Essays on the Economics of Education

by

Zachary Tobin

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2022 by the
Graduate Supervisory Committee:

Esteban Aucejo, Chair
Daniel Silverman
Alvin Murphy

ARIZONA STATE UNIVERSITY

May 2022

ABSTRACT

This dissertation explores the effect of school competition on the human capital accumulation of students. Policies that expand the scope for school choice have become increasingly popular largely due to the belief that this will create incentives for low-performing, incumbent schools to improve academic outcomes. However, there is a general lack of empirical support for these positive academic spillover effects in most contexts.

In the first chapter, I demonstrate that if schools respond to competition through channels not typically considered in standard arguments in favor of school choice, it means that these policies may lead to negative, unintended consequences for academic achievement. I find that increasing the number of schools serving a given market can have a negative effect on test scores through creating incentives for schools to increase the provision of non-academic services that do not contribute to academic preparation, and through the creation of excess costs in the public school system. I use an empirical strategy designed to address strategic location decisions by new entrants as well as student selection across schools to show that entry of a new charter middle school during a recent large-scale charter expansion in North Carolina decreased average traditional public middle school test scores across a school district.

The second chapter considers the extent to which policymakers have tools available to them that can improve the ability of competition to generate the increases in test scores at incumbent schools that they have prioritized. I show that the efficacy of school choice can be improved by providing short-term, partial reimbursements to public school districts for increases in charter school enrollment by resident pupils. I also demonstrate that these effects occur not only due to the direct increase in district revenue associated with reimbursements, but also because the presence of this aid reduces the incentives of school administrators to compete for students through

non-academic channels. The empirical strategy that I use to generate these results leverages plausibly exogenous cutoffs for aid eligibility induced by a unique policy in the state of New York.

ACKNOWLEDGEMENTS

I am very grateful to Esteban Aucejo for his guidance and support throughout my graduate studies. I would also like to thank Dan Silverman and Alvin Murphy for their time and advice. This work has benefited greatly from insights provided by seminar participants at Arizona State University, London School of Economics, the Western Economic Association International Annual Conference, the Society of Labor Economists Annual Meeting, and the Association for Education Finance and Policy Annual Conference. I would also like to acknowledge the valuable support provided by the National Academy of Education/Spencer Dissertation Fellowship program. Finally, I would like to thank the other faculty members at Arizona State University as well as my fellow students for the helpful conversations and insights that they provided throughout this process.

TABLE OF CONTENTS

CHAPTER Page

LIST OF TABLES

Table Page

LIST OF FIGURES

Figure Page

Chapter 1

COMPETITION IN THE NON-PROFIT SECTOR: EVIDENCE FROM A CHARTER SCHOOL EXPANSION IN NORTH CAROLINA

1.1 Introduction

Policies that facilitate the expansion of school choice have enjoyed considerable popularity in the United States over the past few decades, and recent debates have become increasingly focused on the expansion of charter schools. These schools are publicly funded institutions that are granted considerable autonomy over teaching methods and budgeting in exchange for heightened accountability standards. A standard argument in support of charters, and school choice in general, states that an increase in the number of schooling options can improve student outcomes through the exertion of competitive pressure on existing public schools (Hoxby, 2003). In 2016, 25 years after the first charters were granted, these schools had grown to serve around three million students across 43 states (National Center for Educational Statistics, 2019). However, despite this rapid growth, evidence of the spillover effects of charter competition on the traditional public school sector is limited, and results are inconclusive (Bettinger, 2005; Imberman, 2011).

Existing models of school competition often model incumbent public schools as stand-ins that produce education but do not take any strategic decisions.[1] Other studies model them as profit maximizers, where profits can only be used for perquisite consumption (Hoxby, 2003; Neilson, 2020). These theories typically predict that competition among schools will increase efficiency in terms of education produced

[1]For examples, see Epple and Romano (1998), Nechyba (2000), and Ferreyra (2007).

per dollar of expenditure.[2] However, there are a number of key ways in which non-profit organizations differ from their for-profit counterparts, and this suggests that models of competition borrowed from the for-profit sector may not accurately capture the incentives faced by public schools.

The defining characteristic of non-profit organizations is that they are legally barred from distributing net earnings to individuals that control them (Hansmann, 1980). In many economic applications, the non-pecuniary motives of these organizations are modeled through objective functions defined directly over quantity and/or quality of output (Lakdawalla and Philipson, 2006), which can lead to very different implications relative to theories in which firms have purely pecuniary incentives. Weaker profit-maximizing incentives serve as a signal of higher product quality to consumers in sectors where quality is largely non-contractible and the threat of ex-post expropriation of profits is thus relatively high, such as in markets for schooling (Glaeser and Shleifer, 2001). Also, because non-profit organizations are not accountable to any owners, production is much more likely to be influenced by employee preferences and organizational mission statements (Glaeser, 2002).

A large class of non-profit organizations, such as charities, rely on donations for most or all of their funding (Hansmann, 1980), and typically attract donors through fundraising behavior, which does not contribute directly to production. Beginning with Rose-Ackerman (1982), questions regarding whether non-profit organizations engage in socially optimal levels of fundraising have been a major focus of the literature. Theories of non-profit competition predict that an increase in the number of competitors that serve a given market will increase the share of total expenditures de-

[2]Some notable exceptions demonstrate how competitive incentives may be undermined by household preferences for peer quality (Barseghyan *et al.*, 2019), capacity constraints (Cardon, 2003), and any scope for schools to engage in rent-seeking behavior (McMillan, 2004).

voted to fundraising (Castaneda *et al.*, 2008; Aldashev and Verdier, 2010), and is also associated with a shift of influence from workers to donors (Glaeser, 2002). In other words, these theories predict that competition induces non-profit organizations to behave more like their for-profit counterparts, and less like they would in the absence of competition. However, a general lack of consensus exists in the empirical literature regarding the relationship between competition and fundraising. For example, Feigenbaum (1987) finds that competition increases fundraising expenditure among medical research charities, but Thornton (2006) finds a negative relationship between competition and fundraising upon examination of a larger set of organizations. While public schools do not actually earn the majority of their revenue through donations, a clear analogy for fundraising behavior exists in markets for education.

In most contexts across the country, the vast majority of public school revenue is allocated on a per pupil basis, and therefore households effectively act as "donors" to schools when they make enrollment decisions. Most standard economic models consider education an investment, but recent work has found that at the post-secondary level, schools also attract students through the direct provision of consumption amenities, such as athletic programs and other student activities (Pope and Pope, 2009; Jacob *et al.*, 2018). Less evidence of this behavior has been documented at the primary or secondary school level, although in a private school setting, it has been shown that parents value teachers that not only improve test scores but also directly improve student satisfaction (Jacob and Lefgren, 2007). Therefore, analogous to the fundraising behavior of charities, when the degree of competition in a market for education increases, schools may compete for households through an increase in the provision of consumption amenities, even if this is associated with an opportunity cost in terms of academic achievement. This seems especially relevant given that a wave of recent literature finds that household preferences do not vary systematically with school effectiveness (Beuermann *et al.*, 2019; MacLeod and Urquiola, 2019; Abdulkadiroglu *et al.*, 2020).[3]

In this study, I test the implications of a simple model of school competition inspired by the non-profit literature in which schools can provide consumption amenities that do not directly improve academic performance, and which are associated with a time cost in terms of test scores. The primary purpose of the model is to develop and formalize the intuition regarding how increasing the number of schools serving a given market for education can have a negative effect on academic achievement in the market. I also argue that even if policies that facilitate an increase in the number of competitors serving a given market for schooling have adverse implications for academic achievement, it does not necessarily mean that these policies are not socially desirable. This is because they may give households the opportunity to take advantage of increased provision other types of services that they value, which may also facilitate the development of valuable skills that are not measured by standardized math and reading exams. I then provide empirical support for the qualitative implications of this model by leveraging a rapid, large-scale charter sector expansion that took place in the state of North Carolina following the 2011 removal of a cap that limited the number of charter schools across the state to 100. Figure 1.1 provides a means to visualize how the spatial distribution of charter students changed across the state during the first three years of this expansionary period.

The primary empirical results presented in this study utilize a two-way fixed effects (TWFE) strategy, which is essentially a standard difference-in-differences framework extended to a setting with variation in treatment timing. Under the assumption

[3]School effectiveness may play a relatively more important role in household enrollment decisions in contexts in which households receive specific information and guidance regarding the academic quality of public schools located in their area (Campos and Kearns, 2021).

Figure 1.1: Percentage of Public School Students Enrolled in Charter Schools

(a) 2011-2012

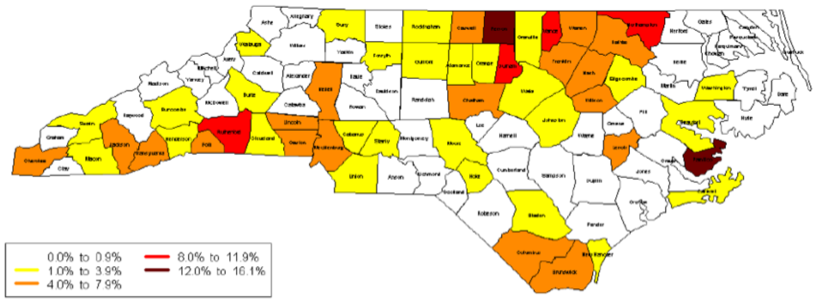

(b) 2014-2015

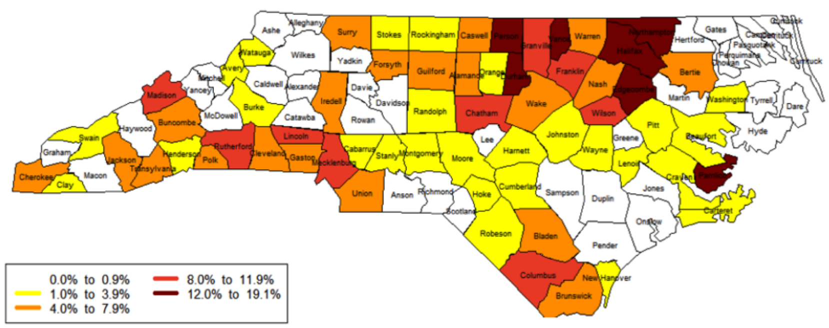

This figure displays the share of public school students enrolled in the charter sector by North Carolina county in the school year prior to the removal of the charter school cap (2011-12) as well as three years later (2014-15). Adapted from *Maps Percentage Of Public School Students In Membership At Charter Schools,* Division of School Business, North Carolina Department of Public Instruction. Retrieved from https://www.dpi.nc.gov/documents/fbs/resources/data/maps—percentage-of-public-school-students-in-membership-at-charter-schools of common counterfactual trends across treatment cohorts, this strategy identifies a variance-weighted average of treatment effects from all possible difference-in-difference comparisons across cohorts and time periods (Goodman-Bacon, 2019).[4] The sample I use to generate these results consists of traditional public middle school students enrolled in schools that serve grades 6-8.[5] I find that charter school competition decreased math standardized test scores for traditional public middle school students by roughly 4% of a standard deviation on average across an LEA. Based on familiar estimates from the teacher value added literature, the magnitude of this effect estimated is comparable to a decrease in teacher ability by about a quarter of a standard deviation (Chetty *et al.*, 2014a), which I interpret as moderately large. It is also worth emphasizing that this estimate represents the effect of competition at the market level, which further distinguishes this study from other recent work on the spillover effects of charter competition that focus on a highly localized treatment (Slungaard Mumma, 2020; Gilraine *et al.*, 2021).

Two challenges that must be addressed when estimating spillover effects of school competition on student achievement are the endogenous location choices of competitors and the non-random selection of students into schools. In order to address potentially strategic location decisions by new charter entrants, I have collected all applications to open a charter school that were submitted to the state of North Carolina during the sample period. In this context, a prospective charter operator must

[4]In Section 1.5, I also demonstrate the robustness of these results by using two alternative empirical strategies, based on different identifying assumptions, through which I obtain very similar estimates.

[5]This is by far the most common grade configuration for middle schools in North Carolina, and exceptions exist primarily in rural areas (Macartney, 2016). Over 90% of all traditional public school students during this time period began middle school in sixth grade and began high school in ninth grade.

apply to a singe state-level authorizer, which then makes a decision to accept or reject the application. These applications contain a variety of detailed information about the proposed school, including the LEA in which it plans to locate.[6] In light of recent evidence regarding strategic behavior of charter schools (Singleton, 2019), it seems very unlikely that all potential charters randomly select the LEA specified on their applications, or that this decision is not directly related to characteristics of the students located in a given market. However, I argue that the accept/reject decision of the state-level authorizer is exogenous to unobservable characteristics of an LEA that are related to student achievement. Therefore, the actual effect of competition can be identified by conditioning on the level of proposed competition, and all empirical results include controls for the total number of applications submitted to the state that specify a given LEA.

One measure I take to address the issue of student selection is to control for lagged test scores across multiple subjects whenever possible in order to address heterogeneity in academic ability. I also go a step further by reproducing all primary results using only the portion of students in my sample that had already completed sixth grade in a traditional public middle school prior to the opening of a new charter middle school in their district. Therefore, these students would not have had the opportunity to enroll in a new charter at the start of middle school. While over 4% of sixth graders in the expansionary period of my sample switch to the charter sector at the start of middle school, less than 1% of seventh and eighth grade students switch after completing at least one year of a traditional public middle school. Therefore, this subsample of students is much less likely to exhibit selection behavior, and if the estimated effect of competition considering only these students is qualitatively similar to the corresponding estimate using the full sample, this suggests that the primary

[6] LEA stands for "Local Education Agency" and is synonymous with "school district."

effects of interest are not being driven by selection alone.

Regarding the mechanisms that are driving the negative relationship that I document between charter competition and traditional public school achievement, I find that an increase in competition leads to an increase in the provision of non-academic services by traditional public schools. The incentives of educators in the state of North Carolina, as well as in many other states across the country, are based heavily on the standardized test performance of the students to which they are assigned. However, households may not have a way to effectively measure the academic quality of schools, and may also place a relatively high value on other types of services. Consistent with predictions from the non-profit literature that competition is associated with a shift of influence from employees to donors (Glaeser, 2002), I find that an increase in the degree of competition in a market for education is associated with a lower level of teacher empowerment and a higher level of household influence in traditional public schools by using teacher responses to a working conditions survey. I also find that charter competition is associated with a decrease in the total share of teacher salaries within a school that are devoted to teachers specializing in subjects tested by the state (math and English/language arts), and I interpret these results as evidence that providing consumption value to households is an important channel through which public schools respond to competition. This is consistent with the evidence documented in Kofoed and Fawson (2020), which finds that public schools in Utah responded to a charter school expansion through an increase in expenditure on total capital outlay and on improvements to existing structures, but did not increase expenditure on classroom instruction.

Another potentially important channel through which charter competition can impact students that remain enrolled in traditional public schools is through changes in instructional expenditure. Competition from the charter sector can potentially generate a negative fiscal externality associated with the fixed costs of school operation (Ladd and Singleton, 2020). The presence of fixed costs implies that when a public school loses revenue associated with lower enrollment due to competition, it will not be able to scale costs downward proportionally. Instead, more revenue per pupil will have to be allocated to covering these fixed costs, which may have adverse implications for academic performance given recent evidence that financial resources are an important determinant of student outcomes (Jackson *et al.*, 2016). That being said, to the extent that schools are spending some resources unproductively (i.e. allocating revenue to perquisite consumption that benefits employees directly but does not benefit students), competition potentially creates incentives for administrators to reallocate revenue toward classroom instruction. Thus, ex-ante the overall effect of charter competition on instructional expenditure is not obvious. Empirically, I find that the negative fiscal effect of charter competition dominates in this context in the sense that charter competition has a negative effect on the total salary that an LEA can offer a teacher, and is also associated with an increase in the share of per pupil expenditure allocated to school administration, which is an example of a cost that is likely fixed in the short-term. Overall, these results are consistent with findings in other contexts which demonstrate that the presence of charter schools may generate excess costs in the public school system (Bifulco and Reback, 2014).

A particularly interesting feature of the charter sector is that schools have the autonomy to offer different bundles of services to students than one might typically find in traditional public schools. For example, charter schools are not required to participate in school nutrition programs or to offer bus transportation, and they may elect to offer a curricular specialty, such as STEM.[7] Some evidence suggests

[7]Participation in school nutrition programs is defined as whether a school offers free/reduced-price lunches to eligible students.

that charter schools prefer to locate in diverse areas in which there is unlikely to be sufficient horizontal differentiation in the traditional public sector to satisfy all household preferences (Glomm *et al.*, 2005). The types of services being offered by charter competitors in a given market may have important implications regarding the marginal student for which schools are competing (Bau, 2019), and theory predicts that if the marginal consumer is different from the average consumer in a given market, competition may not be welfare-improving (Spence, 1975). Recent evidence suggests that the magnitude of the response to school competition depends on the degree of substitutability between the products offered by charter and traditional public schools serving a given market (Gilraine *et al.*, 2021).

One key characteristic that differentiates charter schools is school governance, or how a school is managed. These schools may be run by local members of the community or a non-profit organization, but they may also contract with a for-profit education management organization (EMO).[8] For-profit charters potentially have very different objectives relative to their non-profit counterparts, as they potentially have weaker incentives to improve student achievement beyond the accountability standard set by the state. These schools have been shown to exhibit stronger strategic behavior in terms of location choice (Singleton, 2019). They also spend more per pupil on non-instructional budget items (Singleton, 2017), which may be indicative of an increased propensity to provide consumption amenities relative to purely academic services. I find that traditional public schools subject to competition with entering for-profit charters experienced the largest decreases in teacher empowerment, and consistent with the theory these are also the schools which experienced the largest decreases in

[8]It should be noted that even charters managed by for-profit EMOs are still themselves non-profit organizations, but for the sake of brevity I will sometimes refer to these schools as "for-profit charters."

test scores.

Gilraine *et al.* (2021) also studies spillover effects of school competition in the context of the same charter sector expansion in North Carolina, and finds that proximity of the residential locations of elementary school students to new charter entrants is positively associated with math test scores, although only in cases where the product offered by the charter school is a reasonably close substitute for nearby public schools (i.e. when the charters are not offering completely different curricula or instructional philosophies). Therefore, it is worth noting some key distinctions between these two studies. First, Gilraine *et al.* (2021) focus on estimating an overall policy effect which includes both charter and traditional public school students. While their approach has the benefit of mitigating some concerns about selection across sectors, the focus on traditional public school students in this study is well suited to analyze the channels through which these schools respond to competition. Second, Gilraine *et al.* (2021) estimate a very local effect of competition, and my study instead focuses on comparing the effects of competition across markets. Even if the students within a market who live extremely close to charter schools experience positive effects on exam performance, this does not imply that the average effect of competition across a market is positive relative to a similar market that did not experience any increase in competition. Finally, Gilraine *et al.* (2021) focus on elementary school students while this study focuses on middle schools, and this distinction is potentially important because the scope for responding to competition through the consumption channel is likely greater beyond the primary-school level. For example, middle schools offer athletic programs and elective courses that can not be found at public elementary schools, and therefore there may be more opportunities for middle schools to compete for students through non-academic channels.

While this study focuses primarily on the spillover effects of competition on traditional public school students, it is closely related to the larger literature on charter schools. First, many studies have investigated the causal effect of charter attendance on student outcomes (Abdulkadiroglu *et al.*, 2011; Angrist *et al.*, 2012; Angrist *et al.*, 2016), and typically find at least moderately positive results.[9] However, these effects are often identified from random admission lotteries using samples restricted to charter applicants, which calls into question the external validity of results (Walters, 2018). More recent studies have used richer modeling techniques to explore equilibrium effects of charter entry (Mehta, 2017; Ferreyra and Kosenok, 2018; Walters, 2018), although these studies tend to focus primarily on demand-side effects. Jackson (2012) examines the effect of charter openings on teacher labor markets, but finds limited effects overall with the exception of "difficult to staff" schools,[10] for which there were declines in the number of new teachers hired as well as in measured teacher quality.

In the remainder of the paper, I will first present a stylized model of school competition. Then, I will describe the various sources of data that I have compiled, as well as the empirical strategy used to generate the primary results. Finally, I will discuss these results as well as the underlying mechanisms before concluding.

1.2 Model

This section presents a stylized model of school competition that is inspired by theories of competition among charities in the literature on non-profit organizations. In existing models of school competition, public schools often do not take strategic

[9]These references are intended as a few examples of this literature, as opposed to an exhaustive list. Other examples include Booker *et al.* (2011), Angrist *et al.* (2012), Angrist *et al.* (2013) and Dobbie and Fryer Jr. (2015). For a more complete list, please see Epple *et al.* (2015).

[10]The author defines these schools as those with sufficiently high shares of low-income and minority students.

decisions, or are modeled as maximizers of net revenue that can be used for perquisite consumption (Hoxby, 2003). However, these models may not accurately capture the relevant incentives that public schools face when the degree of competition increases in the market that they serve. Non-profit organizations such as schools differ from for-profit firms in a number of important ways because these organizations have weaker profit-maximizing incentives. This means that we often think of non-profit managers as maximizing different objective functions, and also that workers tend to be able to exert greater influence over production relative to their for-profit counterparts due to a reduced risk of ex-post expropriation.

A large class of non-profit organizations rely almost entirely on donations for revenue, which requires them to engage in fundraising behavior that is designed to attract donors, but that does not contribute to production (Hansmann, 1980). While public schools themselves do not rely on donations for most of their revenue, in this context households effectively act as "donors" because school revenue is allocated at the per pupil level, and therefore households directly determine school revenue with their enrollment decisions. An important insight from the literature on non-profits is that a decrease in revenue (e.g. due to increased competition) is associated with a shift of influence within these organizations from workers to donors (Glaeser, 2002). In other words, competition creates incentives for non-profits to behave more like for-profit firms and less like they would in a counterfactual scenario without competition.

In the context of education, competition may create incentives for incumbent schools to appeal directly to the preferences of households in order to retain revenue. This is relevant because while the incentives of teachers tend to be designed heavily around standardized test scores, households may not have an accurate way to measure or interpret school effectiveness, and may also place a relatively high value on other services that a school can provide. Education contributes to overall societal welfare in addition to benefiting individual students (Friedman, 1955), but because households do not internalize these social benefits, it may be the case that they partially subvert the ability of schools to maximize educational production. For example, a math teacher may believe that students require a substantial quantity of rigorous assignments in order to master a challenging topic, but households may prefer a lighter and less challenging workload in order to give students more time to focus on other endeavors.

Consider a market for education which is represented by a unit circle, as in Salop (1979). There is a mass of otherwise identical households located at each point on the circle which is normalized to one, and locations are indexed by $i \in [0, 1)$. Each household is assigned one student which they can choose to enroll in any one of a given number of schools $j = 1, ..., J$, which are also located around the perimeter of the circle. The circle in this context represents product space, such that different locations are interpreted as different varieties of schooling. A school produces education, which is characterized by a level of per pupil instructional expenditure and a level of consumption amenities, both of which are assumed to be valued by households. These consumption amenities are defined as activities that do not directly improve performance on statewide standardized exams, and as a result they are associated with an opportunity cost in terms of achievement on these exams. For example, this might include time spent on athletic programs and other student organizations. School administrators can also choose to allocate some revenue toward perquisite consumption ("perks"), which provide direct satisfaction to the administrator (e.g. a nicer office), but which do not benefit households in any way.

A key implication of the model is that as the number of schooling options increases in a given market, schools may choose to cater more to the preferences of households by increasing the provision of consumption amenities despite this being associated with a cost in terms of test scores. However, even if they do induce a reduction in test score performance, policies that facilitate an increase in the number of schooling options serving a given market may still be welfare-improving for households because more competition allows them to take advantage of increased consumption amenity provision and shorter distances to school.[11] There is no uncertainty in the model, so school administrators understand exactly how their decisions will impact enrollment.

Each household will choose the school from the set of J options that maximizes the following utility function:

$$\bar{u} + \nu I_j + \eta X_j - \tau d_{ij} \tag{1.1}$$

In this specification, $\bar{u}$ represents the constant utility associated with school enrollment,[12] and households also have preferences over per pupil expenditure on instruction (I_j), consumption amenities (X_j), and distance from school (d_{ij}). The fact that households care only about per pupil expenditure as opposed to total expenditure is a departure from models that focus on effect of competition among other types of non-profits. Most of these studies focus on the specific case of charities, in which the beneficiaries of charity production are not the same agents as the potential donors. Therefore, these studies often choose to model donors as having preferences over the output of all charities in the market (Aldashev and Verdier, 2010), or as a single representative agent that chooses a fraction of income to donate to each charity (Castaneda *et al.*, 2008), which simplifies the analysis. However, in this context households not only act as "donors" through their enrollment decisions, they also

[11]Recall that this is interpreted as the distance between the most preferred variety of schooling of a given household and the variety being offered by the school in which it enrolls its student.

[12]This parameter is assumed to be sufficiently high such that every household will prefer enrollment in some school versus the outside option of no enrollment.

reap benefits from their chosen school. Households only value the portion of total instructional expenditure that benefits their individual student, but the marginal cost of per pupil expenditure for the school administrator is increasing in the number of pupils.

