The Asymmetric Effects of School Facilities on Academic Achievement: Evidence from Texas Bond Votes

Keywords: School Facilities, Student Achievement, Asymmetric Effects

Abstract: Exploiting quasi-experimental variation in levels of capital funding, we provide new evidence on the asymmetric effects of investments in educational facilities on academic achievement. We merge data from 4 million 3rd-8th graders in Texas with school bond outcomes and school openings. We find meaningful differences in the impact across the top and bottom quintiles of the pre-treatment achievement distribution, a result supported by the peer effects literature and a "reshuffling' of higher performing students into new schools. The mean impact of bond passage on reading and math scores is positive – at roughly 1/10th of a standard deviation. However, the effect on the lowest scoring students is four times as large, while the effect reverses in direction for the highest performing students.

I. Introduction

The effect of funding for public education on student achievement has long been a relationship of interest among social scientists. Education is viewed as an important micro level determinant of human capital and macro level determinant of economic growth (e.g., Krueger & Lindahl, 2001; Hanushek, 2013). Widespread access to public education has also been shown to reduce income inequality through enhanced social mobility among lower and middle income households (Sylwester, 2002). Scholars and policy makers consistently note that quality public education provides pathways towards success among disadvantaged youth (Ladd, 2012), and recent work has even demonstrated that school finance reforms helped children from the poorest districts earn higher wages and experience lower poverty rates in the long run (Jackson, Johnson, & Persico, 2016). Put another way, both efficiency and equity related social goals seem to be enhanced through better public education.

At the same time, public education is costly and primarily financed by state and local governments. Sub-national governments in the United States are responsible for all the costs of labor, roughly 90% of the ongoing costs associated with facility maintenance/repair/utilities, and essentially all of the costs associated with building new facilities. Many economists and policy makers have avoided embracing resource-based educational policies due, at least in part, to the striking results contained in the Coleman Report (1966), which found small or even zero effects of additional school spending on student achievement. Motivated by these tradeoffs, numerous studies have investigated the linkages between education spending and student outcomes. A challenge to this literature is that widespread legal reforms have given most US states funding equalization formulas that break the mechanistic linkage between local property taxes and spending on local public education, creating an environment where per pupil spending in many states is flat by design (Hoxby, 2001). Another potential shortcoming of most studies in this area is that estimations almost exclusively focus on the mean/overall effect of higher levels of spending on student outcomes, as opposed to directly investigating that possibility that enhanced levels of resources may have asymmetric effects on the highest and lowest achieving students. Our study seeks to address both of these deficiencies.

Specifically, we turn to a mechanism that local communities still retain autonomy over: voting for property tax bonds used to fund the construction of new schools and/or expand and upgrade existing facilities. Even within a fiscal landscape dominated by equalization; local education infrastructure bonds capture *marginal* autonomous spending that can be related to student achievement outcomes. We use a modified regression discontinuity design (RDD) approach combined with student level fixed effects models to estimate the effect of quasi-random infusions of capital expenditures in Texas school districts. This allows us to test whether or not students at the highest and lowest portions of the student achievement distribution are affected uniformly by new capital funding, or alternatively, if they experience meaningfully different effects. Our findings provide consistent support for the latter story, suggesting the lowest performing students actually gain the most from exposure to higher levels of school funding.

Our treatment and control groups are formed by comparing student test performances after exposure to narrowly passed/failed bond votes, relative to their previous achievement levels. We acknowledge that Tiebout sorting and other political/economic factors lead to circumstances where some school districts are more or less likely to pass educational bonds, leading to concerns regarding the endogeneity of bond passage to pre-existing student characteristics (Schlaffer, 2018). On the one hand, districts with more (less) adequate facilities at a given point in time may be less (more) likely to propose and/or support new bonds, based simply on the satiated (deprived) state of their current needs. On the other hand, some communities may simply fund education to a higher degree that others – meaning new facilities often come into privileged environments where their impacts are likely to be less significant. Regardless of which bias is stronger, it would be hard to interpret the results of studies that did not control for differences in the likelihood of bond passage across communities.

Fortunately, the frequency of school bond votes in Texas combines with the massive size of the student population to allow our modified RDD approach to overcome these selection issues. Of the 1,018 school districts in our panel, 103 districts (10.1%) experienced a narrow bond passage; whereas 99 districts (9.7%) experienced narrow failures. We define 'narrow' as falling within 3% of the margin of determination after validating that the pass/fail comparisons were also valid at more/less narrow windows. We also validate the similarity of our passage/failure districts on a number of observable dimensions, and verify that our voting data passes the standard McCrary (2008) density test for use in regression discontinuity studies. Based on the existing literature, we assume the effects of expanding school facilities on student achievement could take several years to register. Studying the effect of education spending on home prices, Cellini, Ferreria, and Rothstein (2010) find that construction expenditures rose over baseline levels in the second through fifth years following the passage of bonds in California – clearly suggesting that bond passage does not lead to *immediate* changes in school facilities. As such, the causal effects of the events taking place in the several years that follow passage (e.g., new construction, renovation/expansion, reshuffling students and/or teachers within and across districts, changes in class size and composition, etc) are captured by our estimations.

We use data from Texas, a large and diverse environment, to investigate how capital spending affects the performance of students in different portions of pre-existing achievement distribution. The differences between Texas and other states actually make it an ideal choice. Texas is the *second largest and second fastest growing state*. Between the 2000 and 2010 census, the population aged eighteen or younger grew by about 1 million (a 16.5% increase). This rapid population growth places a strain on education infrastructure, which must then expand to accommodate growth. Unsurprisingly, Texas experienced high levels of school overcrowding during this period, especially in urban population centers. Additionally, according to a recent survey of school facilities, the average age of permanent education structures in Texas was 24 years old, while roughly 50% of permanent structures were at least 27 years old (Taylor et al., 2005). Given that the life expectancy of a public school building has been estimated as 50 years (Bureau of Economic Analysis, 2003), many school districts in Texas have been playing catch-up. Local property tax bonds are the primary mechanism to accomplish this.

In the absence of expanding facilities, other solutions to dealing with overcrowding have been advanced. These include the use of portable classroom space, adopting multi-track year-round calendar systems, and redrawing service boundaries when districts try to smooth distributions between over and under populated schools (Graves, 2010). Still, these steps are often characterized as a reaction to inadequately providing traditional facilities due to budgetary constraints. From this perspective, the implicit assumption is that an expanded stock of school facilities would be preferable if cost were no obstacle (i.e., the monetary costs *are* the deterrent). To be fair, other studies have documented certain negative effects of building new schools. These include disrupting the classroom environment, shuffling peer groups, higher turnover of faculty, and higher rates of absenteeism (Engberg et al., 2012). As such, measuring the impact of expanded school facilities on important sub-groups of students is a worthwhile endeavor. Importantly, many of the 'transition costs' identified by the literature are only experienced by students moving into a new school. Our data suggests that in Texas school districts experiencing facility expansion, roughly 77% of students simply 'stay put' in their existing school. The effect of new facilities on this massive sub-population of students is of importance for policy makers and school officials who should protect the interests of parents who may or may not have a child that will ever attend a new school.

To extend the literature by exploring this previously unanswered question, we focus on how exposure to narrowly passed/failed bonds impacts the reading and math scores of students who stay at the same school following those close votes. Since we observe test scores before and after bond passage, we estimate student fixed effects models that allow us to evaluate the potentially asymmetric effects of bond passage on the highest and lowest performing students. While our baseline models use the full sample, we estimate specifications using only higher performing students and lower performing students (defined using pre-treatment test scores), finding meaningful asymmetric effects between the two groups. Understanding the impacts of expanded school facilities on high and low performers is a valuable tool for policy makers. For instance, poorly performing students are at higher risks for dropping out early (Cairns, Cairns, & Neckerman, 1989). These effects have been shown to exhibit path dependency, with poor performance in elementary school still carrying a positive correlation with the probability of dropping out even after controlling for more recent/current grades (Ensminger & Slusarcick, 1992). Therefore, a policy maker looking to increase retention rates may wish to heavily weight benefits to lower performing students. Conversely, officials seeking the attention of high profile awards given to top performers may follow the opposite approach.

