* is strictly positive.*

Proof. $\lim_{e_{ch} \rightarrow 0} \frac{\partial v_{ch}(\mathbf{e})}{\partial e_{ch}} = \infty$ due to the Inada conditions on the test score production function, because there will always be a positive measure of students attending the charter school due to preference shocks. \square

Define $v_{ch}^+ = \max\{v_{ch}, 0\}$. Note that v_{ch}^+ is strictly quasi-concave, due to the strict concavity of v_{ch} when it is above 0.

Lemma 2. *The effort set $E = [\underline{e}, \bar{e}]$ is compact.*

Proof. Let $\underline{e} = 0$. Given any allowable vector of parameters θ there exists \bar{e}_θ such that $v_{ch}(\hat{e}) < 0, \forall \hat{e} > \bar{e}_\theta$. Let $\bar{e} = \max_\theta \bar{e}_\theta$. It exists, so the set is not empty. \square

Lemma 3. *$e_{ch}^*(e_{tps})$ is continuous.*

Proof. Berge's Maximum Theorem requires a continuous objective v_{ch}^+ , and compact and upper-hemicontinuous (UHC) constraint set. Note first that the constraint set, E , is a fixed connected interval, so it is trivially UHC. v_{ch}^+ is continuous, so the Maximum Theorem says the resulting correspondence which is the argmax of v_{ch}^+ is UHC. Because v_{ch}^+ is strictly quasi-concave there is a unique argmax to v_{ch}^+ , which means that $e_{ch}^*(e_{tps})$ is a continuous function. \square

Analogous reasoning applies to $e_{tps}^*(e_{ch})$.

Lemma 4. *There exists an equilibrium to the entry subgame.*

Proof. $\Gamma(e_{tps}, e_{ch}) = (e_{ch}^*(e_{tps}), e_{tps}^*(e_{ch}))$ is a continuous self map on the compact and convex domain E^2 , so there exists an equilibrium by Brouwer's Fixed Point Theorem. \square

B.2 Uniqueness of Equilibrium I do not prove uniqueness of equilibrium in the entry subgame but can rule out multiplicity of the charter school entry decision, given a unique equilibrium in the ensuing entry subgame.

Lemma 5. *There is no multiplicity in the charter school entry decision given uniqueness of equilibrium in the entry subgame.*

Proof. The charter school only receives one shock, \underline{v}_m , when it considers its entry decision, which it knows. It enters if and only if $E_{\nu^e} [v_{ch,m}(\mathbf{e}^*)] \geq \underline{v}_m$, where $E_{\nu^e} [v_{ch,m}(\mathbf{e}^*)]$ is known since under the assumption of the lemma there is a unique equilibrium in chosen effort levels of the entry subgame. \square

Searches for more than one equilibrium for a wide range of parameter values and have never returned more than one equilibrium in a market. Intuitively, there will not be multiple equilibria in the entry subgame so long as schools are not too responsive to each other, which would be satisfied if effort cost were sufficiently convex. In the presence of multiple equilibria for the entry subgame both schools are assumed to play the same equilibrium.

B.3 Data Construction Details

B.3.1 Construction of Effort Variable Individual students were asked how much time they spent on homework per week in an end-of-grade survey. The possible responses were:

Code	Response
A	No homework is ever assigned by all their teachers
B	Less than one hour each week
C	Between 1 and 3 hours
D	More than 3 but less than 5 hours
E	Between 5 and 10 hours
F	More than 10 hours
G	Has homework, but does not do it.

Effort is assumed to be uniformly distributed within each interval to relate the above categorical answers to the continuous effort measure in the model. This student-level measure of homework done is converted to a cardinal measure by taking the midpoint of each interval, where the top category F is mapped to the maximum of 10, and category G is mapped to 0, the latter being consistent with the treatment of effort within the model (i.e., there is no input if no homework has been done). The school-level measure is then created by averaging

this measure over all students attending that school in that year. Substituting the expected value of a random variable is most defensible when the random variable enters the model linearly and the CES technologies being estimated here are unlikely to be perfectly linear. However, defining the model input this way has the advantage that classical reporting errors within an interval would average out.

B.3.2 Construction of Capital Variable The model requires capital levels for charter schools, even in markets where charter schools do not exist. The most natural variable would be per-pupil funding, which has been studied extensively in the economics of education literature (see, e.g., Hanushek, 2003; Hoxby, 2003). However, because charter schools in North Carolina do not receive funding for expenditures on buildings, and other facilities, they tend to have lower per-pupil funding than traditional public schools. Ideally, prospective per-pupil expenditures would be available for all charter schools. In this case, if one assumed that individual schools did not have bargaining power over inputs such as computers, books, or teachers, the relevant per-pupil expenditure level would enter as per-pupil “capital” for charter schools in the model. The problem is that these prospective per-pupil expenditures were not in the data, as they are a somewhat complicated function of the per-pupil capital levels of surrounding public schools. Therefore, to solve the model it is necessary to come up with an adjustment for public school funding to apply to charter schools.