Enrollment at a given school can be found by considering the marginal household located on either side of the school. The distance between school j and the household indifferent between schools j and $j+1$ can be found using the following condition:

$$\bar{u} + \nu I_j + \eta X_j - \tau d_{ij} = \bar{u} + \beta \nu I_{j+1} + \eta X_{j+1} - \tau \left(\frac{1}{J} - d_{ij} \right) \tag{1.2}$$

Solving for d_{ij} and adding this to the corresponding value for the student indifferent between schools j and $j-1$ gives the total enrollment of school j as a function of inputs. If we let $I^{'}$ and $X^{'}$ represent the input choices of the competitors of school j (which school j takes as given), then enrollment at school j as a function of school inputs is given as follows:

$$S(I_j, X_j) = \frac{\nu}{\tau}(I_j - I^{'}) + \frac{\eta}{\tau}(X_j - X^{'}) + \frac{1}{J} \tag{1.3}$$

The test score of each individual student enrolled in school j is given by the following equation:

$$Y(I_j, X_j) = \alpha I_j - \psi X_j \tag{1.4}$$

In this context, α is a parameter that represents the efficiency with which schools convert expenditure to test scores. The test score cost associated with a marginal increase in consumption amenities is represented by ψ, which can be thought of as the opportunity cost of time spent on activities other than preparation for exams. The utility function of a school administrator is given by the following:

$$e^{(1-\phi)M}(Y(I,X))^{\theta}(S(I,X))^{1-\theta} \tag{1.5}$$

This function takes a Cobb-Douglas form with share parameter θ and the test score and enrollment as inputs. This is intended to capture the notion that administrators are evaluated based on both the number of students they enroll as well as their academic performance. This objective function is consistent with the non-pecuniary motives that drive non-profit behavior across a wide variety of industries (Lakdawalla and Philipson, 2006). The coefficient that multiplies the Cobb-Douglas component depends on the parameter $\phi \in [0,1]$, which is intended to capture the degree of altruism of the administrator, and expenditure on perquisite consumption M. Expenditure of this type can be thought of as unproductive because it does not contribute to the ability of a school to fulfill its objective of educating students. More altruistic school administrators derive lower marginal value from an increase in perquisite consumption. Since schools are legally barred from distributing net revenue to employees or other controlling parties, education production will be subject to the following non-distribution constraint:

$$\zeta + M + I \times S(I,X) = \rho S(I,X) \tag{1.6}$$

In this equation, ζ represents the level of fixed cost associated with operating a school, and ρ is the constant level of revenue per pupil. This means that the total revenue allocated to any given school depends on the level of enrollment, which is a realistic feature of markets for education (Ladd and Singleton, 2020).

1.2.1 Equilibrium

Following many other studies that utilize circular city-style frameworks (Salop, 1979; Aldashev and Verdier, 2010), I solve this model to find the symmetric equilibrium in which the enrollment at each school in the market will be $\frac{1}{J}$, where J is the total number of competitors. Depending on the relationship between key parameters, there are two important cases to be considered which have very different implications for the effect of competition on test scores: one in which spending on perquisite consumption by school administrators is strictly positive, and one in which there is no scope for perquisite consumption, and in which administrators allocate all revenue toward classroom instruction. The equilibrium level of perquisite consumption M^* is given by the following:

$$M^* = \begin{cases} \frac{\psi\tau}{J^2(\alpha\eta+\psi\nu)} - \frac{1-\theta}{1-\phi} - \zeta & \text{if } \frac{\psi\tau}{J^2(\alpha\eta+\psi\nu)} - \frac{1-\theta}{1-\phi} > \zeta \\ 0 & \text{otherwise} \end{cases} \tag{1.7}$$

When the level of altruism or the fixed cost of school operation is sufficiently high, the school administrator will choose to allocate all revenue net of fixed costs directly to classroom instruction. If this is not the case, than the level of expenditure on perquisite consumption will be given by the first expression above. The equilibrium level of perquisites in this case is strictly decreasing in the number of competitors J. This is because an increase in the number of competitors creates the incentive for administrators to reallocate expenditure from areas that are not valued by households toward classroom instruction in order to compete for students. The two cases above lead to contrasting implications regarding the relationship between competition and instructional expenditure per pupil in equilibrium, which can be seen in the expression below:

$$I^* = \begin{cases} \rho + \frac{1-\theta}{1-\phi}J - \frac{\psi\tau}{J(\alpha\eta+\psi\nu)} & \text{if } M^* > 0 \\ \rho - \zeta J & \text{if } M^* = 0 \end{cases} \tag{1.8}$$

In the case in which perquisite consumption is zero and all revenue net of fixed costs is allocated toward instruction, the equilibrium level of I is clearly decreasing in the number of competitors J.[13] A greater number of competitors decreases equilibrium enrollment at each school, which implies a loss of total revenue. Therefore, net of fixed costs there is less revenue per pupil left over for classroom instruction, which has adverse implications for test scores. In the case in which perquisite consumption is strictly positive, the level of instructional expenditure is strictly increasing in the number of competitors because increased competition creates incentives for administrators to spend less revenue on perquisites that do not benefit households.

Another important implication of this model has to do with the relationship between competition and the provision of consumption amenities by the school, which households value but which are costly in terms of test scores. The equilibrium expression for the level of consumption amenities is as follows:

$$X^* = \begin{cases} \frac{\alpha}{\psi}\rho - \frac{\alpha\tau}{J(\alpha\eta+\psi\nu)} - \frac{(\theta\psi\gamma-\alpha\eta(1-2\theta))}{(1-\phi)\eta\psi}J & \text{if } M^* > 0 \\ \frac{\alpha}{\psi}\rho - \frac{\theta\tau}{(1-\theta)\eta J} + \frac{(\theta\psi\nu-\alpha\eta(1-2\theta))}{(1-\theta)\eta\psi}\zeta J & \text{if } M^* = 0 \end{cases} \tag{1.9}$$

The additional revenue associated with a marginal increase in consumption amenities becomes more valuable when the number of competitors in the market increases because each school enrolls fewer pupils. If the distance cost from the household utility function τ is sufficiently high, then consumption amenities in equilibrium are increasing in J regardless of whether perquisite consumption is strictly positive.[14] The marginal cost of I is increasing in the number of enrolled pupils because this implies that a school will have to increase expenditure on all currently enrolled students, as well as spend revenue on the new students induced to enroll by this change. How-

[13] Throughout this analysis I implicitly assume that ρ is sufficiently high such that instructional expenditure per pupil is strictly positive in equilibrium.

[14] Recent evidence suggests that distance plays an important role in household enrollment decisions (Abdulkadiroglu *et al.*, 2020).

ever, this marginal cost increases faster when τ is sufficiently high, because fewer new students are induced to switch (and bring in additional revenue) by the increase in I when it is relatively less valuable to households. The decrease in the marginal cost of instructional expenditure per pupil associated with a reduction in enrollment becomes larger as the distance cost τ increases, which increases the benefit of additional revenue and generates a greater level of consumption amenities in equilibrium.

Finally, the relationship between test scores and the number of competitors in the market J is summarized by the following expression:

$$Y(I^*, X^*) = \alpha I^* - \psi X^* = \begin{cases} \frac{\theta(\alpha\eta + \psi\nu)}{(1-\phi)\eta} J & \text{if } M^* > 0 \\ \frac{\theta\tau}{(1-\theta)\eta}\left(\frac{\psi}{J} - \frac{(\alpha\eta+\psi\nu)\zeta J}{t}\right) & \text{if } M^* = 0 \end{cases} \tag{1.10}$$

In the case in which perquisite consumption is zero, test scores are strictly decreasing in the number of competitors. The intuition behind this result is fairly straightforward, because in this case instructional expenditure per pupil I^* is decreasing in J while consumption amenities X^* are increasing in J. In contrast, when perquisite consumption is non-zero, the effect of competition on test scores is strictly positive. In this case, even if the provision of consumption amenities increases, the increase in test scores associated with a higher level of per pupil instructional expenditure will more than offset this cost.

In the case in which perquisite consumption is strictly positive, an increase in the number of competitors in the market will clearly improve the welfare of households. This is because competition increases both instructional expenditure per pupil and the level of consumption amenities provided by schools, and also decreases the distance to school for each household. However, the empirical evidence that I document in the context of a large-scale charter sector expansion in North Carolina is ultimately consistent with the case in which there is no scope for perquisite consumption, and both instructional expenditure per pupil and test scores are strictly decreasing in the number of competitors. That being said, a policy which increases the number of competitors in a market for education (such as the removal of a cap on the number of charter schools) may still be welfare-improving for households overall. This is because the cost associated with a decrease in instructional expenditure per pupil may be more than offset by the benefits of greater provision of consumption amenities and shorter distances to school. Typically, the incentives of public school educators are designed entirely around the performance of their students on standardized exams, but the observation that households can potentially be made better off through policies that induce a reduction in test scores raises questions about the fundamental objectives of a public school. In Appendix A, I discuss more explicitly the welfare implications of a policy that removes a cap on the number of schools in a given market for education in the context of an alternative model which is simplified in the sense that it does not allow for perquisite consumption.[15]

1.3 Data

The primary dataset for this analysis focuses on the time period that begins with the 2009-2010 school year, and ends with the 2014-2015 school year. This corresponds to three years prior to the policy change, and three years after the statewide cap on charter schools was lifted. I rely on a large administrative dataset provided by the North Carolina Education Research Data Center (NCERDC) that contains a variety of observable characteristics about students, teachers, and schools across the public education system.[16] In particular, the dataset will consist of standardized test scores

[15]Recall that in the empirical setting discussed in this paper, all evidence is consistent with the implications of the model in the case in which this channel is not important.

[16]This includes charter schools.

and demographic characteristics of traditional public middle school students enrolled in grades 6-8, as well as characteristics of the schools and LEAs in which they were enrolled. Sixth grade is the most common time for students in North Carolina to switch into the charter sector because it is the year in which the majority of students start middle school, and therefore I expect the effect of competition from the charter sector to be especially important at the middle school level.[17]

In order to open a charter school in North Carolina, prospective operators must apply to the State Board of Education, which may then choose to either accept or reject the application. These applications are publicly available on the Department of Public Instruction website, and I have collected all charter applications submitted to the state during my sample period, whether they were ultimately accepted or rejected. Each of the applications contains detailed information about the LEA in which the school plans to open, the grade levels that the school will serve, and the bundle of services that the school will provide to students. Ultimately, I will use these applications to infer information about the LEAs for which they were submitted, which helps me to address the potentially endogenous location decisions of new entrants. Throughout the study, I define a market for education to be an LEA. In other words, I consider traditional public schools to be in direct competition only with charter entrants serving the same grade levels in the same LEA.[18]

[17]While students also switch from middle school to high school at the start of ninth grade, there were fewer charter high school options available during this time period.

[18]One limitation of this approach is that students are technically not restricted to attend charter schools located in the same LEA into which they are zoned. In fact, charter schools are required to accept applications (and offer admission when not oversubscribed) from any student located anywhere in the state. However, consistent with evidence that distance to school is an important determinant of enrollment decisions (Abdulkadiroglu *et al.*, 2020), the vast majority of charter switchers in my sample remain in the same LEA. Therefore, I believe that in this context, this

1.3.1 Student Characteristics

In order to construct the primary dataset used for this analysis, I first restrict the sample of traditional public school students to those enrolled in middle schools that serve grades 6-8. Table 1.1 compares the baseline year (2011-2012) mean characteristics of those enrolled in LEAs for which at least one charter school application was submitted during my sample, to those enrolled LEAs which were not specified on a single application. Clearly, these two types of LEAs differ in a variety of important ways. First, relative to students enrolled in LEAs that received no applications, students in application districts scored significantly higher on prior-year exams in both math and reading. These students are also significantly less likely to be economically disadvantaged (EDS), and significantly more likely to be black.[19] Motivated by these stark differences, I restrict my primary sample to only LEAs in which at least one new sixth grade charter school applied to open.[20] The primary empirical results rely on comparisons between districts that actually had new charter schools open to districts for which no applications for new charter schools had yet been approved.[21]

Among all LEAs included in my primary sample, Table 1.2 displays mean student characteristics before and after any new middle school opened.[22] For the most part,

---

assumption is reasonable.

[19]The data center defines economic disadvantage as having household income below 185% of the federal poverty line.

[20]I drop a few additional LEAs that had charter schools open during the pre-expansionary period of my sample, which leaves 38 in total.

[21]It is also worth noting that while the LEAs in Wake County and Mecklenburg County are substantially larger in terms of total enrollment than any other LEA in the state, and thus could potentially differ from other LEAs in a variety of important ways, all results in this study are robust to excluding students from these LEAs

[22]All observations that correspond to students in the never-treated LEAs in my sample are included in the "Not Treated" column.

Table 1.1: Summary Statistics (2011-2012)

	Application	No Application	Difference
Previous Year Math Score	0.07	0.03	0.04***
Previous Year Reading Score	0.06	0.03	0.02***
Female	0.50	0.50	0.00
EDS	0.50	0.57	-0.07***
Black	0.29	0.17	0.12***
Hispanic	0.12	0.13	-0.00
N	174,135	50,615	

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of mean characteristics for North Carolina public middle school students in the year prior to the charter expansion, which was the 2011-12 school year. The first column corresponds to LEAs in which at least one charter school offering sixth grade applied to open between 2012-13 and 2014-15, while the second column corresponds to LEAs for which this was not the case. The third column displays the result from a standard differences-in-means test.

it is clear that new charter entry did not induce major changes in the composition of middle school students in these LEAs. One exception is the previous-year math test score, which actually increased on average after the entry of a charter competitor.[23] This increase relative to the average student in the state is likely driven by differences between the application and non-application districts during my sample period. If anything, the fact that previous-year test scores are not decreasing relative to the statewide average after charter entry somewhat mitigates concerns about bias from a cream-skimming effect, although I will still be explicit about how I address potential

[23] Note that this is the previous-year exam score, which will be used as a control in the primary analysis but is not itself an outcome of interest.

selection bias in Section 1.4. The last row of the table represents county-level private school enrollment collected from the Private School Universe survey administered by the NCES.[24] On average, private school enrollment decreases after charter entry, and this decrease is not statistically significant, which suggests that my primary results are not being conflated with growth in the private school sector.

Table 1.2: Traditional Public Characteristics by Treatment Status

	Not Treated	Treated	Difference
Previous Year Math Score	0.07	0.12	0.05***
Previous Year Reading Score	0.07	0.08	0.01***
Female	0.50	0.50	-0.00
EDS	0.48	0.47	-0.01***
Black	0.27	0.30	0.03***
Hispanic	0.12	0.16	0.04***
Private School Enrollment	2097.66	1536.88	-560.79
N	668,490	282,324	

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of mean characteristics for North Carolina middle school students enrolled in LEAs in which at least one charter school applied to open during my sample period. The first column corresponds to years prior to any new charter middle school opening, while the second column corresponds to years after the first new middle school entered the market. The third column displays the result from a standard differences-in-means test.

[24]Note that this survey is administered biennially, so this data only reflects the mean across years that end in an even number. Further, the vast majority of LEA boundaries in North Carolina are drawn exactly along county lines, so this can effectively be thought of as LEA-level private school enrollment.

1.4 Empirical Strategy

The primary results of this study are estimated using a two-way fixed effects model, a commonly utilized extension of a difference-in-differences framework to settings in which treatment timing is staggered. In other words, these models are often used when a researcher might otherwise employ a standard difference-in-differences design, but when there exist multiple treatment cohorts. Goodman-Bacon (2019) shows that under assumptions of common counterfactual trends across timing groups and no dynamic treatment effects, a TWFE strategy will identify a variance-weighted average of all possible comparisons across treatment cohorts and time. The regression specification used to generate the primary results is as follows:

$$y_{igjdt} = \beta_0 + \beta_1 Z_{dt} + \beta_2 X_{igjdt} + f(App_{dt}) + \gamma_j + \delta_t + \epsilon_{ijgdt} \qquad (1.11)$$

In this equation, y_{ijgdt} represents an outcome of interest such as a test score, for student i enrolled in grade level g school j in district d during year t. The right-hand side variable of interest, Z_{dt}, is an indicator which is equal to one if a new charter school offering sixth grade has opened in the LEA to which a given student is assigned. This effectively represents whether the incumbent public schools in the district have been exposed to competition from expansionary charters for the first time.[25] Therefore, the coefficient β_1 is interpreted as the effect of charter school competition. The regression also includes a set of student, school, and LEA-level controls X_{ijgdt}. γ_j and δ_t represent school and year fixed effects, respectively, and all standard errors are clustered at the LEA level.

[25] In Section 1.5.1, I also present results using an alternative strategy in which the treatment variable is the charter share of enrollment among all publicly funded schools, which allows me to consider whether the effects of competition are larger in LEAs in which a greater share of students switched to the charter sector.

The term $f(App_{dt})$ represents a function of the cumulative number of sixth grade charter school applications that have specified LEA d since the beginning of the expansion.[26] One of the primary challenges facing any study that estimates spillover effects of school competition involves the potentially endogenous location choices of competitors. In order to address this issue, I use data on all applications that were submitted to the state-level authorizer, whether they were ultimately accepted or rejected. When available, I have also collected decision rubrics for each application that contain more detailed information about the reasons why reviewers made a particular decision.

In cases where applications are rejected, reviewers nearly always cite reasons such as an unsustainable financial plan or incomplete information in a given section of the application. In contrast, these rubrics rarely cite reasons for rejection related to characteristics of the specific LEA in which the charter school had applied to open. In other words, while it seems likely that charter schools would exhibit strategic behavior when choosing a particular LEA for which to apply (Singleton, 2019), the accept/reject decision made by the state-level authorizer on a given application does not appear to be related to particular characteristics that might affect test score performance across markets. In further support of this point, I observe a number of cases in which a prospective charter school application was initially rejected, but in which the applicant then incorporated the feedback into a new application and ultimately received approval at a later date, suggesting that the reasons for rejection

[26]In my preferred specification this function is fully flexible in that it is a set of indicators that represent any possible number of applications that a district could have been specified on. It is also worth noting that these results are robust to specifications which consider different groups of applications based on observable characteristics as opposed to considering all applications the same (e.g. whether or not a given application expresses the intent to participate in school nutrition programs or offer bus transportation).

cited in the decision rubrics are actually the relevant ones.[27] I argue that in this case it is reasonable to assume that the accept/reject decision made by the state-level authorizer is exogenous to the unobservable characteristics of an LEA that are related to student achievement, and therefore I will compare students in LEAs that generated the same amount of interest from prospective charter school operators, but which differ in whether a new school has actually opened. Under this assumption, the effect of actual competition from charter schools can be identified by conditioning on proposed competition.[28]

In order to provide further support for this assumption, Table 1.3 presents results from estimating the following application-level regression specification:

$$Y_{kdt} = \beta_0 + \beta_1 Accepted_{kdt} + f(App_{kdt}) + \omega_d + \delta_t + \epsilon_{kdt} \tag{1.12}$$

In this equation, Y_{kdt} represents a given characteristic of the district d which was specified on charter school application k submitted in year t, and the primary covariate of interest is an indicator variable that represents whether or not the application was accepted by the state-level authorizer. The specification also includes controls for the total number of applications that had specified a given district, as well as district and year fixed effects.[29] The purpose of this exercise is to demonstrate that condi-

[27]Specifically, out of twelve cases where I see prospective schools that were initially rejected but reapplied to open in the same LEA during the expansionary portion of my sample, six of them were ultimately approved.

[28]This bears some similarity to the identification strategy utilized in Slungaard Mumma (2020), which compares traditional public schools that are located in a very small distance band around actual charters to schools located just as close to other locations where the same charter considered opening (which tend to be in close proximity). A key distinction between this strategy and the one I employ in this study is that I use information on the market location choices of all proposed charter schools, whether they ultimately opened or not.

[29]The bins that represent the number of applications that have specified a given LEA are less

tional on the application controls and fixed effects, whether or not a given application is approved by the state advisory board does not have the ability to predict any observable characteristics of the LEA specified on the application. In Table 1.3, we see that none of the displayed coefficients on the indicator that represents acceptance are close to statistical significance, which is suggestive that the acceptance decision for a charter school application conditional on these controls is not importantly related to the characteristics of the proposed LEA. The dependent variable in the final column represents whether or not the state had already allowed another charter school to enter the same LEA since the beginning of the expansion, and the sample size is slightly smaller because it does not include the applications submitted during the first year.

Table 1.3: Application Status and LEA Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Mean Lagged Math	Mean Lagged Reading	Pct. Black	Pct. Hisp.	Pct. EDS	Prior Entrant
Accepted	0.003	0.005	0.000	-0.001	0.004	-0.088
	(0.005)	(0.004)	(0.001)	(0.001)	(0.009)	(0.070)
R-squared	0.990	0.996	0.999	0.998	0.948	0.910
N	118	118	118	118	118	94

* p<0.10, ** p<0.05, *** p<0.010

This table explores the extent to which an indicator variable that represents whether or not an application was accepted has the ability to predict LEA characteristics conditional on fixed effects that represent LEA, year, and number of applications received. The dependent variable in each column corresponds to a different characteristic. The sample includes all charter applications that indicated an intent to offer sixth grade. "Prior Entrant" is an indicator variable which takes on the value of one if any charter school had already entered the same LEA since the start of the expansion, and this column has fewer observations because it does not include the applications from the first year as no expansionary charter could have opened prior to this point in time. Standard errors are clustered at the LEA level.

Another key empirical challenge in estimating spillover effects of school competition involves student selection behavior across schools. For example, if students that leave a traditional public school for a charter sector alternative possess higher average academic ability than the students that remain in the traditional public sector, this

disaggregated in this specification relative to the one used to generate student-level results due to the small number of applications submitted for some LEA-years.

could potentially bias the estimated spillover effect of charter competition downward. One important way that I address student-level heterogeneity in ability is by including previous-year standardized math and reading end-of-grade (EOG) test scores in all student-level regression specifications. However, I also go a step further by replicating all primary results using only the subsample of students that had already completed sixth grade in a traditional public middle school prior to the opening of any new charter school offering sixth grade in their LEA.[30] This strategy is motivated by the fact that in the North Carolina context, sixth grade is the most common time for a student to switch into the charter sector. During the expansionary period of this sample, 4.1% of students switched from a TPS to the charter sector for sixth grade. This stands in stark contrast to the sample of seventh and eighth graders that had already completed at least a year of traditional public school, in which less than 1% of students switched. Therefore, this subsample is much less likely to exhibit selection behavior because there was no new charter option available in their LEA prior to the completion of sixth grade. If results are qualitatively similar whether or not I restrict the sample in this manner, it suggests that these results are being driven by an actual public school response to charter competition rather than selection alone.

The identification strategy that underlies these results relies importantly on the assumption of parallel counterfactual trends across different treatment cohorts. In other words, it must be the case that trends in test score performance across treatment cohorts and the never-treated group would have been similar in the absence of competition. I explore the plausibility of this assumption in the context of an event study framework, and generate Figure 1.2 using estimates from the following

[30]In other words, this means that I exclude all sixth grade students, as well as any seventh or eighth graders that had the option in sixth grade to attend an expansionary charter school in their LEA.

regression specification:

$$y_{ijgdt} = \beta_0 + \sum_{l=-4}^{-2} \xi_l D_{it}^l + \sum_{l=0}^{2} \xi_l D_{it}^l + \beta_1 X_{ijgdt} + f(App_{dt}) + \gamma_j + \delta_t + \epsilon_{ijgdt} \quad (1.13)$$

In this case, y_{ijgdt} represents the standardized math test score. The model is similar to Equation 1.11, except that instead of a single treatment indicator, the treatment is now represented by a set of relative period indicators D_{it}^l, where l represents the number of periods since treatment. These variables are equal to one if unit i is l periods away from treatment at time t, and zero otherwise. A common test of the parallel trend assumption involves testing whether ξ_l is statistically distinguishable from zero for $l < -1$.[31] While Sun and Abraham (2021) cautions that this test may not always be valid in the presence of dynamic treatment effects, I interpret the relatively flat line prior to the omitted baseline period in this figure as supportive of the assumption of parallel counterfactual trends across cohorts.

1.5 Results

The simple model presented in Section 1.2 predicts that when perquisite consumption by school administrators is strictly positive, test scores will be increasing in the number of competitors serving a given market. However, when there is no scope for perquisite consumption, an increase in school competition leads to a decrease in academic achievement both because more revenue per pupil must be used to cover fixed costs and because of the opportunity cost associated with increased provision of consumption amenities. The purpose of this section is to document the average effect of competition on test scores in the context of the large-scale North Carolina charter sector expansion.

[31] See He and Wang (2017) for an example.

Figure 1.2: Math Test Event Study

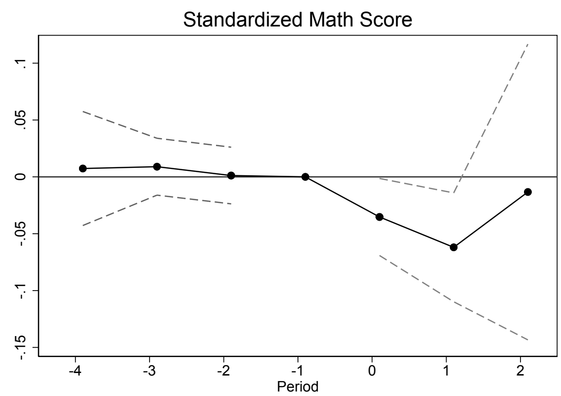

This figure displays a set of coefficients on relative period indicators from an event study framework in which the dependent variable is the standardized EOG math test score. Period 0 represents the first period in which a new charter school has opened in the district of a given student, and the coefficient on period -1 is normalized to zero. The dashed lines represent 95% confidence intervals.

Table 1.4 displays results from the TWFE model described in Section 1.4 where the dependent variable is the standardized end-of-grade (EOG) math test score.[32] The coefficient of interest is on the indicator for whether any new charter school offering sixth grade has opened in the LEA to which a given student is assigned. The regression model to which the first column corresponds contains only the application controls, school and year fixed effects, and previous year math and reading EOG test scores in order to address heterogeneity in academic ability across students. In this column, we see that the coefficient on the competition indicator is negative and statistically significant.

In the second column, which represents the preferred specification, I include a large set of controls at the individual, school, and LEA levels. The quantitative

[32]Only math test results are included in this section. For analogous results using the reading EOG test score as the dependent variable, please see Appendix C. Reading results display very similar qualitative patterns, but are almost always smaller in magnitude.

Table 1.4: Standardized Math Test Score

	(1)	(2)	(3)
New Entrant	-0.040**	-0.043**	-0.031*
	(0.017)	(0.017)	(0.018)
Other Controls	No	Yes	Yes
7th/8th Grade Only	No	No	Yes
R-squared	0.746	0.748	0.750
N	950,814	950,814	575,376

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized math end-of-grade test score. The first column corresponds to a regression specification which includes prior year test scores in math and reading, application controls, and school and year fixed effects. The second column augments this specification wth a large vector of other controls, including gender, race, economic disadvantage status (EDS), grade level, school share black, school share Hispanic, school share EDS, school enrollment, and LEA enrollment. The third column replicates the result from the second column using only the subsample of students that had already begun seventh or eighth grade prior to the opening of a new charter middle school in their LEA, and thus are much less likely to exhibit selection behavior across schools. Standard errors are clustered at the LEA level.

interpretation of the coefficient in this column is that the entry of a competing charter school in the LEA to which a given traditional public middle school student is assigned is associated with an average decrease of 4.3% of a standard deviation in the math EOG exam score. To put this result in perspective, the literature on teacher value added suggests that this is comparable to roughly a quarter of a standard deviation decrease in math teacher ability (Chetty *et al.*, 2014a). In the final column of Table 1.4, I reproduce the results from the preceding column using only the sample of seventh and eighth grade students that had already completed at least a year of a traditional public middle school prior to any new charter school opening in their LEA, which substantially reduces the sample size. Nonetheless, we see that the estimated effect of competition is negative, statistically significant, and qualitatively similar to the full-sample result, suggesting that the primary result is not being driven primarily by student selection.