We find additional school facilities in the district, proxied by narrow passage, registers a small but significantly positive mean effect on both reading and math scores. Once the dynamic impact is fully registered, both scores increase by about one-tenth of one standard deviation. When focusing on the highest and lowest performers, larger effects (of opposing directions) surface. We document a small negative effect on the highest performing students, but a much larger increase in the scores of the lowest performing students. Specifically, students in the bottom quintile gain nearly four-tenths of one standard deviation in their scores. Hence, exposure to more district level facilities seems to help struggling students the most. We also provide evidence that the top students are more commonly 'resorted' into newly built schools following bond passage, complementing the literature on student peer effects and residential income segregation.

Section II presents an overview of the literature on school construction and student achievement. Section III describes our data. Section IV outlines the two empirical approaches we utilize. We discuss our results in Section V and then conclude with Section VI.

II. Background and Theory

A. Education Spending, School Construction and Student Achievement

Three related literatures frame our study. The first broad literature investigates the effects of education spending on student achievement. Dating at least back to the Coleman report in 1966, which concluded that government spending had limited impacts on student achievement, and to the work of Hanushek (1997), which highlighted the lack of a systemic relationship between spending levels and student achievement – one theme in the education literature is that student achievement is not very sensitive to the level of local spending. However, even within this environment, other studies have demonstrated positive effects of school spending on student achievement (e.g., Card & Krueger, 1996; Holmlund, McNally, & Viarengo, 2010), including work indicating the strongest positive effects accrue to the lowest income school districts (Lafortune, Rothstein, & Schanzenbach, 2018). Moreover, work focusing on the effects of school finance reforms provides evidence that enhanced funding leads to higher educational attainment (Hyman, 2017) as well as higher wages and lower poverty rates in the long run (Jackson, Johnson, & Persico, 2016).

The branch of this literature that relates directly to our study investigates the effects of educational spending on capital facilities. For instance, Bowers and Urick (2011) consider the level of upkeep for capital facilities, arguing the physical condition of the facility itself directly influences the learning experiences of students. They found no significant effects of better conditions on student achievement. Another particularly relevant contribution comes from Neilson and Zimmerman (2014), who examine the effect of a comprehensive school construction program in New Haven, Connecticut on standardized math and reading exam scores. The authors use student level fixed effects, a method we also employ, to control for unobservable student level characteristics. They find a significant positive effect on reading scores, but see little evidence that math scores improved. The improvement in reading scores is experienced in the year immediately following new school construction and grows slightly after longer exposure to the new facility. The cumulative gain to reading scores by year six is about one-seventh of a standard deviation. We extend their findings by directly exploring how capital spending impacts lower performing students, relative to higher performing students. These meaningful asymmetric effects are not only found by our estimations to be statistically significant, but we believe the asymmetry is also quite meaningful from an economic and political perspective as well.

Another related study comes from Welsh et al. (2012), who examine a large scale school building and renovation effort in Los Angeles – one of the most overcrowded school districts in the nation. The authors examine the effect of moving to a new school on the moving students, finding significant effects. They also examine the students left behind; arguing those students experience less crowded learning environments. While they find no significant effect on math scores, they do find positive effects for reading. Again, we extend their work by providing a more representative analysis of the school construction occurring in a wide variety of urban, suburban, and rural areas, with differing degrees of facility inadequacy, and directly explore the previously ignored potential for asymmetric effects.

Two other important recent studies come from Martorell, Stange, and McFarlin (MSM) (2016) as well as Hong and Zimmer (2016). Similarly to our study, both use regression discontinuities in voting outcomes to explore the effects of education bonds on student achievement, using data from Texas and Michigan, respectively. MSM verify bond passage is associated with substantial increases in capital expenditures and find that new schools most commonly open their doors two to three years following bond passage. They also document the effects of bond passage on student performance. Although their point estimates of effects on tests scores are at times positive and significant, they are small. Moreover, in the cases where actual school openings are specifically examined, they explain their findings as precise zero estimates of mean achievement effects. Hong and Zimmer examine school district level performances, using a traditional RDD approach to analyze data from 577 school districts in Michigan over the period 1996-2009. Their achievement outcomes of interest were statewide proficiency exam rates. They find no evidence for short run (1-2 year) impacts on student achievement, but do see longer run gains of roughly two to six percent, depending on the specification.

So while the existing literature carries many insights, more remains to be learned. Even when focusing only on recently well placed studies, unresolved issues and disagreements are still present. Importantly, we view our results as largely consistent with the findings in existing studies – particularly MSM – but we also view them as supplementary in at least two important ways. First, we show the asymmetric effect differentially influencing the top and bottom of the student performance distribution. Secondly we are able to speak directly to the issue of how the top students are disproportionately shifted into the newly built schools in these districts. Importantly, from the existing literature we know more about how new facilities impact aggregate (mean) performances across school districts (students) than we know about how they impact the performances of various sub-groups of students.

Additionally, to our knowledge, the best existing work uses all bond votes, whereas we narrow our focus to *only votes that generate new and/or significantly expanded school facilities*. The main difference being our exclusion of busing related bonds, a massive group of outcomes that may not correspond to changes to the learning environment, as they have little impact on the buildings, peers, and teachers that students are exposed to. This is not to say they are meaningless, as students may well prefer newer/nicer buses to older ones, but one could easily argue those effects are different (and likely much less consequential) than the construction/expansion of education facilities.

B. Facility Characteristics, Class Size, and Student Achievement

A second related literature comes from work considering the effects of class size, school overcrowding (a related but not identical topic), and facility quality on student achievement. Perhaps due to the growing number of datasets reflecting student-level randomizations, class size has been investigated far more than other factors. Interest in class size effects is longstanding; for example Glass and Smith (1979) already had the chance to review a number of relevant studies four decades ago. Most would say the modern literature on class size effects began in 1990, when Finn and Achilles explored the Tennessee STAR natural experiment. Their study finds positive and significant benefits for students from a reduction in class size. Other studies using quasi-experimental designs come to similar conclusions. Using a regression discontinuity design facilitated by forced classroom splits following a policy change, Angrist and Lavy (1999) also find positive effects from reductions in class size. However, investigating a similar ‘split-classrooms’ natural experiment, Hoxby (2000) finds no significant relationship between classroom size and student achievement. Moreover, a recent study exploring the effects of class size using the number of local births in a given year as an IV for the endogenous class size variable of interest finds a much smaller gain in achievement than the original Finn and Achilles result (Cho, Glewwe, & Whitler 2012).

A review of class size effects could be pushed much further, as a large number of excellent studies exist. However, we instead summarize this large literature by noting that the estimates of class size effects have varied considerably. While most studies find smaller classes enhance student achievement, many well executed papers also find a null effect. Unfortunately, our student level panel does not directly link students to their classrooms, so we are not able to contribute to this debate. However, since our results are best interpreted as capturing the general treatment effect associated with exposure to the quasi-random infusion of additional capital spending, it is likely that reductions in class size play a role in this environment.

C. Classroom Peer Effects and Student Achievement

A final related literature comes from work investigating classroom peer effects. Recall that we are focusing on the dominant group of students who stay in their school through bond passage. Relative to investigating the entire population, as previous studies have already done, this attention is only merited under two conditions: 1. Selection out of older schools into newly built schools is non-random with respect to student ability, and 2. Peer effects matter. We later provide evidence from our data to support the first claim. Regarding the second, a large literature on peer effects has developed over the past few decades. Findings in this literature are diverse: studies have shown peer spillovers can positively or negatively affect student outcomes (Kindermann, 1993; DuPaul et al,, 1998).

If peer effects are monotonic and stable in magnitude across the achievement distribution, then any given change in the distribution would produce an aggregate change of zero (i.e., a classic zero sum game). However, it is plausible that peer effects are non-monotonic and that they are not stable in magnitude across the distribution. Suppose an average student was harmed *both* by the presence of higher levels of outstanding students *and* struggling students. For example, this would occur if student learning was primarily a function of time spent interacting with the teacher in an environment where both star students and struggling students captured disproportionately large shares of instructor time/resources.