The following algorithm computes a level of capital for both charter and public schools given information that is always observable for a market: i) Convert measures (computers/pupil, teachers/pupil, experienced teachers/pupil) to percentiles (using the same distribution for both charter and public schools), ii) average these percentiles into one index for each school, iii) regress this index on inflation-adjusted per-pupil expenditures for the public school in each market, using separate regressions for charter and public schools,³⁴ iv) use the predicted value from the above regression as the capital measure for that school type in that market.³⁵

An advantage of using linear regressions is that they ensure that capital in the model is an affine function of per-pupil expenditures; that is, the R^2 of model capital on per-pupil expenditures is 1 for both public and charter schools. The manner in which per-pupil expenditures enter is consistent with linear cost functions, that is, inputs to capital are

³⁴The use of separate regressions allows charter and public schools to face different input prices which may be relevant, say, for hiring teachers. Details are available upon request.

³⁵The last step obviates integrating over the errors in the cost functions when solving the charter school’s entry problem. Also, it precludes charter schools from making entry decisions based on unobserved information – that is, the predicted per-pupil capital levels are no different in expectation in duopoly and monopoly markets with the same level of per-pupil expenditures. Although it is possible that such variation may play a role in charter school entry, it is likely second order in understanding charter school entry patterns.

perfect substitutes where a dollar more spent on computers means a dollar less spent on some other input, such as experienced teachers. Another check for whether the capital forecasting model makes intuitive sense is to compare this paper’s findings with those in the literature. The structural estimation finds that capital has a very low marginal product at traditional public schools, a finding consistent with Hanushek (2003), but has a much higher marginal product at charter schools. One potential reason could be that charter schools can hire teachers outside the regular teacher labor market, assign bonuses, etc.

B.3.3 Details about Distance Distributions The model accounts for within-market charter school locations through market-level distance distributions for charter and traditional public schools. These distance distributions do not necessarily correspond to physical locations within markets but do take into account the fact that charter schools are often further from students than traditional public schools.

As with capital, distance distributions are required for all markets in order to calculate the value the charter school would expect to obtain upon entry, which then enters the expression for probability of entry. This section first explains how distance distributions were created and then presents descriptive statistics about distances between students and schools.

Creation of Distance Distributions: Distance distributions are created according to the following algorithm: i) discretize the continuous distance distribution for each market, creating the vectors \vec{r}_{ch} , \vec{r}_{tps} , and $\vec{\mu}$, ii) using markets with charter schools, regress elements of \vec{r}_{ch} , \vec{r}_{tps} , and $\vec{\mu}$ on a two-bin distribution (fraction of students within median [of that market] distance to the public school and further than median distance to the public school, and average the distance for within-median and beyond-median students) for public schools, iii) use this regression to create a predicted fraction of students in each bin, iv) normalize the elements of the predicted fraction of students in each bin to sum to one for each market; this is the distance distribution used in the model. Further details are available upon request.

Descriptive Statistics about Distances Because the distance distributions are a bit unwieldy and many students do not have distance data for both charter and traditional public schools (for example, if a student is only observed in a charter school, the home address may never have been given to the NCERDC), Table 12 presents several moments of simulated distance distributions. The first and second columns provide quantiles of the marginal distributions of distances of students to public and charter schools, respectively. We can see that that marginal distribution of distance from charter schools dominates that for public schools.

The third column shows quantiles of the net distance to public schools. By looking at the net distance from public schools in the third column, which depends on the joint distributions

of distance from public and charter schools, we can see that students are usually closer to the public school. In fact, charter schools would only be closer to students about 20% of the time. This makes sense because markets are designed around traditional public schools; charters enter afterwards and are therefore less likely to be in the center of the population mass.

The fourth column presents quantiles of the market-level average of the net distance of households from the public school, while the fifth column presents quantiles of the share of students in each market who are closer to the public school. These last two columns show that students are typically closer to the public school in most markets. At the same time, there is substantial variation across markets in the typical net distance of students from the public school: there are markets in which the average student would be closer to the charter school as well as markets in which the majority of students would be closer to the charter school. These last columns provide variation to help identify parameters of the test score production function.

Table 12: Moments of Simulated Distribution of Distances (km) of Students from Public and Charter Schools

	r_{tps}	r_{ch}	$r_{tps} - r_{ch}$	$\overline{r_{tps} - r_{ch}}$	$\overline{\mathbb{1}\{r_{tps} < r_{ch}\}}$
1%	0.39	0.51	-5.52	-2.06	0.44
25%	1.79	1.90	-1.46	-1.11	0.82
50%	3.98	5.23	-0.87	-0.85	0.92
75%	8.28	9.48	-0.12	-0.63	0.94
99%	31.18	31.09	2.59	0.10	0.97

B.4 Identification of Test Score Production Function This section shows how distance distributions help to achieve identification of the key parameters in the test score production functions. Though the technology parameters are estimated using full information maximum likelihood, I use a multi-step argument to illustrate the intuition behind how identification is achieved when maximizing the likelihood. I set aside the test score shock because though it does not prevent identification, it complicates the identification argument without adding any intuition. For this argument, assume that market-level ability distributions are known; in practice, they are recovered during estimation.