In summary, consistent with the case of my model in which administrators are not spending school revenue on perquisite consumption, I find that an increase in the degree of competition from charter schools caused a decline in public school test scores. In Appendix D, I further restrict the sample to students that I can match to a particular classroom, and then to the students assigned to teachers for whom I can estimate an out-of-sample measure of value added. This is done in order to demonstrate that the estimated effect of competition is robust to controlling for class size, peer ability, and teacher value added. In Appendix E, I use the conventional method of estimating school value added described in Angrist *et al.* (2017) to show that results are qualitatively consistent when they are estimated at the school level. In Appendix F, I also consider the effect of charter competition on measures of student discipline, in particular changes in the probability or length of suspensions for disciplinary infractions, and I find evidence that traditional public middle schools increased the length of suspensions in response to competition from the charter sector.

1.5.1 Alternative Strategies

A number of recent studies have pointed out potential issues with interpreting an estimate from a TWFE model as a causal parameter in the presence of dynamic treatment effects.[33] Recall that under certain assumptions, the coefficient on the

[33] Examples include Borusyak and Jaravel (2018), Athey and Imbens (2018), Goodman-Bacon (2019), Callaway and Sant'Anna (2019), de Chaisemartin and D'Haultfoeuille (2020), and Sun and Abraham (2021).

treatment indicator in the TWFE model is equal to a weighted average of treatment effects across all two-by-two treatment cohort and time period comparisons (Goodman-Bacon, 2019). Basically, some of these comparisons use already-treated units as the control group during time periods in which they are not actually changing treatment status, and if the treatment effect is not constant over time this could introduce bias into the estimate. Therefore, I complement the primary results using a stacked difference-in-differences approach that mitigates this potential source of bias by facilitating comparisons between treated and not-yet-treated units.

Another potential limitation of the results presented in the previous section is that in estimating the average effect of competition with a single treatment indicator, the strategy does not account for the fact that the effect of competition with new charter schools may depend on the number of charter switchers relative to the total size of the traditional public LEA. In order to address this, I utilize an instrumented difference-in-differences (DDIV) strategy, which corresponds to an instrumental variable framework where I use an indicator which takes a value of one if any new charter middle school has opened in a given LEA as an instrument for the share of students in all publicly-funded schools that are enrolled in the charter sector.[34] In other words, the first stage corresponds to a TWFE regression in which the dependent variable is the charter share.

Stacked Difference-in-Differences

In order to address the issue of bias introduced by dynamic treatment effects, an alternative to a standard TWFE model is a stacked difference-in-differences approach that facilitates comparisons between groups that change treatment status and never or not-yet-treated units. While this procedure has the attractive feature of removing

[34] This is similar to the strategy used in a similar context in Ridley and Terrier (2018).

a potentially important source of bias from the estimate, it comes at the cost of interpretation. Using this strategy, the researcher can only identify a very short-term effect of treatment, because control units may become treated eventually. I follow an approach similar to Deshpande and Li (2019), and create separate datasets for each of the three treatment cohorts.[35] Each dataset includes all LEAs treated in a given year, which are assigned to the treatment group, and a control group of all LEAs which were specified on applications during the expansionary period, but which either did not, or had not yet had a new charter middle school actually open. All observations in a given dataset are assigned the timing of the treatment cohort (e.g. in the dataset corresponding to the 2013 treatment cohort, both the treatment and the control group are assigned a treatment date of 2013). The final sample is created by appending these three datasets together. I then estimate the following regression model:

$$y_{ijgdt} = \beta_0 + \beta_1 G_d + \beta_2 post_t + \beta_3 G_d * post_t + \beta_4 X_{ijgdt} + f(App_{dt}) + \gamma_j + \delta_t + \epsilon_{ijgdt} \quad (1.14)$$

In this case, G_d is an indicator for whether a given LEA is in the treatment group, and $post_t$ is an indicator for whether the observation belongs to the period after treatment. The regressor of interest is the interaction between these two terms. These regressions will also include a set of student, school, and LEA-level controls X_{ijgdt}, and a function $f(App_{dt})$ of the cumulative number of applications that had specified district d since the beginning of the expansion. Note that the treatment group indicator is identified separately from the school fixed effects, because the same school can appear multiple times as both a treated unit and a control unit for

[35] In this case a "cohort" refers to the timing of the treatment. The three groups correspond to the LEAs in which the first new charter school offering sixth grade opened in 2013, 2014, or 2015.

another treatment cohort. Table 1.5 displays results generated from the regression model described by equation 1.14, where the dependent variable is the standardized EOG math test score. This coefficient is negative and statistically significant, which is consistent with the TWFE results, although it is somewhat smaller in magnitude. Recall that relative to the TWFE results, this estimate may potentially suffer less from bias introduced by dynamic treatment effects, but it can only be interpreted as a very short-term effect of competition.

Table 1.5: Standardized Math Test Score: Stacked Difference-in-Differences

	(1)
Treatment Group*Post	-0.022*
	(0.011)
R-squared	0.743
N	2,011,364

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized math end-of-grade test score. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative applications received. Standard errors are clustered at the LEA level.

Instrumented Difference-in-Differences

The instrumented difference-in-differences strategy will identify the local average treatment effect (LATE) of an increase in the charter share associated with the opening of a new competitor under parallel counterfactual trend assumptions on both the treatment (the charter share) and any outcomes of interest (de Chaisemartin, 2010; de Chaisemartin and D'Haultfoeuille, 2018). While this assumption on the EOG

Figure 1.3: Charter Enrollment Event Study

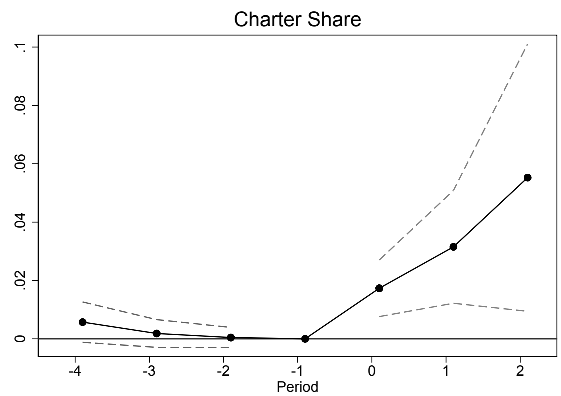

This figure displays a set of coefficients on relative period indicators from an event study framework in which the dependent variable is the share of students in publicly-funded schools in a given LEA that are enrolled in the charter sector. Period 0 represents the first period in which a new charter school has opened in the district of a given student, and the coefficient on period -1 is normalized to zero. The dashed lines represent 95% confidence intervals.

math test score was discussed in Section 1.4, Figure 1.3 explores the plausibility of this assumption on the charter share. I interpret the relatively flat line at zero prior to treatment in this event study generated from an LEA-level sample as evidence that there were no differential pre-trends in the charter share across LEAs included in my primary sample. In order to implement the DDIV strategy, I use two-stage least squares to estimate the following:

$$y_{igjdt} = \beta_0 + \beta_1 CS_{dt} + \beta_2 X_{ijgdt} + f(App_{dt}) + \gamma_j + \delta_t + \epsilon_{igjdt} \tag{1.15}$$

In this equation, y_{ijgdt} corresponds to the standardized math test score. The variable CS_{dt} represents the sixth grade charter share in district d at time t, for which we will use an indicator that represents entry of a new charter school in my primary sample as an instrument. Consistent with the TWFE results, the coefficient on the charter share displayed in Table 1.6 is negative and statistically significant. For the sake of interpretation, the charter share in treated LEAs during the expansionary period of my sample increased by an average of roughly 0.0204. Along with the coefficient in Table 1.6, this implies that the average effect of competition on math scores in treated LEAs was a decrease of roughly 5.68% of a standard deviation, which is qualitatively similar to the TWFE result. Overall, this result suggests that the effect of competition from new charter entrants was greater in areas in which new entrants induced a larger shift in students relative to the size of the LEA.

Table 1.6: Standardized Math Test Score: Instrumented Difference-in-Differences

Charter Share	-2.783*
	(1.475)
N	950,814

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized math end-of-grade test score. The treatment variable of interest is the charter share of an LEA, for which we use a treatment indicator as an instrumental variable. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative applications received. Standard errors are clustered at the LEA level.

1.5.2 Charter Competition by Governance

Relative to their traditional public counterparts, charter schools have considerable autonomy over many important school characteristics. One important source of heterogeneity within the charter sector has to do with governance, or how the school is managed. In this context, charter schools fall broadly into three categories: those run independently by members of the local community, those managed by non-profit charter management organizations (CMOs) which are typically guided by organizational mission statements, and those managed by for-profit education management organizations (EMOs). This third group is of particular interest because these for-profit managers potentially have very different incentives than their non-profit counterparts. On the applications submitted to the state during my sample period, applicants are directly asked whether they intend to contract with a for-profit EMO, and I consider those that answer in the affirmative to be "for-profit" for the sake of this analysis.[36]

Table 1.7: Charter Application Characteristics

	Total	For-Profit	Not For-Profit	Difference
Share Accepted	0.38	0.40	0.37	0.03
Participating in School Nutrition Programs	0.54	0.83	0.49	0.34***
Offering Bus Transportation	0.42	0.16	0.46	-0.31**
Offering Innovative Curriculum	0.25	0.20	0.26	-0.06
Projected Year 1 Enrollment	300.79	490.85	262.00	228.85***
N	119	20	99	

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of mean characteristics for charter school applications relevant to the primary analysis. The first column corresponds to all applications, while the second and third columns correspond to non-profit and for-profit applications, respectively. The fourth column displays the results of a standard difference-in-means test between the two groups.

Table 1.7 summarizes descriptive statistics for the charter school applications relevant for my primary sample of students, and also compares characteristics between the for-profit schools and their non-profit counterparts. The probability of actually receiving approval to open is similar across the two types of schools. However, relative to non-profits, the for-profit schools are more likely to state that they will participate in school nutrition programs, but are significantly less likely to state that they

[36]Note that even schools with for-profit managers are still non-profit organizations themselves, it is not the case that these managers can keep state funding intended for education as profit.

will offer bus transportation to school.[37] I consider applications that state the intent to offer a particular curricular focus (e.g. STEM or performing arts) as "offering an innovative curriculum," and while non-profit applications are slightly more likely to be of this type, this difference is not statistically significant. Finally, relative to non-profits, for-profit charter schools plan to enroll significantly more students during their first year of operation.

In this section, I explore the extent to which the spillover effects of charter school competition on TPS test scores depend on whether or not incumbent schools are competing with charters managed by for-profit EMOs. I argue that for-profit charter schools may be more willing to compete for students through the consumption amenities channel due to differences in objectives between for-profit managers and their non-profit counterparts. For example, EMOs may have less incentive to improve academic achievement beyond the accountability standards set by the state relative to non-profits governed by mission statements that involve improving the academic performance of students as much as possible. In order to provide support for this statement, I consider information from the first year of the proposed budget plan for all charter applications offering sixth grade that were submitted to the state during my sample period.[38] Table 1.8 displays results from the following regression specification:

$$y_{kdt} = \beta_0 + \beta_1 ForProfit_{kdt} + \beta_2 Accepted_{kdt} + \beta_3 TotExp_{kdt} + \beta_4 ProjEnroll_{kdt} + \omega_d + \delta_t + \epsilon_{kdt} \tag{1.16}$$

In this specification, y_{kdt}, represents an inflation-adjusted measure of proposed

[37]In fact none of the for-profit schools that actually opened in this sample offer bus transportation.

[38]Note that because this information comes directly from the charter school applications, this sample only includes potential charter schools, and not traditional public schools.

expenditure for application k to open a school in district d that was submitted in year t. The primary regressor of interest is $ForProfit_{kdt}$, which is an indicator variable equal to one if the application expresses the intent to contract with a for-profit EMO. I also control for whether the application was ultimately accepted, the total levels of proposed expenditure and projected enrollment in the first year of operation, as well as proposed LEA and year effects.

In the first column of Table 1.8, the dependent variable is the proposed level of expenditure on marketing. I focus on this category in particular because this is a clear example of expenditure that is designed to attract households, but that does not contribute to the production of education in the classroom. As expected, I find that for-profit schools plan to spend substantially more on marketing than their non-profit counterparts. In the second column, the dependent variable of interest is total inflation-adjusted expenditure on all budget items not directly related to classroom instruction (excluding marketing). Consistent with evidence from the state of Florida (Singleton, 2017), I find that conditional on total proposed expenditure, schools managed by for-profit EMOs allocate a larger amount of their budget to non-instruction. The model outlined in Section 1.2 predicts that test scores will be lower in markets where the level of altruism is relatively low, such as potentially in markets served by for-profit competitors. If traditional public schools respond to for-profit charter competition by increasing their own provision of non-academic services, then markets subject to this type of competition might be where we would expect to find the largest negative spillover effects of competition on traditional public school test scores.

Table 1.9 summarizes differences between for-profit and non-profit charter school students. In this context, students attending schools managed by for-profit EMOs score substantially higher on lagged math and reading tests relative to both non-

Table 1.8: Projected Charter Expenditure (100s of Dollars)

	(1)	(2)
	Marketing	Total Non-Instruction
For-Profit	236.907***	1566.645**
	(76.880)	(736.541)
Accepted	69.587**	935.638**
	(29.810)	(379.991)
Projected Total Expenditure	0.019***	0.845***
	(0.005)	(0.067)
R-squared	0.750	0.976
N	112	112

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variables are inflation-adjusted categories of proposed charter school expenditure in the first year of operation, which come from the applications submitted to the state. The first column corresponds to expenditure on marketing, while the second column corresponds to total non-instructional expenditure (excluding marketing). Both specifications include controls for whether the application was accepted and the total level of proposed expenditure in the first year, as well as projected enrollment and both LEA and year fixed effects. Standard errors are clustered at the LEA level.

profit charter students as well as the average traditional public schools student in my sample. Further, for-profit charter students are much less likely to be racial minorities or economically disadvantaged relative to non-profit students, although these non-profit students are still less likely to be racial minorities or economically disadvantaged relative to the traditional public students in my sample. These patterns are consistent with other studies in the North Carolina context which find that charter entry has contributed to increased racial segregation in schools across the state (Bifulco and Ladd, 2007; Ladd *et al.*, 2017; Ladd *et al.*, 2018).

Table 1.9: Charter Statistics by Governance

	TPS	For-Profit	Not For-Profit	Difference
Previous Year Math Score	0.08	0.36	0.04	0.32***
Previous Year Reading Score	0.06	0.37	0.11	0.26***
Female	0.50	0.51	0.49	0.03
EDS	0.49	0.11	0.25	-0.15***
Black	0.28	0.12	0.22	-0.11***
Hispanic	0.14	0.07	0.04	0.03***
N	480,041	954	3,343	

* $p<0.10$, ** $p<0.05$, *** $p<0.010$

This table displays a set of mean characteristics for students in grades 6-8 that attend different types of North Carolina schools. The first column displays a set of characteristics corresponding to the traditional public students in my sample from 2012-2013 through 2014-2015. The second column corresponds to charter schools that indicated an intent on their application to contract with a for-profit management organization, while the third column corresponds to charter schools that did not. The fourth column displays the result from a standard differences-in-means test between for-profit and non-profit charter school students.

While the statistics displayed in Table 1.9 may suggest that charters contracting with for-profit managers attempt to locate in very different areas than their non-profit counterparts, Table 1.10 displays the results of an exercise to mitigate concerns about selection that this might raise. Similar to Table 1.3, the estimates displayed in this table demonstrate that conditional on controls for number of applications and both LEA and year effects, an indicator variable that represents whether a charter application expressed the intent to contract with a for-profit manager does not have the ability to predict whether the application was accepted or any of a set of LEA-level characteristics.

Table 1.10: Application and LEA Characteristics by Governance

	(1) Accepted	(2) Lagged Math Mean	(3) Lagged Reading Mean	(4) Pct. Black	(5) Pct. Hisp.	(6) Pct. EDS	(7) Prior Entrant
For-Profit	0.025	0.002	0.004	-0.000	0.000	-0.003	-0.103
	(0.193)	(0.003)	(0.003)	(0.001)	(0.001)	(0.005)	(0.080)
R-squared	0.019	0.990	0.996	0.999	0.998	0.947	0.909
N	118	118	118	118	118	118	94

* p<0.10, ** p<0.05, *** p<0.010

This table explores the extent to which an indicator variable that represents whether or not a charter application expressed the intent to contract with a for-profit management organization has the ability to predict whether the application was accepted or LEA characteristics conditional on fixed effects that represent LEA, year, and number of applications received. The dependent variable in each column corresponds to a different characteristic. The sample includes all charter applications that indicated an intent to offer sixth grade. "Prior Entrant" is an indicator variable which takes on the value of one if any charter school had already entered the same LEA since the start of the expansion, and this column has fewer observations because it does not include the applications from the first year as no expansionary charter could have opened prior to this point in time. Standard errors are clustered at the LEA level.

Given additional concerns about student selection raised by the statistics displayed in Table 1.9, I will also compare the academic spillover effects of competition with for-profit charters to those associated with another type of charter school that displays a very similar pattern of student selection: schools with a STEM focus.[39] STEM charters enroll similar students to for-profit charters, but offer a very different and potentially more academically rigorous product with which incumbent public schools must compete. Therefore, if I were to find much larger negative spillover effects from competition with for-profit charters, this would suggest that the results are being driven by a response to the types of services with which traditional public schools are competing as opposed to selection behavior alone.

Table 1.11 compares mean characteristics of students enrolled in charter schools with a STEM focus to students enrolled in charters that do not have a STEM focus. The first column corresponds to the primary sample of traditional public school students, the second column corresponds to STEM schools, the third column corresponds

[39]No charter schools in my sample are both STEM-focused and managed by a for-profit organization.

to non-STEM schools, and the fourth column to a standard difference-in-means test. The comparison between these two groups of students look remarkably similar to those seen in Table 1.9, which compared for-profit and non-profit charter school students. STEM charter students have significantly higher lagged test scores than their non-STEM counterparts, although the difference in reading test scores is not quite as large as in Table 1.9. STEM students are also significantly less likely to be minorities or economically disadvantaged.

Table 1.11: Charter Statistics by STEM Focus

	TPS	STEM	Non-STEM	Difference
Previous Year Math Score	0.08	0.38	0.05	0.34***
Previous Year Reading Score	0.06	0.32	0.13	0.20***
Female	0.50	0.46	0.50	-0.04*
EDS	0.49	0.11	0.25	-0.14***
Black	0.28	0.08	0.23	-0.15***
Hispanic	0.14	0.01	0.06	-0.04***
N	480,041	837	3,460	

* $p<0.10$, ** $p<0.05$, *** $p<0.010$

This table displays a set of mean characteristics for students in grades 6-8 that attend different types of North Carolina schools. The first column displays a set of characteristics corresponding to the traditional public students in my sample from 2012-2013 through 2014-2015. The second column corresponds to charter schools that indicated an intent on their application to focus on a STEM curriculum, while the third column corresponds to charter schools that did not. The fourth column displays the result from a standard differences-in-means test between STEM and non-STEM charter school students.

The regression specification used to generate the results in Table 1.12 is very similar to the specification discussed in Section 1.4 except in that I include an additional variable which represents the share of all sixth grade charter switchers enrolled in schools that have a particular characteristic.[40] The first two columns correspond to for-profit competition, while the second two columns correspond to competition with STEM charters. In the first column, the coefficient on new charter entry is now interpreted as the effect of competition in LEAs in which no sixth-grade students are switching into for-profit charters, and the effect for a given LEA is represented by this coefficient plus the properly scaled coefficient on the for-profit share. More specifically, the interpretation of the coefficients in this column is that in LEAs that do not compete with a for-profit charter, the average effect of new entry on math test scores is a decrease of 2% of a standard deviation, and is not statistically significant. However, in LEAs in which all switchers to new charter schools attend those managed by for-profit EMOs (a value that actually occurs in the data), new charter entry is associated with a decrease of 12% of a standard deviation, which is substantially larger than the average effect estimated in the previous section. In other words, the negative effect of charter school competition on traditional public test scores is larger when incumbent schools are competing with charters managed by for-profit EMOs.

The second column replicates the results from the first using only the sample of seventh and eighth-grade students that had already completed at least a year of traditional public middle school prior to any new charter school opening in their LEA. While the coefficient on the for-profit share in this column is smaller in magnitude relative to the first, the fact that it is still large, negative, and statistically significant further supports the idea that charter governance is an important source of heterogeneity in driving spillover effects of competition. In contrast, the correspond-

[40] I also include separate sets of controls that represent groups of applications based on whether or not they exhibit this characteristic (e.g. for-profit and non-profit), as opposed to one set of controls that treats all applications the same.

Table 1.12: Standardized Math Test Score by Charter Type

	For-Profit		STEM	
	(1)	(2)	(3)	(4)
6th Grade For-Profit Charter Share	-0.095***	-0.070*		
	(0.031)	(0.037)		
6th Grade STEM Charter Share			0.023	0.025
			(0.048)	(0.048)
New Entrant	-0.025	-0.013	-0.022	-0.011
	(0.020)	(0.019)	(0.018)	(0.019)
R-squared	0.748	0.750	0.748	0.750
N	950,814	575,376	950,814	575,376

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized math end-of-grade test score. The covariate of interest is the share of 6th grade students enrolled in charter schools that opened after the policy change attending schools that indicated the intent to contract with a for-profit EMO, or to offer a STEM focus. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, and LEA enrollment. Each specification also includes cumulative applications received for each charter type (i.e. separate controls for for-profit and not-for-profit applications in columns 1 and 2). The second column of each panel reports results when the sample is restricted to the seventh and eighth grade students that attended sixth grade prior at a traditional public middle school prior to any new charter school opening in their LEA. Standard errors are clustered at the LEA level.

ing coefficients in columns three and four for STEM competition are positive and not statistically significant, despite similar selection patterns across the two types of schools. This demonstrates that the larger negative effect on test scores associated with for-profit charters is driven by a response of traditional public schools to the types of services that they offer, as opposed to selection behavior alone.

1.6 Mechanisms

This section summarizes evidence regarding the mechanisms underlying the negative spillover effect of competition from the charter sector on the academic achievement of traditional public middle school students. This includes evidence of the shift of influence from teachers to households at the middle school level, as well as a fiscal effect that occurs at the LEA level.

1.6.1 School Characteristics

Beginning in 2002, North Carolina has administered a biennial working conditions survey to educators across the state.[41] In order to explore whether charter competition is associated with an increase in the provision of non-academic services in incumbent public middle schools, I construct two measures using teacher responses to this survey. First, I use principal component analysis to generate an index of teacher empowerment from a set of survey items. For each of these items, teachers assign a score of one to four depending on how much of a role they feel that they play in some aspect of school management, such as selecting instructional materials and resources. These scores are averaged at the school level for each item, and the empowerment index is

[41] Due to the biennial nature of the survey, this information is only available for years that end in an even number.

the first principal component of these school-level measures.[42]

I also create a measure which corresponds to the share of teachers within a school that either agree or strongly agree with the statement "Parents/guardians are influential decision makers in this school." Previous literature has documented the fact that post-secondary institutions cater to the consumption value of households (Jacob *et al.*, 2018), but little evidence of this behavior has been documented at the primary or secondary school levels. While the incentives of educators tend to be based heavily on the performance of their students on standardized tests, households may have difficulty measuring school effectiveness, and may also place a relatively high value on other types of services that schools may offer. One result from non-profit theory states that competition creates incentives for these organizations to cater relatively more to donors and relatively less to employees (Glaeser, 2002). Along these lines, I expect to find that competition from the charter sector is associated with a decrease in the influence of teachers, and an increase in the influence of households within incumbent public schools.

Table 1.13 reports middle school-level results from a TWFE specification in which the dependent variable is either the teacher empowerment index or the share of teachers that believe households are influential decision makers. The regression specification used to produce these results is the same as the one described in equation 1.11, except in that it only includes controls at the school and LEA levels. The first column corresponds to the teacher empowerment index, where higher scores indicate a greater level of empowerment. The coefficient displayed in the first column is negative and statistically significant, and the interpretation of this result is that charter competition in this context reduces teacher influence within incumbent mid-

[42]The definitions of specific survey items and associated weights can be found in Appendix B.

Table 1.13: School Characteristics

	Teacher Empowerment	Household Influence
New Entrant	-0.522*	0.066**
	(0.297)	(0.032)
R-squared	0.498	0.788
N	669	669

* $p<0.10$, ** $p<0.05$, *** $p<0.010$

This table reports estimates of coefficients from regressions in which the dependent variable is a school-level characteristic. The first column corresponds to the teacher empowerment index, while the second column corresponds to the share of teachers within a school that either agree or strongly agree with the statement "Parents/guardians are influential decision makers in this school." Controls include school demographics and enrollment, LEA enrollment and average teacher experience, and cumulative charter applications received. Standard errors are clustered at the LEA level.

dle schools.[43] The dependent variable in the second column represents the share of teachers within a middle school that either agree or strongly agree with the statement "Parents/guardians are influential decision makers in this school." The coefficient displayed in this column is positive, statistically significant, and represents almost 10% of the mean level of household influence in this sample. Overall, these results are consistent with the implications from the model in Section 1.2, in which an increase in the number of competitors serving a given market for schooling induced increased provision of non-academic services by these schools.