It is also possible that peer effects are more nuanced than this simple time-use driven example. Instead of assuming higher performing peers are always better for achievement, it is possible that peers of a given level of achievement could have different effects on different types of students. Interesting cases may occur where subpopulations of students benefit more from average, or even from below average peers. This possibility is noteworthy, since it facilitates instances where average test scores could increase for both groups within the framework of a reallocation of students and the changing of peer groups. Burke and Sass (2013) demonstrate that low achieving students experience larger gains from exposure to students with achievement in the middle of the distribution, as compared to greater exposure to the highest achieving students. For example, this would occur if it were more likely that study pairs formed when the size of the initial achievement gap between students was smaller or similarly if teachers were more likely to focus on advanced material that left struggling students behind in classrooms that had more high achievers.

Importantly then, previous research suggests it is possible to *increase* the performance of low achieving students by *reducing* their exposure to high performing students. So for example, high performing students may exhibit positive spillovers onto one another, but evidence that they exhibit a uniformly positive effect on all students is less obvious (and perhaps even lacking). Of direct relevance to our eventual findings, Hoxby and Weingarth (2005) find evidence in favor of the 'Boutique model' of spillover peer effects. According to the Boutique model, it is possible to produce benefits for both high and low performers through more homogenous achievement groupings - with high performers improving other high performers and lower performers similarly helping other lower performers. For example, this would occur if teachers possess heterogeneous skill sets regarding learning facilitation for students situated in different portions of the achievement distribution, or if the choices teachers made over classroom curriculum were better suited to situations where students in the class had smaller differences in previously learned material (i.e., more homogenous current knowledge) than it was to cases with more heterogeneity. Therefore, as the opening of a new school will result in elevated student mobility, the 'reshuffling' of students and corresponding peer effects could move in a number of directions.

III. Data

To investigate the relationship between educational facilities and student achievement we use three distinct panel data sources that overlap in terms of geography (all are from Texas) as well as their timing. They are: 1. student level data containing math and reading test scores, 2. voting data from school district level bond results, and 3. school campus level activation data. We obtained test scores for the population of students from the Texas Education Agency. They contain reading and math scores on state exams taken in grades three through eight, beginning with the 2003-2004 school year and running through the 2013-2014 school year. Beyond reporting the school of attendance for each student in each year, unique student IDs allow us to track students test scores over multiple years – an advantage that allows the use of student level fixed effects. Since we do use student level fixed effects in our estimations, we drop any student in the original data where we see only a single year of test scores reported. An observation count of 2,025,898 students is what remains *after* this initial filter requiring at least two distinct years seeing test scores is applied. Our achievement data spans two different standardized tests taken by Texas students. The Texas Assessment of Academic Skills (TAAS) was in place up through 2010-2011. Texas switched to the State of Texas Assessment of Academic Readiness (STAAR) in 2011-2012. Since these tests use different scoring/scales, we normalize the distribution of all math and reading scores for each year. As such, each test score is best thought of as a level of performance within a very large distribution of outcomes in the state of Texas for a given year.

Our bond data contains the entire population of results for Texas school district level bonds voted on between 1995 and 2014. As mentioned before, an important aspect of our study is that we limit our sample to the 2,197 initiatives that directly mentioned school facility construction/expansion and/or major land acquisitions. Unfortunately, the bond description comes from a simple labeling system that does not allow for precise determination of which bonds lead to *new* campuses (i.e., as opposed to expanding existing facilities). This means that our analysis of the effects of quasi-random bond passage on student achievement cannot be *precisely* linked to the building of a new school. However, we expect passage of the bonds we consider to be strongly correlated with facility expansion, and we see no strong reason why expanding facilities within an already existing campus would not potentially carry similar effects to the new school case. For example, reductions in student-teacher ratios and significant student reshuffling may still occur, it is now just between locations/classrooms in the same location as opposed to different locations.

Table 1 provides strong evidence supporting the straightforward expectation that districts experiencing a narrow passage do in fact see more new schools. School districts experiencing passage of a bond (by any margin) are significantly more likely to experience a new school opening in both the 3 year and 5 year window following passage, when compared to the sample at large. Rather than dive even more deeply into this issue, we refer interested readers to MSM (2016), who examine this exact issue using very similar data to our own.

While the vast majority of the 2,197 bonds we examine passed by relatively large margins of victory (e.g., more than 3%), our empirical approach of course relies very heavily on the 103 bonds (4.7% of votes) that passed by less than 3% and the 99 bonds (4.5% of votes) that failed by less than 3%. Although the number of 'narrow' margins is relatively small, at right around 10% of total votes, the subset of students that were exposed to a close vote of either type is still massive, at 859,693. For this reason, our approach focuses directly on the quasi-experimental variation in school funding levels that was experienced by students falling into this group of students exposed to a 'narrow' margin vote. A McCrary density test was performed over the bond votes to investigate the distribution of electoral support outcomes (McCrary, 2008). We pass the McCrary test with an estimated P-value of 0.058 and standard error of 0.148, as the test fails to reject the discontinuity hypothesis. Figure 1 shows the results of the McCrary test as a graph of the distribution of bond votes by the percentage of "yes" votes received. Two expected results quickly surface visually. First, the peak of the density function occurs well past the 50% passage threshold. This was expected, as the simple reality of the data is that most bonds succeed. Additionally, the upward sloping portion of the density function that travels through the 50% threshold does not seem to experience any noticeable structural breaks at or near the passage threshold. Again, we expected this result as the technique has also been used with similar data by other previous studies.

We additionally explored alternative bandwidths that were larger and smaller than the 3% margin. Many similarities in the resulting distributions emerged from this comparison, and the treatment/control formation seems to be valid in a number of alternative choices. At the end of the day, we selected the 3% margin because it produced a large number of votes in both the treatment and control groups (i.e., 103 narrow passages and 99 narrow failures), easily passed the McCrary test, maintained group comparability on a number of important dimensions, and fell comfortably within the range of choices made by previous studies using similar identification strategies. The statistics reported in Figures 2, 3, and 4 have been specifically grouped into districts experiencing marginal bond failure, marginal bond passage, and all school districts (for reference). These comparisons give additional evidence that the formation of our treatment/control groups of students is valid. For example, Figure 2 shows the narrow passage/failure districts have strikingly similar age distributions, including strong similarities for the age ranges involving school age children. Similar evidence comes from Figure 3, which reports the median district's percent of the population falling within different income brackets. This data comes from the American Communities Survey; the five year estimates from 2010. Again, the difference between the marginal passage and the marginal failure districts is trivial. Finally, Figure 4 reports the median district's percentage of the population split among urban and rural residents, again displaying similar to one another.

The final component of our data is the Texas Directory of Schools, a data set that is maintained by the Texas Education Agency. This facility level data contains the most recent date each campus became categorized as an active school. The school district is also reported, allowing straightforward merges with our student and bond data. For the majority of facilities, this file contains all the information we need. However, a handful of rare circumstances led to cases where the recorded 'activation' date was not the date the school actually opened. If a previously operating campus experiences a temporary shutdown for renovation, the Texas Directory of Schools data will then reflect the date of re-opening, rather than the original opening. It is conceivable that a major campus renovation/re-opening has a different effect on student achievement than a newly built campus. Unfortunately, our data cannot distinguish between these two cases.

Each school district/year observational pairing is coded with a series of binary variables: year 0 (opening year), year 1 (1 year after opening), year 2 (2 years after opening), and so on. This process restarts at 0 each time a new facility is activated within the district. An advantage of this approach is that the timing of campus activations provides an easily formed control group for our regression analysis. Collectively, student level observations taken from district/year pairings *without* concurrent or recent campus activations serve as the control group for students in district/year observations that do. When paired with student level fixed effects that control for unobserved differences across students, our regression analysis isolates the variation in student performance that is specifically associated with campus activation. While this approach has a number of intuitive advantages, a drawback is that it does not directly address the issue of multiple school openings occurring close to one another in time (i.e., each opening is treated as an isolated occurrence). This could overstate the effect of interest, as a campus opening could be 'stacked' near/with another that occurred shortly before it. Fortunately, this outcome is extremely rare, allowing some confidence this is not affecting our results. A final drawback is that it limits the number of instances of observations representing long spans between new campus activations. We address this issue by creating a single variable to account for long lagged observations to increase these instances, and to reduce the possibility that a small number of observations drive our eventual results.

Only school districts with at least one school opening during the investigated period are included in the later school activations models. From the 170 districts experiencing one or more activations during our panel, we have 4,704 distinct campuses. This accounts for nearly two-thirds of all of the schools in Texas that were active over the same time period.