Recall that the test score for a student at school s can be written as $a_i \bar{a}_{smt}^{\theta_s} E y_s(1, 1, e_{smt}, k_{smt})$. This can also be written as $g_s(a_i, \bar{a}_{smt}) h_s(e_{smt}, k_{smt})$, i.e., separated into components based on own and peer ability, $g_s(\cdot)$, and based on observable school effort and capital, $h_s(\cdot)$. Let $\mathbf{r}_m \in \mathcal{R}$ denote the pair of distance distributions for market m , where \mathcal{R} is the set of feasible distance distributions. Assume full support of distance distributions over \mathcal{R} . Such full

support assumptions are commonly invoked in such identification arguments.

Define a set of markets and time periods m_{ch} which all have distance distributions where there is not a positive measure of students who would choose to attend the public school because the students are all too far from it. A key to this set being nonempty is that households have a positive cost of commuting, which is consistent with the estimation results. Analogously, let m_{tps} define the set of markets where distance distributions are such that there is not a positive measure of students who would choose to attend the charter school, were it to enter.³⁶ The full support assumption from earlier must be augmented with “dense-ness” (or continuity of the distribution of distance distributions) assumptions which ensure that m_{ch}, m_{tps} are sufficiently populated (e.g., not just one market).

Consider m_{ch} first. Because all students in such markets attend the charter school, we can write the expected test score for charter schools in such markets as $g_s(\bar{a}_m, \bar{a}_m)h_s(e_{stm}, k_{stm})$, where \bar{a}_m is the mean of the market-level ability distribution. The two arguments in g_s are the same because there is a continuum of students in each market (as there is in the model). By considering the same market over different time periods in which a charter school entered, we can recover estimates of $h_{ch}(e_{ch,t,m}, k_{ch,t,m})$ by differencing the log of $g(\bar{a}_m, \bar{a}_m)h_s(e_{stm}, k_{stm})$ and running market-period regressions of log mean test scores on the difference. Although $h_s()$ are in principle non-parametrically identified, this paper adopts a relatively flexible CES production function due to the known data demands of non-parametric identification. Therefore, we could then run non-linear least squares to decompose $h_{ch}(\widehat{e_{ch,t,m}}, \widehat{k_{ch,t,m}})$ into $Ey_s(1, 1, e_{stm}, k_{stm})$, recovering the charter school technology parameters $(\omega_{ch}, \alpha_{ch}, \beta_{ch}, \tau_{ch})$. Intuitively, variation in both schools’ capital levels as well as the effort productivity shocks create variation in school effort and capital, even when there all students within a market attend one school. Analogously, we can use the set of markets m_{tps} to recover the analogous parameters for the public school.

It still remains to recover the coefficient on peer quality at charter schools, θ_{ch} . As with estimating the parameters of the CES technology, to estimate this parameter (and θ_{tps} , for peer quality at public schools), I use full-information maximum likelihood. As before, I will sketch the remainder of this argument because it is illustrative to talk about the intuition behind identification. First consider how to identify θ_{ch} . To estimate θ_{ch} , first use the estimates of $h_{ch}(\widehat{e_{ch,t,m}}, \widehat{k_{ch,t,m}})$ to get estimates of $g_{ch}(\widehat{a_i}, \widehat{a_{ch,t,m}})$, using the residuals. Let $\check{y}_{i,ch,t,m} = (y_{i,ch,t,m}/h_{ch}(\widehat{e_{ch,t,m}}, \widehat{k_{ch,t,m}}))$ be the test score expunged of observed school inputs.

Again, consider a market and time period $m, t \in m_{ch}$, where all students attend the charter school due to the distance distribution in that market that period. The mean expunged

³⁶Though not necessary for this identification argument, in principle m_{tps} could be augmented with markets in those years where there is no observed charter school entry.

test score in the charter school in that market would then be:

$$\begin{aligned}\bar{y}_{ch,t,m} &= (\bar{a}_{ch,t,m})^{1+\theta_{ch}} \\ &= (\bar{a}_{t,m})^{1+\theta_{ch}} = (\bar{a}_m)^{1+\theta_{ch}},\end{aligned}$$

where the first line follows from there being a continuum of students, the first part of the second line follows from the fact that all students in m, t chose the charter school (by assumption), and the second part of the second line follows from the invariance of market ability distributions over time. Because the market ability distribution is known (again, by assumption for this argument), we can compute \bar{a}_m and then solve for θ_{ch} or estimate it by considering all charter schools collected in m_{ch} . Perform the analogous exercise to estimate θ_{tps} for public schools.