I also explore whether charter competition has an effect on the share of total teacher salaries within a school that are allocated to teachers that specialize in

[43]The standard deviation of the teacher empowerment index is roughly 1.65, which means that the coefficient in the table corresponds to a decrease in teacher empowerment of over 30% of a standard deviation.

math and English/language arts, the two subjects most commonly tested by the state. Given evidence that task-specific human capital is an important determinant of teacher performance (Ost (2014)), it seems reasonable that a teacher would be most effective at improving student test scores in a given subject when they are able to focus all of their attention on the associated curriculum (e.g. a math teacher may be most effective at generating high math test scores when they only have to teach math courses, as opposed to courses across multiple subjects). However, if a school responds to competition by increasing the provision of services that do not directly improve exam scores, the school may have to allocate more revenue toward hiring other types of teachers, and may also have to broaden the responsibilities of the teachers that it already employs. Therefore, if charter competition induces an increase in the provision of consumption amenities at incumbent public schools, I expect this to be associated with a decrease in the share of salaries devoted to teachers who teach only math or only English/language arts courses.

While some schools in this sample devote over 50% of total salaries to teachers that specialize in these subjects, roughly a quarter of them do not employ any teachers that specialize completely in math or English/language arts, and therefore a typical linear model is not particularly well-suited for this analysis. Instead, I utilize a fractional response model, which estimates partial effects on the expected value of a fractional response (Papke and Wooldridge, 1996; Papke and Wooldridge, 2008). This framework is especially well suited to deal with a corner solution at zero, and has the advantage of much less restrictive assumptions relative to a Tobit approach when one is primarily interested in average partial effects, because one is not required to impose a parametric model on the conditional density of the response variable. I follow Wagner (2003) and estimate a fractional response logit specification with school-specific intercepts by Bernoulli quasi-maximum likelihood estimation.

Table 1.14 reports both the estimate of the parameter on new charter entry from the model as well as the average partial effect. The mean share of teacher salaries devoted to math and English/language arts specialists in this sample is roughly 11%, and therefore the estimate in Table 1.14 represents a decrease of roughly 14% relative to the mean. Overall, these results demonstrate that new charter entry had a negative and statistically significant effect on the share of teacher salaries allocated to teachers that completely specialize in tested subjects, which I interpret as further support of the idea that charter competition induced an increase in the provision of consumption amenities by incumbent public schools.

Table 1.14: Math/ELA Teacher Salary Share

	Parameter	APE
New Entrant	-1.061**	-0.015**
	(0.490)	(0.007)
N	1,338	1,338

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates from a fractional response logit model in which the dependent variable is the share of all teacher salaries that are allocated to teachers that specialize in tested subjects (i.e. math and English/language arts). The estimate in the first column corresponds to the parameter on the indicator for whether a new charter middle school has entered the LEA, while the estimate in the second column represents the average partial effect. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative applications received. Standard errors are clustered at the LEA level.

Table 1.15 is designed to explore whether public schools respond differently to competition based on the types of services being offered to students by competing charters. The first four columns explore the effect of for-profit charter schools, while the second four columns correspond to STEM charters. In addition to the teacher empowerment index and the share of teachers that believe households are influential decision makers, I include two additional dependent variables. First, from the working conditions survey I construct an index of preparation time constructed from items related to the amount of time that teachers spend on activities related directly to preparation for local or state-level standardized tests.[44] I also consider the natural log of total courses offered by a school in a given year.[45] Students are required to take courses in core subjects as determined by the state curriculum, but also typically select elective courses throughout middle school. If schools respond to competition through channels not directly associated with improvements in state-level math and reading exam scores, one way that they may appeal directly to household preferences is by broadening the set of courses offered to students.

The dependent variable in the first column of the first panel of Table 1.15 is the teacher empowerment index, and the estimated coefficient on the for-profit share is negative and large in magnitude. The result in column three demonstrates that specifically in the for-profit case, competition was associated with a large decrease in the exam preparation time index.[46] Finally, the result in column four demonstrates that competition is associated with an increase in the total number of distinct courses offered by a school, which can be interpreted as further evidence that these schools were catering to household tastes in a manner not directly related to improving standardized test scores in math and reading. The quantitative interpretation of this result is that in districts in which all charter entrants were managed by for-profit EMOs, the number of course offerings increased by nearly 30%, which corresponds

[44]The definitions of specific survey items and associated weights can be found in Appendix B.

[45]This is measured as the number of distinct four-digit course codes matched to the school.

[46]The standard deviation of the time index is 1.26, and therefore the magnitude of the coefficient suggests that in districts in which incumbent schools faced competition with only for-profit charters, the index of preparation time decreased on average by over a standard deviation.

Table 1.15: School Characteristics by Type of Competition

	For-Profit				STEM			
	Tchr. Emp.	HH Inf.	Time Ind.	Log Courses	Tchr. Emp.	HH Inf.	Time Ind.	Log Courses
6th Grade	-1.003*	0.032	-1.343***	0.303***				
For-Profit Share	(0.548)	(0.036)	(0.237)	(0.065)				
6th Grade					0.223	0.051	0.245	-0.181
STEM Share					(0.618)	(0.076)	(0.668)	(0.145)
New Entrant	-0.414	0.053	-0.039	-0.019	-0.616*	0.029	-0.274	0.031
	(0.403)	(0.034)	(0.209)	(0.061)	(0.313)	(0.042)	(0.272)	(0.073)
R-squared	0.506	0.787	0.509	0.993	0.500	0.785	0.492	0.992
N	669	669	669	1,338	669	669	669	1,338

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is a school-level characteristic. The covariate of interest is the share of 6th grade students enrolled in charter schools that opened after the policy change attending schools that indicated the intent to contract with a for-profit EMO, or to offer a STEM focus. The first column of each panel corresponds to the teacher empowerment index, while the second column of each panel corresponds to the share of teachers within a school that either agree or strongly agree with the statement "Parents/guardians are influential decision makers in this school." The dependent variable in the third column of each panel is an index of time spent by teachers on activities related to exam preparation, and the dependent variable in the fourth column is the natural log of the total number of distinct courses offered by a middle school. Controls include school demographics and enrollment, LEA enrollment and average teacher experience, and cumulative applications received. Each specification also includes cumulative applications received for each charter type (i.e. separate controls for for-profit and not-for-profit applications in columns 1 and 2). Standard errors are clustered at the LEA level.

to roughly an additional 4.5 distinct courses relative to the mean number of offerings by schools in this sample. In contrast, none of the coefficients on the STEM share in the second panel of Table 1.15 are statistically significant, which suggests that the specific type of product being offered by charter schools with which they are in direct competition is an important determinant of the response of incumbent public schools.

Table 1.16 explores heterogeneity in the effect of charter competition on the share of teacher salaries devoted to specialists in math and English/language arts based on whether the specific charters with which an incumbent public school is competing exhibit a particular characteristic. I use a fractional response logit model similar to the one previously described in this section, except I include an additional variable which represents the share of sixth graders in new charter schools that attend a charter school of a given type (e.g. for-profit). The left panel corresponds to for-profit charter

Table 1.16: Math/ELA Teacher Salary Share by Charter Characteristic

	For-Profit		STEM	
	Parameter	APE	Parameter	APE
6th Grade For-Profit Share	-0.054	-0.000		
	(0.809)	(0.001)		
6th Grade STEM Share			-0.026	-0.000
			(0.809)	(0.003)
New Entrant	-1.297**	-0.018**	-0.917	-0.013
	(0.578)	(0.008)	(0.785)	(0.011)
N	1,338	1,338	1,338	1,338

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates from a fractional response logit model in which the dependent variable is the share of all teacher salaries that are allocated to teachers that specialize in tested subjects (i.e. math and English/language arts). The covariate of interest is the share of 6th grade students enrolled in charter schools that opened after the policy change attending schools that indicated the intent to contract with a for-profit EMO, or to offer a STEM focus. The estimate in the first column corresponds to the parameter on the indicator for whether a new charter middle school has entered the LEA, while the estimate in the second column represents the average partial effect. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, and LEA enrollment. Each specification also includes cumulative applications received for each charter type (i.e. separate controls for for-profit and not-for-profit applications in columns 1 and 2). Standard errors are clustered at the LEA level.

schools, while the right panel corresponds to their STEM counterparts. Overall, these results suggest that while charter competition may be associated with a significant decrease in the share of teacher salaries allocated to math and English/language arts specialists, there does not appear to be substantial heterogeneity in this effect across different types of competition.

1.6.2 LEA Expenditures

Charter school entry may also have a negative fiscal impact on incumbent public schools because some costs of school operation are fixed, at least in the short-term (Ladd and Singleton, 2020). Revenue is allocated to public school districts at the per pupil level, but when an LEA loses students to the charter sector, it may not be able to scale costs downward proportionally. Basically, the decline in enrollment implies that the share of per pupil expenditure required to cover fixed costs will increase, leaving the LEA with less per pupil revenue to spend on instruction.

One implication of the model presented in Section 1.2 in the case in which competition has a negative effect on test scores (as documented in Section 1.5) is that competition is associated with lower per pupil expenditure on classroom instruction. Teachers in North Carolina are paid according to a salary schedule entirely based on observable characteristics such as years of experience and highest degree type. However, in order to compete for high quality teachers, LEAs may also offer a supplement to this base salary which is financed by locally-collected revenue. Therefore, variation in the level of this supplement across LEAs in close proximity may have important implications for the quality of classroom instruction, and there is substantial heterogeneity in supplement levels across the state. For example, in 2014-2015, the average annual supplement offered to teachers within a LEA ranged from $0 in seven cases to $6,892 in Chapel Hill/Carrboro City. As an example of a cost with direct implications for the quality of instruction, I focus on these local supplements as a means to explore the effect of charter competition on instructional expenditure in this context. I have collected the average local supplement offered to teachers across an LEA from the North Carolina Department of Public Instruction website.[47]

Table 1.17: LEA Expenditure

	(1)	(2)
	Local Teacher Supp.	PP Expenditure on School Admin.
New Entrant (Any Grade Level)	-1.559*	0.394**
	(0.885)	(0.180)
R-squared	0.958	0.749
N	228	228

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is a measure of LEA-level expenditure. The first column corresponds to the average local teacher supplement offered in a given LEA, while the second column corresponds to per pupil expenditure on school administration, both scaled in terms of hundreds of dollars. Controls include LEA demographics and enrollment, as well as cumulative applications received. Standard errors are clustered at the LEA level.

Table 1.17 summarizes the set of LEA-level results regarding fiscal spillover effects of charter competition. The explanatory variable of interest is an indicator for whether any new charter school has opened in a given LEA.[48] The dependent vari-

[47]One limitation of this data is that I only observe the average level of the supplement across an LEA even though it is not necessarily the case that all teachers within an LEA receive an identical amount.

[48]Note that this is slightly different than the treatment variable that was considered in the previous two sections, which was whether a new charter school that expressed an intent on its application to offer sixth grade has opened. This is because only LEA-level expenditure information is available, and the negative fiscal impact associated with losing students to the charter sector affects an LEA regardless of the grade levels at which the competition occurs. Appendix G reproduces the

able in the first column of Table 1.17 is the inflation-adjusted average local teacher salary supplement (measured in hundreds of dollars), and the estimated coefficient is negative and statistically significant. The interpretation of this estimate is that competition is associated with a $155.90 decrease in the level of the average annual supplement, which represents roughly 5% of the mean value of the average teacher supplement in this sample. While this effect may seem relatively small in magnitude, it represents a non-trivial share of income for some teachers, especially early in their careers.[49] Since not all teachers necessarily receive the same local supplement within a given district, it is also possible that some teachers were disproportionately affected, although this can not be tested directly with the available data.

In the second column, the dependent variable is the per pupil level of LEA expenditure on school administration. Regardless of whether an incumbent public school loses some of its students to a nearby charter entrant, it will still require a principal and associated support staff in order to operate.[50] Therefore, this measure is intended as an example of a fixed cost that can not be easily scaled with enrollment in the short term. The coefficient in the second column is positive, statistically significant, and represents roughly 7% of the average level of per pupil expenditure on school administration in this sample. This suggests that after charter entry, LEAs are allocating a greater share of per pupil expenditure toward covering fixed costs, which leaves less left over for the direct interaction of students and teachers in the classroom. These results are consistent with the implications of the model in the case in which there is no scope for perquisite consumption, and in which charter competition is

results from this section using the sixth grade treatment from the previous sections, and qualitatively the results are very consistent.

[49]For example, in 2017 new teachers without a graduate degree earned an annual income of $35,000.

[50]School principal salaries are also determined by a state-level salary schedule.

associated with a reduction in both instructional expenditure per pupil and academic achievement.

1.6.3 Explaining the Effect of Competition

Table 1.18: Standardized Math Test Score Mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
New Entrant	-0.043**	-0.032*	-0.032*	-0.028*	-0.026	-0.026*	-0.028*	-0.031*	-0.009
	(0.017)	(0.016)	(0.016)	(0.015)	(0.017)	(0.015)	(0.015)	(0.016)	(0.014)
Household Influence			0.009						-0.047
			(0.032)						(0.038)
Teacher Empowerment				0.013***					0.014***
				(0.003)					(0.003)
Local Teacher Supp.					0.001				0.002
					(0.001)				(0.001)
PP Exp. on School Admin.						-0.014			-0.017*
						(0.012)			(0.010)
Math/ELA Salary Share							0.039**		0.028
							(0.019)		(0.021)
Preparation Time Index								0.009	0.007
								(0.005)	(0.005)
R-squared	0.748	0.748	0.748	0.748	0.748	0.748	0.748	0.748	0.748
N	950,814	473,194	473,194	473,194	473,194	473,194	473,194	473,194	473,194

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized math end-of-grade test score. The first column displays the primary result from Table 1.4, while the second column replicates this result restricting the sample to years in which working conditions survey data is non-missing. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative applications received. Columns 3-8 each include one of the variables that represent proposed mechanisms associated with competition, and the ninth column includes all of these variables jointly. Standard errors are clustered at the LEA level.

The purpose of Table 1.18 is to provide evidence that the school and LEA-level mechanisms proposed in the previous sections can at least partially explain the documented spillover effects of charter school competition on traditional public school test scores. The first column displays the effect of charter competition on end-of-grade math test scores estimated from the full sample as in Table 1.4, while the second column reproduces the primary results using only the subsample for which the working conditions survey data is available (i.e. from years that end in an even number).

The estimated coefficient in the second column is qualitatively similar to the primary result, although it is slightly smaller in magnitude.

In columns three through eight, I explore the extent to which conditioning on variables that represent the proposed mechanisms reduces the coefficient on new charter entry by separately including a covariate that represents household influence, teacher empowerment, the average local teacher salary supplement, per pupil expenditure on school administration, the share of total salaries devoted to teachers that specialize in tested subjects, and the index of preparation time, respectively. Conditioning on each of these variables separately at least somewhat reduces the magnitude of the estimated coefficient on new charter entry. In column nine, we ultimately see that when we condition on all of these measures jointly, the estimated coefficient on new charter entry is reduced by over 70% and is no longer statistically significant. It is also worth noting that as expected, the coefficients on the teacher empowerment index, the local teacher supplement, the math or English/language arts specialist salary share, and the preparation time index are positive, and the coefficients on household influence and per pupil expenditure on school administration are negative. In the ninth column, only the estimates on teacher empowerment and per pupil expenditure on school administration are individually statistically significant, but I can reject the null hypothesis of an F-test of joint significance for the six covariates with a p-value of 0.0004. Overall, I find that consistent with the implications of the model in the case in which there is no scope for perquisite consumption by school administrators, a negative fiscal impact and the shift of influence within these schools from teachers to households are important mechanisms underlying the spillover effects of charter competition.

1.7 Conclusion

Despite the rapid expansion of the charter sector across the United States over the past three decades, the spillover effects of charter competition on students that remain enrolled in traditional public schools are still not well understood. Existing theories often do not allow public schools to take strategic decisions, or model them as maximizers of net revenue, which may not accurately capture the incentives faced by public schools when they experience competition. The literature on non-profit organizations predicts that among other things, competition is associated with higher expenditure on fundraising and a shift of influence from employees to donors. Motivated by theories of competition for donors among charities, I consider a model in which schools can provide consumption amenities to households that do not directly improved academic achievement, and I use this framework to derive some important implications of school competition. In this analogy, households play a similar role to donors because they directly determine the revenue of schools through their enrollment decisions. The public school system in the state of North Carolina provides an interesting setting to test these implications empirically due to a recent, large-scale expansion of the charter sector.

Using a two-way fixed effects strategy, I find a negative effect of competition from the charter sector on the test scores of public middle school students. I focus on two important channels through which traditional public school districts respond to entry from nearby charters. First, I find that charter competition is associated with a decrease in local teacher salary supplements and an increase in per pupil expenditure on school administration. This is consistent with the implications of the model in the case in which administrators do not spend school revenue on perquisite consumption, and thus in which there is a limited scope for administrators to reallocate existing school resources in a more productive manner. If a school loses some students to a new competitor, the associated reduction in revenue is roughly proportional to the number of switchers. However, because some costs of school operation are fixed at least in the short term, this implies that a greater share of revenue per pupil must be allocated toward covering these fixed costs, and thus less revenue is left over for classroom instruction.

The model also predicts that an increase in the degree of competition in a market for education leads to an increase in the level of consumption amenities provided by schools, despite the fact that this is associated with a cost in terms of test scores. Consistent with this prediction, I find that traditional public schools subject to competition from new charter entrants experienced a shift of influence from teachers to households, and that this shift can explain part of the negative overall effect of competition on test scores. I also explore some heterogeneity in these effects based on the types of charters with which incumbent public schools are competing, and find the largest negative effects on test scores in districts being served by charter schools managed by for-profit organizations, which relative to their non-profit counterparts may have weaker incentives to focus on academic quality beyond meeting accountability standards.

These findings have important implications for our understanding of policies that expand the scope for school choice, as well as policies that facilitate increased competition among non-profit organizations in general. While competition is likely to elicit a response from incumbent organizations, they may respond along margins that have adverse implications for product quality. These results raise important questions not only about whether facilitating competition among schools is a useful policy tool when the objective is to improve student achievement, but also about the primary objectives of a school.

Chapter 2

TRANSITIONAL AID AND THE EFFICACY OF SCHOOL COMPETITION

2.1 Introduction

Over the past few decades, many states across the US have implemented policies that facilitate competition among schools with the intention of creating incentives for schools to improve their effectiveness. At the encouragement of the federal government (e.g. Race to the Top grants), charter schools have become one of the most prevalent alternatives to the traditional public schools that students are zoned into based on their residential location. Charter schools are publicly funded, which means they can not charge tuition or screen students, but they operate outside of the traditional public system, which means they have considerable autonomy over curriculum and personnel relative to a traditional public school. Supporters of charter schools argue that these institutions improve student outcomes both directly, as they potentially improve the match quality between students and schools, as well as indirectly, because competition creates incentives for low-performing incumbent schools to improve.

While improving academic outcomes is often a central motivation behind policies that expand access to charter schools, in the first chapter I demonstrate that incumbent public schools may respond to competition through channels not typically considered by basic arguments in support of school choice, and that this may have unintended consequences for the test scores of traditional public school students. For example, competition may lead to increased provision of non-academic services that provide direct consumption value to households (e.g. athletic programs) but do not contribute to preparation for end-of-year exams, even if this behavior is associated with a cost in terms of exam scores. Also, most revenue is distributed to public school districts at the pupil level, or in other words total revenue scales approximately proportionally with enrollment. However, at least in the short-term, some of the costs of school operation are fixed, and thus can not be scaled proportionally with enrollment. Therefore, if a public school district loses some resident pupils to a new charter entrant, it will have less revenue per remaining pupil that it can allocate toward the classroom net of fixed costs (Ladd and Singleton, 2020; Bifulco and Reback, 2014). This is significant given evidence that school district expenditure is an important determinant of student outcomes (Jackson *et al.*, 2016), and raises questions about the efficacy of school choice as a policy tool for improving student achievement.

In order to mitigate fiscal concerns associated with charter school expansions, a small number of states have experimented with policies that fully or partially reimburse LEAs for revenue appropriated by charter schools. The motivation behind these policies is that while some costs of school operation may be fixed in the short-term, over a longer time horizon public schools can adjust their cost structure in a manner that reduces the excess costs created by charter school entry (e.g. through consolidating multiple public schools with reduced enrollment due to charter competition). Therefore, transitional aid payments typically take the form of reimbursements that phase out over the first few years after a given increase in charter enrollment occurs.

One channel through which transitional aid can improve student achievement is through a direct increase in the total amount of revenue that a given district can spend on classroom instruction. However, transitional aid can also improve test scores because it reduces the value to public school administrators of competing for students through the provision of non-academic services that may be associated with an opportunity cost in terms of academic performance. In this paper, I present a stylized model of school competition that illustrates how transitional aid payments can mitigate or reverse a negative relationship between school competition and test scores. I then leverage a unique transitional aid policy in the state of New York to provide an empirical test for the qualitative implications of the model.

Given that the majority of school district revenue is intended to be distributed at the pupil level, in practice LEAs are responsible for making tuition payments for resident pupils enrolled in charter schools that depend on the level of per pupil expenditure in the LEA. However, eligible LEAs in the state of New York receive transitional aid payments in the form of partial tuition reimbursements for increases in charter enrollment for three years after a given increase occurs. More specifically, LEAs are reimbursed for 80% of tuition payments in the first year after the enrollment increase, 60% of tuition payments in the second year after the enrollment increase, and 40% of tuition payments in the third year after the enrollment increase. However, New York LEAs are only eligible to receive transitional aid if at least 2% of all resident pupils are enrolled in the charter sector, or if at least 2% of total expenditure consists of tuition payments to charter schools.[1] I leverage the plausible exogeneity of aid eligibility in the neighborhood around these cutoffs to implement a regression discontinuity strategy extended to settings with multiple possible assignment variables (Choi and Lee, 2018; Wong *et al.*, 2013; Reardon and Robinson, 2012). Importantly, I find that transitional aid eligibility is associated with a significant increase in instructional expenditure per pupil as well as in average LEA performance on statewide exams for students enrolled in grades 3-8.

Another interesting feature of the New York transitional aid policy is the one-year delay between a given increase in charter school tuition payments for an aid-eligible LEA and the actual receipt of the first reimbursement. In other words, when a given LEA experiences an increase in resident pupils enrolled in the charter sector, its ad-

[1]LEAs in New York City are also specifically excluded from aid eligibility.

ministrators know whether or not they will ultimately receive transitional aid for a year prior to any potential payments. This provides an interesting opportunity to test whether the effect of aid eligibility on student achievement operates through additional channels besides the direct increase in LEA revenue. For example, if it were the case that transitional aid reduced the incentives of school administrators to respond to competition through an increase in the provision of consumption amenities, we might expect aid eligibility to improve exam scores even before any changes in revenue. I use an event study framework to demonstrate that among LEAs which became eligible for transitional aid for the first time from 2013-2017, average exam scores increased during the year prior to the first reimbursement, although they increased further once aid payments began. This suggests that a direct increase in LEA revenue is not the only mechanism through which the effect of transitional aid eligibility operates.

The existing study most similar to this one is Mann and Bruno (2020), and therefore it is worth highlighting some of the key distinctions between these analyses. Mann and Bruno (2020) leverage a policy change in Pennsylvania and implement a difference-in-differences strategy to demonstrate that transitional aid reimbursements are associated with increased instructional expenditure and at least partially mitigate the negative effect of increased charter enrollment on student achievement in the traditional public sector. First, in a different context and under very different identifying assumptions, I find largely consistent results regarding the effect of transitional aid eligibility on LEA expenditure behavior, and therefore in some sense these results can be viewed as complementary. More importantly, the primary contribution of this study is to highlight the fact that transitional aid may affect student outcomes through additional channels besides changes in LEA revenue because it reduces the value of competing for students through non-academic channels. I highlight this indirect channel both in the model of school competition as well as in the of empirical results.

The large literature on the causal effect of charter school enrollment on student outcomes has largely relied on random admissions lotteries for identification.[2] More recently, a number of studies have utilized richer models to explore equilibrium effects of charter school expansions (Walters, 2018; Mehta, 2017), as well as entry and exit decisions within the charter sector (Singleton, 2019; Singleton, 2017), but the primary contribution of this study is to the growing literature on the spillover effects of charter school competition on traditional public schools and their students (Slungaard Mumma, 2020; Gilraine *et al.*, 2021; Imberman, 2011). Motivated by a recent wave of evidence that household enrollment decisions do not necessarily appear to be importantly related to school effectiveness (Abdulkadiroglu *et al.*, 2020; MacLeod and Urquiola, 2019; Beuermann *et al.*, 2019), this paper is closely related to studies which find that LEAs may respond to competition through channels that do not directly contribute to academic achievement (Arsen and Ni, 2012; Cook, 2018; Kofoed and Fawson, 2020). This study also represents a contribution to the literature on the fiscal effects of charter school competition, which often finds a negative fiscal impact of charter competition on LEAs (Ladd and Singleton, 2020; Bifulco and Reback, 2014), although not necessarily in contexts in which transitional aid is available (Ridley and Terrier, 2018).

In the remainder of the paper, I first introduce a stylized model of school competition in Section 2.2 to develop intuition, while Section 2.3 describes the features of the institutional setting for the empirical analysis as well as the sources of data that I have compiled. Section 2.4 describes the empirical strategy used to generate the primary results, and Section 2.5 provides a discussion of these results before Section

[2]For a review of this literature, please see Epple *et al.* (2015).

2.6 concludes.

2.2 Model

The purpose of this section is to describe a stylized model of school competition that demonstrates how increased competition among schools can potentially lead to a decrease in academic performance across a market, but also how the provision of transitional aid can reduce or even reverse this effect by decreasing the value of competing for students through the provision of non-academic services.[3] Consider a school district that consists of a mass λ of households, a traditional public school, and possibly a charter school.[4] A school produces education, which is characterized by a level of per pupil instructional expenditure and a level of consumption amenities, the latter of which is assumed to be directly valued by households.

Each household in this market is assigned one student for which it makes an enrollment decision, and each one of these households is identical except for a heterogeneous cost of enrolling in the charter school k_i, which is uniformly distributed on $\left[-\bar{k}, \bar{k}\right]$. Charter schools can not legally charge tuition, so k represents non-pecuniary costs of charter enrollment. This might include time and effort spent on gathering information and filling out applications, or the inability to take advantage of certain services that are only offered in the public sector in some contexts (e.g. bus transportation to school). I also allow k_i to take negative values in order to capture the fact that some households may place a relatively high value on services that are only offered in the

[3]Other theoretical studies of school competition also show how competition can fail to bring about improvements in student achievement in the presence of peer effects (Barseghyan *et al.*, 2019), capacity constraints (Cardon, 2003), or any scope for schools to engage in rent-seeking behavior (McMillan, 2004).