IV. Empirical Methodology

To investigate how additional resources for educational infrastructure influence students' performance, we follow two distinct approaches. Each approach has advantages and disadvantages relative to the other. Both use scores on standardized reading and math exams to measure student achievement and employ student level fixed effects. Both approaches also cluster standard errors at the school/facility level to account for unobserved factors that may also influence student achievement. Most importantly, both approaches produce largely similar evidence with respect to the effect of school facilities on student achievement.

Note that before a new school is built, several strategic decisions have already been made by the district. Decisions concerning whether or not to ask voters for funds to build a new campus, to renovate/expand a current campus, or instead to supplement an existing facility with portable buildings have already been made. The population of districts deciding to hold a bond vote, as opposed to pursuing other alternatives or simply making due with existing resources, may not be representative of all districts. Moreover, school districts whose bonds pass may not reflect the same learning environments as the group of districts whose bonds fail.

We assume both of these processes are non-random. Hence, we mitigate potential selection bias by limiting our sample to students exposed to a bond vote falling within 3% of the required 50% passage threshold (i.e., allowing exposure to the voting discontinuity to form the sample). From the previously mentioned McCrary test, we find evidence to support continuity/randomness around the forcing threshold (i.e., the 50% level of support needed for passage) and observe similar characteristics for districts falling into the narrow passage/failure groups. Additionally, the simple observation that 103 close votes pass and 99 close votes fail is reassuring.

For student i, attending school campus k, located in district j, in year t, we estimate our main regression discontinuity model as:

$$(1) \qquad Y_{i,k,j,t} = \beta_1 MP_{j,t} + \beta_2 MF_{j,t} + s_i + \varepsilon_{i,k,j,t}$$

where $Y_{i,k,j,t}$ represents a standardized test score (reading, math, or their sum), $MP_{j,t}$ is a vector of dummy variables denoting transitions around a bond's marginal passage, $MF_{j,t}$ is a vector of dummy variables denoting the same around a bond's marginal failure, s_i is a vector of student level fixed effect dummies, and $\varepsilon_{i,k,j,t}$ is a residual term clustered at the school campus level that is assumed to be normally distributed. We note that the continuous variable for vote share is not needed in this regression, as we have already pre-selected our sample by using only students exposed to a narrow passage or a narrow failure.

At first glance, by including vectors of dummy variables for both narrow passage and narrow failure, the equation seems to include an exhaustive group of outcomes. However, this is not the case, as three distinct types of observations are present in the regression. A given observation could have exposure to a recent narrow passage, could have exposure to a recent narrow failure, or could be an observation where a student has not yet been exposed to a narrow vote, but later in the sample will experience exposure to a close vote that then places them into one of the first two groups. We frame this estimation choice as being akin to common experimental design procedure where subjects provide baseline choices in a pre-treatment exposure environment, and are then randomly assigned into the treatment/control groups that influence the subsequent rounds of choices.

As such, the comparison of passing post-vote observations to the control group of pre-vote observation years constitutes one treatment/control pairing, while comparing failing post-vote observations to the control group can be viewed as another. We acknowledge other models exploiting the same underlying natural experiment are feasible given the nature of our data. We selected this version for ease of discussing the dynamic nature of our main result of interest, as it relates to the length of exposure to the bond passage, and because we eventually find that *close failure outcomes* actually carry a negative effect on achievement relative to the pre-vote control category. We find this result intuitive given the likely reactions of teachers, administrators, students, and parents in the school districts suffering from a negative outcome to a hotly contested vote.

We form the vectors $MP_{j,t}$ and $MF_{j,t}$ to include eight years of lagged dummy variables. We made the choice of eight years after early explorations revealed that beyond this point the overall/total effect of exposure to a close vote ‘flattened out’ (i.e., the marginal effect of additional exposure went to zero). So while the overall impact was found to be permanent in level effects, the size of the effect stops changing beyond the 7-8 year out mark.

The main advantage of using this modified RDD approach is that it makes a strong case for identification of any causal effects of bond passage versus failure. The weakness is that the underlying outcome of interest – the district receiving new and/or improved educational facilities – is measured imprecisely. Put another way, these models produce relatively precise estimates of exposure to the close bond votes, an outcome that serves as a strong proxy for the construction, expansion, and/or major renovations that follow.

Our second empirical approach reverses the advantages and disadvantages of the first, making it a natural complement. These additional models directly measure new school openings, but introduce the cost of using a weaker identification strategy. The ideal approach to constructing complementary regressions would be to link the activation of each particular campus with the passage of its ‘parent’ bond, and to then use the same voting discontinuity at narrow passage/failure outlined above. This would offer stronger identification and improved precision. Unfortunately, our data does not offer direct linkages between bond passage and new school activations. Specifically, the complementary models are estimated as follows:

$$(2) \qquad Y_{i,k,j,t} = \beta_1 A_{k,t} + \beta_2 X_{j,t} + s_i + \varepsilon_{i,k,j,t}$$

where $Y_{i,k,j.t}$ remains the standardized test scores for reading, math, and summed total scores, $A_{k,t}$ is a vector of dummy variables denoting instances of a new school opening in the district, as well as lagged values of new openings from previous years, $X_{j,t}$ contains variables for school district size and the amount of recent activity the district has experienced regarding school activation, s_i is again the student level fixed effect, and $\varepsilon_{i,k,j.t}$ is a residual term clustered at the school campus level that is assumed to be normally distributed.

All variables connecting to new school activations are generated from the Texas Education Agency directory of campuses regarding the year that a campus became active. Because the regression uses student level fixed effects, our ability to control for changes in school campuses is limited. However, any unobserved differences across campuses or districts that students experience uniformly over time are effectively controlled for by this approach. The models include the year a school becomes active, one year after activation, two years after activation, and three years after activation, all as independent dummy variables. All other observations (i.e., cases where the newest campus in the school district is at least 4 years old) are grouped together as a single reference category. Early explorations indicated this cutoff worked well (i.e., behaved similarly to slightly longer and shorter cut-offs). Given the established result that capital spending takes multiple years to increase following bond passage, and the fact that construction of a new facility takes time before new students could actually attend, we view our choice here as largely consistent with the implied timing of the previous approach.

The percentage of the population of schools in the district activated in the previous three years is also included in the regression. The expected sign of this variable is ambiguous. It is possible that districts experiencing a larger percent of activated campuses are benefitting from increased resources devoted to education, which may result in better test scores for students. Conversely, it is possible that many newly activated campuses signals greater disruption in the district; possibly causing an unstable learning environment for students.

Fortunately, the size of our student level data set and our ability to include student level fixed effects in our estimations facilitates a unique opportunity to stratify the data, allowing us to focus on various segments of the student achievement distribution. Importantly, we are able to determine if any asymmetric effects of educational facilities on student achievement are present. We estimate all of our main models using the full sample, as well as the top and bottom quintile of pre-activation achievers. It is possible that high and low performing students experience different effects from adding new facilities, particularly given previous literature suggesting students at different achievement levels may respond asymmetrically to the changing composition/dynamics found within their existing classrooms. In particular, the non-random patterns of student and teacher mobility following passage could modify the composition of classrooms. After exploring our data, we found strong evidence that students who moved from an existing campus to a newly active campus have test scores that, on average, are significantly higher (i.e., seen in Table 2). This creates the possibility for the effect of a new school opening to have a differential effect on students through peer effects, which we later discuss in the context of our results.

V. Results

From our primary results presented in Table 3, we find the quasi-random passage of a bond has a (delayed) significantly positive mean impact on student achievement. Relative to the years leading up to the close bond vote, students perform slightly worse than average during the year of the close vote for both the close passage and close failure groups. The negative 'vote year' effect on students in the passage group is much smaller, but both are statistically significant. While there is no reason to think a close bond vote (for either a yes or no outcome) would immediately change the physical characteristics of the environment in a district, it is reasonable to expect the behavior of teachers, administrators, parents, and students all could be subject to effects associated simply with taking a bond vote. Consider a simple conjecture that teachers and administrators may use more classroom time educating their students about the relative merits of the bond itself, potentially at the opportunity cost of spending less time reviewing the scholastic curriculum. Additionally, teachers may have their overall outlook/morale towards teaching negatively affected by a failed outcome. Given the outcomes both occur within the same year – it is natural to ask whether or not the time-order requirement for causality is met in this case. We believe it is. For example, suppose a student takes a standardized test in April following a narrowly failed bond vote in January.