At first pass, it may seem odd that within-market distance distributions play no role in the above identification argument. There is an intuitive reason this oft-used source of identifying variation is not necessary for this “identification at infinity” type argument: by considering markets with “perfect” sorting and the fact that there is another restriction implied by integration, i.e., $\bar{a}_{ch,t,m} = \int a f_{ch,t,m}(a) da$, where $f_{ch,t,m}(a)$ is the density function for student ability at the charter school in market m in period t , the relationship between own and peer ability is known. If the relationship between own and peer ability were not known the above argument would not go through. However, within-market variation in distance from charter and public schools, though not necessary for identification, helps to more precisely estimate the technologies. Intuitively, the further all students in a market are from the charter school, the less sorting on ability there would be in that market.

B.5 Sample Selection Table 13 compares selected variables for the full sample and estimation samples. The means of most variables for public schools in the full and estimation samples are similar. In both the full and estimation samples, markets with charter schools have higher fractions of Black and Hispanic students, yet charter schools have lower fractions of both types of students relative to public schools in such markets. In both samples, female students comprise a smaller share of students at charter schools than they do for both types of public schools, and students attending charter schools are much more likely to have had at least one parent who has attended at least some college than students at either type of public school (in the estimation sample, 75% for charter schools versus 60% for public schools in markets with charters and 43% for public schools in markets without charters).

Table 14 shows how sample restrictions affect the test score distribution. In particular, removing students who attended a public school outside their market increases the average

test score for students attending charter schools.³⁷ About 20% of excluded students are from one charter school.

³⁷Note that the extent of the bias induced by using a subset of charter school students cannot be derived from simple comparisons of the mean test scores of charter and public schools in duopoly markets. What matters is the distribution of ability among excluded households.

Table 13: Summary Statistics for Full and Estimation Samples

Ethnicity	Summary Statistics for Full Sample						Summary Statistics for Estimation Sample									
	All Markets			Mkts. w. Charters.			All Markets			Mkts. w. Charters.						
	%	N	All Schools	%	N	Public	%	N	All Schools	%	N	Public	%	N	Public	
White	63.45	644300	61.75	4689	53.52	38214	64.22	601396	62.82	92080	68.40	3359	56.07	35443	67.85	53120
Black	29.44	298955	30.93	2349	37.89	27054	28.78	269552	30.02	44000	25.23	1239	36.47	23052	25.17	19709
Hispanic	3.08	31274	1.25	95	4.31	3077	3	28103	2.74	4016	1.34	66	3.26	2064	2.40	1882
Other	4.03	40925	6.07	461	4.28	3056	4	37464	4.43	6492	5.03	247	4.20	2657	4.58	3583
Female	49.09	1015511	46.93	7594	49.53	71402	49.07	936515	49.62	146588	48.61	4911	50.33	63216	49.08	78294
% Parent College Grade	47.92	1015365	70.11	7591	56.92	71373	47.06	936401	51.48	146563	75.22	4911	60.07	63216	43.04	78276
% Attending Charter	6.97	1015511	6.81	7594	6.94	71402	6.97	936515	6.89	146588	6.80	4911	6.97	63216	6.84	78294
% With Charter in Market	0.75	1015511	100.00	7594	0.00	71402	0.00	936515	3.37	146588	100.00	4911	0.00	63216	0.00	78294
	7.78	1015511	100.00	7594	100.00	71402	0.00	936515	46.59	146588	100.00	4911	100.00	63216	0.00	78294

Table 14: Test Score Distribution for Full and Estimation Samples by Market and School Type

	Duopoly Markets		Full Sample		Estimation Sample	
	Mean	SD	Mean	SD	Mean	SD
Charter	2.975	0.974	7376	0.938	4941	0.938
Public	3.032	0.995	67219	0.990	63216	0.990
Monopoly Markets						
Public	3.002	0.933	893413	3.0278	0.915	65085

B.6 Discussion of Market Definition Markets need to be defined in order to solve the model. Unfortunately, there is no natural (i.e., institutional) definition of market boundaries because charter schools may draw from more than one traditional public school. The least restrictive assumption would be to treat the entire state as one market. Such an assumption would not be a good fit for this analysis, which studies charter schools across the entire state of North Carolina. Therefore, there has to be some smaller definition of a market to head off the somewhat absurd scenario wherein schools in the state compete for students hundreds of miles away.

Any feasible definition of markets must be applicable even to locations without charter schools; otherwise, a charter school's value of entry, hence its probability of entry, cannot be computed. In the end, a tradeoff between fidelity to the institutional environment and computational complexity determined which feasible alternative worked best. Because the route taken excludes some households, I conducted several exercises to examine how robust the main findings were in terms of this assumption.

I start by discussing potential alternative definitions of markets. I then discuss in further detail the robustness exercises carried out, in particular, re-estimation of the model excluding markets with higher shares of students violating market definitions.