[4]I will consider a baseline case with no competition as a benchmark before analyzing a scenario in which the public school and charter school compete for students.

charter sector (e.g. a particular curricular focus or educational philosophy). Households choose to enroll their student in the school, either public school p or charter school q, that maximizes the following utility function:

$$\bar{u} + \eta X_j - \mathbb{1}(j)k_i, \ j \in \{p, q\} \tag{2.1}$$

In this case, $\bar{u}$ is a parameter chosen to be sufficiently high such that enrollment in some school will always dominate the outside option of no enrollment. X_j is the level of consumption amenities offered by school j, which represents services valued by the household that do not contribute directly to preparation for end-of-year exams (e.g. athletic programs or other student organizations). The third term represents a relative utility cost (or benefit if $k_i < 0$) that is only incurred if the household enrolls its student in the charter school. For the sake of tractability, in this formulation households do not directly value instructional spending or academic quality. While this may seem like a strong assumption, it is consistent with recent evidence that household enrollment decisions are not importantly related to school effectiveness (Abdulkadiroglu *et al.*, 2020). Each school maximizes the following objective function, which represents the total amount of education it produces:

$$S_j Y_j, \ j \in \{p, q\} \tag{2.2}$$

In this case, S_j represents total enrollment at school j, while Y_j represents a test score, which is given by the following:

$$Y(I_j, X_j) = \alpha_j I_j - \psi_j X_j, \ j \in \{p, q\} \tag{2.3}$$

The school chooses a level of instructional expenditure I_j and a level of consumption amenities X_j to maximize total education production. In the above expression, α_j is a parameter that determines the efficiency with which school j can convert instructional expenditure to academic achievement. The score is also decreasing in the level of consumption amenities offered by school j with constant marginal cost ψ_j. This is intended to capture the idea that time spent on the provision of services not directly related to test score preparation are associated with an opportunity cost in terms of performance. Therefore, school administrators must balance the cost of consumption amenities in terms of test scores with the potential benefits in terms of enrollment given the value of these amenities to households.

2.2.1 Baseline: No Competition

We will begin by analyzing a market in which there is only a public school with no charter competition, which will serve as a useful benchmark case when we ultimately consider the effects of competition. In this case, all households will enroll their student in the public school. The public school administrator will maximize education production subject to the following budget constraint:

$$\zeta_p + \lambda I_p = \lambda \rho \tag{2.4}$$

Total school expenditure includes the fixed cost ζ_p as well as total expenditure on instruction, which is equal to instructional expenditure per pupil I_p times enrollment λ. This may not exceed total school revenue $\lambda\rho$. This formulation captures a realistic feature of markets for public education in that revenue is allocated at the per pupil level, and therefore total school revenue scales proportionally with enrollment.

In the absence of competition, the school has no incentive to provide consumption amenities, and will therefore choose $X_p = 0$, which means that all households receive utility $\bar{u}$. The school will allocate all revenue to classroom instruction, which means that in equilibrium, instructional expenditure per pupil will be given by the following expression:

$$I_p^* = \rho - \frac{\zeta_p}{\lambda} \tag{2.5}$$

This implies that the test score for each student will be as follows:

$$Y_p^* = \alpha_p \left(\rho - \frac{\zeta_p}{\lambda} \right) \tag{2.6}$$

These values will serve as a useful comparison in the following section, when we consider the effect of competition on academic achievement as well as the role of transitional aid.

2.2.2 Competition and Transitional Aid

We will now consider the effect of competition in the context of a static game in a market with both a public and a charter school in operation, and in which the administrators of these schools make input decisions simultaneously. In this case, households will have a choice in which type of school to enroll their student. The level of enrollment at each school in this setting can be found by considering the marginal household in terms of the non-pecuniary charter enrollment cost, or in other words the cost which would make a household indifferent between enrollment in either school. The marginal level of the charter enrollment cost $\hat{k}$ is given by the expression below:

$$\hat{k} = \eta(X_q - X_p) \tag{2.7}$$

Recall that X_q represents consumption amenities provided by the charter school, and X_p represents consumption amenities provided by the public school. All households with a charter enrollment cost below this marginal value will enroll in the charter school, which implies the following level of charter school enrollment:

$$S_q = \frac{\lambda}{2\bar{k}}\left(\eta(X_q - X_p) + \bar{k}\right) \tag{2.8}$$

All households with a charter enrollment cost greater than $\hat{k}$ will enroll in the public school, which gives us the following expression for public school enrollment:

$$S_p = \frac{\lambda}{2\bar{k}}\left(\bar{k} - \eta(X_q - X_p)\right) \tag{2.9}$$

Given these levels of enrollment, the charter school will maximize Equation 2.2 subject to the following budget constraint:

$$\zeta_q + I_q S_q = \rho S_q \tag{2.10}$$

Total charter school expenditure is given by the fixed cost ζ_q and total instructional expenditure, which is equal to enrollment S_q times instructional expenditure per pupil I_q. Total charter school revenue is equal to the constant level of per pupil revenue ρ times enrollment S_q. In a significant departure from the benchmark case without competition, the budget constraint of the public school is given by the following:

$$\zeta_p + I_p S_p = \rho\left(S_p + \mu S_q\right) \tag{2.11}$$

Similarly to the charter school, total public school expenditure consists of the fixed cost ζ_p plus total instructional expenditure $I_p S_p$. The public school also receives per pupil revenue ρ for each of the S_p students that it enrolls. However, this school also receives transitional aid in the form of a reimbursement for a share of the revenue that it would have received in a counterfactual scenario without competition, where the share is given by $\mu \in [0, 1]$. In other words, $\mu = 0$ corresponds to a scenario with no transitional aid, and $\mu = 1$ corresponds to a situation in which the public school is fully reimbursed for any decrease in revenue associated with charter enrollment. The purpose of this analysis will be to perform a simple comparative statics exercise in which we consider how equilibrium levels of inputs and test scores in this market depend on the level of the reimbursement share μ. Solving for the equilibrium levels of instructional expenditure and consumption amenities at each school yields the following proposition:

Proposition 1 *In equilibrium, I_q^* is strictly increasing in μ and X_q^* is strictly decreasing in μ, which implies that Y_q^* is strictly increasing in μ. X_p^* is strictly decreasing in μ, but the effect of μ on I_p^* is ambiguous, and therefore the effect of μ on Y_p^* is also ambiguous. If ζ_p is sufficiently low, I_p^* and Y_p^* are strictly increasing in μ.*

The full expressions that represent I_j^* and X_j^* for $j \in \{p, q\}$ can be found in Appendix I. The clearest way to explain the intuition behind each statement in the above proposition is to begin with the equilibrium level of consumption amenities provided by the public school X_p^*. When the share of tuition to be reimbursed μ increases, it reduces the marginal benefit of providing consumption amenities. This is because a higher reimbursement share implies that the school will lose less revenue for each student that enrolls in the charter school, and therefore this generates a lower level of amenities in equilibrium.

When the public school offers a lower level of consumption amenities, it means that the charter school can also afford to offer a lower level of amenities while remaining competitive for students. Therefore, X_q^* is also decreasing in μ. However, the difference in amenity provision across schools $X_q^* - X_p^*$ is increasing in μ, which implies that in equilibrium charter enrollment is increasing in μ and public enrollment is decreasing in μ. Therefore, a higher reimbursement share implies that the charter school will have greater total revenue net of the fixed cost of operation ζ_q, which means that instructional expenditure per pupil I_q^* is increasing in μ. When considered along with the associated decrease in X_q^*, this implies an unambiguous increase of the charter school test score Y_q^* in μ.

There are two offsetting effects of μ on equilibrium instructional expenditure per pupil in the public school I_p^*. First, the direct increase in revenue from the reimbursement allows the school to spend more on instruction. However, the reduction in public school enrollment associated with an increase in μ also means that more expenditure per pupil will have to be allocated toward covering the fixed cost ζ_p, reducing the amount that can be allocated to classroom instruction. If the fixed cost of public school operation ζ_p is sufficiently low, the reimbursement channel will dominate and both I_p^* and Y_p^* will be strictly increasing in μ.

Relative to the baseline case with a single school in operation, households are strictly better off under competition due to the provision of consumption amenities. However, the effect of competition on test scores in the market is less clear. On one hand, test scores may be lower due to economies of scale (as both schools incur a fixed cost) and also due to the opportunity cost associated with amenity provision. However, if the tuition reimbursement share μ is sufficiently high, and the charter school fixed cost ζ_q is sufficiently low, then it could also be the case that competition increases academic achievement. Overall, this simple analysis demonstrates how competition among schools may create the incentive for administrators to increase the provision of consumption amenities at the cost of academic achievement, but also how this effect can be mitigated or reversed through the provision of transitional aid for decreases in revenue associated with charter enrollment. This intuition is central in motivating the empirical analysis introduced in Section 1.4.

2.2.3 Optimal Reimbursement Share

This section provides a discussion of the optimal reimbursement share in a given market from the perspective of social welfare, which is highly relevant in designing effective programs that provide transitional aid to LEAs for increases in charter enrollment. Under a reasonable parametrization of the model, transitional aid reimbursements unambiguously increase academic performance in both the charter and public schools. However, this comes at both the direct cost of providing the reimbursement as well as a loss of household welfare associated with lower provision of consumption amenities. In order to analyze the optimal reimbursement share μ^*, I consider a social welfare function W which consists of the sum of utility for both the public and charter school administrators, the utility of all households in the market, and the cost of providing the reimbursement. Assuming that $W(\mu)$ admits a well-defined interior optimum, an application of the Implicit Function Theorem then gives the following proposition:

Proposition 2 *If η and $\bar{k}$ are sufficiently high, then the optimal reimbursement share μ^* is decreasing in λ.*

The above proposition states that if the marginal value to households of an increase in consumption amenities is sufficiently high, and if an LEA is sufficiently heterogeneous in terms of preferences for charter schooling as represented by $\bar{k}$, then the optimal reimbursement share is strictly decreasing in the total size of the market. This result suggests that transitional aid may be most beneficial in smaller LEAs in terms of total resident pupils. Marginal increases in revenue associated with aid reimbursements in large districts may be less valuable than the consumption amenities that would be provided in the absence of aid as long as there is sufficient heterogeneity in tastes for charter schooling to guarantee that these changes in revenue can not generate an unreasonably large shift in enrollment across schools. However, in LEAs with fewer students and lower total revenue, competition creates greater incentives for administrators to increase the provision of consumption amenities at the cost of academic performance. Therefore, the presence of transitional aid reimbursements in this type of market may be particularly important regarding whether competition is able to generate the positive academic effects that policymakers often have in mind when considering policies that expand the scope for school choice.

2.3 Institutional Context and Data

The charter sector in the state of New York, which began granting charters in 1998, has grown to serve over 150,000 students across over 300 schools. The state originally limited the number of charter schools in operation to 100, but has subsequently raised this cap twice, first to 200 in 2007, and then to the current, non-binding level of 460 in 2010. This state provides a particularly interesting setting in which to study spillover effects of charter school competition because of the unique manner in which transitional aid payments are provided to public school districts for increases in charter enrollment.[5] New York is one of a small group of states to have experimented with policies that partially reimburse LEAs for losses in revenue due to increased enrollment of resident pupils in the charter sector.[6] However, unlike other contexts in which transitional aid payments are made for all increases in charter enrollment, New York LEAs only become eligible to receive aid once either charter enrollment or payments made to charter schools exceed a given threshold.

The state of New York partially reimburses eligible LEAs for revenue appropriated

[5]LEAs in New York City are prohibited from receiving transitional aid for increases in charter enrollment, and therefore New York City is excluded entirely from this analysis.

[6]Examples of other states that have implemented some form of transitional aid include Massachusetts and Pennsylvania.

by the charter sector for three years after an increase in charter enrollment occurs. Revenue is approximately distributed at the pupil level, and therefore in practice LEAs are required to make tuition payments that depend on their level of per pupil expenditure for each resident pupil enrolled in a charter school. LEAs are reimbursed 80% of increased tuition payments to charter schools in the year following a given increase in charter enrollment (known as "Transitional Aid A"), as well as 60% of these payments during the second year following the increase ("Transitional Aid B"), and 40% of these payments during the third year following the increase ("Transitional Aid C"). However, while the magnitude of the aid payment is proportional to the change in charter enrollment, LEAs are only eligible to receive aid if either at least 2% of resident pupils are enrolled in the charter sector, or if payments to charter schools comprise at least 2% of total expenditure.

In the vast majority of LEAs, only the enrollment threshold is binding, and therefore this is the only threshold that will be considered for the sake of this analysis.[7] Therefore, LEAs in the sample are eligible to receive any transitional aid if either the first, second, or third lag of the charter share exceeds the 2% threshold. Data on charter enrollment, transitional aid eligibility, and transitional aid payments has been generously provided by the New York State Education Department (NYSED). I supplement this information with district-grade-year level mean test scores and demographic characteristics available on the NYSED website, as well as LEA revenue and expenditure data from the LEA Finance Survey conducted by the National Center for Education Statistics.

Many LEAs across the state of New York did not have any resident pupils enrolled

[7]Three LEAs are dropped from the primary sample because they meet the eligibility threshold for either Transitional Aid A, Transitional Aid B, or Transitional Aid C by the expenditure criterion but do not meet the enrollment criterion.

Figure 2.1: Charter Share

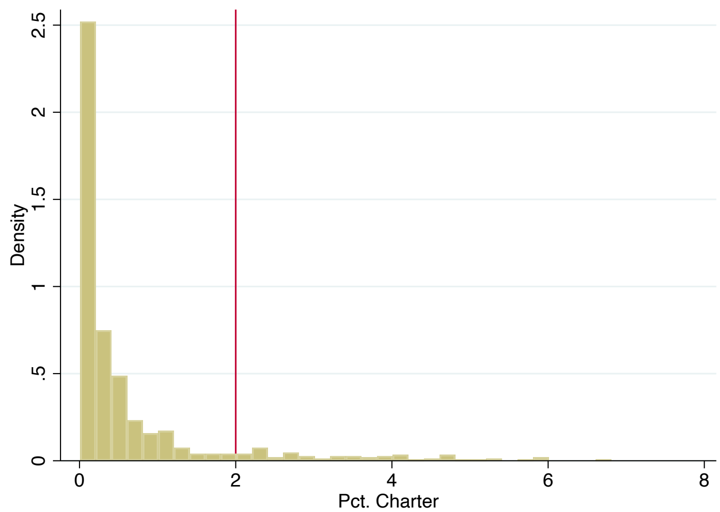

This figure displays a histogram of the charter share of New York LEAs from 2013-2017. The vertical line is at the threshold (2%) beyond which an LEA becomes eligible to receive transitional aid for increases in charter enrollment. The sample used to generate this figure does not include LEAs with a charter share of zero, or the top 5% of LEAs in terms of the charter share.

in the charter sector during the time period upon which the primary empirical analysis is focused, which runs from 2013-2017. Figure 2.1 displays the density of charter shares across all LEAs during this time for which at least one resident pupil was enrolled in a charter school.[8] The vertical line at 2% represents the threshold beyond which LEAs are eligible to receive transitional aid for increases in tuition payments to charter schools. This figure demonstrates that a relatively small mass of districts are located around this cutoff value, as the majority of LEAs have charter shares well below 1%. Further, there is no visual discontinuity at the threshold value, which suggests that LEAs are not able to manipulate enrollment numbers to become eligible for aid payments.

[8] Due to the presence of some extreme outliers in terms of charter enrollment, the top 5% of LEAs in terms of charter share were also excluded from the figure.

2.4 Empirical Strategy

This section provides details regarding the regression discontinuity and event study frameworks used to generate the primary results of this analysis.

2.4.1 Regression Discontinuity

The first objective of this empirical analysis is to document the effect of transitional aid eligibility on LEA expenditure behavior. If it were the case that any LEA in New York received transitional aid for an increase in charter enrollment among resident pupils, then it would be difficult to identify this effect due to endogeneity concerns. This is because there are potentially many important observable and unobservable differences between LEAs categorized by whether or not the share of resident pupils enrolled in charter schools is increasing. However, some of the unique features of the New York transitional aid policy allow me to mitigate this identification challenge through the implementation of a regression discontinuity design based on thresholds for transitional aid eligibility. More specifically, LEAs in this analysis are only eligible to receive transitional aid for increases in charter enrollment if at least 2% of resident pupils are enrolled in the charter sector.[9] An important assumption required for this strategy to identify the effect of aid eligibility on LEA expenditure is that eligibility is exogenous among the subsample of LEAs with charter shares in the neighborhood of the eligibility threshold.

An additional feature of this setting is that instead of a single assignment variable and threshold, as in a standard regression discontinuity design, there are multiple assignment variables and thresholds. An LEA in this context is eligible to receive

[9] As discussed in Section 1.3, LEAs are also technically eligible for aid if at least 2% of total expenditure consists of tuition payments made to charter schools, but this alternative threshold is not binding for any of the LEAs in the sample.

any transitional aid if at least one of the previous three lags of the charter share exceeds 2%. However, consider the case of an LEA with a current charter share above the eligibility threshold, but for which the first three lags of this share are not. While this LEA would not be eligible to receive transitional aid in the current period, administrators would potentially be aware that the district would be eligible to begin receiving aid in the following period. While any effect of future aid eligibility on expenditure behavior could not be driven by an increase in revenue due to aid payments, we still might expect to find an effect if this creates incentives for school administrators to reallocate existing revenue across expenditure categories. Overall, this leaves us with four possible thresholds (i.e. eligibility cutoffs for aid eligibility in the following period, Transitional Aid A, Transitional Aid B, and Transitional Aid C), as well as four different assignment variables (i.e. the current charter share and its first three lags) to consider.

Given complications associated with multiple assignment variables, and the fact that relative to a traditional regression discontinuity design it is possible to define many different potential treatments of interest, the literature on multiple-score regression discontinuity suggests a number of possible estimation strategies that are relevant in this context (Choi and Lee, 2018; Wong *et al.*, 2013; Reardon and Robinson, 2012). In this analysis I follow the approach described in Choi and Lee (2018), which allows me to estimate the treatment effect of each type of aid eligibility within the same regression model. This means that I can explicitly test for differences in the effect of aid eligibility associated with a given increase in charter enrollment both prior to and across each of the first three years after this increase occurs. Relative to other multiple-score regression discontinuity strategies, the frontier approach is associated with increased efficiency in estimation, but relies on the assumption that the response surface is modeled correctly. The regression specification that I use to estimate the effect of transitional aid eligibility on LEA expenditure behavior is as follows:

$$\begin{aligned} y_{dt} = \beta_0 + \beta_1 FA_{dt} + \beta_2 TA_{dt} + \beta_3 TB_{dt} + \beta_4 TC_{dt} + \beta_5 S_{dt} \\ + f(CS_{dt}, CS_{dt}^{-1}, CS_{dt}^{-2}, CS_{dt}^{-3}) + \omega_d + \delta_t + \epsilon_{dt} \end{aligned} \tag{2.12}$$

In the above equation, y_{dt} represents expenditure per pupil on a given budget category (e.g. instructional expenditure per pupil). The regressors of interest are FA_{dt}, TA_{dt}, TB_{dt}, and TC_{dt}. These variables represent eligibility for aid in the following period, eligibility for Transitional Aid A, eligibility for Transitional Aid B, and eligibility for Transitional Aid C, respectively. Therefore, the primary coefficients of interest are β_1, β_2, β_3, and β_4. It is possible for an LEA to be eligible for any combination of these types of transitional aid, and thus these coefficients can be separately identified, which provides an interesting opportunity to explore dynamic effects of aid eligibility. While we would not expect the effect of future aid eligibility on LEA expenditure behavior to be driven by changes in total LEA revenue (as no reimbursements from the state had yet been received), we still might expect to find a non-zero effect if future aid eligibility creates incentives for school administrators to reallocate revenue across expenditure categories.

This regression specification also includes a quadratic of each of the four assignment variables that determine aid eligibility, more specifically the current charter share in period t and the first three lags of this variable. I also control for total district enrollment S_{dt}, district fixed effects, and year fixed effects in order to address heterogeneity in some other important determinants of LEA revenue. The literature on optimal bandwidth selection in regression discontinuity analysis has not been extended to settings with multiple assignment variables, so as a rule-of-thumb I include observations within 0.4 percentage points around any of the enrollment thresholds in the preferred specification. Standard errors throughout the empirical analysis are clustered at the LEA level. I also demonstrate that the primary results of this study are qualitatively robust to alternative choices of bandwidth.

In addition to LEA expenditure behavior, one of the primary concerns of a policy-maker considering the provision of transitional aid should be the effect of aid eligibility on student achievement. Therefore, I also use this regression discontinuity strategy to analyze the effect of transitional aid eligibility on average LEA performance on statewide exams. Using data on mean exam scores at the district-grade-year level, I estimate a regression specification of the following form:

$$\begin{aligned} y_{gdt} = \beta_0 + \beta_1 FA_{dt} + \beta_2 TA_{dt} + \beta_3 TB_{dt} + \beta_4 TC_{dt} + \beta_5 S_{dt} \\ + \beta_6 X_{gdt} + f(CS_{dt}, CS_{dt}^{-1}, CS_{dt}^{-2}, CS_{dt}^{-3}) + \omega_d + \delta_t + \epsilon_{gdt} \end{aligned} \tag{2.13}$$

Aside from the unit of the sample, this specification is very similar to the one used to estimate the effect of aid eligibility on LEA expenditure behavior. In this case, y_{gdt} represents a mean test score at the district-grade-year level. These scores are standardized by the mean and standard deviation of the student-level distribution of exam scores, such that the statewide mean is zero and units can be interpreted as a percentage of a standard deviation. The transitional aid indicators provide an interesting opportunity to explore the dynamic effects of aid eligibility on exam scores, and we might expect to find a positive effect of eligibility for future aid if the effect of aid eligibility on student achievement operates through additional channels besides the direct increase in LEA revenue. This specification also includes a set of controls at the district-grade-year level, which includes various demographic characteristics, the total number of students tested in a given unit, and grade-level fixed effects.

Table 2.1: Demographic Characteristics

	(1) Pct. Black	(2) Pct. Hispanic	(3) Pct. EDS	(4) Students Tested
Eligibility for Future Aid	-0.050*	0.052**	0.026	-11.200
	(0.026)	(0.019)	(0.032)	(44.605)
Aid Eligibility A	-0.029	0.020	0.009	33.503
	(0.024)	(0.017)	(0.011)	(25.298)
Aid Eligibility B	0.009	0.026	0.023	53.087**
	(0.014)	(0.021)	(0.014)	(20.498)
Aid Eligibility C	-0.010	0.012	0.006	65.359
	(0.016)	(0.024)	(0.020)	(46.748)
R-squared	0.026	0.156	0.208	0.244
N	306	306	306	306

* $p<0.10$, ** $p<0.05$, *** $p<0.010$

This table explores characteristics of LEA-grade level-year triplets in the neighborhood of the four transitional aid thresholds. The first column corresponds to the share of black students, the second corresponds to share of Hispanic students, the third corresponds to the share of EDS students, and the fourth corresponds to the total number of students tested. Controls include a second-degree polynomial of the charter share and the previous three lags total district enrollment across all grade levels, as well as district, year, and grade level fixed effects.

In order to support the assumption that aid eligibility is exogenous in the sample of LEAs with a charter share in the neighborhood of at least one of the eligibility thresholds, it is helpful to demonstrate that there are not important differences between LEAs located just on either side of these thresholds in terms of observable characteristics. A common test performed in the large literature on regression discontinuity designs involves the use of observable characteristics as dependent variables in the primary specification in order to test whether there are any discontinuities in these measures (Imbens and Lemieux, 2008). Table 2.1 displays a set of regression results generated from Equation 2.4.1 in which each column represents a different demographic characteristic at the LEA-grade-year level. The ideal scenario would be if there were no statistically significant coefficients across any characteristic or type of aid eligibility. However, given the relatively small sample size and the fact that the model includes multiple assignment variables and thresholds, it is not surprising that there are a few characteristics for which the coefficient that represents one type of aid eligibility is statistically significant. Overall, the coefficients in Table 2.1 are small in magnitude and mostly insignificant in the statistical sense. Therefore, the results displayed in this table can be interpreted as evidence that there are not systematic differences across LEAs with charter shares just on either side of an eligibility threshold. I will also use each of these LEA characteristics as control variables in the exam score analysis in order to address potential small sample biases.

Descriptive Statistics

Table 2.2 displays a set of summary statistics that characterize the public education system across New York, as well as specifically in the sample of LEAs that will be used to generate the primary results, which have charter shares in the neighborhood of at least one of the thresholds that determine transitional aid eligibility. Since I restrict attention to a set of LEAs that may not be representative of the average LEA across the state, proper interpretation of the empirical results requires a characterization of differences between LEAs that are included in my sample and those that are not. Table 2.2 displays a set of sample means at the district-grade-year level for grade levels 3-8 from 2013-2017. The first column corresponds to all district-grade-years across the state during this sample period that had a non-zero charter share which was not in the neighborhood of at least one of the eligibility thresholds (excluding New York City), the second column corresponds to the LEAs included in the primary sample considered in the empirical analysis, and the third column displays the results of a standard difference-in-means test. The primary sample clearly differs from other LEAs across the state along a variety of important dimensions. Relative to observations not included in the primary sample, LEAs with charter shares around the aid eligibility thresholds enroll significantly higher shares of minority students and economically disadvantaged students (EDS). They are also significantly larger on average in terms of number of students tested at the district-grade-year level. In light of these differences, it is important to exercise caution in extrapolating the primary results to LEAs in other areas that appear to be very different in terms of observable characteristics.