Timing plays as important role in interpreting our findings. The initial impacts (i.e., first few years after passage) are minimal at best. With the exception of two coefficients on reading scores (one that carries over into total scores) in years two and three, the estimated effects are insignificant during the first four years. This finding complements MSM (2016), who show capital spending levels following bond passage in Texas districts does not increase during the first two years after the vote. The construction of new schools is most likely to begin and/or be in progress during the years we find no positive effect, but that facility completion likely triggers the significant positive effects we see for both math and reading scores in years five through eight. Of course, some of the projects in our data are clearly expansions and/or major renovations, which may take slightly less time to register their impact. All told, we estimate the mean long-run effect of exposure to district level bond passage to be slightly less than one-tenth of a standard deviation when focusing on total test scores. So for example, a student at the 50th percentile at the time of passage would improve to the 53rd percentile eight years after the vote. While this effect is statistically significant, it is certainly not very large.

Importantly, the effect of a negative bond vote (again, compared to the pre-vote year control group) show an even more interesting pattern. A failed bond vote quickly registers a negative impact on student test results. For example, using math or reading scores (or the combination), we find that lower scores are experienced immediately, and that this loss is retained in the long run. The negative vote does not produce a greater intensity of effect as time passes, as was the case for the close positive votes, but rather an effect that begins at its largest and then seems to dissipate slightly over time. A possible explanation for this pattern is that frustrated school districts may turn to lesser preferred solutions to the overcrowding problems they face. For example, a district may respond by adding temporary portable capacity. Alternatively, they may re-propose a subsequent bond with different characteristics; potentially that carries a higher likelihood of success. For example, possibly asking for less money or including other projects that are popular to voters.

To determine if re-proposal following failure was biasing our results, we estimated a restricted version of the estimation where districts with a successful bond vote (passing at any margin) occurring within three years after an initial close failure were dropped. The results carry the same sign, magnitude, and significance for the coefficients associated with marginal failure as those shown in Table 3. We find this reassuring.

The finding that the "close yes" and "close no" outcomes display the opposite dynamics when it comes to comparing short run versus longer run impacts, in ways that make sense given the institutional setting, offers additional reassurance that our selected quasi-experimental variation is identifying a causal effect associated with the increased levels of capital spending associated with bond passage. Perhaps due to our choice to focus on the large group of students who remain attending the same school, we find that the effect on math scores is actually larger than the effect on reading scores. Other investigations (e.g., Neilson & Zimmerman, 2014) find the effect of expanded educational facilities to be slightly larger for reading scores.

Table 4 presents the results of equation (2), which offers the direct investigation of new school openings. Consistent with the primary results, the effect of opening a new school in the district is positive for all three measures of achievement. Again, we find the point estimates are slightly larger for math than for reading, but both exam results are highly significant across all five dummy variables for exposure (at different lag lengths) to a new school within the district. The weakest benefit comes in the initial year of activation, but strengthens as students experience additional years beyond activation. The smaller initial effect may suggest that the activation itself carries elements of disruption that may be short lived. Alternatively, school officials may "learn by doing" in terms of better utilizing new facilities after they gain experience in the new school. We also note that, in general, larger school districts seem to have lower tests scores, implied from the negative coefficient on number of school in the district.

While we prefer the models including this variable over those that do not (i.e., to control for the differences between large (urban/suburban) and small (rural) districts), all the main results concerning our variables of interest are not sensitive to this choice. Furthermore, although very statistically significant, the magnitude of this effect is small relative to the effect of activations.

Tables 5 and 6 provide our main extensions focusing on the top and bottom quintiles of the pre-existing achievement distribution. Recall that we are interested in exploring the possibility that capital funding impacts high and low performers' asymmetrically. We consider one of the main contributions of our study to be *the finding that asymmetries matter* – with the lowest quintile of students retaining the positive effect, now with a much larger magnitude (typically in the range of three to five times as large). However, the highest quintile group experiences a negative impact on all three measures of academic achievement, with a stronger negative effect on reading scores than for math scores.

Additionally, we find highly consistent results when we revisit our original regression discontinuity approach (i.e., equation (1) presented in Table 3). Using equation (1) and restricting the sample to the highest achieving quintile, we obtain coefficients on marginal bond passage variables that flip in sign to negative effects for both math and reading scores – the opposite of the mean effects seen in Table 3. On the other hand, the lowest quintile achievers are subject to effects in the same direction, but that increase dramatically in terms of magnitude; displaying a similar increase in magnitude as the jump seen when moving from Table 4 to Table 6. As such, the asymmetric effect surfaces regardless of the approach used.

Of course, an important question is what factors may help explain these asymmetric impacts. If the asymmetry is being driven primarily by changes in classroom sizes, our findings would be difficult to align with other convincing findings in the literature. A starting point would be to assume that both high and low performing students experience roughly similar changes in class size. While we have no direct evidence of this, again because our measures do not place students into specific classrooms, we contend this is a reasonable assumption. In this case, since Table 5 suggests the highest performing students experience a negative effect of school activation, we would then have to believe that reductions in class size carried a negative effect on the performance of top students – a narrative that runs against convincing evidence from studies like Ding and Lehrer (2011), who find the best students experience *larger* benefits from reductions in class size than average and lower performing students. Therefore, we are not able to explain our asymmetry using this explanation. In fact, any potentially asymmetric effects of class size reductions may likely work against the pattern of our findings.

However, as was discussed earlier, we know the pre-activation test scores of students that eventually move to a newly opened school are significantly higher than the scores of students who do not move into the new school. Therefore, upon activation of a new campus, the fraction of high achieving students in pre-existing schools in the district must fall, while the fraction of lower performing students must, in turn, rise. This stylized fact combines with findings from the peer effects literature to form at least one strong explanatory channel for our asymmetry. Specifically, our results align closely with the Boutique model of peer effects attributed to Hoxby and Weingarth (2005), who argue students from all performance levels increase their achievement when they are surrounded by other students near their own current performance level. Here, the removal of higher achievement students is expected to generate a negative effect on high scoring students (i.e., who lose similar ability peers), but also to positively impact the lower scoring students (i.e., who gain a larger fraction of similar ability peers).

Put another way, when a new campus is activated, the existing schools experience more performance homogeneous classrooms, as the higher performing students are more likely to leave. This compositional shift seems to benefit the lowest scoring students while simultaneously harming the highest scoring students. On net, since the results for the highest achiever quintile run counter to the class size story, but do fall in line with the Boutique model, we argue our findings are likely being driven by peer effects. Of course, the effects 'compound' when it comes to explaining the lowest achievement quintile, in the sense that smaller class sizes and larger fractions of classmates at similar ability levels would *both* tend to increase achievement for this group – perhaps explaining why their gain is much larger in magnitude.

VI. Conclusion

Exploiting a public voting process that generated quasi-experimental variation in school facility investments in Texas school districts, we investigate the effects of enhanced levels of capital spending on student achievement. Across empirical methodologies displaying complementary advantages and disadvantages, the same story surfaces. We find additional school facilities increase achievement on average, by a small amount (roughly one-tenth of a standard deviation in performance on standardized test scores). Opening a new school in the district was found to improve performances from students *that did not attend it*, an important result to be aware of for policy makers who must convince local resident voters that raising property taxes to fund infrastructure expansions is desirable. So for example, this result could be helpful in combating concerns from parents in cases where they (incorrectly) expect no benefit will accrue to their child from a new school their child will not eventually attend.

Importantly, we find that not all students are winners. The expansion of school facilities seems to carry asymmetric effects on the students in a manner that depends on their pre-existing level of achievement. Students scoring in the bottom quintile of the distribution experience larger positive effects, approaching nearly one half of a standard deviation by the time the full/cumulative impact plays out, compared to students at the top of the distribution who actually experience negative impacts. We highlight that this effect is limited to the group of students who 'stay put' in their school following bond passage – an important caveat given that we demonstrate that high achieving students are more likely to attend the new schools.