B.6.1 Potential Alternative Market Definitions I considered using both feeder schools and public school districts to define markets. I also considered aggregating public schools into "synthetic schools" comprising larger markets than single-school catchment areas. I will now discuss in detail the issues encountered when defining markets in these ways and why, unfortunately, none provides an obviously better definition for this analysis.

Alternative definition 1, Feeder schools: One way to determine which students may have been choosing between their middle school and a particular charter school would be to first determine which primary schools feed middle schools competing with each charter school, and then assign any students in any of those primary schools to be in the same market as the charter school.

The problem with using feeder schools to define markets is that they cannot be used to define markets in areas charter schools have not entered. Consider an isolated area with two middle schools and no charter schools anywhere nearby. Because we have not observed charter school entry in this area, the feeder school algorithm cannot be used to determine whether these middle schools should comprise one or two markets. It is necessary to define markets in such areas, however, to calculate the value of entry to the charter school in areas where there are no charters, which drive the estimates of the entry cost distribution. One option would be to assign each middle school in this example to be in its own market.

Unfortunately, such a rule would treat markets observed with charters (i.e., those where we grouped middle schools into one market, based on feeder schools) and markets not observed with charters differently.

Identifying feeder schools in markets without charter schools is easier when considering geographically isolated traditional public schools, because there are no “markets” to potentially group together. Interestingly, the robustness exercise which re-estimates the model using only markets with low shares of market crossers focuses on similar markets. That is, it effectively uses the markets where the feeder school definition could reasonably be applied even without having observed charter school entry, because there are no other public school options nearby.

A related approach would aggregate public schools into markets based on observed patterns of competition. For example, if a charter school has many students in attendance who transferred from two nearby public schools, one might combine the two public schools into one market. However, similar to feeder schools, this method does not provide a rule to combine public schools in areas where there are no charter schools, which is necessary in order to compute equilibria for all markets. By using distance to assign charter schools to markets, I provide a definition of markets that can be applied even in the absence of charter schools.

Alternative definition 2, School districts: Another plausible alternative to using distance to create markets would use an administrative boundary, such as the school district in which a traditional public school is located. However, creating markets this way would be problematic for three reasons:

1. Extremely high computational burden: If every school could compete with all the schools in their district, it would be necessary to solve the game for every possible combination of charter school entry outcomes, which would tremendously increase the computational burden.³⁸ There is no natural upper limit on the number of charter schools that might enter a district. The upper bound on the number of potential charter school entrants should be large enough so as not to mechanically fit entry probabilities by limiting the number of entrants. For example, if there could be as many charter schools as there are public schools, this would entail computing equilibria for over 2 million games for the largest district, which has 21 public middle schools, clearly an infeasible exercise. Additionally, the vast majority of these games (> 99.999%) would involve more than one charter school, and each game would be more computationally expensive to solve because it would pit up to 42 schools against each other, as opposed to the current setup where there is only one school competing with

³⁸A district with up to five charter schools would have 31 ($=\sum_{k=1}^5 \binom{5}{k}$) possible games. A district with up to ten charter schools would have 1,023 possible games.

any other school in each game.³⁹

Unfortunately, existing results from work in related areas do not easily extend here. Seim (2006)'s method for estimating an entry game where firms compete with multiple firms cannot be used because the payoff to entry is endogenous in the current setting. Allowing multiple charter schools to compete with a public school would require calculating the value of entry for a subgame corresponding to every possible combination of charter schools, in every market and time period. Seim (2006) was able to implement this because she assumed that the matrix characterizing how the presence of firms affected other firms' payoffs was a primitive object and, therefore, policy invariant. One contribution of this paper is that it studies interactions between charter and traditional public schools, which would not be possible to study if the degree of interdependence between school payoffs was a model primitive.

Ferreira and Kosenok (2015) allow students to attend all charter schools in a large market (Washington D.C.). A single market is a good fit for their environment because it corresponds to an administrative unit, and also implies a reasonable geographic size for a student to commute within.⁴⁰ They assume that multiple charter schools might compete with the same public school, but that public schools only react to charter schools through reduced-form policy functions. This assumption allows them to avoid solving a game where every charter and public school optimizes inputs, taking into account every other charter and public school, which would necessarily result in the environment studied in the current paper, due to endogenous public school effort choices. Again, because a focus of the current paper is on public school reactions to charter school entry, and investigating how those reactions would look in markets that have not yet had charter schools, I would lose a key and novel element of this paper's analysis by treating public school policy functions as fixed. That, and the marked difference in geographic scopes, causes me to view my paper and Ferreira and Kosenok (2015) as complementary.

2. It is not clear where within districts charter schools should be located: In general, it is not

³⁹Having more than two schools compete would also make calculating each school's objective more difficult. As seen in the equilibrium characterization in eq. (22), the model emits a simple expression for the probability that a student of a certain ability will choose the charter school when there are two schools in the student's choice set. If, say, a multinomial logit were used instead the model would no longer emit such a simple expression for the measure of students attending each school when solving for school best responses.