2.4.2 Event Study

This section describes an additional test to determine whether the effect of transitional aid eligibility on student achievement operates through an indirect effect (e.g. a redistribution of time or other resources from non-academic to academic services), in addition to the direct effect of increased revenue associated with reimbursements. This is done by leveraging the delayed timing of transitional aid payments in the New York context. Eligibility for aid is determined by the charter share of enrollment in a

Table 2.2: Summary Statistics

	Non-Primary Sample	Primary Sample	Difference
Pct. Black	0.09	0.25	-0.16***
Pct. Hispanic	0.13	0.20	-0.08***
Pct. EDS	0.41	0.70	-0.29***
Number Tested	234.70	449.02	-214.32***
Observations	4,035	306	4,341

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of descriptive statistics at the LEA-grade-year level. The first column displays a set of mean characteristics for LEAs that are not included in the primary sample for the regression discontinuity analysis, the second column displays a set of mean characteristics for units that are included in the primary sample, and the third column displays the results of a standard difference-in-means test.

given LEA, and the size of aid payments is proportional to increases in the number of resident pupils enrolled in charter schools. However, LEA administrators potentially become aware of reimbursements a year before the first of these payments actually occurs. If aid eligibility reduces the incentives of school administrators to compete for students through the provision of non-academic services, then we might expect to find a positive effect of aid eligibility on average exam scores even prior to the receipt of aid payments. In order to test for this indirect effect of aid eligibility, I utilize the following event study framework to analyze the set of LEAs that became eligible to receive transitional aid for the first time from 2013-2017:

$$y_{gdt} = \beta_0 + \sum_{l=-6}^{-3} \xi_l D_{dt}^l + \sum_{l=-1}^{3} \xi_l D_{dt}^l + \beta_1 X_{gdt} + \omega_d + \delta_t + \epsilon_{gdt} \tag{2.14}$$

In this case, y_{gdt} represents a standardized average exam score at the district-grade-year level. This model includes a set of relative period indicators D_{it}^{l} to test for dynamic effects of aid eligibility, where l represents the number of periods since treatment. These variables are equal to one if unit i is l periods away from treatment at time t, and zero otherwise. The first relative period in which an LEA is eligible to receive current aid is represented by $l = 0$, and I normalize the coefficient on the relative period represented by $l = -2$ to zero. Particular attention will be focused toward the coefficient on the relative period represented by $l = -1$, as this is the period during which administrators are potentially aware that an LEA is eligible to receive future aid, but has not yet received any reimbursements. A positive and statistically significant coefficient on this indicator would suggest that the direct increase in district revenue is not the only channel through which transitional aid eligibility can affect student achievement. This specification also includes district and year effects, as well as controls for demographic characteristics at the district-grade-year level and a quadratic of the current charter share of enrollment. I bin all of the observations that correspond to relative periods prior to $l = -5$ together, and I bin all of the observations that correspond to relative periods after $l = 2$ together.

Similar to a standard difference-in-differences framework, identification of the treatment effects of interest relies on the assumptions of parallel counterfactual trends across treatment cohorts and no anticipatory effects prior to treatment.[10] A common test of these assumptions in the literature involves testing whether ξ_l is statistically distinguishable from zero for $l < -1$.[11] Identification in the dynamic context also

[10]Note that in this case I am considering the period in which LEAs are eligible for future aid but have not yet received any reimbursements as part of the treatment. Therefore, while effects on exam scores during this period could be interpreted as occurring in anticipation of future aid receipt, this does not violate the assumption.

[11]See He and Wang (2017) for an example.

requires the assumption that each treatment cohort follows the same dynamic path of treatment effects (Sun and Abraham, 2021).

2.5 Results

This section provides a discussion of the primary results from the empirical analysis. More specifically, this includes regression discontinuity results regarding the effect of transitional aid eligibility on expenditure and mean statewide exam performance at the district-grade-year level. I also discuss an event study which demonstrates that aid eligibility has a positive effect on average exam scores even prior to the receipt of any reimbursements, suggesting that this effect operates through additional channels besides the direct increase in LEA revenue.

2.5.1 LEA Expenditures

Table 2.3 displays results from the estimation of Equation 2.4.1 regarding the effect of transitional aid eligibility on LEA expenditure behavior at the per pupil level among LEAs within 0.4 percentage points of at least one of the enrollment thresholds. The dependent variables that correspond to each column in the table are categories of LEA expenditure per pupil in units of 100s of dollars. The rows in the table correspond to estimates that represent the effect of eligibility for each type of transitional aid (i.e. future aid, A, B, or C), as determined by the current charter enrollment share and its first three lags. These coefficients can be identified separately from one another because a given LEA may be eligible for any possible combination of aid types, and therefore this environment provides an interesting setting to explore the dynamic effects of aid eligibility.[12] The dependent variable in the first column

[12] I have also explored specifications that allow for interactions between the different types of aid, but I did not find significant effects beyond those displayed in this section.

is current expenditure per pupil, which is defined as total expenditure on all day-to-day school district operations.[13] The coefficient on future aid eligibility in the first column is not statistically distinguishable from zero, which is expected given that this coefficient is identified from LEAs that did not actually receive any aid payments. However, the coefficients on each of the three indicators for current aid eligibility are positive, although only the effects of eligibility for Transitional Aid B and Transitional Aid C are statistically significant.

The second and third columns of Table 2.3 break down current expenditure into two relevant components. These are namely instructional expenditure per pupil, which is defined as expenditure related to the direct interaction between students and teachers in the classroom (including teacher salaries and benefits), and expenditure on other support services. From the estimates in the second row of the table, we see that most of the average increase in current expenditure per pupil was allocated directly toward instruction, which we might expect to be a particularly important determinant of student achievement. The coefficient in the second row of the second column is statistically significant at a 95% confidence level, and the quantitative interpretation of this estimate is that eligibility for Transitional Aid A increases instructional expenditure per pupil by over $200.[14] The results displayed in rows three and four suggest that LEAs allocate the larger increases in current expenditure per pupil associated with Transitional Aid B and Transitional Aid C toward both instruction and support services, although none of these coefficients are individually statistically significant.

The total support services expenditure category that corresponds to the third column of Table 2.3 consists of a wide variety of subcategories. For example, this

[13]This excludes debt repayments and capital outlays.

[14]This represents about 1.5% of average instructional expenditure per pupil in this sample.

Table 2.3: Expenditure Per Pupil

	(1) Current	(2) Instructional	(3) Service	(4) Business/Other
Eligibility for Future Aid	-1.145	-0.159	-1.232	-1.267*
	(4.754)	(2.544)	(2.191)	(0.680)
Aid Eligibility A	2.516	2.081**	0.588	-0.078
	(3.150)	(0.816)	(2.814)	(0.464)
Aid Eligibility B	5.304*	2.649	3.264	-0.095
	(2.677)	(1.505)	(2.037)	(0.374)
Aid Eligibility C	6.141*	2.411	3.771	-0.481
	(3.134)	(2.146)	(2.821)	(0.549)
R-squared	0.822	0.890	0.553	0.441
N	51	51	51	51

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of LEA-grade-year results regarding the effect of transitional aid eligibility on various categories of LEA expenditure per pupil. The dependent variable in the first column corresponds to total expenditure per pupil on day-to-day LEA operations, while the dependent variable in the second column corresponds to expenditure per pupil on the direct interaction between students and teachers in the classroom. The dependent variable in the third column represents per pupil expenditure on non-classroom support services, and the dependent variable in the fourth column is per pupil expenditure on business/other expenses. Total district enrollment is included as a control, and the specification also includes district and year fixed effects. Standard errors are clustered at the LEA level.

includes pupil support services (e.g. counseling and nursing services) that we might expect to be non-trivial determinants of household enrollment decisions. However, an example of a subcategory of total support services that is less likely to have direct implications for household decisions is classified as business/central/other expenditures by the LEA expenditure survey, and this is the dependent variable for the results displayed in the fourth column. Examples of services that fall into this category include fiscal services, evaluation services, and information services. The estimated coefficient on the future aid eligibility indicator displayed in the fourth column of Table 2.3 is negative and statistically significant at a 90% confidence level. I interpret this as evidence that LEAs compensated for the lack of actual increases in revenue associated with this treatment category by decreasing expenditure per pupil on a category that did not have direct implications for households. In contrast, the coefficient on the future aid eligibility indicator in the column that corresponds to instructional expenditure per pupil is almost zero. The estimated coefficients on each of the current aid eligibility indicators in the fourth column are also negative, although they are smaller in magnitude and not statistically significant. This suggests that LEAs did not increase business/central/other expenditure per pupil even when they were eligible to receive reimbursements for increased tuition payments to charter schools.

Given that procedures to compute optimal bandwidths in the regression discontinuity literature have not been extended to settings with multiple assignment variables, it is especially important to demonstrate that these results are robust to alternative choices of bandwidth. This is accomplished in Figure 2.2, which displays results for each of the dependent variables included in Table 2.3 across bandwidths ranging from 0.3-0.55 percentage points around each charter enrollment eligibility threshold. The values plotted at 0.4 on the x-axis of each figure correspond exactly to the coefficients displayed in Table 2.3. These figures clearly demonstrate that the results displayed

Figure 2.2: Expenditure per Pupil ($100s) Bandwidth Plots

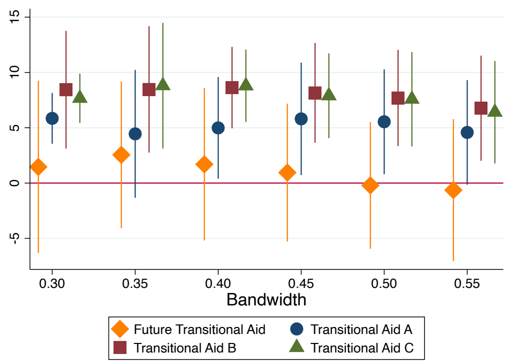

(a) Current

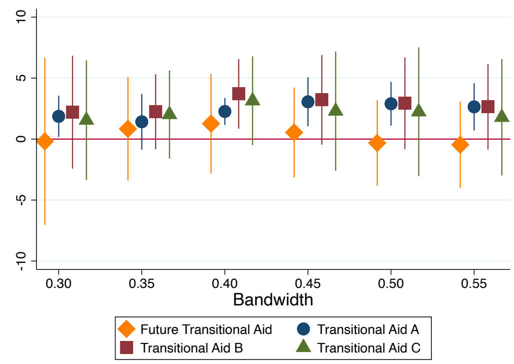

(b) Instructional

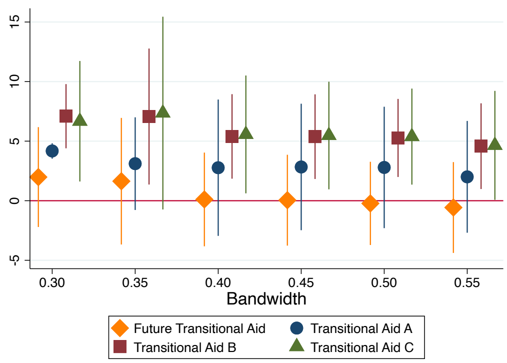

(c) Total Support Service

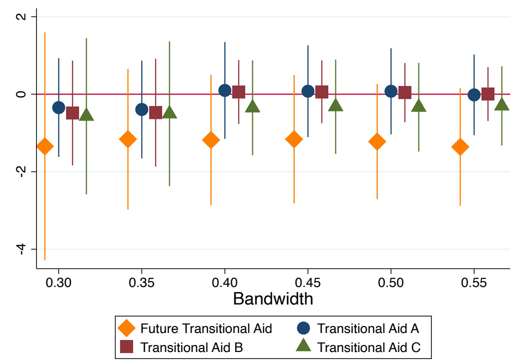

(d) Business/Other

These figures display the coefficients and 95% confidence intervals on indicators that represent each type of transitional aid eligibility on various categories of expenditure in regression models that include observations located within various bandwidths around the threshold value of charter enrollment required to be eligible for transitional aid (2%). The bandwidths are displayed on the x-axis, while the level of expenditure per pupil measured in $100s of dollars is displayed on the y-axis.

in Table 2.3 are qualitatively robust to variation in the bandwidth used to generate them, which provides further support for the insights drawn from this analysis.

2.5.2 Student Achievement

Given the effect of aid eligibility on LEA expenditure behavior, it is natural to wonder whether increases in instructional expenditure due to transitional aid are associated with meaningful improvements in academic outcomes. Table 2.4 displays the coefficients on the indicators that represent each type of aid eligibility estimated from Equation 2.4.1, where the dependent variable is average performance on a statewide math exam at the district-grade-year level. The sample includes the same LEAs used to generate the expenditure results, but the sample size is larger because mean exam performance is reported at the grade level for grades 3-8. The most important takeaway from the table is that the estimated coefficient on the indicators that represent each of the three types of current aid eligibility are positive and statistically significant at a 95% confidence level. These estimates are also moderately large in magnitude, especially given that they represent a district-grade-year level mean. Relative to familiar estimates from the teacher value added literature, each of these coefficients corresponds to an increase of over two-thirds of a percentage point in annual earnings at age 28 (Chetty *et al.*, 2014b). Overall, this suggests that despite the costly nature of transitional aid policies from the perspective of a state government, these policies are potentially capable of generating substantial improvements in average student performance on statewide exams.

The fact that these coefficients on different types of aid eligibility are separately identified also allows us to consider dynamic effects of eligibility. While none of the coefficients are statistically distinct from one another, the magnitude of the estimates increases as we consider aid eligibility further removed from a given increase in the

Table 2.4: Average Math Score

Eligibility for Future Aid	0.020
	(0.039)
Aid Eligibility A	0.066***
	(0.021)
Aid Eligibility B	0.095**
	(0.039)
Aid Eligibility C	0.098**
	(0.042)
R-squared	0.668
N	306

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of LEA-grade-year results regarding the effect of transitional aid eligibility on the average math test score across an LEA. Controls include a second-degree polynomial of the current charter share as well as each of the first three lags, the total number of students tested in a given district-grade-year unit, and total district enrollment across all grade levels. Controls for demographic characteristics also include variables that represent the shares of tested students that are black, Hispanic, and economically disadvantaged. The specification also includes district, year, and grade level fixed effects. Standard errors are clustered at the LEA level.

charter share. This result is somewhat surprising given that the magnitude of reimbursements decreases in periods farther from a given increase in charter sector enrollment.[15] I interpret this as evidence that over time, transitional aid eligibility allows LEAs to adjust their cost structures in a manner conducive to generating further improvements in academic outcomes. Finally, the coefficient on eligibility for future aid is also positive, although it is smaller in magnitude and not statistically significant. Given that this represents a treatment category in which LEA administrators have not actually received any aid payments, this is suggestive that the effect of transitional aid eligibility on student achievement may operate through additional channels besides the direct increase in LEA revenue associated with reimbursements, although I provide further support for this idea in Section 2.5.2.

Figure 2.3 provides a means to visualize whether the effects of transitional aid eligibility on average math performance at the district-grade-year level are robust to alternative choices of bandwidth around the baseline choice of 0.4 percentage points that was used to generate the primary results displayed in Table 2.4. It is clear that estimates are qualitatively similar for each of the bandwidths displayed along the x-axis of the figure, in the sense that I always find positive effects of each of the three types of current aid eligibility. This provides further support for the insights drawn from consideration of the primary results. Results regarding the effect of transitional aid eligibility on average English/Language Arts exam scores at the district-grade-year level can be found in Appendix J.

[15] Recall that for example, Transitional Aid A consists of a reimbursement for 80% of increased tuition payments to charter schools, while Transitional Aid C only consists of a reimbursement for 40% of increased tuition payments to charter schools.

Figure 2.3: Average Math Score Bandwidth Plot

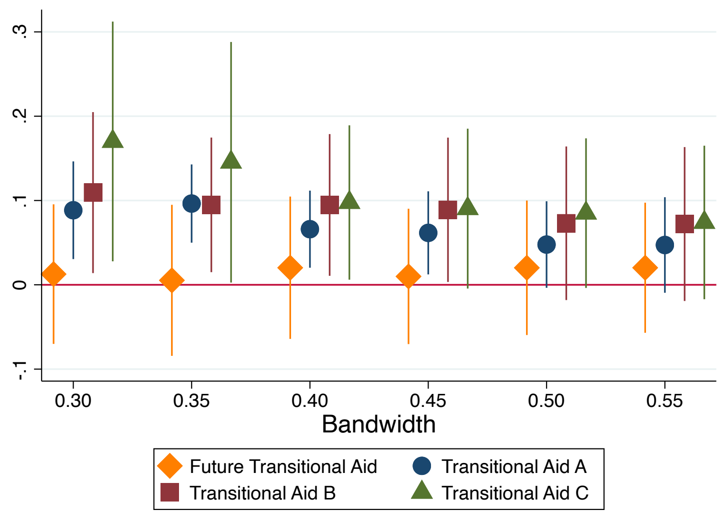

This figure displays the estimated effect and 95% confidence interval of each type of transitional aid at various bandwidths around the enrollment thresholds for aid eligibility (2%). The dependent variable is average LEA math exam performance at the district-grade-year-level, and bandwidths are displayed on the x-axis.

Future Aid Eligibility

This section presents evidence that the overall effect of transitional aid eligibility on student achievement operates through channels in addition to the direct increase in LEA revenue. Recall that eligible New York LEAs receive transitional aid payments for increases in the number of resident pupils enrolled in the charter sector, but these payments begin one year after a given increase occurs. In other words, administrators of an eligible LEA are potentially aware of eventual transitional aid payments for a year before the first of these reimbursements actually occurs. If I find an increase in test scores associated with aid eligibility during the year in which an LEA has not actually received an aid payment yet, this can be interpreted as evidence of an indirect effect of aid eligibility associated with changes in the incentives of school administrators.

Figure 2.4: Average Math Score Event Study

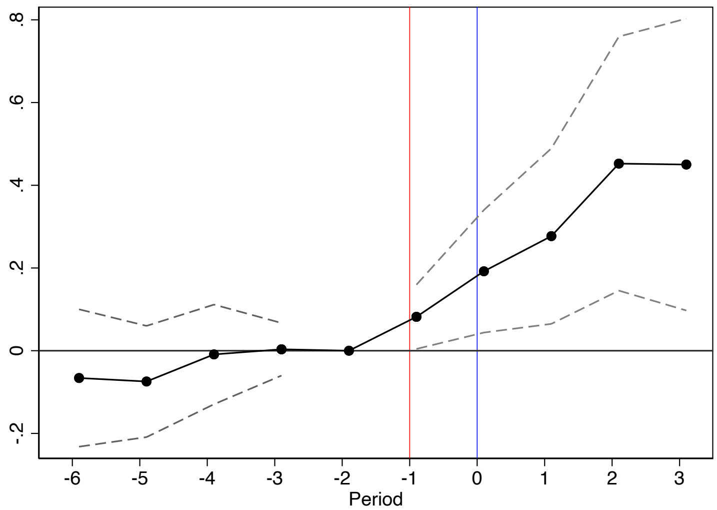

This figure displays a set of coefficients on relative period indicators from an event study framework in which the dependent variable is average LEA math exam performance at the district-grade-year-level. Period 0, marked by the blue line, represents the first period in which an LEA is actually eligible to receive transitional aid payments. Period -1, marked by the red line, represents the period in which an LEA becomes aware that it is eligible to receive transitional aid in the following period. The coefficient on period -2 has been normalized to zero, and the dashed lines represent 95% confidence intervals.

Figure 2.4 displays results from estimating Equation 2.4.2, where the dependent variable is mean performance on statewide math exams at the district-grade-year level. The sample includes all LEAs that became eligible for transitional aid for the first time from 2013-2017, and units on the y-axis can be interpreted as a percentage of a standard deviation of the aggregate distribution of test scores. Period 0, denoted by the blue line, represents the first period in which an LEA is eligible to receive current aid. Period -1, denoted by the red line, represents the period in which LEA administrators are potentially aware that they will be eligible to receive aid in the following period, but have not yet received any actual reimbursements.

The positive coefficient on period -1, which is statistically significant at a 95%

confidence level, is evidence that transitional aid eligibility has a positive effect on student achievement even prior to the receipt of any aid payments. I interpret this as evidence that the effect of transitional aid on academic outcomes does not operate solely through changes in LEA revenue, but also because aid reduces the incentives of school administrators to compete for students through non-academic channels. That being said, the effect of aid eligibility on statewide math exam performance increases further at the blue line and beyond, suggesting that increases in revenue and changes to cost structure over the medium to long-term may also be important mechanisms through which the effect of aid eligibility operates. The coefficients on relative periods prior to period -2, which are very small in magnitude and statistically indistinguishable from zero, are commonly interpreted as evidence in support of the parallel trends assumption required for identification of the treatment effects, although Sun and Abraham (2021) cautions that this test may not be valid in the presence of heterogeneous treatment effects across treatment cohorts.[16] The analogous event study figure for English/Language Arts exam scores can be found in Appendix J.

2.6 Conclusion

Policies that expand the scope for school choice, and charter school expansions in particular, have enjoyed tremendous popularity across the United States over the past two decades. One of the most standard arguments in support of school choice is that competition for students should incentivize incumbent public schools to improve their academic quality. However, if incumbent schools respond to competition through channels not typically considered in these arguments (e.g. through an increase in

[16]It is worth noting that results are qualitatively consistent when I use the alternative estimator proposed in Sun and Abraham (2021), which relaxes this treatment effect homogeneity assumption.

the provision of non-academic services), this may have unintended consequences for student achievement. Further, charter school entry may create excess costs in the public school system due to the fact that some costs of school operation are fixed at least in the short-term, which can further decrease the efficacy of competition as a tool for improving academic outcomes. A small number of states have experimented with policies that provide transitional aid for increases in charter school enrollment to traditional public school districts in order to relieve some of the fiscal burden created by growth in the charter sector. I show in the context of a stylized model of school competition how the presence of transitional aid can increase test scores in both public and charter schools, and can diminish the incentive for school administrators to compete for students through the provision of consumption amenities that do not help students prepare for statewide exams.

I provide empirical support for the qualitative implications of this model by leveraging plausibly exogenous eligibility cutoffs for transitional aid in the state of New York to demonstrate that aid eligibility is associated with an increase in instructional expenditure per pupil, as well as in the average LEA performance of students in grades 3-8 on statewide math exams. Finally, I demonstrated that aid eligibility was associated with improved academic performance even prior to the receipt of any actual payments. Importantly, these results suggest that transitional aid payments made to LEAs for increases in charter enrollment may improve academic performance not only due to the direct increase in LEA revenue, but also because they change the incentives of school administrators in a manner that makes competition a more effective tool for improving academic outcomes. Therefore, while transitional aid policies may be costly for a given state to implement, there are some contexts in which we should expect them to be socially beneficial.

REFERENCES

Abdulkadiroglu, A., J. Angrist, S. Dynarski, T. J. Kane and P. Pathak, "Accountability and Flexibility in Public Schools: Evidence from Boston's Charters and Pilots", Quarterly Journal of Economics **126**, 2, 699–748 (2011).

Abdulkadiroglu, A., P. A. Pathak, J. Schellenberg and C. R. Walters, "Do Parents Value School Effectiveness?", American Economic Review **110**, 5, 1502–1539 (2020).

Aldashev, G. and T. Verdier, "Goodwill Bazaar: NGO Competition and Giving to Development", Journal of Development Economics **91**, 1, 48–63 (2010).

Angrist, J. D., S. R. Cohodes, S. M. Dynarski, P. A. Pathak and C. R. Walters, "Stand and Deliver: Effects of Boston's Charter High Schools on College Preparation, Entry, and Choice", Journal of Labor Economics **34**, 2, 275–318 (2016).

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**, 4, 837–860 (2012).

Angrist, J. D., P. D. Hull, P. A. Pathak and C. R. Walters, "Estimating the Effects of Private School Vouchers in Multidistrict Economies", Quarterly Journal of Economics **132**, 2, 871–919 (2017).

Angrist, J. D., P. A. Pathak and C. R. Walters, "Explaining Charter School Effectiveness", American Economic Journal: Applied Economics **5**, 4, 1–27 (2013).

Arsen, D. and Y. Ni, "The Effects of Charter School Competition on School District Resource Allocation", Education Administration Quarterly **48**, 1, 3–38 (2012).

Athey, S. and G. Imbens, "Design-based Analysis in Difference-In-Differences Settings with Staggered Adoption", Working Paper (2018).

Barseghyan, L., D. Clark and S. Coate, "Peer Preferences, School Competition, and the Effects of Public School Choice", American Economic Journal: Economic Policy **11**, 4, 124–158 (2019).

Bau, N., "Estimating an Equilibrium Model of Horizontal Competition in Education", Working Paper (2019).

Baude, P. L., M. Casey, E. A. Hanushek and S. G. Rivkin, "The Evolution of Charter School Quality", Economica **87**, 345, 158–189 (2020).

Bettinger, E. P., "The Effect of Charter Schools on Charter Students and Pubic Schools", Economics of Education Review **24**, 133–147 (2005).

Beuermann, D. W., C. K. Jackson, L. Navarro-Sola and F. Pardo, "What is a Good School and Can Parents Tell? Evidence on the Multidimensionality of School Output", IDB Working Paper Series (2019).

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**, 1, 31–56 (2007).

Bifulco, R. and R. Reback, "Fiscal Impacts of Charter Schools: Lessons from New York ", Education Finance and Policy **9**, 1, 86–107 (2014).

Booker, K., T. R. Sass, B. Gill and R. Zimmer, "The Effects of Charter High Schools on Educational Attainment", Journal of Labor Economics **29**, 2, 377–415 (2011).

Borusyak, K. and X. Jaravel, "Revisiting Event Study Designs", Working Paper (2018).

Callaway, B. and P. H. Sant'Anna, "Difference-in-Differences with Multiple Time Periods", Working Paper (2019).

Campos, C. and C. Kearns, "The Impact of Neighborhood School Choice: Evidence from Los Angeles' Zones of Choice", Working Paper (2021).

Cardon, J. H., "Strategic Quality Choice and Charter Schools", Journal of Public Economics **87**, 729–737 (2003).

Castaneda, M. A., J. Garen and J. Thornton, "Competition, Contractibility, and the Market for Donors to Nonprofits", The Journal of Law, Economics, and Organization **24**, 1, 215–246 (2008).

Chetty, R., J. N. Friedman and J. E. Rockoff, "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates", American Economic Review **104**, 9, 2593–2632 (2014a).

Chetty, R., J. N. Friedman and J. E. Rockoff, "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood", American Economic Review **104**, 9, 2633–2679 (2014b).