In terms of public campaigns for education bonds, we argue that in addition to groups of voters who have little interest in investing in the human capital of students in the school system, one group holding a legitimate set of concerns associated with opening new schools could be parents of high achieving students that live in the parts of the school district that will 'stay put' (i.e., perhaps the older and lower income areas relative to higher income areas with newer housing). At the same time, strong students from newer and potentially wealthier neighborhoods are more likely to leave to attend the new campus. Collectively, our findings complement the boutique model of peer effects advanced by Hoxby and Weingarth (2005). Their work suggests that creaming off high performing students to new schools may simultaneously help the lowest performing students, while harming the higher performing students who stay at the original school. We also complement recent work suggesting the strongest potential for gains in student achievement comes from schools/districts that are initially performing poorly (e.g., Carlson & Lavertu, 2018), an intuitive result.

Further study of peer effects as they relate to the expansion of school facilities is merited. Do teachers systematically modify curriculum in response to the changing composition of ability in their classrooms? Do peer-to-peer interaction dynamics change when fewer top performing students are present? Does teacher mobility play a role? These and other questions remain outside of our current contribution.

Importantly, we find similar results across two complementary empirical approaches, perhaps because our lengthy panel data allows all of our estimated models to include student level fixed effects – which is arguably the most important step we take to mitigate bias in a data environment where we observe student achievement, but do not observe many student level home/family characteristics that would certainly correlate to student achievement. On net, our results suggest policy makers should not assume a 'one size fits all' effect accurately characterizes the impact of new educational facilities on student achievement, and that predictable groups of winners and losers may be formed as school districts expand their educational facilities.

References

Angrist, J., & Lavy, V. (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *The Quarterly Journal of Economics*, 114(2), 533-575.

Bowers, A., & Urick. A. (2011). Does high school facility quality affect student achievement? A two-level hierarchical linear model. *Journal of Education Finance*, 37(1), 72-94.

Bureau of Economic Analysis. (2003). Fixed assets and consumer durable goods in the United States, 1925-99. Washington, DC: U.S. Government Printing Office.

Burke, M., & Sass, T. (2013). Classroom peer effects and student achievement. *Journal of Labor Economics*, 31(1), 51-82.

Cairns, R., Cairns, B., & Neckerman, H. (1989). Early school dropout: Configurations and determinants. *Child Development*, 60(6), 1437-1452.

Card, D., & Krueger, A. (1996). School resources and student outcomes: An overview of the literature and new evidence from North and South Carolina. *Journal of Economic Perspectives,* 10(4), 31-50.

Carlson, D., & Lavertu, S. (2018). School improvement grants in Ohio: Effects on student achievement and school administration. *Education Evaluation and Policy Analysis,* 40(3), 287-315.

Cellini, S., Ferreira, F., & Rothstein, J. (2010). The value of school facilities investments: Evidence from a dynamic regression discontinuity design. *The Quarterly Journal of Economics,* 125(1), 215-261.

Cho, H., Glewwe, P., & Whitler, M. (2012). Do reductions in class size raise students' test scores? Evidence from population variation in Minnesota's elementary schools. *Economics of Education Review*, 31(3), 77-95.

Coleman, J. (1966). Equality of educational opportunity. Washington, DC: U.S. Government Printing Office.

Ding, W., & Lehrer, S. (2011). Experimental estimates of the impacts of class size on test scores: robustness and heterogeneity. *Education Economics*, 19(3), 229-252.

DuPaul, G., Ervin, R., Hook, C., & McGoey, K. (1998). Peer tutoring for children with attention deficit hyperactivity disorder: Effects on classroom behavior and academic performance. *Journal of Applied Behavior Analysis*, 31(4), 579-592.

Engberg, J., Gill, B., Zamarro, G., & Zimmer, R. (2012). Closing schools in a shrinking district: Do student outcomes depend on which schools are closed? *Journal of Urban Economics*, 7(2), 189-203.

Ensminger, M. & Slusarcick, A. (1992). Paths to high school graduation or dropout: A longitudinal study of a first-grade cohort. *Sociology of Education*, 65(2), 95-113.

Finn, J., & Achilles, C. (1990). Answers and questions about class size: A statewide experiment. *American Educational Research Journal,* 27(3), 557-577.

Glass, G., & Smith, M. (1979). Meta-analysis of research on class size and achievement. *Educational Evaluation and Policy Analysis*, 1(1), 2-16.

Graves, J. (2010). The academic impact of multi-track year-round school calendars: A response to school overcrowding. *Journal of Urban Economics,* 67(3), 378-391.

Hanushek, E. (1997). Assessing the effects of school resources on student performance: An update. *Educational Evaluation and Policy Analysis,* 19(2), 141-164.

Hanushek, E. (2013). Economic growth in developing countries: The role of human capital. *Economics of Education Review,* 37(1), 204-212.

Holmlund, H., McNally, S., & Viarengo, M. (2010). Does money matter for schools? *Economics of Education Review,* 29(6), 1154-1164.

Hong, K., & Zimmer, R. (2016). Does investing in school capital infrastructure improve student achievement? *Economics of Education Review*, 53(1), 143-158.

Hoxby, C. (2000). Does competition among public schools benefit students and taxpayers? *American Economic Review,* 90(5), 1209-1238.

Hoxby, C. (2001). All school finance equalizations are not created equal. *The Quarterly Journal of Economics,* 116(4), 1189-1231.

Hoxby, C., & Weingarth, G. (2005). Taking race out of the equation: School reassignment and the structure of peer effects (No. 7867). Working paper.

Hyman, J. (2017). Does money matter in the long run? Effects of school spending on educational attainment. *American Economic Journal: Economic Policy,* 9(4), 256-280.

Jackson, K., Johnson, R., & Persico, C.. (2016). The effects of school spending on educational and economic outcomes: Evidence from school finance reforms, *The Quarterly Journal of Economics,* 131(1), 157-218.

Krueger, A., & Lindahl, M. (2001). Education for growth: Why and for whom? *Journal of Economic Literature,* 39(4), 1101-1136.

Kindermann, T. (1993). Natural peer groups as contexts for individual development: The case of children's motivation in school. *Developmental Psychology,* 29(6), 970-977.

Ladd, H. (2012). Education and poverty: Confronting the evidence. *Journal of Policy Analysis and Management*, 31(2), 203-227.

Lafortune, J., Rothstein, J., & Schanzenbach, D. (2018). School finance reform and the distribution of student achievement. *American Economic Journal: Applied Economics*, 10(2), 1-26.

Martorell, P., Stange, K., & McFarlin, I. (2016). Investing in schools: Capital spending, facility conditions, and student achievement. *Journal of Public Economics,* 140(1), 13-29.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, Vol. 142(2), 698-714.

National Clearinghouse for Educational Facilities. (2016). http://www.ncef.org. Accessed 4/26/2016.

Neilson, C., & Zimmerman, S. (2014). The effect of school construction on test scores, school enrollment, and home prices. *Journal of Public Economics*, 120(1), 18-31.

Schlaffer, J. (2018). Financing public education facilities: The role of elderly populations and geographic mobility. *Social Science Quarterly,* 99(1), 118-135.

Sylwester, K. (2002). Can education expenditures reduce income inequality? *Economics of Education Review,* 21(1), 43-52.

Taylor, L., Texas Department of Education. (2005). Meeting needs? A survey of school facilities in the state of Texas. Committee on public school finance, http://bush.tamu.edu/research/faculty/txschoolfinance/papers/facilitiesreport.pdf.

Texas Education Agency. (1999). School size and class size in Texas public schools. Policy Research Report. http://tea.texas.gov/acctres/Spec_PRR_12_1999.pdf

Welsh, W., Coghlan, E., Fuller, B, & Dauter, L. (2012). New schools, overcrowding relief, and achievement gains in Los Angeles-- Strong returns from a $19.5 billion investment. P.B. 12-2. *Policy Analysis for California Education*, PACE (NJ1).