⁴⁰Even the work of Ferreira and Kosenok (2015), which allows for quite rich school choice sets (i.e., the entire DC school district), does not allow households from outside DC to attend DC charter schools, even though in reality it is possible for households outside DC to attend DC charter schools. See the Washington DC School Reform Act, Section 38-1802.06 for details (<http://www.dcpccb.org/sites/default/files/report/School%20Reform%20Act.pdf>). This simply reinforces the point that any market definition for the study of charter schools requires making a tradeoff between fidelity to the institutional environment and computational complexity.

clear where within a district a charter school should be located if the market were the school district boundary. This matters because districts would be much larger than markets as they are currently defined, increasing the importance of considering within-market location decisions. The current algorithm computes a market-level household distance distribution for students from charter schools, using the distance distribution of students from public schools (see Appendix B.3.3 for details). If there were multiple public schools in a market they would likely have preferences over location within the market. Allowing for this would require modeling the charter school location decision within a market, which would be considerably expensive computationally, if a pure strategy location equilibrium existed at all. If, to simplify computation, charter school locations within districts were assumed to be exogenous, this would eliminate the ability of a charter school to compete with a particular public school, which the current definition of markets allows for. Additionally, it is entirely possible that modeling within-district locations might produce outcomes similar to those produced under the current definition of markets, where charter schools are assigned to compete with particular traditional public schools.

3. Some students would still violate market boundaries: Though not common, students may cross school district boundaries to attend charter schools. Therefore, some students would need either to be excluded from the analysis in the manner implemented in this paper or dealt with some other way.

B.7 Robustness of Results to Market Definition This section explains how the choice of modeling markets does not compromise the validity of estimates of model parameters and inferences about charter school effects.

Charter school students in the estimation sample may have higher average ability than those in the overall population because the algorithm used to assign charter schools to markets excludes some students observed attending charter schools, and the excluded students on average have lower test scores. This could in principle affect the distribution of student ability at charter schools, which could affect both estimates of production function parameters and the distribution of estimated treatment effects.

I recover unbiased parameters for test score production technologies for charter and public schools by developing an identification strategy that separates unobserved student ability from the effectiveness of school inputs, even when students could sort on ability due to differences in school inputs (details in Appendix B.4). Therefore, estimates of test score production technologies are not biased by the sample restrictions induced by the market creation algorithm. After establishing that parameters have been consistently estimated, the estimated model can be used to quantify the extent to which excluding these households

might affect the simulated distribution of treatment effects – in particular the direct effect of charter school entry. I implement this by allocating excluded charter school students to a synthetic market and solving for equilibrium charter school effort levels in this synthetic market, given the ability distribution of excluded charter school students, and estimating treatment effects as in Section 6. I also re-estimated the model excluding markets where more than 5% of students crossed market boundaries to attend charter schools, and found that estimated parameter values and treatment effects changed very little (see Appendix B.7.1).

B.7.1 Robustness of Estimation Results to Sample Selection This paper uses a geographic rule to construct markets, where markets based on closest traditional public school. The advantage of this rule is that it works in all areas and is also tractable enough to afford analysis of endogenous effort choices. To examine how this definition may have affected this paper’s findings, I re-estimated model parameters excluding markets where more than 5% of the students were observed ever crossing boundaries to attend a charter school. I then calculated the distribution of treatment effects, finding that the results are quite similar.

First, it is important to note that few students violate the definition of markets. About 3% of students in a market cross the market boundary at some point to attend a charter school, which is a very low number. In only five (out of 496) markets does the share of students crossing boundaries to attend charter schools exceed 5%. This group comprises 4.2% of student-year observations from the estimation sample. Table 15 shows summary statistics for student test scores and market shares in this subsample. It shows that the students who violate the market definitions are similar to those who do not violate the market definitions, seen in Table 4 and this is not causing any major bias. The excluded markets contain some of the larger charter schools with lower test scores. This could be due to both inputs being provided or student abilities; the treatment effects recomputed for the subsample (presented below) suggest that it is more so the latter.

Table 15: Sample Means of Test Score and Market Share for Subsample Excluding Markets Where More than 5% Students Cross Boundaries to Attend Charter School

Duopoly Markets	Test Scores	Market Share
Charter	3.200	0.055
Public	3.073	0.945
Monopoly Markets		
Public	3.030	1.000

I now discuss the re-estimated model parameters on this subsample. The results described below are very similar to the baseline estimation sample results. In fact, using the standard

errors computed under the baseline specification, which includes all markets, none of the re-estimated parameters are statistically significantly different from those estimated from the baseline estimation sample. The main parameters driving the direct and spillover effects of charter schools are essentially unchanged, with three minor differences:

1. The student cost parameters that were most affected by this exclusion criterion were the student effort cost parameter c_e (which increased by 36%, from -3.869 to -2.494) and the interaction between student ability and the psychic cost of attending a charter school, $c_{ch,a}$ (which increased by 71%, from -0.361 to -0.106). The latter may stem from the higher test scores of the subsample.
2. The standard deviation of the school preference shock, σ_η , decreases by about 5% when these markets were excluded (from 7.985 to 7.585).
3. Excluding markets with more than 5% of students crossing boundaries disproportionately excludes markets where charter schools had entered. Therefore, estimates of charter school technology parameters change slightly, while from traditional public schools remain unchanged. The substitution parameter for the charter school, β_{ch} , decreases by about 1.5% from -0.675 to -0.686. The returns to scale parameter at charter schools, τ_{ch} , decreases by about 1% from 0.953 to 0.943.