Choi, J.-y. and M.-J. Lee, "Regression Discontinuity with Multiple Running Variables Allowing Partial Effects", Political Analysis **26**, 258–274 (2018).

Cook, J. B., "The Effect of Charter Competition on Unionized District Revenues and Resource Allocation", Journal of Public Economics **158**, 48–62 (2018).

de Chaisemartin, C., "A Note on Instrumented Difference-in-Differences", Working Paper (2010).

de Chaisemartin, C. and X. D'Haultfoeuille, "Fuzzy Differences-In-Differences", Review of Economic Studies **85**, 2, 999–1028 (2018).

de Chaisemartin, C. and X. D'Haultfoeuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects", American Economic Review **110**, 9, 2964–2996 (2020).

Deshpande, M. and Y. Li, "Who Is Screened Out? Application Costs and the Targeting of Disability Programs", American Economic Journal: Economic Policy **11**, 4, 213–48 (2019).

Dobbie, W. and R. G. Fryer Jr., "The Medium-Term Impacts of High-Achieving Charter Schools on Non-Test Score Outcomes", Journal of Political Economy **123**, 5, 985–1037 (2015).

Epple, D. and R. Romano, "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects", American Economic Review **88**, 33–62 (1998).

Epple, D., R. Romano and R. Zimmer, "Charter Schools: A Survey of Research on Their Characteristics and Effectiveness", NBER Working Paper Series **21256** (2015).

Feigenbaum, S., "Competition and Performance in the Nonprofit Sector: The Case of US Medical Research Charities", The Journal of Industrial Economics **35**, 3, 241–253 (1987).

Ferreyra, M. M., "Estimating the Effects of Private School Vouchers in Multidistrict Economies", American Economic Review **97**, 3, 789–817 (2007).

Ferreyra, M. M. and G. Kosenok, "Charter School Entry and School Choice: The Case of Washington, D.C.", Journal of Public Economics **159**, 160–182 (2018).

Figlio, D. and S. Loeb, "School Accountability", Handbook of the Economics of Education **3**, 383–421 (2011).

Friedman, M., "The Role of Government in Education", Economics and the Public Interest (1955).

Gilraine, M., U. Petronijevic and J. D. Singleton, "Horizontal Differentiation and the Policy Effect of Charter Schools", American Economic Journal: Economic Policy **13**, 3, 239–276 (2021).

Glaeser, E. L., "The Governance of Not-For-Profit Firms", NBER Working Paper **8921** (2002).

Glaeser, E. L. and A. Shleifer, "Not-For-Profit Entrepreneurs", Journal of Public Economics **81**, 99–115 (2001).

Glomm, G., D. Harris and T.-F. Lo, "Charter School Location", Economics of Education Review **24**, 451–457 (2005).

Goodman-Bacon, A., "Difference-in-Differences with Variation in Treatment Timing", Working Paper (2019).

Hansmann, H. B., "The Role of Nonprofit Enterprise", The Yale Law Journal **89**, 5, 835–901 (1980).

Hanushek, E. A., "Assessing the Effects of School Resources on Student Performance: An Update", Educational Evaluation and Policy Analysis **19**, 2, 141–164 (1997).

Hanushek, E. A., J. F. Kain, S. G. Rivkin and G. F. Branch, "Charter School Quality and Parental Decision Making with School Choice", Journal of Public Economics **91**, 823–848 (2007).

He, G. and S. Wang, "Do College Graduates Serving as Village Officials Help Rural China?", American Economic Journal: Applied Economics **9**, 4, 186–215 (2017).

Hoxby, C., "School Choice and School Productivity. Could School Choice Be a Tide that Lifts All Boats?", The Economics of School Choice pp. 287–342 (2003).

Imbens, G. W. and T. Lemieux, "Regression Discontinuity Designs: A Guide to Practice", Journal of Econometrics **142**, 2, 615–635 (2008).

Imberman, S. A., "The Effect of Charter Schools on Achievement and Behavior of Public School Students", Journal of Public Economics **95**, 850–863 (2011).

Jackson, C. K., R. C. Johnson and C. Persico, "The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms", The Quarterly Journal of Economics **131**, 1, 157–218 (2016).

Jackson, K. C., "School Competition and Teacher Labor Markets: Evidence from Charter School Entry in North Carolina", Journal of Public Economics **96**, 431–448 (2012).

Jacob, B., B. McCall and K. Stange, "College as Country Club: Do Colleges Cater to Students Preferences for Consumption?", Journal of Labor Economics **36**, 2, 309–348 (2018).

Jacob, B. A. and L. Lefgren, "What Do Parents Value in Education? An Empirical Investigation of Parents' Revealed Preferences for Teachers", Quarterly Journal of Economics **122**, 4, 1603–1637 (2007).

Kinsler, J., "Understanding the Black-White School Discipline Gap", Economics of Education Review **30**, 6, 1370–1383 (2011).

Kofoed, M. S. and C. Fawson, "A Neighborly Welcome? Charter School Entrance and Public School Competition on the Capital Margin", Public Choice (2020).

Ladd, H. F., C. T. Clotfelter and J. B. Holbein, "The Growing Segmentation of the Charter School Sector in North Carolina", Education Finance and Policy **12**, 4, 536–563 (2017).

Ladd, H. F., C. T. Clotfelter and M. Turaeva, "Choosing Charter Schools in North Carolina: What Do Parents Value?", CALDER Working Paper **196-0618-1** (2018).

Ladd, H. F. and J. Singleton, "The Fiscal Externalities of Charter Schools: Evidence from North Carolina", Education Finance and Policy **15**, 1, 191–208 (2020).

Lakdawalla, D. and T. Philipson, "The Nonprofit Sector and Industry Performance", Journal of Public Economics **90**, 8-9, 1681–1698 (2006).

Macartney, H., "The Dynamic Effects of Educational Accountability", Journal of Labor Economics **34**, 1, 1–28 (2016).

MacLeod, W. B. and M. Urquiola, "Is Education Consumption or Investment? Implications for School Competition", Annual Review of Economics **11**, 563–589 (2019).

Mann, B. A. and P. Bruno, "The Effects of Charter School Enrollment Losses and Tuition Reimbursements on School Districts: Lifting Boats or Sinking Them?", Educational Policy **5**, 1–30 (2020).

McMillan, R., "Competition, Incentives, and Public School Productivity", Journal of Public Economics **88**, 1871–1892 (2004).

Mehta, N., "Competition in Public School Districts: Charter School Entry, Student Sorting, and School Input Determination", International Economic Review **58**, 4, 1089–1116 (2017).

National Center for Educational Statistics, "Public Charter School Enrollment", The Condition of Education 2019 pp. 1–4 (2019).

Nechyba, T. J., "Mobility, Targeting, and Private-School Vouchers", American Economic Review **90**, 1, 130–146 (2000).

Neilson, C., "Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students", Working Paper (2020).

Ost, B., "How Do Teachers Improve? The Relative Importance of Specific and General Human Capital", American Economic Journal: Applied Economics **6**, 2, 127–151 (2014).

Papke, L. E. and J. M. Wooldridge, "Econometric Methods for Fractional Response Variables with an Application to 401(k) Plan Participation Rates", Journal of Applied Econometrics **11**, 6, 619–632 (1996).

Papke, L. E. and J. M. Wooldridge, "Panel Data Methods for Fractional Response Variables with an Application to Test Pass Rates", Journal of Econometrics **145**, 1-2, 121–133 (2008).

Pope, D. G. and J. C. Pope, "The Impact of College Sports Success on the Quantity and Quality of Student Applications", Southern Economic Journal **75**, 3, 750–780 (2009).

Reardon, S. and J. Robinson, "Regression Discontinuity with Multiple Rating-Score Variables", Journal of Research on Educational Effectiveness **5**, 83–104 (2012).

Ridley, M. and C. Terrier, "Fiscal and Education Spillovers from Charter School Expansion", NBER Working Paper Series **25070** (2018).

Rose-Ackerman, S., “Charitable Giving and Excessive Fundraising”, Quarterly Journal of Economics **97**, 2, 193–212 (1982).

Sacerdote, B., “Peer Effects in Education: How Might They Work, How Big Are They, and How Much Do We Know Thus Far?”, Handbook of the Economics of Education **3**, 249–277 (2011).

Salop, S. C., “Monopolistic Competition with Outside Goods”, The Bell Journal of Economics **10**, 1, 141–156 (1979).

Segal, C., “Classroom Behavior”, The Journal of Human Resources **43**, 4, 783–814 (2008).

Singleton, J. D., “Putting Dollars Before Scholars? Evidence from For-Profit Charter Schools in Florida”, Economics of Education Review **58**, 43–54 (2017).

Singleton, J. D., “Incentives and the Supply of Effective Charter Schools”, American Economic Review **109**, 7, 2568–2612 (2019).

Slungaard Mumma, K., “The Effect of Charter School Openings on Traditional Public Schools in Massachusetts and North Carolina”, Working Paper (2020).

Spence, A. M., “Monopoly, Quality, and Regulation”, The Bell Journal of Economics **6**, 2, 417–429 (1975).

Sun, L. and S. Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects”, Journal of Econometrics **225**, 175–199 (2021).

Tedin, K. L. and G. R. Weiher, “Does Choice Lead to Racially Distinctive Schools? Charter Schools and Household Preferences”, Journal of Policy Analysis and Management **21**, 1, 79–92 (2002).

Thornton, J., “Nonprofit Fund-Raising in Competitive Donor Markets”, Nonprofit and Voluntary Sector Quarterly **35**, 2, 204–224 (2006).

Wagner, J., “Unobserved Firm Heterogeneity and the Size-Exports Nexus: Evidence from German Panel Data”, Review of World Economics **139**, 161–172 (2003).

Walters, C., “The Demand for Effective Charter Schools”, Journal of Political Economy **126**, 6, 2179–2223 (2018).

Wong, V. C., P. M. Steiner and T. D. Cook, “Analyzing Regression-Discontinuity Designs with Multiple Assignment Variables: a Comparative Study of Four Estimation Methods”, Journal of Educational and Behavioral Statistics **38**, 2, 107–141 (2013).

APPENDIX A

ALTERNATIVE MODEL: WELFARE IMPLICATIONS

This section presents a simplified version of the model discussed in Section 1.2, with the purpose of discussing the welfare implications of a policy that increases the number of competitors in a market for education, such as the removal of a cap on charter schools. In order to do so, I include an entry stage prior to administrators making input decisions, and demonstrate that in the unique symmetric equilibrium of this model, the number of competitors that enter the market may be higher or lower than the socially optimal level. The primary distinction between this version and the model in Section 1.2 is in the objective function of the school administrator, which in this case can be seen below:

$$S(I, X) \times (\alpha I - \psi X) \tag{A.1}$$

Recall that the term in parentheses represents the test score of an individual student, and as this is multiplied by enrollment, this function can be thought of as representing total education production. This is a simplification because the administrator can not spend on perquisite consumption, which the results in Section 1.5 suggest was not particularly important in the North Carolina context. In solving for the unique symmetric equilibrium, the decision rules for the per pupil level of instructional expenditure I^* and the level of consumption amenities X^* have the same qualitative relationship with the number of competitors J as in the case of the model in Section 1.2 in which perquisite consumption was zero.[1] Suppose that $\bar{w}$ represents the fixed utility associated with an outside option that a potential school administrator would have to forego in order to open a school. Then, in order to find the number of schools that enter the market in a competitive equilibrium, I substitute the equilibrium decision rules back in to the objective function and solve for the number of competitors J^* that satisfies the following condition:

$$S(J^*)\left[\alpha I(J^*) - \psi X(J^*)\right] = \bar{w} \tag{A.2}$$

This equation has a unique solution, which is given by the following:

$$J^* = \sqrt{\frac{\tau}{\zeta\nu\psi + \eta(\alpha\zeta + \bar{w})}} \tag{A.3}$$

The production of education provides value to society (Friedman, 1955), and this model predicts that an increase in school competition leads to lower test scores. However, because competition also allows households to take advantage of increased provision of consumption amenities and shorter distances to school, the overall welfare effect of competition is not obvious. This effect will depend on whether the existing number of schools in the market is greater than or less than the socially optimal level. Therefore, I solve for the socially optimal number of schools in a given market for education by maximizing the total welfare function with respect to the number of schools. This function is given by the following:

$$U(J) + J(V(J) - \bar{w}) \tag{A.4}$$

[1]Recall that consumption amenities were increasing in the number of competitors, and both instructional expenditure per pupil and test scores were decreasing in the number of competitors.

In this equation, $U(J)$ represents utility summed over all households on the circle, and $V(J)$ represents the objective function of a school administrator in equilibrium as a function of the number of competitors J. The welfare function is globally concave under the following assumption regarding the household preference parameter on consumption amenities and the corresponding cost parameter from the objective function of the administrator:

$$\eta > \frac{4}{5}\psi \tag{A.5}$$

Solving for the socially optimal level of schools J_s, and comparing this value to the competitive equilibrium level of schools J^*, gives us the following proposition:

Proposition 3 $\frac{J_s}{J^*} = \sqrt{5\eta - 4\psi}$

This result states that if the value that households place on consumption amenities is sufficiently high (low) relative to the cost of their provision, then the socially optimal number of schools will be greater (less) than the competitive equilibrium level. The reason that the equilibrium number of schools may not be equal to the social optimum in this model is because of the business-stealing effect. When a new school chooses to enter the market, it steals students from incumbent schools, but the new entrant does not internalize this cost when making an entry decision. This is why there is "too much" variety in equilibrium in standard versions of circular city models. However, in this case, the decision of a new school to enter the market not only makes competitors worse off, but it also affects the utility of the households that enroll their students at these competing schools. These households experience a cost through a decrease in instructional expenditure per pupil, but this may be more than offset by increased provision of consumption amenities and shorter distances to school. Since these benefits to households can potentially dominate overall, the relationship between the competitive equilibrium and socially optimal level of schools depends on the relationship between η and ψ.

This result gives us an interesting way to frame policies such as the 2011 removal of the charter school cap in North Carolina. The policy change led to an increase in the number of charter schools in operation across the state, which implies that under the cap, the number of schools must have been constrained to a level lower than the competitive equilibrium. Under these circumstances, if $\eta > \frac{4}{5}\psi + \frac{1}{5}$, or in other words if the marginal value of consumption amenities to the household is sufficiently high relative to the marginal cost to the school of providing them, then the removal of the cap will lead to an unambiguous welfare improvement despite the fact that this implies a lower level of test scores across the market. This result is driven by the shorter distances and greater level of consumption amenities that households enjoy as a result of increased competition. However, if $\frac{4}{5}\psi < \eta < \frac{4}{5}\psi + \frac{1}{5}$, then the effect on overall welfare is ambiguous and depends on the relationship between the socially optimal, competitive equilibrium, and cap-constrained number of schools. While this stylized framework is unlikely to accurately capture the full social cost of a decrease in academic achievement, the observation that a policy which facilitates an increase in the number of schooling options could be welfare improving despite an associated decrease in test scores raises important questions about the fundamental objectives of a public school.

APPENDIX B

WORKING CONDITIONS SURVEY ITEMS

Tables B.1 and B.2 list the WCS survey items used to generate the teacher empowerment index and the index of preparation time, respectively, as well as their weights and definitions.

Table B.1: Teacher Empowerment Survey Items

Item	**Weight**	**Definition**
empowerment_b1	0.5239	Selecting instructional materials and resources
empowerment_b3	0.4493	Setting grading and student assessment practices
empowerment_b4	0.5132	Determining the content of in-service professional development programs
empowerment_b5	0.2533	The selection of teachers new to this school
empowerment_b6	0.4429	Establishing student discipline procedures

Table B.2: Preparation Time Survey Items

Item	**Weight**	**Definition**
time_i10	0.0958	Individual planning time
time_i11	0.2670	Collaborative planning time
time_i13	0.5388	Preparation for required federal, state, and local assessments
time_i14	0.5524	Delivery of assessments
time_i15	0.5693	Utilizing results for assessments

APPENDIX C

NORTH CAROLINA READING RESULTS

This section contains tables which summarize the effect of charter competition on standardized reading test scores using both the primary and alternative empirical strategies discussed in the main text in the context of math test scores. In general, the large literature that studies the effect of school accountability systems on student achievement almost always finds more pronounced effects for math scores than for reading scores (Figlio and Loeb, 2011). Table C.1 is analogous to Table 1.4 in the main text, where the dependent variable is now the EOG reading test score. In the first column, we see that while the estimated effect of charter competition on reading scores is negative, the coefficient is smaller in magnitude than the corresponding estimate for math, and it is not statistically significant. When we include a large vector of control variables in the second column, the estimate is statistically significant, but is roughly half the size of the negative effect of competition on math test scores. In the third column, when we restrict the sample to only the seventh and eighth graders that had already begun a traditional public middle school prior to the opening of any new charter middle school in their LEA, the estimated coefficient is negative, but not statistically significant.

Table C.1: Standardized Reading Test Score

	(1)	(2)	(3)
New Entrant	-0.020	-0.021*	-0.015
	(0.012)	(0.012)	(0.012)
Other Controls	No	Yes	Yes
7th/8th Grade Only	No	No	Yes
R-squared	0.704	0.708	0.707
N	950,814	950,814	575,376

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized EOG reading test score. The first column corresponds to a regression specification which includes prior year test scores in math and reading as well as application controls. The second column augments this specification wth a large vector of other controls, including gender, race, economic disadvantage status (EDS), grade level, school share black, school share Hispanic, school share EDS, school enrollment, and LEA enrollment. The third column replicates the result from the second column using only the subsample of students that had already begun seventh or eighth grade prior to the opening of a new charter middle school in their LEA. Standard errors are clustered at the LEA level.

Table C.2 explores heterogeneity in the effect of charter competition on reading test scores based on the type of charter schools with which incumbent public schools are competing. The regression specification includes an additional covariate which represents the share of all sixth grade students enrolled in new charter schools of a particular type. The first two columns correspond to for-profit charter schools while the second two columns correspond to charter schools with a STEM focus. The result in the first column of each panel utilizes the full sample, while the second column of each panel restricts the sample to seventh and eighth graders that had already begun a traditional public middle school prior to the opening of any new charter middle school in their LEA. Unlike the math results, we see that in this case the share of charter students attending a for-profit school does not seem to be importantly related to spillover effects of charter competition on reading test scores.

The purpose of the results displayed in Table C.3 is to explore the relationship

Table C.2: Standardized Reading Test Score by Charter Type

	For-Profit		STEM	
	(1)	(2)	(3)	(4)
6th Grade For-Profit Charter Share	-0.011	0.010		
	(0.033)	(0.037)		
6th Grade STEM Charter Share			-0.006	-0.028
			(0.029)	(0.031)
New Entrant	-0.019	-0.014	-0.024	-0.010
	(0.012)	(0.013)	(0.015)	(0.016)
R-squared	0.708	0.707	0.708	0.707
N	950,814	575,376	950,814	575,376

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized EOG reading test score. The covariate of interest is the share of 6th grade students enrolled in charter schools that opened after the policy change attending schools that indicated the intent to contract with a for-profit EMO, or to offer a STEM focus. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment. Each specification also includes cumulative applications received for each charter type (i.e. separate controls for for-profit and not-for-profit applications in columns 1 and 2). The second column of each panel reports results when the sample is restricted to the seventh and eighth grade students that attended sixth grade prior to any new charter school opening in their LEA. Standard errors are clustered at the LEA level.

between variables that represent proposed mechanisms discussed in Section 1.6 and spillover effects of charter school competition on reading test scores. The first column displays the effect of charter competition on end-of-grade reading test scores estimated from the full sample as in Table C.1, but in the remaining columns the sample is restricted to only observations with non-missing data from the working conditions survey (i.e. from years that end in an even number). In the second column, the coefficient on new charter entry is not statistically significant, and therefore there is no significant effect of competition on reading scores to explain. However, it is worth noting that teacher empowerment, the local teacher salary supplement, and the share of total teacher salaries devoted to teachers that specialize in tested subjects appear to be positively related to reading test scores.

C.1 Alternative Strategies

This section displays reading test score results using the instrumented difference-in-differences and stacked difference-in-differences strategies discussed in section 1.4. Table C.4 displays the estimated effect of charter school competition on standardized reading test scores using the stacked difference-in-differences strategy. The estimate is statistically significant, and as with many of the other results in this section, it has the same sign as the corresponding result for math although it is smaller in magnitude.

Table C.5 displays the DDIV result. While the estimated effect of competition has the same sign as the corresponding result for math, it is smaller in magnitude and not statistically significant.

Table C.3: Standardized Reading Test Score Mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
New Entrant	-0.021*	0.009	0.006	0.012	0.015	0.011	0.012	0.009	0.023
	(0.012)	(0.016)	(0.016)	(0.016)	(0.015)	(0.018)	(0.016)	(0.017)	(0.015)
Household Influence			0.039*						0.009
			(0.021)						(0.023)
Teacher Empowerment				0.009***					0.008***
				(0.002)					(0.002)
Local Teacher Supp.					0.002				0.002*
					(0.001)				(0.001)
PP Exp. on School Adm.						-0.006			-0.009
						(0.009)			(0.008)
Math/ELA Salary Share							0.038**		0.033**
							(0.014)		(0.014)
Preparation Time Index								0.004	0.003
								(0.003)	(0.003)
R-squared	0.708	0.710	0.710	0.710	0.710	0.710	0.710	0.710	0.710
N	950,814	473,194	473,194	473,194	473,194	473,194	473,194	473,194	473,194

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized reading end-of-grade test score. The first column displays the primary result from Table C.1, while the second column replicates this result restricting the sample to years in which working conditions survey data is non-missing. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative applications received. Columns 3-8 each include one of the variables that represent proposed mechanisms associated with competition, and the ninth column includes all of these variables jointly. Standard errors are clustered at the LEA level.

Table C.4: Standardized Reading Test Score: Stacked Difference-in-Differences

Treatment Group∗Post	-0.013*
	(0.006)
R-squared	0.704
N	2,011,364

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized reading end-of-grade test score. Baseline controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, and LEA enrollment. Each regression also includes a set of indicator variables that represent the cumulative number of applications that the LEA to which a school is assigned has received since the beginning of the expansion. Standard errors are clustered at the LEA level.

Table C.5: Standardized Reading Test Score: Instrumented Difference-in-Differences

Charter Share	-1.385
	(0.933)
N	950,814

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized reading end-of-grade test score. The treatment variable of interest is the charter share of an LEA, for which we use a treatment indicator as an instrumental variable. Controls include gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative applications received. Standard errors are clustered at the LEA level.

APPENDIX D

CLASSROOM-LEVEL ROBUSTNESS CHECK

The purpose of this section is to explore whether the estimated effect of charter school competition is robust to conditioning on class size, peer ability, and teacher value added. Policies designed to improve academic achievement through a reduction in class sizes have enjoyed considerable popularity across the United States (Hanushek, 1997), and it seems plausible that losing students to charter competition might allow traditional public schools to offer smaller classes. Further, a number of studies have found peer quality to be an important determinant of academic performance (Sacerdote, 2011). To the extent that there is selection on academic ability into the charter school sector in this context, changes in peer ability could also be a channel through which spillover effects of charter school competition are operating. In order to address these issues, I restrict my sample to only the individuals who can be matched to a classroom.[1]

Also, a large literature has found that teacher quality is an important determinant of student achievement (Chetty *et al.*, 2014a). Therefore, if an expansion of the charter sector substantially changes the composition of teachers employed at traditional public schools, for example because some teachers switch to charters or leave for employment opportunities in other areas, this could have important implications for academic achievement. In order to investigate this possibility, I further restrict the classroom-matched samples to only those classrooms with teachers that were also employed by the North Carolina Department of Public Instruction from 2007-2009 (the two school years prior to the beginning of my sample). I construct a measure of teacher value added by first estimating the following student-level regression model on observations from only these two out-of-sample years:

$$y_{igjdt} = \beta_0 + \beta_1 X_{igjdt} + \delta_t + \epsilon_{ijgdt} \tag{D.1}$$

In this case, y_{ijgdt} represents a standardized test score for student i in grade g enrolled in school j in district d at time t. The vector X_{ijgdt} contains a large set of individual, school, and district characteristics, most importantly including lagged test scores to account for heterogeneity in academic ability. I then compute residuals for each student, and take means at the teacher level. Finally, I standardize this teacher-level measure such that units can be interpreted as a percentage of a standard deviation.

Table D.1 summarizes the results described above when the dependent variable is the standardized math test score. In the first column, the estimated effect of competition using only the classroom-matched sample is qualitatively similar but larger in magnitude than the result presented in the main text. In the second and third columns, conditioning on class size and both the mean and standard deviation of lagged math test scores for all other students assigned to the classroom does not substantially change the coefficient on charter entry. The fourth column reproduces this result using only the sample of students matched to teachers for whom I can estimate out-of-sample value added. Finally, in the fifth column when teacher value added is also included as a control, we see that the estimated coefficient on competition only decreases by about 16%. Overall, these results suggest that reductions in class size, changes in peer ability, and changes in the composition of public school

[1] This is done separately for math and English/language arts, so the two subsamples do not necessarily include the same students.

Table D.1: Standardized Math Test Score: Classroom and Teacher

	(1)	(2)	(3)	(4)	(5)
New Entrant	-0.061***	-0.058***	-0.055***	-0.058***	-0.049***
	(0.019)	(0.019)	(0.019)	(0.016)	(0.016)
Math Class Size		0.002***	0.001***	0.001	0.001
		(0.000)	(0.000)	(0.001)	(0.001)
Peer Math Mean			0.098***	0.092***	0.092***
			(0.009)	(0.008)	(0.007)
Peer Math SD			0.076***	0.090***	0.082***
			(0.021)	(0.025)	(0.023)
Math Value Added					0.117***
					(0.010)
R-squared	0.703	0.704	0.706	0.708	0.714
N	699,819	699,819	699,819	312,009	312,009

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized math end-of-grade test score. All regressions include baseline controls, school and year fixed effects, and controls for cumulative applications. Results displayed in the first three columns are generated from the portion of students in my main sample which can be matched to a math classroom. The second column augments the first by including a control for math class size, while the third also contains controls for the mean and standard deviation of the prior-year math test score for the other students assigned to the classroom. The results displayed in the fourth and fifth columns are generated from the subsample of students that can be matched to classrooms with teachers for which the out-of-sample value added measure can be estimated. The fifth column augments the fourth by adding a control for teacher value added. Standard errors are clustered at the LEA level.

teachers are unlikely to be the primary mechanisms driving the spillover effects of charter competition on student achievement in the traditional public sector.