Figure 1: McCrary Distribution Continuity Test

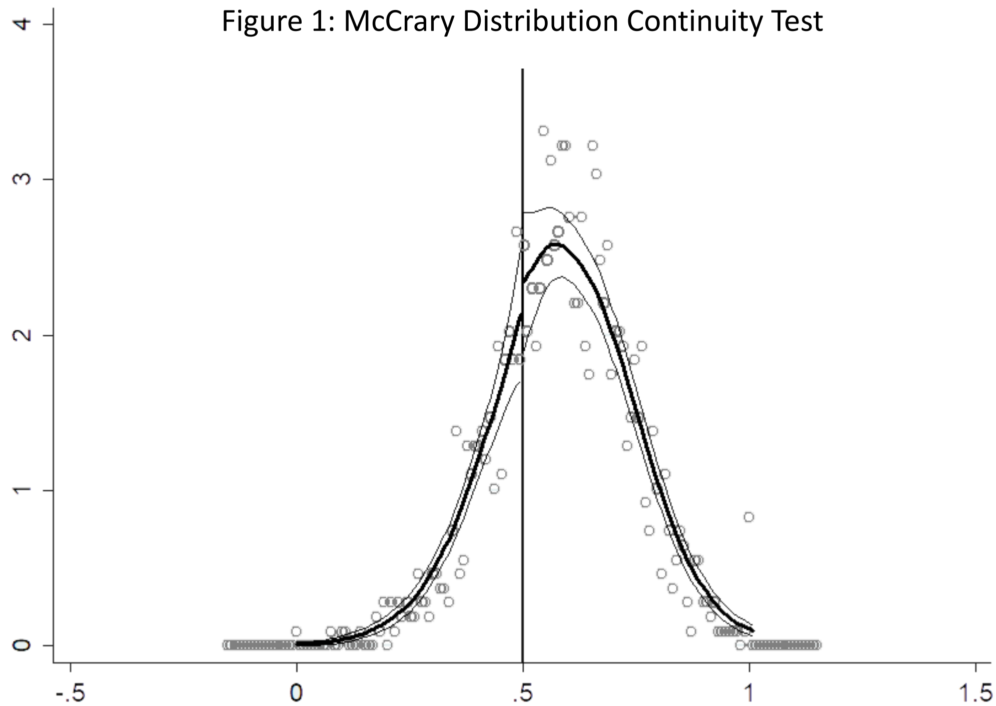

Percent of “Yes” votes.   P-value: 0.05

Figure 2: Age distribution of residents by RD group

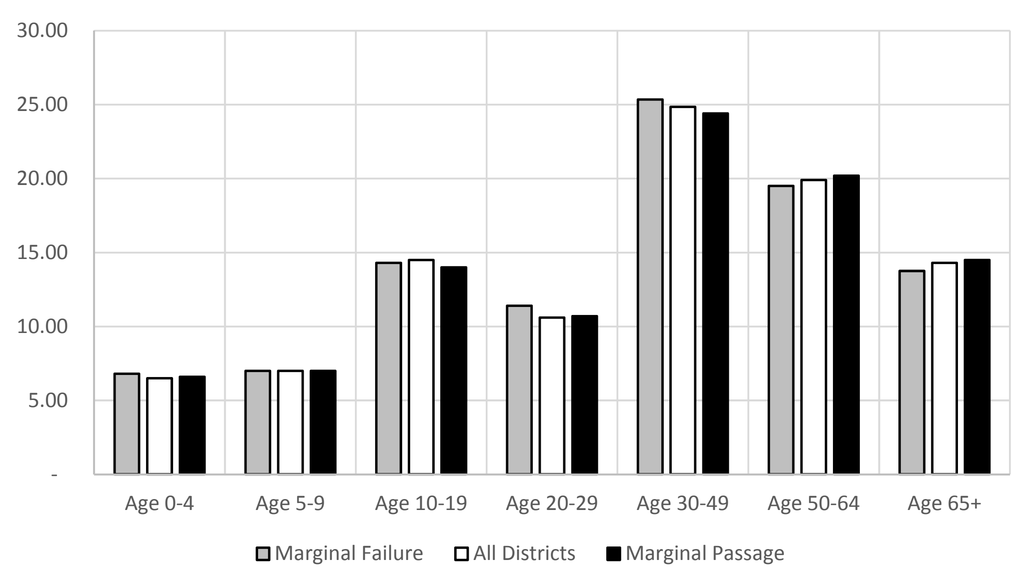

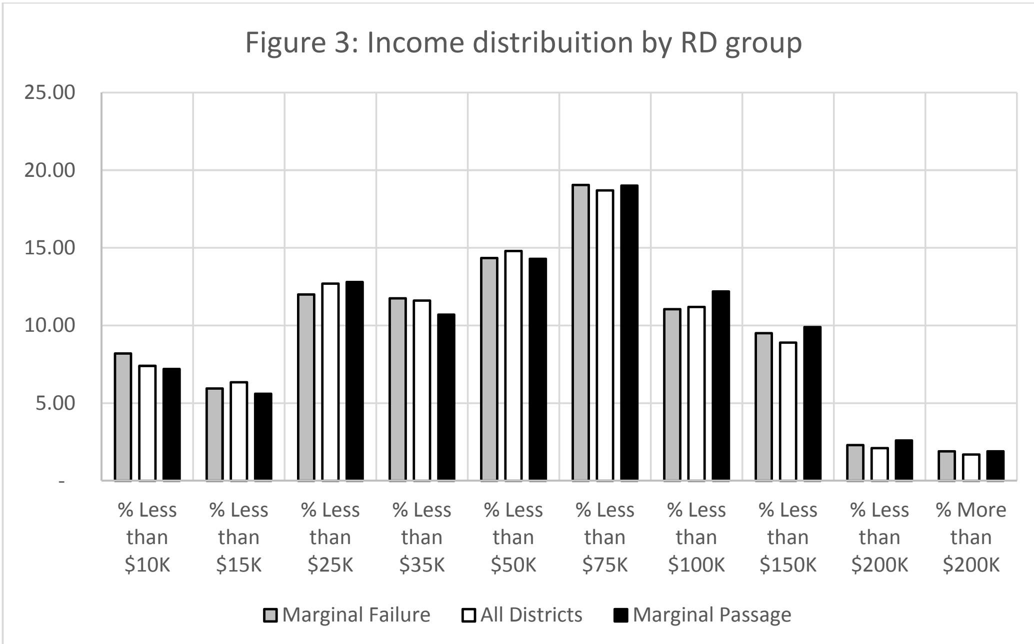

Figure 3: Income distribuition by RD group
25.00
20.00
15.00
10.00
5.00
-
% Less than $10K
% Less than $15K
% Less than $25K
% Less than $35K
% Less than $50K
% Less than $75K
% Less than $100K
% Less than $150K
% Less than $200K
% More than $200K
Marginal Failure
All Districts
Marginal Passage

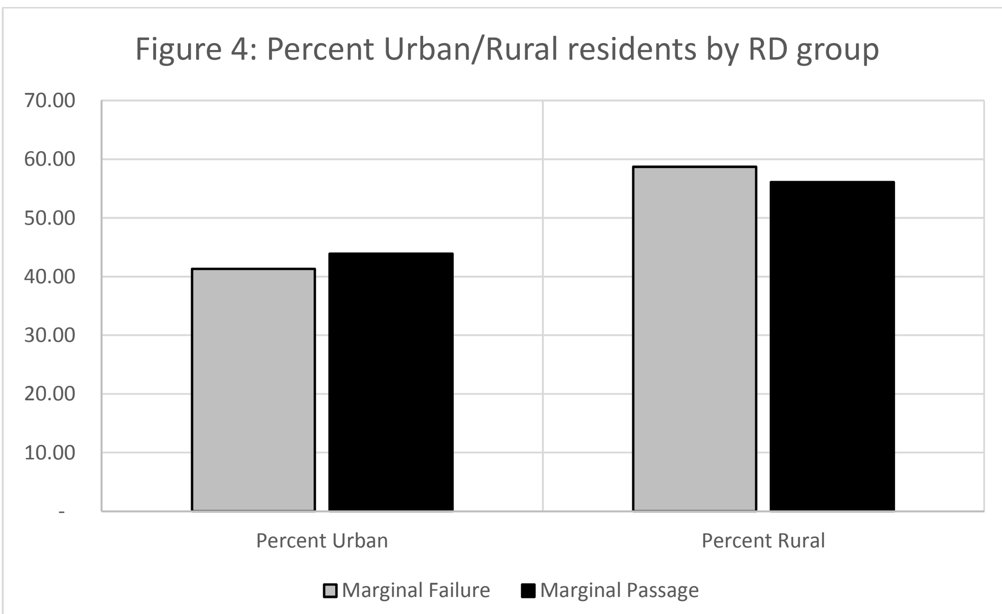

Figure 4: Percent Urban/Rural residents by RD group
70.00
60.00
50.00
40.00
30.00
20.00
10.00
-
Percent Urban
Percent Rural
Marginal Failure
Marginal Passage