Using these re-estimated parameters and the corresponding subsample of the data, I also re-computed the distribution of treatment effects to check whether the quantitative findings were substantially changed. As might have been expected given the similar parameter estimates, these are also very similar for the estimation sample and the subsample excluding markets. Specifically, Table 16 recomputes treatment effects presented in Table 11, using the re-estimated parameters and excluding the five markets where more than 5% of the students were observed crossing boundaries to attend charter schools. Two results of interest, the effect of treatment on the treated and the spillover effect, remain basically the same: the estimated mean of the direct effect $\bar{\Delta}^{direct}$ is 12% of a standard deviation (sd), up from 11% sd in the baseline, and the estimated mean spillover effect $\bar{\Delta}^{spill}$ is smaller, down to 1% sd from 2% sd in the baseline.

B.7.2 Using the Model to Correct for Potential Sample Selection Bias In this section I provide an approximation for the bias induced on the direct treatment on the treated (TOT), i.e. TOT for students attending charter schools, by excluding students who attend charter schools outside their designated markets.

Table 16: Robustness Results for Mean Direct and Spillover Treatment Effects by School Choice in Duopoly and Monopoly Markets

Duopoly Markets	$\overline{\Delta}^{direct}$	$\overline{\check{\Delta}}^{direct}$	$\overline{(\Delta^{direct} - \check{\Delta}^{direct})}$	$\overline{\Delta}^{spill}$
TOT (Attend charter)	0.116	0.113	0.003	0.003
TOU (Attend public)	-0.002	-0.012	0.010	0.010
ATE	0.007	-0.003	0.010	0.010
Monopoly Markets: TOU				
Attend charter	0.013	0.016	-0.002	-0.002
Attend public	-0.092	-0.097	0.005	0.005
ATE	-0.084	-0.089	0.004	0.004

Recall that allowing these excluded students in the model would imply interdependence between markets, violating the assumption that the markets form a partition of North Carolina. This interdependence would render the model computationally intractable: An equilibrium where all schools were competing with all other schools would be far more complex than the two-school version used in this paper because it would require solving equilibrium outcomes for not only a large number of schools but also *for every possible configuration of charter school entry decisions*.

One concern with this restriction is that excluding charter school students who on average have lower test scores might affect mean ability of students at charter schools. As previously discussed, this does not in principle induce a bias in estimates of test score production parameters at charter or public schools: those parameters are consistently estimated so long as student ability can be controlled for. If school effort choices and effective school inputs to test score production have been consistently estimated, changes in the student ability distribution will change the scale of the direct TOT but not alter its direction, so long as mean ability and the difference between effective inputs at charter and monopoly public schools for students attending charter schools is positive.⁴¹ In this scenario, a lower mean of the charter school student ability distribution would only diminish the direct treatment effect, resulting in upwards bias from excluding lower ability students.

I first approximate what the direct TOT would have been for the excluded students, and then compute an average of direct TOT for the included and excluded samples, weighed by the sample sizes. To capture the fact that excluded charter school students have lower average test scores than included charter school students, I assume that the excluded charter

⁴¹ “Effective inputs” are the expected test score at a school for student with ability 1, i.e. $\omega_s \bar{a}_s^{\theta_s} (\alpha_s (e_s)^{\beta_s} + (1 - \alpha_s) k_s^{\beta_s})^{\tau_s / \beta_s}$. Differences in student ability exacerbate differences in effective inputs between schools.

school students all reside in one synthetic market, denoted $m = ex$.⁴² Let I_{ex} denote the set of excluded students, $n_{ex} = |I_{ex}|$ the number of excluded students, and F_{ex} the distribution function for ability of excluded students. The direct effect of treatment on the treated in market ex is

$$(28) \quad \overline{\Delta}_{ex}^{direct, TOT} = \int \Delta_{i,ex}^{direct} f_{ex}(a_{i,ex}) da_{i,ex}$$

where $f_{ex}(a_{i,ex})$ is the density of ability for excluded students and the direct effect for i , $\Delta_{i,ex}^{direct}$, is the difference in test scores between the synthetic charter and monopolist public schools in market ex :

$$\Delta_{i,ex}^{direct} = Ey_{ch}(a_i, \bar{a}_{ch,ex}, e_{ch,ex}, k_{ch,ex}) + \nu_{i,ch,ex}^y - Ey_{tps}(a_i, \bar{a}_{ex}, e_{tps,ex}^{mono}, k_{tps,ex}) - \nu_{i,tps,mono,ex}^y$$

where $k_{s,ex}$ and $e_{s,ex}$ denote capital and effort levels for school s in market ex .