Table D.2 summarizes the reading test results, and similar to the other results presented throughout this study, these effects tend to be smaller in magnitude relative to those associated with math. In the first column, we see that the estimated effect of charter competition on reading scores in the classroom-matched sample is negative and statistically significant. The coefficient of interest does not change substantially in the second and third columns when we include controls for class size and peer ability. When the classroom-matched sample is restricted to only students assigned to teachers with non-missing value added in the third column, we no longer find a significant effect of charter competition, and this does not change when we include teacher value added as a control.

Table D.2: Standardized Reading Test Score: Classroom and Teacher

	(1)	(2)	(3)	(4)	(5)
New Entrant	-0.021*	-0.020	-0.021*	-0.010	-0.010
	(0.012)	(0.012)	(0.012)	(0.015)	(0.015)
ELA Class Size		0.001***	0.001***	0.001***	0.001***
		(0.000)	(0.000)	(0.000)	(0.000)
Peer Reading Mean			0.070***	0.068***	0.068***
			(0.002)	(0.003)	(0.003)
Peer Reading SD			-0.003	0.006	0.007
			(0.009)	(0.011)	(0.011)
Reading Value Added					0.033***
					(0.005)
R-squared	0.704	0.704	0.705	0.703	0.703
N	862,838	862,838	862,838	362,690	362,690

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is the standardized reading end-of-grade test score. All regressions include baseline controls, school and year fixed effects, and controls for cumulative applications. Results displayed in the first three columns are generated from the portion of students in my main sample which can be matched to an English/Language Arts classroom. The second column augments the first by including a control for English/Language Arts class size, while the third also contains controls for the mean and standard deviation of the prior-year reading test score for the other students assigned to the classroom. The results displayed in the fourth and fifth columns are generated from the subsample of students that can be matched to classrooms with teachers for which the out-of-sample value added measure can be estimated. The fifth column augments the fourth by adding a control for teacher value added. Standard errors are clustered at the LEA level.

APPENDIX E

NORTH CAROLINA SCHOOL VALUE-ADDED

In this section, I explore the effect of charter school entry on school-level measures of both math and reading value-added. I create these measures in a style similar to the conventional method described in Angrist *et al.* (2017). First, I run individual-level regressions of the following form:

$$y_{ijgdt} = \beta_0 + \beta_1 X_{igt} + \theta_{jdt} + \epsilon_{ijgdt} \tag{E.1}$$

In this case, y_{ijgdt} represents an individual test score in either math or reading, X_{igt} represents a vector of controls that includes prior-year test scores in multiple subjects, and θ_{jdt} is a school-year effect. Next, I take these school-year effects and standardize them within each year such that units can be interpreted as a percentage of a standard deviation. Finally, I use these standardized measures as the dependent variable in the following school-level regression specification:

$$\theta_{jdt} = \beta_0 + \beta_1 Z_{dt} + \beta_2 X_{jdt} + f(App_{dt}) + \epsilon_{jdt} \tag{E.2}$$

Table E.1 summarizes the results. As I would expect given the patterns documented in the Section 1.5, the effect of charter competition on school effectiveness is negative and statistically significant in both subjects. In math, charter entry is associated with a 22.6% of a standard deviation decrease in school value-added, and in reading a decrease of 18.7% of a standard deviation.[1] Overall, these results demonstrate that when we analyze the effect of charter competition on effectiveness at the school level, we find results that are qualitatively similar to those discussed in Section 1.5.

Table E.1: School Value-Added

	(1) Math	(2) Reading
New Entrant	-0.226**	-0.187*
	(0.107)	(0.107)
R-squared	0.534	0.396
N	1,338	1,338

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variables is a standardized measure of school value-added. The first column corresponds to math, while the second column corresponds to reading. Controls include school demographics and enrollment, LEA enrollment and average teacher experience, and cumulative applications received. Standard errors are clustered at the LEA level.

[1] I also find that after the entry of a new charter school serving sixth graders, the median value added of pre-expansion charters serving sixth graders in the same district increases by 0.5665 in math and by 0.1173 in reading. In contrast, median value added of the new charter schools in math (-0.1231) is much lower than either traditional public schools or pre-expansion charters, although it is worth noting new charter value added in reading (0.6128) is comparable to pre-expansion charters and much higher than in traditional public schools. This is somewhat consistent with studies in other contexts which find that charter schools initially perform worse than public schools after entry, but that their performance increases over time as the lowest performers exit the market and the schools that remain improve (Hanushek *et al.*, 2007; Baude *et al.*, 2020).

APPENDIX F

STUDENT DISCIPLINE

Previous work has documented that student safety and discipline are important factors in household enrollment decisions (Tedin and Weiher, 2002), and that measures of student behavior serve as a proxy for non-cognitive skills (Segal, 2008). Imberman (2011) finds that charter competition is associated with a decrease in the reported number of disciplinary infractions, although it is unclear whether this reflects actual changes in student behavior or simply changes in enforcement. In this section, I explore whether there is any effect of charter competition on student discipline in the North Carolina context. One limitation of this data is that schools are only required to report infractions that result in out-of-school suspensions (OSS) to the state, which tend to be very serious in nature (most schools also report less serious infractions, but some do not). Therefore, for the purposes of this exercise I consider only disciplinary incidents that occurred at schools that report at least some less serious disciplinary infractions during each year of my sample.

The results summarized in Table F.1 are generated from a sample at the level of the disciplinary infraction, and correspond to the following regression specification:

$$y_{qigjdt} = \beta_0+\beta_1 Z_{dt}+\beta_2 Previous_{qigjdt}+\beta_3 Serious_{qigjdt}+\beta_4 X_{qigjdt}+f(App_{dt})+\gamma_j+\delta_t+\epsilon_{qigjdt} \tag{F.1}$$

In this specification, y_{qigjdt} is an outcome variable related to disciplinary incident q, for student i in grade g, enrolled in school j, located in district d, during year t. The primary independent variable of interest Z_{dt} is an indicator for whether a new charter school has opened in district d by time t. I also include controls for how many previous disciplinary incidents the student has committed in the same school year, as well as an indicator for whether the offense is especially serious.[1] The dependent variable in the first column is an indicator which takes the value of one if the incident resulted in OSS or an expulsion. The coefficient in the first column is positive, suggesting that new charter entry may have induced schools to hand out more serious consequences for a given infraction, but the estimate is not statistically significant. In the second column, the dependent variable is the natural log of the number of days suspended for a given incident, and the sample is restricted to incidents that resulted in short-term suspensions of 10 days or less.[2] The coefficient in the second column is positive and statistically significant, suggesting that schools at least somewhat increased the length of suspensions handed out for similar infractions after experiencing an increase in charter competition. Overall, these results provide suggestive evidence that an increase in charter competition is associated with slightly harsher disciplinary policies in incumbent public schools, but it is difficult to infer whether these measures acted as an effective deterrent against student misbehavior.

[1]The data center provides a code for serious report categories, for example rape or robbery with a deadly weapon. I consider infractions that do not fall into any of these specific categories to be somewhat less serious in nature.

[2]This represents about 99% of all OSS (Kinsler, 2011).

Table F.1: Student Discipline

	(1) Suspension/Expulsion	(2) Log Days Suspended
New Entrant	0.033	0.046*
	(0.034)	(0.025)
R-squared	0.122	0.130
N	616,786	148,822

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variables are measures related to student discipline. The sample is at the disciplinary incident level. The first column corresponds to a specification in which the dependent variable is an indicator for whether a student received an out-of-school suspension or expulsion for a disciplinary incident. The second column corresponds to number of days suspended, and the sample for this column includes only short-term suspensions. Controls include student gender, race, economic disadvantage status, lagged math and reading test scores, grade level, school demographics and enrollment, LEA enrollment, and cumulative number of charter school applications in the same LEA. Standard errors are clustered at the LEA level.

APPENDIX G

NORTH CAROLINA LEA EXPENDITURES: SIXTH-GRADE TREATMENT

The purpose of this section is to confirm that the LEA-level fiscal results are similar whether treatment is defined as entry of any new charter school, as in the LEA results in the main text, or as new entry of charter schools planning to offer sixth grade, which is the primary treatment considered in the individual and school-level results. I utilize a treatment indicator that corresponds to any charter entry for the LEA-level results in the main text because only LEA-level expenditure information is available, and the negative fiscal externality associated with competition from the charter sector is relevant regardless of the grade levels at which competition takes place. Clearly, the total number of charter school entrants in a given LEA during this period is very highly correlated with the number of entrants offering sixth grade, but if I find qualitatively similar results when I use the sixth grade treatment indicator, it lends further support to the idea that these mechanisms are important in driving the test score results presented in Section 1.5. In Table G.1, the estimated effects of new charter entry are very similar to the results displayed in Table 1.17, although the estimate that corresponds to the local teacher salary supplement is no longer statistically significant.

Table G.1: LEA Expenditures: Sixth-Grade Treatment

	(1) Local Tchr. Supp.	(2) Exp. on School Adm.
New Entrant	-1.796	0.387**
	(1.133)	(0.160)
R-squared	0.959	0.747
N	228	228

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is a measure of LEA-level expenditure. The first column corresponds to the average local teacher supplement offered in a given LEA, while the second column corresponds to per pupil expenditure on school administration, both scaled in terms of hundreds of dollars. Controls include LEA demographics and enrollment, as well as cumulative applications received. Standard errors are clustered at the LEA level.

APPENDIX H

ALTERNATIVE STRATEGIES: ANALYSIS OF MECHANISMS

Figure H.2: School-level Event Study Figures

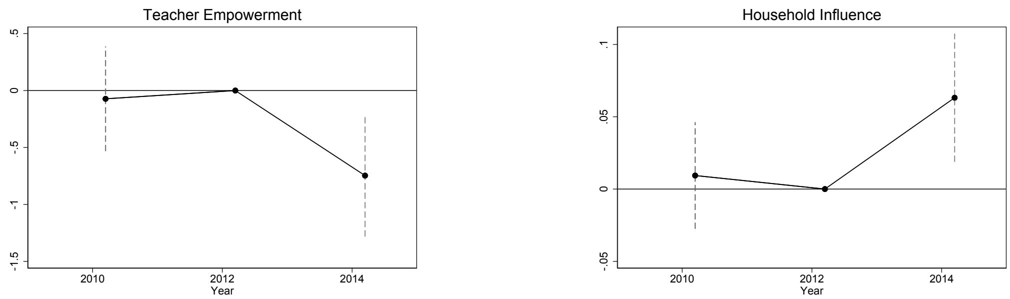

These figures display a set of coefficients on year indicators from a school-level event study framework in which the dependent variable is either the teacher empowerment index or the share of teachers within a school that believe that households are influential decision makers. The coefficient on 2012 is normalized to zero. and the dashed lines represent 95% confidence intervals.

This section contains a discussion of the mechanisms described in Section 1.6 in the context of the instrumented difference-in-differences and stacked difference-in-differences approaches discussed in Section 1.5.

H.1 School

One limitation of the TWFE strategy is that it does not allow us to test the extent to which the magnitude of the effect of charter competition depends on the number of charter switchers relative to traditional public school enrollment in a given district. Therefore, this section summarizes some school-level results from Section 1.6 in the context of the DDIV strategy, specifically the teacher empowerment index and the share of teachers within a school that believe that parents/guardians are influential decision makers.[1] As discussed in Section 1.5, this strategy relies on parallel counterfactual trends assumptions on both the charter share (which was discussed in Section 1.5), and the outcomes of interest. Figure H.2 displays event study figures for each of these outcomes, and I interpret the fact that only the post-treatment coefficient is statistically significant in each figure as evidence in support of this assumption.

Table H.1 displays DDIV results from school-level regressions in which the dependent variable is either the teacher empowerment index or the share of teachers within a school that believe that parents/guardians are influential decision makers. In the first column, we see that the coefficient on the charter share is negative and statistically significant. Scaled by the average change in the charter share of 0.0204, this suggests that the average effect of competition on the teacher empowerment index was roughly -0.5226, which is qualitatively consistent with the TWFE teacher empowerment result. The coefficient in the second column is positive and statistically significant. Scaled by the average change in the sixth grade charter share, this suggests that the share of teachers within a treated school believing that households

[1] I do not report results from the stacked difference-in-differences strategy at the school level because the working conditions survey was only administered in years that ended with an even number. Therefore, there is only one possible period in which a school can be treated, so dynamic treatment effects are not a concern.

Table H.1: School Characteristics: Instrumented Difference-in-Differences

	(1) Teacher Empowerment	(2) Household Influence
Charter Share	-25.616*	3.250**
	(14.897)	(1.627)
N	669	669

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is a school-level characteristic. The first column corresponds to the teacher empowerment factor, while the second column corresponds to the share of teachers within a school that either agree or strongly agree with the statement "Parents/guardians are influential decision makers in this school." The treatment variable of interest is the charter share of an LEA, for which we use a treatment indicator as an instrumental variable. Controls include school demographics and enrollment, LEA enrollment, and cumulative applications received. Standard errors are clustered at the LEA level.

Figure H.4: LEA-level Event Study Figures

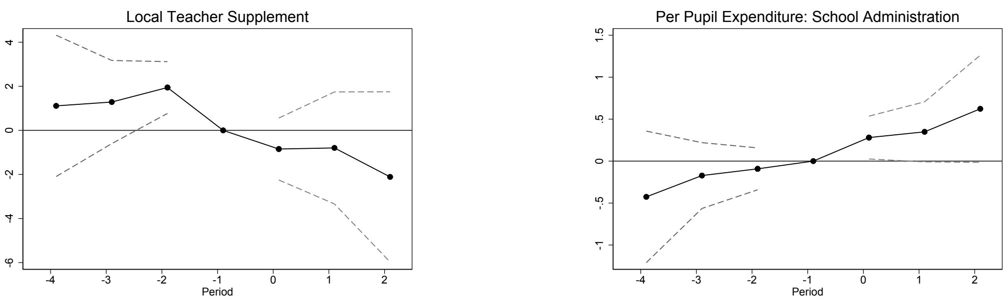

These figures display a set of coefficients on relative period indicators from a LEA-level event study framework in which the dependent variable is either the average local teacher salary supplement or per pupil expenditure on school administration. Period 0 represents the first period in which a new charter school has opened in the district of a given student, and the coefficient on period -1 is normalized to zero. The dashed lines represent 95% confidence intervals.

are influential decision makers increased by an average of 0.0496, which is also consistent with the TWFE result regarding household influence. Overall, these results suggest that the largest shifts in influence from teachers to households within incumbent public middle schools occurred in areas which had the greatest number of charter switchers relative to traditional public enrollment.

H.2 Local Education Agency

This section summarizes LEA-level results regarding the fiscal spillover effects of charter competition. Event study figures that correspond to the average local teacher salary supplement and the level of per pupil expenditure on school administration can be seen in Figure H.4.

Table H.2 displays LEA-level results using the stacked difference-in-differences strategy on a sample created in the same manner as described in Section 1.5. The dependent variable in the first column is the average local teacher salary supplement, while in the second column it is the per pupil level of expenditure on school administration. The coefficient of interest in each column is statistically significant and has

Table H.2: LEA Expenditures: Stacked Difference-in-Differences

	(1) Local Teacher Supp.	(2) PP Expenditure on School Admin.
Treatment Group∗Post	-0.869*	0.234**
	(0.493)	(0.101)
R-squared	0.957	0.773
N	546	546

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is a measure of LEA-level expenditure. The first column corresponds to the average local teacher supplement offered in a given LEA, while the second column corresponds to per pupil expenditure on school administration, both scaled in terms of hundreds of dollars. Controls include LEA demographics and enrollment, as well as cumulative applications received. Standard errors are clustered at the LEA level.

Table H.3: LEA Expenditures: Instrumented Difference-in-Differences

	(1) Local Teacher Supp.	(2) PP Expenditure on School Admin.
Charter Share	-70.426*	17.784**
	(38.229)	(7.196)
N	228	228

* p<0.10, ** p<0.05, *** p<0.010

This table reports estimates of coefficients from regressions in which the dependent variable is a measure of LEA-level expenditure. The first column corresponds to the average local teacher supplement offered in a given LEA, while the second column corresponds to per pupil expenditure on school administration, both scaled in terms of hundreds of dollars. The treatment variable of interest is the charter share of an LEA, for which we use a treatment indicator as an instrumental variable. Controls include LEA demographics and enrollment, as well as cumulative applications received. Standard errors are clustered at the LEA level.

the same sign as the corresponding TWFE estimate, although they are each somewhat smaller in magnitude. Overall, the estimates generated from this alternative strategy are largely consistent with the primary results in Section 1.6.

Table H.3 displays DDIV results for LEA-level regressions in which the dependent variable is either the average local teacher salary supplement or the level of per pupil expenditure on school administration. In the first column, we see that the coefficient on the charter share is negative and statistically significant. Along with the average increase in the total charter share during the expansionary period of 0.0204, this implies a decrease in the average local supplement of $143.67. The coefficient in the second column is positive and statistically significant. This coefficient implies an average increase in per pupil expenditure on school administration in treated LEAs of $36.28, which is consistent with the results discussed in Section 1.6 regarding per pupil expenditure on school administration. These estimates largely support the conclusions that were drawn from the primary results.

APPENDIX I

TRANSITIONAL AID MODEL SOLUTION

This section describes the equilibrium levels of instructional expenditure per pupil and consumption amenities at both the public and charter schools derived from the model in Section 1.2 under the scenario with competition. The equilibrium levels of instructional expenditure per pupil I_q^*, consumption amenities X_q^*, and enrollment S_q^* at the charter school are as follows:

$$I_q^* = \rho - \frac{2\zeta_q \bar{k}}{\lambda \left(\frac{\eta\rho}{3} \left(\frac{\alpha_q}{\psi_q} - (1-\mu) \frac{\alpha_p}{\psi_p} \right) + \bar{k} \right)} \tag{I.1}$$

$$X_q^* = \frac{\rho}{3} \left(2\frac{\alpha_q}{\psi_q} + (1-\mu) \frac{\alpha_p}{\psi_p} \right) - \frac{\bar{k}}{\eta} \tag{I.2}$$

$$S_q^* = \frac{\lambda}{2\bar{k}} \left(\frac{\eta\rho}{3} \left(\frac{\alpha_q}{\psi_q} - (1-\mu) \frac{\alpha_p}{\psi_p} \right) + \bar{k} \right) \tag{I.3}$$

In equilibrium, instructional expenditure per pupil at the charter school q is strictly increasing in the reimbursement share μ, and consumption amenities at the charter school are strictly decreasing in μ. This implies that the test score at the charter school is unambiguously increasing in μ. Charter school enrollment S_q^* is also strictly increasing in μ. The equilibrium levels of instructional expenditure per pupil I_p^* and consumption amenities X_p^* at the public school are as follows:

$$I_p^* = \rho - \frac{2\zeta_p \bar{k}}{\lambda \left(\bar{k} - \frac{\eta\rho}{3} \left(\frac{\alpha_q}{\psi_q} - (1-\mu) \frac{\alpha_p}{\psi_p} \right) \right)} + \mu\rho \frac{\frac{\eta\rho}{3} \left(\frac{\alpha_q}{\psi_q} - (1-\mu) \frac{\alpha_p}{\psi_p} \right) + \bar{k}}{\bar{k} - \frac{\eta\rho}{3} \left(\frac{\alpha_q}{\psi_q} - (1-\mu) \frac{\alpha_p}{\psi_p} \right)} \tag{I.4}$$

$$X_p^* = \frac{\rho}{3} \left(\frac{\alpha_q}{\psi_q} + 2(1-\mu) \frac{\alpha_p}{\psi_p} \right) - \frac{\bar{k}}{\eta} \tag{I.5}$$

$$S_p^* = \frac{\lambda}{2\bar{k}} \left(\bar{k} - \frac{\eta\rho}{3} \left(\frac{\alpha_q}{\psi_q} - (1-\mu) \frac{\alpha_p}{\psi_p} \right) \right) \tag{I.6}$$

In equilibrium, consumption amenities at the public school are strictly decreasing in the reimbursement share μ. If the fixed cost of public school operation ζ_p is sufficiently low, instructional expenditure per pupil at the public school in equilibrium is strictly increasing in μ, which then implies that the equilibrium test score at the public school is also strictly increasing in μ. Public school enrollment S_p^* is strictly decreasing in μ because an increase in the reimbursement share implies a larger decrease in consumption amenities at the public school relative to the charter school.

APPENDIX J

NEW YORK ENGLISH/LANGUAGE ARTS RESULTS

This section describes results regarding the effect of transitional aid eligibility on average LEA exam performance in English/Language Arts for students enrolled in grades 3-8. These estimates are consistent with the larger literature on school accountability measures, which tends to find larger effects on math test scores relative to verbal test scores (Figlio and Loeb, 2011). The regression discontinuity results in Table J.1 demonstrate that consistent with the case of math, transitional aid eligibility is associated with an increase in average LEA performance in English/Language Arts, although none of the estimates are statistically significant. It is worth noting that while the coefficients estimated on each of the three current transitional aid eligibility indicators are smaller in magnitude than the corresponding math estimates, the coefficient on future aid eligibility is actually larger than the corresponding result in math.

Table J.1: Average English/Language Arts Score

Eligibility for Future Aid	0.078
	(0.056)
Aid Eligibility A	0.042
	(0.034)
Aid Eligibility B	0.076
	(0.056)
Aid Eligibility C	0.080
	(0.071)
R-squared	0.702
N	306

* p<0.10, ** p<0.05, *** p<0.010

This table displays a set of LEA-grade-year results regarding the effect of transitional aid eligibility on the average English/Language Arts test score across an LEA. Controls include a second-degree polynomial of the current charter share as well as each of the first three lags, the total number of students tested in a given district-grade-year unit, and total district enrollment across all grade levels. Controls for demographic characteristics also include variables that represent the shares of tested students that are black, Hispanic, and economically disadvantaged. The specification also includes district, year, and grade level fixed effects. Standard errors are clustered at the LEA level.

Figure J.1 explores the sensitivity of these results to the choice of bandwidth around the charter enrollment thresholds for aid eligibility. The estimated coefficients displayed in the figure are all positive, although many of them are not statistically significant. These estimates appear to be relatively stable across alternatives, especially outside of very small bandwidths for which sample size issues become a larger concern.

Figure J.2 displays results from an event study framework in which the dependent variable is average English/Language Arts exam performance at the district-grade-year level, and the sample includes LEAs that became eligible to receive transitional aid for the first time from 2013-2017. The red line in the figure denotes the period in which the charter enrollment share first exceeds the eligibility threshold, while the blue line denotes the period in which the LEA is actually eligible to receive aid payments for the first time. The estimates displayed in the figure are very consistent with the corresponding math results, particularly in that test score improvements

Figure J.1: Average English/Language Arts Score Bandwidth Plot

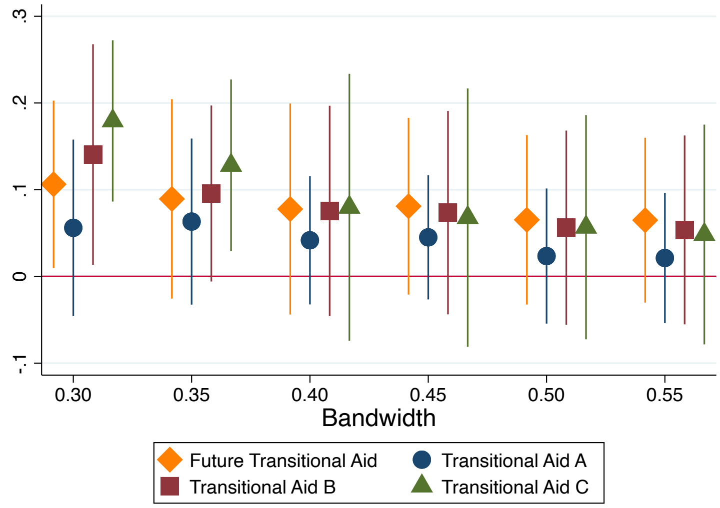

This figure displays the estimated effect and 95% confidence interval of each type of transitional aid at various bandwidths around the enrollment thresholds for aid eligibility (2%). The dependent variable is average LEA English/Language Arts exam performance at the district-grade-year-level, and bandwidths are displayed on the x-axis.

appear to begin prior to the receipt of the first transitional aid payment, but none of the coefficients are individually statistically significant at a 95% confidence level.

Figure J.2: Average English/Language Arts Score Event Study

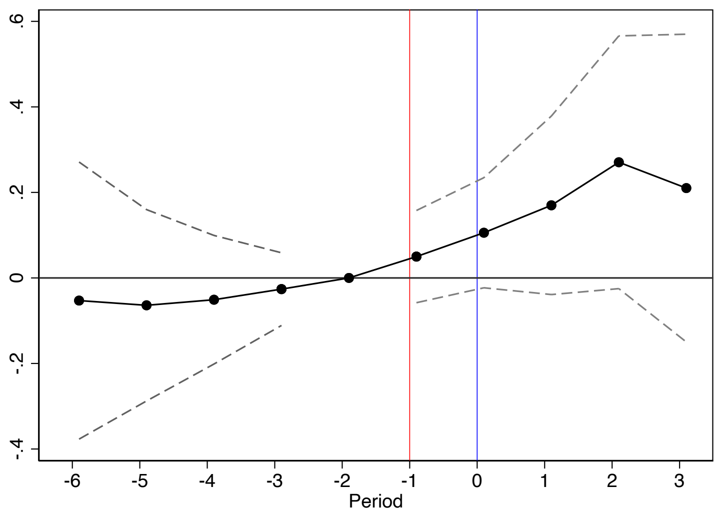

This figure displays a set of coefficients on relative period indicators from an event study framework in which the dependent variable is average LEA English/Language Arts exam performance at the district-grade-year-level. Period 0, marked by the blue line, represents the first period in which an LEA is actually eligible to receive transitional aid payments. Period -1, marked by the red line, represents the period in which an LEA becomes aware that it is eligible to receive transitional aid in the following period. The coefficient on period -2 has been normalized to zero, and the dashed lines represent 95% confidence intervals.