Table 1: Bond Results and New School Activations						
	Total districts	Districts with a new school (%)	Yes vote districts	Districts with a new school (%)	No vote districts	Districts with a new school (%)
3 Years out	834	182 (22%)	574	141 (25%)	216	32 (15%)
5 Years out	834	182 (22%)	525	153 (29%)	166	37 (22%)

Table 2: Average Student Test Scores by Mobility			
VARIABLES	Reading Score	Math Score	Total Score
Mean of Never Movers	-0.006188	-0.0057096	-0.0063974
	(0.0003798)	0.0003805	0.0003816
Mean of Movers	0.0289878	0.0332633	0.0346787
	(0.0018131)	(0.0018014)	(0.0018185)
H_0: Mean of Movers – Mean of Never Movers = 0	-18.7781***	-20.7761***	-21.8245***
Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1			

Table 3: Model (1) – Modified Regression Discontinuity Approach

VARIABLES	Reading Score	Math Score	Total Score
Yes Vote	-0.018*	-0.030	-0.030**
	(0.011)	(0.0191)	(0.015)
Yes Vote, 1 Year Lag	0.009	0.009	0.005
	(0.012)	(0.020)	(0.016)
Yes Vote, 2 Year Lag	0.023*	0.020	0.019
	(0.012)	(0.021)	(0.017)
Yes Vote, 3 Year Lag	0.035***	0.034	0.032*
	(0.013)	(0.021)	(0.018)
Yes Vote, 4 Year Lag	0.019	0.017	0.014
	(0.013)	(0.020)	(0.017)
Yes Vote, 5 Year Lag	0.032**	0.035*	0.032*
	(0.013)	(0.021)	(0.017)
Yes Vote, 6 Year Lag	0.050***	0.076***	0.063***
	(0.013)	(0.021)	(0.017)
Yes Vote, 7 Year Lag	0.072***	0.089***	0.082***
	(0.013)	(0.022)	(0.018)
Yes Vote, 8 Year Lag	0.072***	0.094***	0.087***
	(0.014)	(0.022)	(0.018)
No Vote	-0.100***	-0.160***	-0.137***
	(0.018)	(0.026)	(0.022)
No Vote, 1 Year Lag	-0.088***	-0.150***	-0.125***
	(0.018)	(0.026)	(0.022)
No Vote, 2 Year Lag	-0.069***	-0.135***	-0.106***
	(0.016)	(0.025)	(0.020)
No Vote, 3 Year Lag	-0.059***	-0.129***	-0.098***
	(0.016)	(0.024)	(0.020)
No Vote, 4 Year Lag	-0.043***	-0.094***	-0.069***
	(0.016)	(0.024)	(0.020)
No Vote, 5 Year Lag	-0.024	-0.065***	-0.044**
	(0.015)	(0.024)	(0.020)
No Vote, 6 Year Lag	-0.023	-0.072***	-0.049**
	(0.015)	(0.023)	(0.020)
No Vote, 7 Year Lag	-0.045***	-0.082***	-0.065***
	(0.017)	(0.024)	(0.021)
No Vote, 8 Year Lag	-0.034*	-0.048*	-0.041*
	(0.017)	(0.026)	(0.022)
Constant	0.073***	0.106***	0.099**
	(0.011)	(0.018)	(0.015)
Observations	2,466,558	2,466,558	2,466,558
R-squared	0.777	0.793	0.823

Clustered standard errors in parentheses, Student fixed effects included
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Model (2) – New School Activations

VARIABLES	Reading Score	Math Score	Total Score
Number of Schools in District	-0.000525***	-0.000247**	-0.000427***
	(6.90e-05)	(9.81e-05)	(8.35e-05)
Percent of New Schools Last 3yrs	0.000381**	-0.000632*	-0.000208
	(0.000193)	(0.000323)	(0.000262)
0 Years Since Activation	0.0167***	0.0357***	0.0300***
	(0.00353)	(0.00517)	(0.00432)
1 Years Since Activation	0.0282***	0.0509***	0.0443***
	(0.00379)	(0.00557)	(0.00465)
2 Years Since Activation	0.0342***	0.0586***	0.0519***
	(0.00431)	(0.00652)	(0.00542)
3 Years Since Activation	0.0582***	0.0667***	0.0682***
	(0.00474)	(0.00685)	(0.00573)
4+ Years Since Activation	0.0637***	0.105***	0.0928***
	(0.00654)	(0.00902)	(0.00785)
Constant	-0.00967***	-0.0273***	-0.0203***
	(0.00364)	(0.00505)	(0.00428)

Clustered standard errors in parentheses, Student fixed effects included
*** p<0.01, ** p<0.05, * p<0.1

Table 5: Model (2) – New School Activations – Highest Quintile Scores

VARIABLES	Reading Score	Math Score	Total Score
Number of Schools in District	-8.27e-05	-0.000498***	-0.000139
	(9.08e-05)	(0.000114)	(0.000103)
Percent of New Schools Last 3yrs	-0.00131***	-0.000962**	-0.00151***
	(0.000313)	(0.000388)	(0.000372)
0 Years Since Activation	0.0699***	0.0392***	0.0654***
	(0.00616)	(0.00572)	(0.00578)
1 Years Since Activation	-0.244***	-0.0145**	-0.136***
	(0.00635)	(0.00626)	(0.00614)
2 Years Since Activation	-0.289***	-0.0374***	-0.175***
	(0.00786)	(0.00767)	(0.00781)
3 Years Since Activation	-0.319***	-0.0484***	-0.203***
	(0.0104)	(0.00878)	(0.00967)
4+ Years Since Activation	-0.332***	-0.0377***	-0.214***
	(0.0118)	(0.0117)	(0.0117)
Constant	1.193***	0.895***	1.143***
	(0.00458)	(0.00514)	(0.00483)

Clustered standard errors in parentheses, Student fixed effects included

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 6: Model (2) – New School Activations – Lowest Quintile Scores

VARIABLES	Reading Score	Math Score	Total Score
Number of Schools in	-0.000877***	-0.000938***	-0.000925***
District	(8.92e-05)	(0.000108)	(0.000100)
Percent of New Schools	0.000171	-0.000439	-0.000372
Last 3yrs	(0.000354)	(0.000464)	(0.000410)
0 Years Since Activation	-0.00311	0.0439***	0.0244***
	(0.00582)	(0.00729)	(0.00636)
1 Years Since Activation	0.304***	0.172***	0.256***
	(0.00653)	(0.00818)	(0.00721)
2 Years Since Activation	0.373***	0.216***	0.318***
	(0.00829)	(0.00967)	(0.00891)
3 Years Since Activation	0.403***	0.224***	0.335***
	(0.00827)	(0.0105)	(0.00919)
4+ Years Since Activation	0.411***	0.290***	0.379***
	(0.0101)	(0.0126)	(0.0114)
Constant	-1.280***	-0.964***	-1.218***
	(0.00552)	(0.00666)	(0.00608)

Clustered standard errors in parentheses, Student fixed effects included
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Endnotes

[1] The National Clearinghouse for Educational Facilities has extensive data on facilities in the US.

[2] A review of the literature investigating all forms of public school spending and student outcomes lies beyond the scope of this paper. Interested readers could see Hanushek (1997). Several papers focusing on school facilities and local school bonds are discussed below.

[3] Estimating the average treatment effect for students who do not leave the school district following the bond votes is all our data allows. Estimating the average treatment effect for all students are not identical exercises. We later present evidence showing the characteristics of displaced students following quasi-random bond passage are not representative of the overall student population to help frame our results. A distinguishing feature of our analysis is that, to our knowledge, we provide the first estimates of how bond passage affects the massive group of students who stay put in the same schools – an interesting sub-group and certainly the largest fraction of the student population. Because this group of staying students is so dominantly large relative to the overall population of students, it would be difficult to argue the effects we find do not generalize at least reasonably well to the larger population.

[4] Rather than fully review the large literature on the effects of various school facility traits on student outcomes, we direct interested readers to Bowers and Urick (2011).

[5] While our data is a long panel, there are still challenges associated with estimating how bond passage influences school districts in the long run. One way to interpret our findings is that once the estimated effect 'levels out' (i.e., become stable in size) around year eight, the benefit has been fully accrued but the improved conditions continue to impact student achievement levels.