As was seen in Section 6.2, the treatment effect on the treated depends crucially on endogenous school inputs and student sorting on unobserved ability. To evaluate the above expression one needs the ability distribution for excluded students F_{ex} , which can be approximated using the same method used in estimation (Section 4), but substituting in charter school *average* capital and effort inputs across all charter schools in all markets.⁴³ As before, assume the test score distribution is normal, which means that it is sufficient to recover the mean and standard deviation for the excluded ability distribution.⁴⁴ Mean ability for excluded students is

$$\bar{a}_{ch,ex} = \left(\frac{\bar{y}_{ch,ex}}{y_{ch}(1, 1, \bar{e}_{ch}, \bar{k}_{ch})} \right)^{\frac{1}{1+\theta_{ch}}},$$

where \bar{e}_{ch} and \bar{k}_{ch} denote average charter school capital and observed effort across all markets. The standard deviation of excluded student ability is recovered analogously. After recovering the distribution of ability for excluded students, solve for equilibrium charter and monopoly public school effort levels in market ex by assuming that each school in the synthetic market has the average capital level for that school type across all markets, and that students are the same distance from both the charter and public school.

Finally, one can combine treatment effects for included and excluded students to form

⁴²Observations from 1998-2001 are pooled into the same synthetic market for the current exercise.

⁴³An alternative would be to posit an effort level for the charter school that enters the recovery of the ability distribution for excluded students, $\tilde{e}_{ch,ex}$, solve for the subsequent equilibrium effort level, $\tilde{e}'_{ch,ex}$, and iterate to find a fixed point in the effort level for the synthetic charter school serving excluded students.

⁴⁴Even if market ability distributions were normal, student sorting on ability would generically induce the distribution of abilities for students at charter schools to be non-normal. Nevertheless, this method captures the fact that excluded students have lower average test scores.

an estimate of the extent to which the sample selection procedure may bias estimates of the direct TOT. Let $\overline{\Delta}_{incl}^{direct,TOT}$ denote the expected direct effect of treatment on the treated for students attending charter schools who were retained in the sample, which represents n_{incl} students. The estimate of the overall direct effect of treatment on the treated, taking into account excluded households, is

$$(29) \quad \overline{\Delta}_{all\ markets}^{direct,TOT} = \frac{\overline{\Delta}_{ex}^{direct,TOT} n_{ex} + \overline{\Delta}_{incl}^{direct,TOT} n_{incl}}{n_{ex} + n_{incl}}.$$

The market size of the synthetic market containing excluded students affects charter and monopoly public school inputs, so results are reported for three scenarios: i) the market serving excluded students is three times as large as an average market, ii) four times as large, and iii) five times as large. The mean direct effect on excluded students is 11.1% of a standard deviation (sd) in the first case (5.7% sd and 2.1% sd for cases ii and iii, respectively), which returns a mean direct effect on charter school students to 10.9% sd in the first case (9.7% sd and 7.9% sd for cases ii and iii, respectively).

Another way to characterize the extent to which excluded students might have a different treatment effect than those included in the estimation sample is to use a back of the envelope calculation applying the same effective inputs they would have received at charter schools, adjusting for their different (lower) mean ability. Using this method, I calculate that the direct TOT would be 9.4% sd, which is similar to the result obtained using the previous method. In summary, excluding the students does not substantially affect the mean direct effect of charter school entry on student achievement.

B.8 Sensitivity to Non-pecuniary Cost I now discuss how estimates of direct and spillover effects would change if the cost of attending charters, c_{ch} , was significantly lower. If the cost parameter partially captures binding charter school capacities, reducing it allows us to understand how capacity constraints may affect achievement. I re-estimated the model with the restriction that the charter school non-pecuniary cost parameter c_{ch} to be one-tenth of its value in the baseline specification, which corresponds to about 1.5 sd in test scores. Other than the other cost parameters, the parameter most affected is the standard deviation of the school preference shock σ_{η} , which falls to about one-tenth of its baseline value. Ability sorting is therefore stronger, causing the model to increase the probability that high-ability students attend charter schools.

Although the fit of the restricted model is not as good as that of the baseline, qualitative findings about how charter schools affect student achievement are similar between the restricted and baseline models. Under the restricted model, the effect of attending charter

schools for attendant students is substantial and positive (0.189 sd), and varies quite a bit between charter schools. The mean spillover effect on public school students is -0.01 sd. That is to say, as in the baseline results, charter schools on average increase achievement for attendant students, are quite heterogeneous in how much they do so, and cause relatively small mean spillovers on students at competing public schools.